Short-Input Random Oracles (and PRFs) over Abelian Groups: New Generalisations and Better Bounds

Ritam Bhaumik
TII, Abu Dhabi
bhaumik.ritam@gmail.com

Abstract

At Eurocrypt '26, Bhaumik et al. initiated a systematic study of short-input random oracles built from public random permutations over $\mathbb{Z}_2^n$. While their main focus was on length-preserving random oracles, they also proved that XORP, the expanding construction at the core of CENC (Iwata, FSE '06), is indifferentiable from a random function up to about $N^{2/3}$ queries for a permutation with domain size N ($= 2^n$). We extend this study in several directions. For classifying short-input expanding random oracles, we introduce GXP, a family of constructions that use independent random permutations over an arbitrary abelian group of order N to build a function with several output blocks, each the difference of two masked permutation outputs. GXP covers XORP with one or several permutations, LXoP (Dinur, Eurocrypt '25), and the widely-studied sum of two permutations.

We study the security of GXP in both keyless and keyed settings. In the keyless setting, a recent preprint from Kim has shown that the sum of two permutations over $\mathbb{Z}_2^n$ with a class of simulators called *uniform simulators* achieves a tight indifferentiability from a random oracle up to about $N^{3/4}$ queries. We build on this analysis, and prove that every member of GXP, over every finite abelian group, with arbitrary automorphism masks (under an extra width condition for some singular masks) and with unused permutation inputs, is indifferentiable up to about $N^{3/4}$ queries, for a uniform simulator with three trials and an acceptance coin. For admissible masks we give a sharper bound with explicit constants; other masks only add lower-order terms. The coin is necessary: without it, the simulator is distinguishable at about $N^{2/3}$ queries. Extending Kim's attack, with coset progressions in place of subgroups, we show that the exponent $3/4$ is tight for this simulator, for every family whose construction domain has density bounded below.

In the keyed setting, we extend the Fourier-analytic method put forward by Bhaumik et al. in another recent preprint, and establish the pseudorandomness with advantage $O_w(q/N^{3/2})$ for every rigid-free member of GXP, with explicit constants and over every finite abelian group. Pseudorandomness is thus governed by rigid-freeness, a strengthening of admissibility: rigid pairs and non-admissible masks can both limit it to about q/N, while neither affects the indifferentiability exponent.

The fork constructions ForkEDMD and ForkCENC of Dutta et al., and EDMD, in which the output of one permutation call becomes the input of all the other calls, are not members of GXP. A change of variables on the permutations turns each of them into a member with exactly the same security, so all our results apply to them as well.

Contents

1 Introduction **4**
- 1.1 Motivation . . . 4
- 1.2 Our contributions . . . 5
- 1.3 Technical overview . . . 8
- 1.4 Related work and credits . . . 10
- 1.5 Organisation . . . 11

2 Preliminaries **11**
- 2.1 Indifferentiability and pseudorandomness . . . 12
- 2.2 Divergences . . . 13
- 2.3 Fourier analysis on finite abelian groups . . . 13

3 The LCP family and its subfamily GXP **14**
- 3.1 Linear combinations of permutations . . . 14
- 3.2 The subfamily GXP . . . 16
- 3.3 Same-class pairs, admissibility and rigid pairs . . . 17
- 3.4 Instances . . . 18
- 3.5 Forks and the reducible family . . . 18

4 The coin simulator **19**
- 4.1 Legal sets and the index-free weight . . . 19
- 4.2 Simulator description . . . 22
- 4.3 Equivalent presentations . . . 24

5 Indifferentiability of the GXP family **25**
- 5.1 Standing assumptions and query types . . . 25
- 5.2 Main theorem . . . 26
- 5.3 Hybrid and plan of the proof . . . 27
- 5.4 From the real world to the hybrid . . . 28
- 5.5 From the hybrid to the ideal world . . . 43
- 5.6 Bounds for the instances . . . 48

6 Lower bounds for the coin simulator **52**
- 6.1 The coin is necessary . . . 52
- 6.2 Matching attack on the coin simulator . . . 54

7 Arbitrary masks **59**
- 7.1 Activity of same-class pairs . . . 59
- 7.2 The theorem . . . 60
- 7.3 Plan of the proof . . . 61
- 7.4 Charges for singular pairs . . . 63
- 7.5 Real to hybrid . . . 64
- 7.6 Hybrid to ideal . . . 68
- 7.7 Lower bounds and the meaning of tightness . . . 68

7.8 Proof of Theorem 3 69

8 Pseudorandomness **70**
8.1 Non-admissible masks 70

9 Conclusions **72**
9.1 Open problems 72

References **74**

Index of notation **79**

Glossary of terms **81**

Appendix **82**

A Mathematical background **82**
A.1 Probability and information theory 82
A.2 Finite abelian groups 86
A.3 Fourier analysis on finite abelian groups 88

B Evaluation of the constants **92**
B.1 Numerical steps in Sections 5.4 and 5.5 93
B.2 Numerical steps in Sections 5.6 to 8 95

C Derivations for the ideal world **97**

D Deferred proofs for Sections 4 and 6 **99**
D.1 Proof of Lemma 5 99
D.2 Proof of the matching attack 100

E Pseudorandomness of rigid-free families **109**
E.1 Fixed-support moment 109
E.2 Pullback and reduction to moments 111
E.3 Merging recursion 112
E.4 Summing over levels 114

1 Introduction

Bhaumik et al. [BDD+26] studied how to build a short-input random oracle from public random permutations. Among the constructions with two calls they found that only the sum of permutations and EDMD, both with independent permutations, reach beyond the birthday bound, and they proved that XORP, the expanding construction at the core of CENC, is indifferentiable from a random function up to about $N^{2/3}$ queries. This paper continues that study. We widen the class of constructions, borrowing from Dutta et al. [DGL22] the fork family, in which one permutation output feeds the inputs of several further calls. We also sharpen the bounds by adapting two recently devised proof techniques: Kim's analysis of the sum of two permutations [Kim26], which reaches about $N^{3/4}$ queries with the uniform simulator, and the Fourier-analytic method of Bhaumik et al. [BGGJ26] for pseudorandomness.

Sums of permutations. Many symmetric constructions turn random permutations into a function by adding outputs. Examples include the following:

- the sum of two independent permutations, $\pi_1(x) + \pi_2(x)$ [BKR98] (Figure 1c);
- Iwata's XORP [Iwa06], which calls one permutation at $w + 1$ domain-separated inputs and outputs the w differences with the first call (Figure 1d);
- Dinur's LXoP [Din25], which combines consecutive calls through a linear map σ (Figure 1e);
- the fork constructions ForkEDMD and ForkCENC of Dutta et al. [DGL22] (Figure 1f).

We study all of these through one family, which we call GXP (Figure 1b). The first three are instances of GXP. The fork constructions, in which the output of one call is fed into the inputs of the others, are not, but they reduce exactly to instances of GXP, by a change of variables on the permutations; we call the constructions that reduce in this way GXP$^+$ (Section 3.5).

On an input x, a member F of GXP outputs w blocks, $\mathrm{F}(x) \in G^w$. Each block is a difference $\lambda_j \pi_{c(j)}(\iota_j(x)) - \lambda_0 \pi_{c(0)}(\iota_0(x))$ of masked outputs of k independent permutations of a finite abelian group G of order N, where the masks λ_j are group automorphisms. Several calls may share a permutation, and the calls of one permutation may leave some of its inputs unused. GXP is the subfamily, with differences as outputs, of a more general class LCP of linear combinations of masked permutation outputs (Figure 1a), for which we give a compact notation (Section 3). We study the indifferentiability of GXP from a random function [MRH04], which is the public-permutation notion: an indifferentiable member can stand in for a short-input random oracle, the use that motivates [BDD+26]. We also study its pseudorandomness, against distinguishers that see only the construction.

1.1 Motivation

Beyond the birthday bound. A block cipher is modelled as a pseudorandom permutation, but many of its uses need a pseudorandom function: keystreams in counter mode, nonce-based message authentication [CS16], key derivation, and so on. By the PRP–PRF switch, a random permutation looks like a random function only up to about $\sqrt{N}$ queries. Many lightweight block ciphers have 64-bit blocks, among them PRESENT [BKL+07], PRINCE [BCG+12], SKINNY-64 [BJK+16] and GIFT-64 [BPP+17]. For them the bound is 2^{32} evaluations, which is no meaningful security: the Sweet32 attacks [BL16] used birthday-bound collisions to recover plaintext from 64-bit ciphers in TLS and OpenVPN. With 80-bit blocks the bound of 2^{40} is not meaningful either, and even with 96-bit blocks 2^{48} evaluations are a serious practical concern [BDD+26, Sec. 1].

Sums of permutations were introduced to remove this loss. Bellare et al. [BKR98] proposed the sum of two permutations, and Lucks [Luc00] gave the first verifiable proof that it is secure beyond $\sqrt{N}$ queries, together with the generalisation to k permutations. Bellare and Impagliazzo [BI99], and later Dai et al.

[DHT17], devised statistical techniques that improved the bound. Patarin's mirror theory [Pat08, Pat10] treated the problem combinatorially and claimed security up to about N queries; it was extended to the sum of k permutations [CLP15], and verified proofs of its central theorem appeared only later [DNS22, CDN+23]. Dinur [Din24] then proved a tight bound by Fourier analysis, which Eberhard [Ebe17] had already obtained, for every finite abelian group, in a work little known to cryptographers. Expanding variants were proposed to amortise the cost, producing w output blocks from $w + 1$ calls. The first was XORP, the core of Iwata's CENC mode [Iwa06]; its optimal security was later shown to follow from Patarin's theorem [IMV16]. Dutta et al. introduced the fork constructions [DGL22, DL25], and Dinur introduced LXoP [Din25]. All of this work concerned the pseudorandomness of these constructions for a *secret* permutation.

These constructions are increasingly used where the permutation is not secret, or where the adversary controls the key. Examples include the following [BDD+26, Sec. 1.3]:

- committing security of authenticated encryption, where an attacker who chooses keys must not produce a ciphertext that decrypts validly under two keys;
- short-input random oracles built from a public permutation, for instance in the Fiat–Shamir transform or for a wide GGM tree in MPC-in-the-head protocols. Feneuil and Rivain [FR26] analyse such trees in the random-oracle and ideal-cipher models, and the latest version of FAEST builds its vector commitments from an AES-based generator [BBB+26].

The appropriate security notion for these uses is indifferentiability [MRH04]. If a construction is indifferentiable from a random function, the composition theorem of [MRH04] lets it replace that function in any single-stage security game; multi-stage games need more care [RSS11]. For example, indifferentiability of XORP yields the committing security of CENC [BDD+26, Sec. 4.2].

The state of the art. Several published indifferentiability proofs for sums of public permutations [MPN10, MP15, Lee17, BN18a], including a claim of full N-query security for $\pi_1 + \pi_2$ [BN18a], were later found to rest on an assumption that holds only for the weaker sequential notion [Gun22]; see [BDD+26, Sec. 1.4] for this history. Gunsing et al. [GBJ+23] re-established regular indifferentiability up to about $N^{2/3}$ queries for $\pi_1 + \pi_2$, and gave an attack with about $N^{5/6}$ queries on the natural simulators. Bhaumik et al. [BDD+26] proved about $N^{2/3}$ queries for EDMD [MN17] and about $N^{2/3}/(w+1)^{4/3}$ for XORP with $w + 1$ calls, and showed that before their work no regular indifferentiability result was known for EDMD or XORP. In the sequential notion, Datta et al. [DDMN25] studied EDM and STH. Kim [Kim26] raised $\pi_1 + \pi_2$ to $N^{3/4}$ with the uniform simulator, and showed that this exponent is optimal for that simulator. To our knowledge, the other constructions above have had no indifferentiability bound beyond $N^{2/3}$, and LXoP, ForkEDMD and ForkCENC with $w \geq 2$ had none at all.

1.2 Our contributions

Indifferentiability. Our results extend Kim's upper bound to every member of GXP, with a uniform simulator that makes three trials and tosses an acceptance coin for unused inputs. For *admissible* families, a condition on the masks of calls that share a permutation (Definition 5), Theorem 1 gives

$$\mathbf{Adv}^{\text{indiff}} \leq 10\, n^{3/2}(w+1)^2 \frac{q^2}{N^{3/2}} + 2w(w+1)\frac{q}{N} + 3w^2\frac{\sqrt{q}}{N}.$$

Here n is the largest number of calls to one permutation, and the last two terms appear only in the presence of certain same-permutation pairs. For arbitrary automorphism masks, Theorem 3 gives the same exponent $3/4$ for the same simulator. A pair of calls whose mask difference is singular, with kernel of size $\varkappa$, only adds lower-order terms of orders q/N, $\sqrt{q}/N$, $\sqrt{q\varkappa}/N$ and $q\sqrt{\varkappa}/N^{3/2}$, and construction-only attacks show

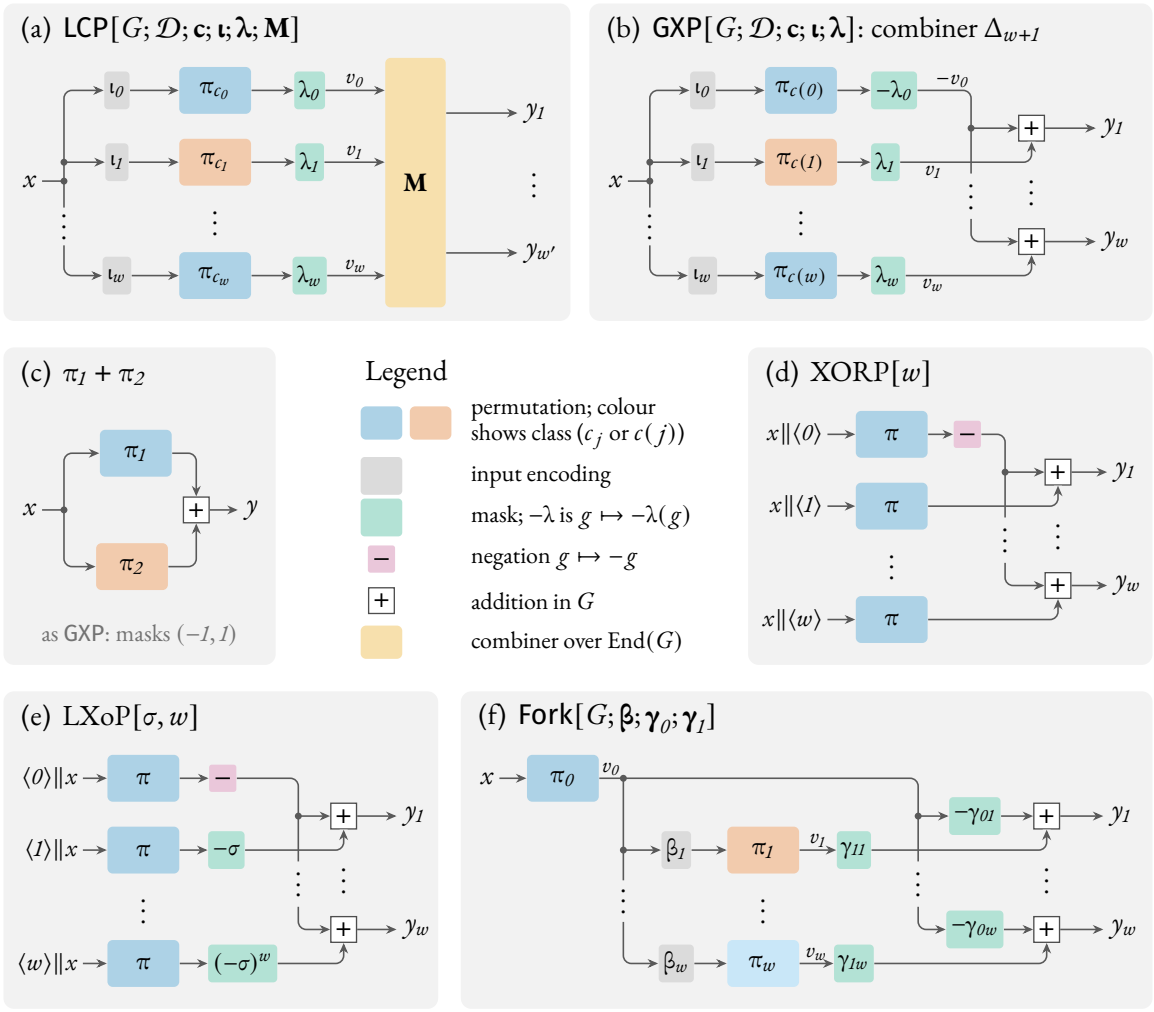


Figure 1: The constructions of this paper as circuits. (a) The general family LCP (Definition 2), with an arbitrary combiner **M**. (b) Its subfamily GXP (Definition 3): every output block is the difference $y_j = v_j - v_0$ of two masked outputs. (c)–(e) Instances of GXP (Table 2): the sum of two permutations; XORP, with $w+1$ tagged calls to one permutation, so that all calls form one class; and LXoP, which adds the masks $(-\sigma)^j$. (f) A fork (Definition 7), in which the output v_0 of the first call feeds the inputs of the others; its output is $y_j = \gamma_{1j} v_j - \gamma_{0j} v_0$. It is not an instance of GXP, but it is in GXP$^+$: by Lemma 18 it reduces exactly to one (Section 3.5).

instance	indifferentiability	tight	pseudorandomness
$\pi_1 + \pi_2$	$40\,q^2/N^{3/2}$	yes	$117\,q/N^{3/2}$
domain-separated sum	$114\,q^2/N^{3/2} + 4q/N$	yes	not covered
XORP, independent perm.s	$10(w+1)^2q^2/N^{3/2}$	yes	$C_w q/N^{3/2}$
XORP$[w]$, one perm.	$10(w+1)^{7/2}q^2/N^{3/2} + 2w(w+1)q/N$	yes	not covered
LXoP$[\sigma, w]$	$10(w+1)^{7/2}q^2/N^{3/2} + 3w^2\sqrt{q}/N$	yes	$C_w q/N^{3/2}$
ForkEDMD, ForkCENC	$10(w+1)^2q^2/N^{3/2}$	yes	$C_w q/N^{3/2}$

Table 1: Bounds for the main instances, from Theorem 1 with $\textsc{Sim}_3$ for $2w \le q \le N/64nw$, and from Theorem 4 for $q \le c_w N$. The columns are explained in Section 1.2.

that such terms cannot be removed entirely (Proposition 7). Admissibility is therefore not needed for the exponent $3/4$. It does matter for efficiency: for admissible families the simulator makes a constant expected number of draws per query, while for arbitrary masks we bound only the advantage, not the running time (Remark 6). The coin cannot be omitted in a theorem that covers free inputs. Without it, a simulator that answers unused inputs as an independent lazy permutation is distinguishable at about $N^{2/3}$ queries on single-permutation families with a constant fraction of free inputs (Proposition 4).

Tightness. The exponent $3/4$ is optimal for this simulator. Section 6.2 extends Kim's attack to every member of GXP over every finite abelian group, with arbitrary automorphism masks, at fixed width and class sizes and with a construction domain of density bounded below. It gives a distinguisher with $O(N^{3/4})$ primitive queries and no construction queries that has advantage $3/4$ against $\textsc{Sim}_3$ (Corollary 5). Through converses of the reductions of Section 5.6, the same holds for ForkEDMD and ForkCENC (Corollary 6). Security beyond $N^{3/4}$ queries therefore needs a different simulator; for $\pi_1 + \pi_2$, Kim's Gyroscope simulator reaches about $N^{4/5}$ [Kim26].

Pseudorandomness. Indifferentiability gives pseudorandomness up to about $N^{3/4}$ queries for every family (Corollary 7). For *rigid-free* families, we prove a sharper bound with explicit constants (Theorem 4, proved in Appendix E):

$$\mathbf{Adv}^{\mathrm{prf}} \le C_w \frac{q}{N^{3/2}} \qquad (q \le c_w N).$$

Rigid-free means that for any two calls j, j' to the same permutation, $\lambda_j^{-1} - \lambda_{j'}^{-1}$ is an automorphism; in particular, the family is admissible and no two such calls share a mask. This is exactly the condition under which the Fourier-analytic method of Bhaumik et al. [BGGJ26], which we extend from LXoP, applies. Distinct masks alone do not suffice. For example, on $\mathbb{Z}_N$ with N even, two calls to one permutation with masks 1 and -1 give a single output at statistical distance exactly $1/2(N-1)$ from uniform. Families in which two calls to one permutation share a mask cannot do better than about q/N, by a simple collision test. Identity-mask XORP on one permutation is such a family. Non-admissible masks also degrade pseudorandomness, by an amount that grows with the kernel size (Section 8.1), although they do not affect the indifferentiability exponent.

Table 1 collects the bounds for the main instances. In the indifferentiability column, n is 1, 2 or $w+1$ as appropriate. The column "tight" records that the exponent $3/4$ is optimal for $\textsc{Sim}_3$ at fixed width (Corollaries 5 and 6). In the pseudorandomness column, $c_1 \ge 1/40$, LXoP is taken with w-admissible σ, and the bound for the forks follows through their rigid-free bases (Proposition 3). The two families marked "not covered" have rigid pairs, which give an $\Omega(q/N)$ attack.

New results. We provide below an overview of the new results presented in this work:

- *A single theorem for the whole family.* Theorem 3 gives Kim's exponent $3/4$ for every member of GXP, with a single simulator and proof, and Theorem 1 sharpens the bound for admissible families. A new construction of this form is covered without any condition on its masks. Compared with previous work:
  - XORP with one permutation and $w+1$ calls improves from $N^{2/3}/(w+1)^{4/3}$ to about $N^{3/4}/(w+1)^{7/4}$ queries, which is better whenever $w+1$ is below about $N^{1/5}$. This compares leading powers only: the bound of [BDD+26] also has a logarithmic factor, and its simulator in effect assumes that $w+1$ is a power of two, since it samples inverse candidates from all unused inputs and reads each as a tagged input;
  - EDMD improves from $N^{2/3}$ to $N^{3/4}$. This also follows by combining Kim's bound for $\pi_1 + \pi_2$ with the equivalence of EDMD and the sum of two permutations in [BDD+26, proof of Theorem 3, App. A.2]; here it is a special case of the forks;
  - to our knowledge, LXoP, ForkEDMD and ForkCENC with $w \geq 2$ receive their first indifferentiability bounds.

  For a 128-bit permutation, the query threshold moves from about 2^{85} to 2^{96}, ignoring constants, logarithmic factors and polynomial factors in the width. For a 64-bit permutation it moves from about 2^{43} to 2^{48}.

- *Every finite abelian group.* All theorems and the attack hold over any finite abelian group, not only $\mathbb{F}_2^d$. So they also cover permutations of $\mathbb{Z}_N$ and of prime fields, the setting of arithmetisation-oriented designs, with masks given by arbitrary group automorphisms. For indifferentiability this is new even for $\pi_1+\pi_2$, which [Kim26] treats over $\mathbb{F}_2^n$. For pseudorandomness, sums of independent permutations were already treated over every finite abelian group [Ebe17, BGGJ26]; what is new there is the families with several calls to one permutation, such as LXoP, which [BGGJ26] treats over $\mathbb{F}_2^n$.

- *Explicit constants and a structural account of the error.* Given this history, we give complete proofs with explicit constants. Each term of the bound has an identifiable source. The $q^2/N^{3/2}$ term is present for every family. For admissible families, the q/N term comes only from calls that share a permutation and a mask, and the $\sqrt{q}/N$ term only from calls that share a permutation with different masks; singular mask differences add the terms of Theorem 3.

- *Rigid-freeness governs pseudorandomness, not the indifferentiability exponent.* Rigid pairs and non-admissible masks add only lower-order terms to the indifferentiability bound, but both can make pseudorandomness strictly worse than the rigid-free bound $q/N^{3/2}$ (Sections 7.7 and 8.1).

1.3 Technical overview

Indifferentiability. We follow Kim's architecture [Kim26]. A hybrid answers every query by a simple explicit law. The real world is compared with the hybrid, and the hybrid with the simulator, by the chain rule for relative entropy. Five difficulties appear beyond the two-permutation case:

1. *Shared calls.* Several outputs of one permutation enter one input. We describe the possible outputs of an input by a single anchor and a *legal set* of anchors (Section 4.1), which carries masks, classes and same-permutation collisions uniformly. The weight that the hybrid gives to fresh inputs turns out to be independent of the call index (Lemma 4), and this is what makes the hybrid simple. The chi-square

of a fresh construction response is computed in Fourier form, where the cross terms between block pairs vanish (Lemma 8).

2. *Rigid pairs.* Two calls to one permutation with the same mask force two blocks to differ, which a random function violates with probability $1/N$. This discrepancy cannot be absorbed in the divergence calculation. It is paid once by a coupling (Section 5.5), and this is the q/N term.

3. *Masks that are merely admissible.* For general masks, the probability that a hidden output collides with a queried one involves coincidences between affine maps. For a particular prefix these can be large (Remark 1). We bound them only *in expectation over the game*. To do so we index everything by *exposure slots*, fixed before any adaptive choice. Charging fixed slots is essential, because an adaptive distinguisher can select large coincidences (Remark 4).

4. *Unused inputs.* The real world treats unused inputs of a permutation like fresh ones. The simulator rejects family candidates with probability about $1 - \rho_b$, so it must reject unused candidates with the same probability. The acceptance coin does this, and it can be implemented with a dummy oracle value (Section 4.2).

5. *The posterior of the oracle.* Each rejection tilts the conditional law of the unrevealed oracle values. We bound the tilt of one value through a threshold form of the log-likelihood (Lemma 16). It carries one factor n per call of the queried permutation.

For arbitrary masks (Section 7), a pair of calls whose mask difference is singular forbids a whole coset of anchors for some values. Values that forbid large cosets are rare; they are removed by a stopped hybrid that halts with exactly the real probability of revealing one, so that only a total-variation cost is paid. Small cosets enter through per-input support sizes, charged to fixed slots as above. The tilt bound of Lemma 16 already tolerates the resulting alignment of forbidden sets.

The matching attack. Kim's distinguisher for $\pi_1 + \pi_2$ fills a subgroup of order about $N^{3/4}$ with outputs of one call. It then makes inverse queries whose candidates the simulator must reject whenever a fixed test event holds for their hidden oracle value. These rejections are invisible, but they tilt the posterior of every unrevealed value towards the test event, and later inverse answers reveal this bias. We extend the attack in three ways:

- Groups without suitable subgroups, such as $\mathbb{Z}_p$, are handled by *coset progressions* $H + \{0, g, \ldots, (L-1)g\}$, which exist at every scale (Lemma 31); the filled progression is twice as long as the target sets, so that one test event is forbidden at every target.
- When the two attacked calls use the same permutation, two random shifts keep the target and test sets apart (Lemma 21).
- With several calls per permutation, the tilt of one value takes the same threshold form as in the upper bound.

Each useful test then has a drift of about r^2/N^2, where r is the size of the target sets, against a noise of about $r/\sqrt{N}$ for the whole statistic, which gives the threshold $r \approx N^{3/4}$ (Section 6.2).

Pseudorandomness. After an output character is pulled back to the permutation outputs, the density of the construction is a product of injection coefficients, one per permutation, and each input imposes one linear relation on its labels (Lemma 45). A merging recursion of Eberhard and Dinur [Ebe17, Din25], applied as in [BGGJ26], bounds second moments of such coefficients, provided each merge can be undone. In our setting, merges stay within one permutation, and undoing a merge needs exactly that the family is rigid-free: for any two calls of one permutation, the difference of the inverse masks must be invertible (Lemma 47).

The leaves of the recursion need the fixed-support moment Θ_k over a general group. We show that Θ_k is a closed-form function of (N, k) alone, with a Gamma-integral representation (Lemma 44).

1.4 Related work and credits

Our proofs build on two lines of work, on indifferentiability and on pseudorandomness, and we credit them here in detail, together with other related results:

- **Kim.** The architecture of the indifferentiability proof is Kim's [Kim26]: the hybrid, the chain rule across query types, and the conditional law of the unrevealed oracle values. Each lemma here that generalises a statement of [Kim26] is re-proved for the family; no result of [Kim26] is used as a black box. Relative to [Kim26], the new ingredients are:
  - the legal-set formulation and the index-free weight;
  - the Fourier computation of Lemma 8;
  - the per-input multiplicities in the bias and failure bounds;
  - the coupling for rigid pairs;
  - the exposure-slot bounds for admissible masks;
  - the acceptance coin;
  - the stopped hybrid and singular-excess charges for arbitrary masks (Section 7).

  The matching attack of Section 6.2 is Kim's distinguisher [Kim26, Sec. 4 and App. E]. The new ingredients there are coset progressions in place of subgroups, separating shifts for calls on one permutation, the threshold form of the tilt for several calls per permutation, and converses of the fork reductions.

- **Bhaumik et al. [BGGJ26].** The pseudorandomness proof follows [BGGJ26]: the reduction to L_2 moments, the pullback of characters, and the recursion with label recovery. They in turn build on Dinur [Din24, Din25] and Eberhard [Ebe17]. They treat sums of independent permutations over every finite abelian group, and LXoP over $\mathbb{F}_2^n$. We extend the argument for LXoP to families with several permutations, each called several times, to unused inputs, to arbitrary automorphism masks and to every finite abelian group. For LXoP, Theorem 4 is an explicit form of [BGGJ26, Theorem 7]. For $\pi_1 + \pi_2$, [BGGJ26, Theorem 2] is sharper below the birthday bound. Choi et al. [CHJ26] also use Fourier analysis, to prove quantum security of the XOR of $r \geq 2$ independent permutations of $\{0,1\}^n$; our pseudorandomness bound is for classical distinguishers.

- **Bhaumik et al. [BDD+26].** They describe every two-call construction over $\mathbb{F}_2^n$ by a matrix of linear maps [BDD+26, Sec. 3]; LCP extends this description to more calls, permutations and groups, but does not allow one call's output to feed another call's input (Section 3.1), and they show that EDMD is equivalent to the sum of two permutations by a change of variables on the permutations [BDD+26, proof of Theorem 3, App. A.2]. Lemma 18 extends this change of variables to forks with any number of branches and to output masks, and the class GXP+ of Section 3.5 is built on it. They also study XORP with a single permutation in the indifferentiability setting, with a two-trial simulator, which we do not analyse.

- **Other work.** Dutta et al. [DGL22] prove pseudorandomness of the fork constructions for *secret* pairwise independent permutations; a revised version, by Dutta and List, is [DL25]. Corollary 4 is an indifferentiability statement for public permutations, which is a different model with a different bound. Truncating a permutation is another way to obtain a function, and Choi et al. [CLL19] study its indifferentiability. Chen et al. [CLM19], who introduced the sum of Even–Mansour ciphers, and

Mennink [Men25] build keyed PRFs from public permutations. That model combines a secret key with primitive queries, unlike ours. The post-quantum security of such sums is studied for the sum of Even–Mansour ciphers in [TGL25] and for the keyed sum of permutations in [DDM+26]. For committing security from explicit collision assumptions on a forkcipher, rather than from indifferentiability, see [AEG+26]. For pseudorandomness of XORP with a single permutation, see [BN18b]. The chi-square method of Dai et al. [DHT17] is a related route to pseudorandomness bounds, which we do not use.

1.5 Organisation

The paper is organised as follows:

- Section 2 fixes notation, defines the two security notions, and states the analytic tools we use; the standard background behind them is recalled in Appendix A.
- Section 3 defines the general family LCP and the subfamily GXP studied here, together with admissibility and rigid pairs, lists the published constructions that are instances of GXP, and defines the class GXP+ of constructions, such as the forks, that reduce to it.
- Section 4 describes the coin simulator, with a worked example, and the moves that change a family's presentation without changing its security.
- Section 5 proves indifferentiability for admissible masks. After the statement of Theorem 1 and the plan of its proof, Sections 5.4 and 5.5 carry out the two comparisons, and Section 5.6 derives the bounds for the instances.
- Section 6 shows that the coin is necessary (Section 6.1) and gives a matching attack, which shows that the exponent $3/4$ is tight for this simulator (Section 6.2); the proofs of the attack and of the equivalence lemma are in Appendix D.
- Section 7 extends the theorem to arbitrary masks (Theorem 3) and collects the lower bounds.
- Section 8 discusses pseudorandomness; the Fourier proof for rigid-free families is in Appendix E.
- Section 9 summarises the results and lists open problems.

The references are followed by an index of notation and a glossary of terms, each linked to the definitions. The appendices contain the mathematical background (Appendix A), the evaluation of the constants, derivations for the ideal world, the deferred proofs of Sections 4 and 6, and the proof of the pseudorandomness theorem.

2 Preliminaries

We assume standard background in three areas: probability and information theory (distributions, total variation, relative entropy and chi-square divergence), finite abelian groups and their endomorphisms, and Fourier analysis on finite abelian groups. All of it is recalled in Appendix A, in Sections A.1 to A.3 respectively, together with the proofs of the lemmas stated in this section. Nothing in that appendix is new; it is included so that the paper is self-contained, and a reader who wants the conventions spelled out should start there. Here we fix notation, give the definitions specific to this paper, and list the known results that we use.

Notation. Throughout, G is a finite abelian group of order N, written additively. Coset progressions, which the matching attack uses, are defined in Definition 26. We write $[k] := \{1, \dots, k\}$ and $(a)_t := a(a-1)\cdots(a-t+1)$ for the falling factorial, with $(a)_0 = 1$. $\mathrm{Unif}(S)$ is the uniform distribution on a finite set S, and $\mathrm{Perm}(G)$ is the set of permutations of G. $\mathrm{End}(G)$ and $\mathrm{Aut}(G)$ are the endomorphisms and the automorphisms of G, and $\ker\varphi$ and $\mathrm{Im}\,\varphi$ are the kernel and the image of $\varphi \in \mathrm{End}(G)$ (Section A.2). For distributions P, Q on a finite set we use the following divergences (Section A.1):

- total variation $\mathrm{TV}(P,Q) = \frac{1}{2}\sum_x |P(x) - Q(x)|$;

- relative entropy $\mathrm{KL}(P\|Q) = \sum_x P(x)\log(P(x)/Q(x))$;
- chi-square divergence $\chi^2(P\|Q) = \sum_x (P(x)-Q(x))^2/Q(x)$.

Logarithms are natural.

2.1 Indifferentiability and pseudorandomness

A *construction* F maps an input $x \in \mathcal{D}$ to an output in G^w, making oracle calls to k permutations $\pi = (\pi_1, \ldots, \pi_k)$ of G. We consider two experiments.

- **Indifferentiability** [MRH04]. In the *real world* $\textsc{Real}$, the permutations are independent and uniform. The distinguisher gets the construction F^π and the primitive interfaces π_b and π_b^{-1} for every $b \in [k]$. In the *ideal world* $\textsc{Ideal}$, a function $\Gamma : \mathcal{D} \to G^w$ is uniform. The distinguisher gets Γ as its construction interface, and a simulator $\textsc{Sim}$ with oracle access to Γ answers its primitive queries. The simulator may answer $\perp$, which marks a failure. The advantage is

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\textsc{Sim}}(q) := \max_{\mathcal{A}} \left|\Pr[\mathcal{A}^{\textsc{Real}} = 1] - \Pr[\mathcal{A}^{\textsc{Ideal}} = 1]\right|,$$

  over distinguishers $\mathcal{A}$ making at most q queries in total, to the construction and to the primitives.

- **Pseudorandomness.** The distinguisher makes at most q construction queries and no primitive queries. It gets F^π in $\textsc{Real}$ and Γ in $\textsc{Ideal}$. Its advantage $\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q)$ is defined in the same way.

Distinguishers are computationally unbounded and, without loss of generality, deterministic. A pseudorandomness distinguisher is an indifferentiability distinguisher that never calls the simulator. So $\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \le \mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\textsc{Sim}}(q)$ for every simulator $\textsc{Sim}$.

The *transcript* of an interaction is the sequence of its queries and responses. A deterministic distinguisher chooses each query as a function of the earlier responses, and its output is a function of the transcript. So its advantage is at most the total variation between the transcript laws in the two worlds (Fact 1), and we bound that total variation throughout.

The constructions of this paper evaluate the permutations at several inputs determined by one construction input x. For the upper bounds we give the distinguisher more information than the ordinary interface does, which only makes it stronger.

Definition 1 (Augmented interface)**.** The permutation inputs evaluated on a construction input x are the *calls* of x, and their outputs together form the *tuple* of x (Definition 3). In the *augmented* interface, a primitive query at any call of x returns the whole tuple of x, and so does a successful inverse query whose answer is a call of x. A query at a *free input*, which the construction never uses (Definition 4), or an inverse query answered by such an input, returns only that input and its output, as in the ordinary interface.

Augmentation costs the analysis nothing. In $\textsc{Real}$, the augmented interface simply evaluates the permutations at all calls of x. In $\textsc{Ideal}$, our simulator assigns all outputs of an input at once (Section 4.2), so it can return the stored tuple. Conversely, a distinguisher for the ordinary interface is a distinguisher for the augmented interface that ignores some of the answers. So, for the same simulator, the augmented advantage is at least the ordinary one, and an upper bound in the augmented game is also an upper bound in the ordinary game. Our upper bounds are proved in the augmented game; our lower bounds state their interface explicitly, since a lower bound in the augmented game does not transfer to the ordinary game. We say that the construction value of x is *revealed* once the transcript contains it, either as the answer to a construction query or through the tuple of x, from which it can be computed.

We measure the running time of a simulator in expectation, assuming that the group operations, the masks and their inverses, the injections ι_j and their inverses can be evaluated efficiently, and that G can be sampled uniformly. For admissible families, our simulator uses rejection sampling with a constant expected number of draws per call, and capping the sampling at c draws gives a worst-case running time at a cost of at most $q\,16^{-c}$ in the advantage (Section 4.2). For arbitrary masks we bound the advantage but not the running time (Remark 6).

2.2 Divergences

We use the following standard facts. Each is recalled, with a proof or a reference, in Section A.1:

- *Chain rule* (Fact 6). For transcripts T_0, T_1 of q responses with $\mathrm{supp}(T_0) \subseteq \mathrm{supp}(T_1)$,

$$\mathrm{KL}(T_0 \| T_1) = \sum_{i=1}^{q} \mathbb{E}_{\tau \leftarrow T_0^{(i-1)}} \left[\mathrm{KL}\big(T_0(\cdot \mid \tau) \,\|\, T_1(\cdot \mid \tau)\big) \right],$$

  where $T^{(i-1)}$ is the prefix of the first $i-1$ responses and $T(\cdot \mid \tau)$ the law of the next response.
- $\mathrm{KL}(P\|Q) \le \chi^2(P\|Q)$ (Fact 2).
- *Pinsker's inequality* (Fact 5). $\mathrm{TV}(P,Q) \le \sqrt{\mathrm{KL}(P\|Q)/2}$.
- *Convexity* (Fact 3). KL and χ^2 are jointly convex in their two arguments.
- *Parseval's identity* on a finite abelian group (Section 2.3).

Two elementary lemmas are used repeatedly. Their proofs are in Section A.1.

Lemma 1. *Let Q be a distribution on a finite set Ω, $g : \Omega \to \mathbb{R}$, and $P(x) := Q(x)e^{g(x)}/\mathbb{E}_Q[e^g]$. Let $d := \max g - \min g$ over $\mathrm{supp}(Q)$. Then $e^{-d} \le P/Q \le e^d$ on $\mathrm{supp}(Q)$ and $\chi^2(P\|Q) \le e^{2d}\,\mathrm{Var}_Q(g)$. Consequently, if U is uniform on a set V and $U_g \propto Ue^g$ with $0 \le g \le A$, then $\chi^2(U\|U_g) \le e^{3A}\mathbb{E}_U[g^2]$. Moreover, $\chi^2(U\|\cdot)$ is convex, so the bound survives averaging the second argument over any mixture.*

So tilting a distribution by a factor e^g costs about $\mathrm{Var}(g)$ in chi-square, as long as g has bounded oscillation.

Lemma 2. *Let P be a distribution on Ω, let Q be a distribution on $\Omega \cup \{\perp\}$ with $Q(\perp) = p < 1$, and let $Q^+ := Q(\cdot)/(1-p)$ on Ω. Then $\mathrm{KL}(P\|Q) = \mathrm{KL}(P\|Q^+) + \log\frac{1}{1-p}$.*

Thus, when Q may fail and P never does, comparing P with Q costs exactly $\log\frac{1}{1-p} \approx p$ more than comparing P with Q conditioned on success.

2.3 Fourier analysis on finite abelian groups

Characters, the dual group, dual endomorphisms and Parseval's identity are recalled in Section A.3. We use the following notation.

Characters and duals. Write the character group $\widehat{G}$ additively, with 0 the principal character, and write $\mathrm{e}(\psi, g) := \psi(g) \in \mathbb{C}$ for the pairing. Then $|\mathrm{e}(\psi, g)| = 1$, and the pairing turns sums in either argument into products: $\mathrm{e}(\psi+\psi', g) = \mathrm{e}(\psi, g)\,\mathrm{e}(\psi', g)$ and $\mathrm{e}(\psi, g+g') = \mathrm{e}(\psi, g)\,\mathrm{e}(\psi, g')$. For $\varphi \in \mathrm{End}(G)$, the dual $\varphi^* \in \mathrm{End}(\widehat{G})$ is defined by $\mathrm{e}(\varphi^*\psi, g) = \mathrm{e}(\psi, \varphi g)$. Then $(\varphi+\varphi')^* = \varphi^* + \varphi'^*$ and $(\varphi\varphi')^* = \varphi'^*\varphi^*$. Moreover φ is an automorphism iff φ^* is, since $\ker\varphi^*$ is the annihilator of $\mathrm{Im}\,\varphi$ (Fact 15). For tuples $\beta \in \widehat{G}^r$ and $y \in G^r$ put $\mathrm{e}(\beta, y) := \prod_i \mathrm{e}(\beta_i, y_i)$ and $\mathrm{supp}(\beta) := \{i : \beta_i \ne 0\}$. The functions $y \mapsto \mathrm{e}(\beta, y)$ form an orthonormal basis of $L^2(\mathrm{Unif}(G^r))$. Hence *Parseval's identity*: $\mathbb{E}[|f|^2] = \sum_\beta |\hat{f}(\beta)|^2$ with $\hat{f}(\beta) := \mathbb{E}_y[f(y)\overline{\mathrm{e}(\beta, y)}]$ (Fact 13).

Injection coefficients. For $0 \le r \le N$ and $\beta \in \widehat{G^r}$ let

$$\hat{\mu}_r(\beta) := \mathbb{E}\left[\overline{e(\beta, Y)}\right], \qquad Y \text{ a uniform injection } [r] \to G,$$

and for $k \ge 1$ let

$$\Theta_k := \sum_{\beta \in (\widehat{G} \setminus 0)^k} |\hat{\mu}_k(\beta)|^2.$$

For the empty tuple we set $\hat{\mu}_0(()) = 1$ and $\Theta_0 = 1$. A uniform injection $[r] \to G$ is a sample of r elements of G without replacement, in random order. The quantity $\hat{\mu}_r(\beta)$ is the Fourier coefficient at β of the density of Y relative to $\mathrm{Unif}(G^r)$ (Fact 14). For sampling *with* replacement, that is, for Y uniform on G^r, every coefficient with $\beta \neq 0$ would vanish; the coefficients $\hat{\mu}_r(\beta)$ therefore measure the dependence created by excluding repeated values. For example, $\hat{\mu}_2(\psi, -\psi) = -1/(N-1)$ for every $\psi \neq 0$ (Example 4). In Appendix E, Θ_k appears as the second moment of these coefficients on one fixed support of size k. The following lemma collects their properties. It is proved in Section A.3.

Lemma 3 (after [Eber17] and [BGGJ26, Sec. 2.4 and App. A])**.** *Let $1 \le r \le N$ and $\beta \in \widehat{G^r}$.*

(a) $\hat{\mu}_r(\beta) = \hat{\mu}_k(\beta|_J)$ *for* $J = \mathrm{supp}(\beta)$ *and* $k = |J|$. *Zero labels may be inserted or deleted.*

(b) $\hat{\mu}_r(\beta)$ *is real,* $|\hat{\mu}_r(\beta)| \le 1$, *and* $\hat{\mu}_r(\beta) = 0$ *unless* $\sum_i \beta_i = 0$. *In particular,* $\hat{\mu}_1(\psi) = 0$ *for* $\psi \neq 0$.

(c) If $\beta_1, \beta_2, \beta_3$ *are nonzero with* $\beta_1 + \beta_2 + \beta_3 = 0$, *then* $\hat{\mu}_3(\beta) = 2/(N-1)(N-2)$.

(d) (Merging.) *If* $\mathrm{supp}(\beta) = J$ *with* $|J| = k \ge 2$ *and* $j \in J$, *then*

$$\hat{\mu}_r(\beta) = -\frac{1}{N-k+1} \sum_{i \in J \setminus \{j\}} \hat{\mu}_r\left(\beta^{(j,i)}\right), \qquad \beta_\ell^{(j,i)} := \begin{cases} 0, & \ell = j, \\ \beta_i + \beta_j, & \ell = i, \\ \beta_\ell, & \text{otherwise.} \end{cases}$$

Merging removes the label of call j and adds it to the label of one other call i. Each merge reduces the support by at least one, at the cost of a factor $(k-1)/(N-k+1)$, which is what makes the coefficients small when k is small compared with N.

3 The LCP family and its subfamily GXP

We first define a general family LCP of constructions that combine masked outputs of several permutations linearly (Section 3.1). We then specify the subfamily GXP studied in this paper, in which every output block is the difference of two masked calls (Section 3.2), and introduce the conditions on the masks that our theorems need (Section 3.3). Section 3.4 lists the published constructions that are instances of GXP, and Section 3.5 defines the larger class GXP^+ of constructions, including the fork constructions, that reduce to GXP exactly.

3.1 Linear combinations of permutations

We start from the simplest construction and generalise it step by step. Each step introduces one of the parameters of the formal definition, so that the definition itself contains nothing new.

From the sum of permutations to XORP. The sum of two independent permutations maps $x \in G$ to $\pi_1(x) + \pi_2(x)$ (Figure 1c). It makes two calls, one to each permutation, at the same input x, and it outputs one block. Iwata's XORP [Iwa06] (Figure 1d) calls a single permutation π of $G = \mathbb{F}_2^d$ several times instead. To keep these calls at distinct inputs, it appends a tag to the input: it evaluates π at $x \| \langle j \rangle_s$ for $j = 0, \ldots, w$, where $\langle j \rangle_s$ is the s-bit encoding of j and $x \in \mathbb{F}_2^{d-s}$, and it outputs the w blocks $\pi(x \| \langle j \rangle_s) - \pi(x \| \langle 0 \rangle_s)$ for $j = 1, \ldots, w$ (over $\mathbb{F}_2^d$, subtraction is XOR). Compared with the sum, XORP differs in three ways: several calls share one permutation; different calls feed different encodings of x to the permutation; and there are several output blocks, each a combination of the outputs of the calls.

Each of these differences becomes a parameter, and a fourth one, masks, covers the remaining published constructions:

- *Which permutation each call uses.* A sequence $\mathbf{c} = (c_0, \ldots, c_w)$ with values in $[k]$ says that call j uses permutation c_j. For the sum, $\mathbf{c} = (1, 2)$; for XORP, $\mathbf{c} = (1, \ldots, 1)$.
- *How each call encodes the input.* An injection $\iota_j : \mathcal{D} \to G$ maps the construction input to the input of call j. Calls to the same permutation must be made at distinct inputs, so their images must be disjoint. For the sum, $\mathcal{D} = G$ and every ι_j is the identity; for XORP, $\mathcal{D} = \mathbb{F}_2^{d-s}$ and $\iota_j(x) = x \| \langle j \rangle_s$.
- *A mask on each output.* Before the outputs are combined, call j applies an automorphism λ_j of G to its output. Dinur's LXoP [Din25] applies powers of a linear map σ to consecutive calls, and the fork constructions [DGL22] lead to masks through a change of variables (Lemma 18).
- *How the masked outputs are combined.* A matrix $\mathbf{M}$ with entries in $\mathrm{End}(G)$, the *combiner*, maps the vector of $w+1$ masked outputs to the output blocks (Section A.2). For the sum, $\mathbf{M} = (1\ 1)$; for XORP, $\mathbf{M}$ is the $w \times (w+1)$ matrix whose row j has -1 in column 0, 1 in column j and 0 elsewhere.

Together with the group G and the domain $\mathcal{D}$, this gives six parameters. Three of them are sequences indexed by the calls, and one is the combiner matrix. We write them positionally, in this order. Figure 1a shows the resulting construction as a circuit.

Earlier matrix descriptions. Describing such constructions by linear maps is not new. Bhaumik et al. [BDD+26, Sec. 3] describe every construction with two calls and one output block over $\mathbb{F}_2^n$ by a 3×3 block matrix of linear maps. Their matrix allows linear maps on the input, the output, and the input of the second call, which may depend on the output of the first call; earlier categorisations of this kind are [CMP21, BCEJ23]. LCP extends this description to any number of calls and output blocks, several permutations each called several times, and every finite abelian group with automorphism masks. It does not allow the input of a call to depend on another call's output; the forks, which need this, are handled through the class GXP^+ of Section 3.5.

Definition 2. Let G be a finite abelian group and $\mathcal{D}$ a finite nonempty set. The construction makes $w+1 \geq 1$ calls, indexed by $j \in \{0, \ldots, w\}$, and outputs $w' \geq 1$ blocks. The parameters are:

- a sequence $\mathbf{c} = (c_0, \ldots, c_w) \in [k]^{w+1}$ that takes every value in $[k]$: call j uses permutation c_j;
- a sequence $\boldsymbol{\iota} = (\iota_0, \ldots, \iota_w)$ of injections $\mathcal{D} \to G$ such that, for each $b \in [k]$, the images $\iota_j(\mathcal{D})$ with $c_j = b$ are pairwise disjoint;
- a sequence $\boldsymbol{\lambda} = (\lambda_0, \ldots, \lambda_w) \in \mathrm{Aut}(G)^{w+1}$ of masks;
- a $w' \times (w+1)$ matrix $\mathbf{M}$ with entries in $\mathrm{End}(G)$, the *combiner*.

Given permutations $\pi_1, \ldots, \pi_k$ of G, the construction $\mathsf{LCP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}; \mathbf{M}]$ maps $x \in \mathcal{D}$ to

$$\mathbf{M} \cdot \left(\lambda_j \pi_{c_j}(\iota_j(x))\right)_{0 \leq j \leq w} \in G^{w'},$$

where $\mathbf{M}$ acts on the column vector of masked outputs.

Integer entries of $\mathbf{M}$ act on G by repeated addition. We use the following shorthands:

- (b^i) abbreviates the constant sequence $(b, \dots, b)$ of length i;
- $\textsc{id}$ is the layout $\mathcal{D} = G$ and $\iota_j = \mathrm{id}$, which requires the c_j to be distinct;
- over $G = \mathbb{F}_2^d$ and $\mathcal{D} = \mathbb{F}_2^{d-s}$, $\textsc{tag}_s$ is the layout $\iota_j(x) = x \| \langle j \rangle_s$, and $\textsc{tag}'_s$ is $\iota_j(x) = \langle j \rangle_s \| x$;
- $\mathbf{I}$ is the sequence of identity masks, and $(\lambda(j))_j$ a sequence given by a formula in j;
- $\mathbf{I}_{w+1}$ is the $1 \times (w+1)$ combiner $(1, \dots, 1)$, so that $w' = 1$.

For example, the sum of k independent permutations is $\mathsf{LCP}[G; G; (1, \dots, k); \textsc{id}; \mathbf{I}; \mathbf{I}_k]$, with $w + 1 = k$ calls and one output block. Constructions with combiners other than the one of Section 3.2, such as this sum for $k \geq 3$, can be named in this notation but are not analysed here.

3.2 The subfamily GXP

The subfamily studied in this paper uses the *difference combiner*

$$\Delta_{w+1} := \left(-\mathbf{I} \,\middle|\, I_w\right) \in \mathbb{Z}^{w \times (w+1)},$$

whose row j produces the difference between the masked outputs of call j and call 0. We write

$$\mathsf{GXP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}] := \mathsf{LCP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}; \Delta_{w+1}],$$

with $w + 1$ calls and $w' = w$ output blocks. Figure 1b shows it as a circuit. Spelled out, with the notation used throughout the paper, it is the following.

Definition 3. An instance of GXP has the following parameters:

- a number $w \geq 1$ of output blocks and $w + 1$ *calls*, indexed by $j \in \{0, 1, \dots, w\}$;
- a number k of permutations and a surjective *class map* $c : \{0, \dots, w\} \to [k]$, the sequence $\mathbf{c}$ of Definition 2 with $c(j) := c_j$; call j uses permutation $c(j)$, the calls that use permutation b form *class* b, and the class sizes are $n_b := |c^{-1}(b)| \geq 1$, with $n := \max_b n_b \leq w + 1$;
- a finite nonempty input domain $\mathcal{D}$ and injections $\iota_j : \mathcal{D} \to G$ such that, for every class b, the images $\iota_j(\mathcal{D})$ with $c(j) = b$ are pairwise disjoint;
- masks $\lambda_0, \dots, \lambda_w \in \mathrm{Aut}(G)$.

Given permutations $\pi_1, \dots, \pi_k$ of G and $x \in \mathcal{D}$, define the *raw outputs* $p_j(x) := \pi_{c(j)}(\iota_j(x))$, the *masked outputs* $v_j(x) := \lambda_j(p_j(x))$, and

$$\mathrm{F}(x) := \big(v_1(x) - v_0(x),\ \dots,\ v_w(x) - v_0(x)\big) \in G^w.$$

Definition 4 (Anchor, free inputs)**.** Write $y = (y_1, \dots, y_w)$ for a construction value and set $y_0 := 0$, so that $v_j = v_0 + y_j$ for every j. The masked output v_0 is the *anchor* of x; given y, it determines all $w + 1$ raw outputs. The inputs of π_b that the construction never uses are the *free inputs*

$$\mathrm{Free}_b := G \setminus \bigcup_{c(j)=b} \iota_j(\mathcal{D}), \qquad |\mathrm{Free}_b| = N - n_b|\mathcal{D}|.$$

The family is *balanced* if every Free_b is empty.

In a balanced family $n_b|\mathcal{D}| = N$ for every b, and all class sizes are equal. The construction never evaluates a permutation at a free input, but a distinguisher may. Free inputs arise naturally: XORP and LXoP with $w + 1$ calls use $\lceil \log_2(w+1) \rceil$ tag bits, and leave $2^{\lceil \log_2(w+1) \rceil} - (w+1)$ tags unused (Section 3.4). Two counts recur. Let $F := \sum_b n_b^2$ be the number of ordered same-class pairs (j, j'), including $j = j'$, and let $\bar{F} := (w+1)n \geq F$. For balanced families $F = \bar{F}$.

Different parameters can describe constructions that are equivalent for both security notions: masks can be absorbed into the permutations class by class, any call can serve as the reference call, and only the size of the domain matters, not the layout. This is Lemma 5 in Section 4.3, which is stated there because its proof also transports the simulator.

3.3 Same-class pairs, admissibility and rigid pairs

A pair of calls $j \neq j'$ with $c(j) = c(j')$ is a *same-class pair*. The raw outputs $p_j(x)$ and $p_{j'}(x)$ are outputs of one permutation at two distinct inputs, so they differ. In terms of the anchor and the construction value this reads

$$\mu_{jj'}(v_0) \neq \lambda_{j'}^{-1}(y_{j'}) - \lambda_j^{-1}(y_j), \qquad \mu_{jj'} := \lambda_j^{-1} - \lambda_{j'}^{-1} \in \mathrm{End}(G). \tag{1}$$

Definition 5. The family is *admissible* if for every same-class pair, $\mu_{jj'}$ is either 0 or an automorphism of G. A same-class pair with $\mu_{jj'} = 0$, that is with $\lambda_j = \lambda_{j'}$, is *rigid*. Let Λ be the set of unordered rigid pairs. An admissible family with $\Lambda = \varnothing$ is *rigid-free*.

For a rigid pair, (1) holds iff $y_j \neq y_{j'}$, whatever v_0 is. For a non-rigid pair of an admissible family, it excludes exactly one value of v_0. Accordingly, we single out the following values.

Definition 6 (Structurally good values)**.** A construction value $y \in G^w$ is *structurally good* if $y_j \neq y_{j'}$ for every rigid pair (with $y_0 = 0$). We write $\mathcal{S}$ for the set of structurally good values.

Each rigid pair removes at most N^{w-1} values, so

$$|G^w \setminus \mathcal{S}| \leq |\Lambda|\, N^{w-1} \leq \tbinom{w+1}{2} N^{w-1}.$$

In the real world $\mathrm{F}(x) \in \mathcal{S}$ always, whereas a uniform function leaves $\mathcal{S}$ with probability at most $|\Lambda|/N$ per value. In the indifferentiability proof this is the only cost of rigid pairs. In pseudorandomness the same property gives an $\Omega(q/N)$ distinguisher (Appendix E).

Reading the conditions off the notation. For an instance $\mathsf{GXP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}]$, the number of permutations k is the number of distinct entries of $\mathbf{c}$, and n_b is the multiplicity of b in $\mathbf{c}$. A same-class pair is a pair $j \neq j'$ with $c_j = c_{j'}$; it is rigid iff moreover $\lambda_j = \lambda_{j'}$. The instance is admissible iff $\lambda_j^{-1} - \lambda_{j'}^{-1} \in \{0\} \cup \mathrm{Aut}(G)$ for every same-class pair, and rigid-free iff in addition no same-class pair has equal masks.

Remark 1. Admissibility controls collisions *within* one input. Two inputs x, x' interact through $p_j(x) = p_{j'}(x')$ for $c(j) = c(j')$. In terms of their anchors a, a' and values y, y' this condition is

$$a = \alpha_{jj'}(a' + y'_{j'}) - y_j, \qquad \alpha_{jj'} := \lambda_j \lambda_{j'}^{-1}.$$

Comparing two such conditions, for ordered same-class pairs $f = (j, j') \neq f'$, involves the endomorphisms $\alpha_f^{-1} - \alpha_{f'}^{-1}$, and these need not be 0 or invertible. One could impose that they are ("strong admissibility"). Then the forward-query and companion lemmas of Section 5.4 would hold pointwise for every prefix. Strong admissibility is strictly stronger than admissibility, however, as the following family shows:

- $G = (\mathbb{F}_{16} \times \mathbb{F}_8)^2$, $w = 3$ and $\lambda_j = \sigma^j$;
- σ multiplying the $\mathbb{F}_{16}$ coordinates by an element of order 5 and the $\mathbb{F}_8$ coordinates by an element of order 7.

This family is admissible, since $I + \sigma^r$ is invertible for $r \leq 3$. But $\sigma^3 + \sigma^{-2} = \sigma^{-2}(\sigma^5 + I)$ has a kernel of size 256. Our proof needs only admissibility: the coincidence counts that strong admissibility would bound pointwise are bounded *in expectation over the game* (Section 5.4).

construction	notation	rigid pairs
$\pi_1 + \pi_2$ [BKR98]	$\mathsf{GXP}[G; G; (1, 2); \mathbf{ID}; (-1, 1)]$	none
sum of k permutations	$\mathsf{LCP}[G; G; (1, \ldots, k); \mathbf{ID}; \mathbf{I}; \mathbf{I}_k]$	—
$\pi(0\|\|x) \oplus \pi(1\|\|x)$	$\mathsf{GXP}[\mathbb{F}_2^d; \mathbb{F}_2^{d-1}; (1^2); \mathbf{TAG}'_1; \mathbf{I}]$	$\{0, 1\}$
XORP$[w]$ [Iwa06, BDD+26]	$\mathsf{GXP}[\mathbb{F}_2^d; \mathbb{F}_2^{d-s}; (1^{w+1}); \mathbf{TAG}_s; \mathbf{I}]$	all
XORP, independent perm.s	$\mathsf{GXP}[G; G; (1, \ldots, w+1); \mathbf{ID}; \boldsymbol{\lambda}]$, any $\boldsymbol{\lambda}$	none
LXoP$[\sigma, w]$ [Din25, BGGJ26]	$\mathsf{GXP}[\mathbb{F}_2^d; \mathbb{F}_2^{d-s}; (1^{w+1}); \mathbf{TAG}'_s; ((-\sigma)^j)_j]$	none†
LXoP, in the form of [BGGJ26]	$\mathsf{LCP}[\mathbb{F}_2^d; \mathbb{F}_2^{d-s}; (1^{w+1}); \mathbf{TAG}'_s; \mathbf{I}; B_\sigma]$	as above
fork constructions [DGL22]	in GXP^+, with a base $\mathsf{GXP}[G; G; (1, \ldots, w+1); \mathbf{ID}; \boldsymbol{\lambda}]$	none

Table 2: Published constructions in the notation of Section 3, with their rigid pairs (Definition 5). Here $s = \lceil \log_2(w+1) \rceil$. †For w-admissible σ. The rows are explained in Section 3.4.

3.4 Instances

Table 2 lists the instances. We comment on the rows as follows:

- **Tagged layouts.** In the rows over $\mathbb{F}_2^d$, each call prefixes the input with an s-bit tag, where $s = \lceil \log_2(w+1) \rceil$ and $w+1$ is the number of calls ($w+1 = 2$ for the domain-separated sum $\pi(0\|x) \oplus \pi(1\|x)$). The unused tags are free inputs.
- **XORP$[w]$.** We write XORP$[w]$ for XORP with w output blocks and $w+1$ calls to one permutation, by analogy with LXoP$[\sigma, w]$; [BDD+26] calls it XORP$_{w+1}$. The family is balanced iff $w + 1 = 2^s$; otherwise the tags $j \geq w+1$ are free inputs.
- **LXoP** (Figure 1e). The output of [BGGJ26, eq. (23)] is $\left(\pi(\iota_{j-1}x) + \sigma\pi(\iota_j x)\right)_{j=1}^{w}$, that is, identity masks with the combiner B_σ, the $w \times (w+1)$ matrix whose row j has 1 in column $j-1$, σ in column j, and 0 elsewhere. The kernel of B_σ, as a map on the vector of raw outputs, equals the kernel of $\Delta_{w+1} \circ \mathrm{diag}((-\sigma)^j)$, so the two outputs differ by an invertible linear map that a distinguisher can apply itself. Since the LCP form has identity masks and a different combiner, its rigid pairs must be read off the GXP form. There $\mu_{jj'} = (-\sigma)^{-j'}\left((-\sigma)^{j'-j} - 1\right)$ is invertible for all $j < j' \leq w$ exactly when σ is w-admissible in the sense of [BGGJ26, eq. (22)], i.e. $I + \sigma^r$ is invertible for $1 \leq r \leq w$. So LXoP is rigid-free precisely for w-admissible σ; for other σ there may be rigid or singular pairs (Section 7).
- **One call per permutation** ($n = 1$). There are no same-class pairs, so the family is rigid-free for any masks. Here the masks can even be absorbed into the permutations (Lemma 6).
- **Sums of k permutations.** For $k = 2$ the combiner $\mathbf{I}_2$ with identity masks gives the same construction as Δ_2 with masks $(-1, 1)$. For $k \geq 3$ the output has one block but k calls, and it is not of the form of Definition 3.
- **Fork constructions.** These are not instances of GXP, because one call's output is fed into the other calls' inputs. They belong to the class GXP^+ of Section 3.5, which reduces them exactly to GXP instances; the base in the table is given by Lemma 18, which also determines the masks $\boldsymbol{\lambda}$.

3.5 Forks and the reducible family

In every instance of GXP, the input of each call depends only on the construction input x. The fork constructions of Dutta et al. [DGL22] (Figure 1f) do not have this form: the output of one call is fed into

the inputs of the others. We define them in general, and then the class of constructions that our analysis covers through exact reductions to GXP.

Definition 7. Let $w \geq 1$ and $\beta_j, \gamma_{0j}, \gamma_{1j} \in \mathrm{Aut}(G)$ for $1 \leq j \leq w$. The *fork* $\mathsf{Fork}[G; \beta; \boldsymbol{\gamma}_0; \boldsymbol{\gamma}_1]$ uses $w+1$ independent permutations $\pi_0, \ldots, \pi_w$ of G and maps $x \in G$ to

$$\mathsf{Fork}^{\pi}(x) := \left(\gamma_{1j} v_j - \gamma_{0j} v_0\right)_{1 \leq j \leq w}, \qquad v_0 := \pi_0(x), \qquad v_j := \pi_j(\beta_j v_0),$$

so the first call is made at x, and call $j \geq 1$ is made at β_j applied to the output of the first call.

EDMD is the case $w = 1$ with $\beta_1 = \gamma_{01} = \gamma_{11} = \mathrm{id}$ over $\mathbb{F}_2^d$. A fork is not an instance of GXP, since no choice of the injections ι_j can make the input of call j depend on another call's output. It is, however, equivalent to one. Replacing π_j by $\pi'_j := \pi_j \circ \beta_j \circ \pi_0$ turns call j into a call of a new permutation at x itself, and the new permutations are again independent and uniform. For $w = 1$ this is the change of variables by which Bhaumik et al. [BDD+26, proof of Theorem 3, App. A.2] show that EDMD is equivalent to the sum of two permutations. In general, together with a translation of the primitive queries, it is Lemma 18 in Section 5.6. Two further reductions in that subsection, Lemmas 19 and 20, remove a permutation that is applied before all others and restrict the construction domain. The constructions reached in this way form the following class.

Definition 8. The class GXP^+ of constructions *reducible to* GXP is the smallest class of constructions, each built from independent uniform permutations of G, with the following properties:

(R1) Every instance of GXP is in GXP^+.

(R2) Every fork (Definition 7) is in GXP^+.

(R3) *Pre-permutation.* If $C' \in \mathsf{GXP}^+$ uses the permutations $\pi'_1, \ldots, \pi'_m$, then the construction C with $m+1$ permutations defined by $C^{\pi_0,\ldots,\pi_m} := C'^{\pi_1 \circ \pi_0, \ldots, \pi_m \circ \pi_0}$ is in GXP^+.

(R4) *Restriction.* If $C \in \mathsf{GXP}^+$ has domain $\mathcal{D}$ and $\emptyset \neq \mathcal{X} \subseteq \mathcal{D}$, then C restricted to $\mathcal{X}$ is in GXP^+.

Every member of GXP^+ arises from a single instance of GXP, its *base*, by finitely many applications of these rules. For a member obtained by (R1), the base is the instance itself. For a fork, it is the instance of Lemma 18, with $k = w+1$ permutations, one call each, full domain G and the masks given there. Rules (R3) and (R4) keep the base of C' and of C. (If a construction can be reached in several ways, any of them may be used.)

Proposition 3 shows that every member of GXP^+ is at least as secure, for both notions, as its base, against a simulator obtained by wrapping the coin simulator. The fork constructions of [DGL22] are members: ForkEDMD$[w]$ is a fork restricted to the inputs $x \| 0^s$; EDMD is a fork with $w = 1$; and ForkCENC$[w]$ is XORP with $w+1$ independent permutations under one pre-permutation (Corollary 4). We do not know whether ForkSTHCENC and ForkEDM-CTR of [DGL22] are members.

4 The coin simulator

This section describes the simulator for which Theorems 1 and 3 are proved. Section 5 proves its indifferentiability for admissible masks, and Section 6 shows that the acceptance coin is necessary and that the exponent $3/4$ cannot be improved for this simulator. Throughout, primitive queries use the augmented interface of Definition 1, in which a query at one call of an input returns the outputs of all its calls.

4.1 Legal sets and the index-free weight

When the simulator answers a primitive query at a call of an input x that it has not yet assigned, it must invent outputs for all $w+1$ calls of x at once. These outputs must satisfy three constraints:

(i) they must produce the construction value, that is, $v_j - v_0 = y_j$ for every j, where $y = \Gamma(x)$ and $v_j = \lambda_j(p_j)$ are the masked outputs;
(ii) each permutation must remain a permutation, so an output of π_b that is already in use cannot be used again;
(iii) two calls of x that use the same permutation are made at distinct inputs (Definition 3), so they must receive distinct outputs.

Recall from Section 3.2 that, once y is fixed, all $w+1$ outputs of x are determined by the single masked output v_0, the anchor: $v_j = v_0 + y_j$ and $p_j = \lambda_j^{-1}(v_j)$. So choosing the outputs of x amounts to choosing one group element v_0, and the three constraints become conditions on v_0. The *legal set* $\mathrm{Leg}(y)$ defined below is the set of anchors that satisfy them. The simulator samples the anchor uniformly from it, and the proof repeatedly needs its size.

Definition 9 (Legal set)**.** Fix *ranges* $R_b \subseteq G$ for $b \in [k]$; in the simulator, R_b is the set of outputs of π_b already in use. Write $\tilde{R}_j := \lambda_j(R_{c(j)})$. For $y \in \mathcal{S}$ the *legal set* $\mathrm{Leg}(y)$ is the set of $v_0 \in G$ such that $v_0 + y_j \notin \tilde{R}_j$ for every j and (1) holds for every non-rigid same-class pair; for $y \notin \mathcal{S}$, $\mathrm{Leg}(y) := \varnothing$. An anchor in $\mathrm{Leg}(y)$ is *legal*.

A tuple $(v_0, \dots, v_w)$ with $v_j = v_0 + y_j$ is a tuple of masked outputs consistent with y, with the ranges and with class-wise distinctness exactly when $v_0 \in \mathrm{Leg}(y)$. For $y \in \mathcal{S}$,

$$|\mathrm{Leg}(y)| = N - \Big|\bigcup_{j=0}^{w}(\tilde{R}_j - y_j)\Big| - e(y), \tag{2}$$

where $e(y)$ counts the distinct non-rigid collision values of v_0 that are not already in the union. If the family is admissible, then $0 \le e(y) \le \binom{w+1}{2}$. This bound on $e(y)$ uses admissibility; for arbitrary masks, $e(y)$ is bounded differently in Section 7. An anchor is lost if one of its $w+1$ outputs is already in use, which removes at most $\sum_j |R_{c(j)}|$ anchors, or if it makes two calls of one permutation collide, which removes at most one anchor per non-rigid same-class pair when the masks are admissible. Rigid pairs need no condition on v_0: for them, distinct outputs are equivalent to $y_j \neq y_{j'}$, which is part of $y \in \mathcal{S}$.

Inverse queries raise the opposite question. An inverse query $\pi_b^{-1}(t)$ asks which input of π_b has output t. A candidate input $\iota_j(x)$ can be the answer only if giving call j of x the output t is consistent with everything else, and whether it is depends only on the construction value y of x: an input x with value y can have $p_j(x) = t$ exactly when the corresponding anchor $\lambda_j(t) - y_j$ is legal.

Definition 10 (Forbidden values)**.** Fix a class b, a *target* $t \notin R_b$ and a call $j \in c^{-1}(b)$. The values *forbidden* for (t, j) are

$$\mathrm{Forb}_{t,j} := \big\{y \in G^w : \lambda_j(t) - y_j \notin \mathrm{Leg}(y)\big\}.$$

For $y \in \mathcal{S}$, membership means that one of two alternatives holds:
(F1) $y_{j'} - y_j \in \tilde{R}_{j'} - \lambda_j(t)$ for some $j' \neq j$;
(F2) some non-rigid same-class pair collides at $v_0 = \lambda_j(t) - y_j$.

The first alternative says that another call j' of x would receive an output already in use. For each j' it is a condition on the single difference $y_{j'} - y_j$, so it describes a set containing a $|R_{c(j')}|/N$ fraction of all construction values (Section A.2). The second alternative is a union of at most $\binom{w+1}{2}$ affine hyperplanes, each containing a $1/N$ fraction of all values (Fact 9). Hence, if $|R_{b'}| \le s_{\max}$ for every class b',

$$|\mathrm{Forb}_{t,j}| \le \phi_0 N^w, \qquad \phi_0 := \frac{w s_{\max} + w^2}{N}, \qquad \frac{|\mathcal{S} \cap \mathrm{Forb}_{t,j}|}{|\mathcal{S}|} \le \phi_{\mathcal{S}} := \phi_0 \frac{N^w}{|\mathcal{S}|} \le \frac{32}{31}\phi_0.$$

The last inequality assumes $w^2 \le N/32$. The bound on $|\mathrm{Forb}_{t,j}|$ holds for all $y \in G^w$, not only on $\mathcal{S}$: outside $\mathcal{S}$ the rigid constraints are also affine hyperplanes of density $1/N$, and rigid and non-rigid pairs together number at most $\binom{w+1}{2} \le w^2$.

Lemma 4. *Let $R_{b'} \subseteq G$ be arbitrary sets and put $M_{b'} := N - |R_{b'}|$. Let b be a class, $t \notin R_b$, and $j \in c^{-1}(b)$. Then*

$$\rho_j(t) := \Pr_{z \leftarrow \mathrm{Unif}(\mathcal{S})}\left[\lambda_j(t) - z_j \in \mathrm{Leg}(z)\right] = \frac{(M_b - 1)_{n_b - 1} \prod_{b' \neq b} (M_{b'})_{n_{b'}}}{|\mathcal{S}|} =: \rho_b.$$

This depends on b, on the range sizes, and on the masks only through the rigid-pair pattern, that is, through $|\mathcal{S}|$. It does not depend on the call $j \in c^{-1}(b)$ or on the target t. If $|R_{b'}| \le s_{\max}$ for all b', then $1 - \rho_b = |\mathcal{S} \cap \mathrm{Forb}_{t,j}|/|\mathcal{S}| \le \phi_{\mathcal{S}}$ and $1 - \rho_b \le w(s_{\max} + n)/N$.

Proof. Let $\mathcal{T}_0$ be the set of tuples $(p_{j'})_{j' \neq j}$ such that:

- $p_{j'} \notin R_{c(j')}$ for every $j' \neq j$;
- completed by $p_j := t$, the tuple is class-wise injective.

For $z \in \mathcal{S}$ put $v_0 := \lambda_j(t) - z_j$, $v_{j'} := v_0 + z_{j'}$ and $p_{j'} := \lambda_{j'}^{-1}(v_{j'})$, so that $p_j = t$. The condition $v_0 \in \mathrm{Leg}(z)$ says two things. First, $p_{j'} \notin R_{c(j')}$ for all j'; for $j' = j$ this is $t \notin R_b$, which holds. Second, $p_{j'} \neq p_{j''}$ for every non-rigid same-class pair. For a rigid pair, $p_{j'} \neq p_{j''}$ is equivalent to $z_{j'} \neq z_{j''}$, which is part of $z \in \mathcal{S}$.

Hence $z \mapsto (p_{j'})_{j' \neq j}$ maps $\mathcal{S} \setminus \mathrm{Forb}_{t,j}$ into $\mathcal{T}_0$. It is injective, because z is recovered from the full tuple by $z_{j'} = \lambda_{j'}(p_{j'}) - \lambda_0(p_0)$, using $v_0 = \lambda_0(p_0)$ and $z_0 = 0$. It is surjective: for an element of $\mathcal{T}_0$ completed by $p_j = t$, the same formula returns some z. This z lies in $\mathcal{S}$, since for a rigid pair $z_{j'} = z_{j''}$ would force $p_{j'} = p_{j''}$. And $\lambda_j(t) - z_j = \lambda_0(p_0)$ lies in $\mathrm{Leg}(z)$, because the range and collision conditions on v_0 are exactly the conditions defining $\mathcal{T}_0$. So $\rho_j(t) = |\mathcal{T}_0|/|\mathcal{S}|$. To count $\mathcal{T}_0$, note that in class b the $n_b - 1$ outputs other than t are distinct elements of $G \setminus (R_b \cup \{t\})$. In every other class b', the $n_{b'}$ outputs are distinct elements of $G \setminus R_{b'}$. This gives the product formula.

Finally, $1 - \rho_b = |\mathcal{S} \cap \mathrm{Forb}_{t,j}|/|\mathcal{S}| \le \phi_{\mathcal{S}}$ by definition. For the second bound, use $(M - 1)_{n_b - 1} \ge (M - n)^{n_b - 1}$ and $(M)_{n_{b'}} \ge (M - n)^{n_{b'}}$, with $M \ge N - s_{\max}$, and $\sum_{b'} n_{b'} - 1 = w$. This gives $\rho_b \ge (1 - (s_{\max} + n)/N)^w$, and Bernoulli's inequality finishes. □

Example 1. We work through these notions for XORP[2] with one unused tag. Let $G = \mathbb{F}_2^d$, $\mathcal{D} = \mathbb{F}_2^{d-2}$, one permutation π, identity masks, and the three calls $\iota_j(x) = x \| \langle j \rangle_2$ for $j = 0, 1, 2$; the inputs $x \| \langle 3 \rangle_2$ are free. Over $\mathbb{F}_2^d$, addition and subtraction are both XOR:

- *Values and anchors.* The construction value is $y = (y_1, y_2)$ with $y_j = \pi(x \| \langle j \rangle_2) + \pi(x \| \langle 0 \rangle_2)$. The anchor is $v_0 = \pi(x \| \langle 0 \rangle_2)$, and the three outputs of x are v_0, $v_0 + y_1$ and $v_0 + y_2$.

- *Structurally good values.* All three pairs of calls share a permutation and a mask, so all are rigid, and $\mathcal{S} = \{y : y_1 \neq 0,\ y_2 \neq 0,\ y_1 \neq y_2\}$, of size $(N - 1)(N - 2)$. A random function gives a value outside $\mathcal{S}$ with probability about $3/N$; this is the source of the coupling term $2w(w + 1)q/N$.

- *Legal sets.* For $y \in \mathcal{S}$, rigid pairs impose no condition on the anchor, so $\mathrm{Leg}(y) = \{v_0 : v_0,\ v_0 + y_1,\ v_0 + y_2 \notin R\}$, where R is the set of outputs of π in use.

- *Forbidden values.* An inverse query at a target $t \notin R$ can be answered by the call-j input of x only if the anchor $t + y_j$ is legal, that is, only if the other two outputs of x, namely $t + y_j + y_{j'}$ for $j' \neq j$ (with $y_0 = 0$), avoid R.

- *The coin.* With $M := N - |R|$, Lemma 4 gives $\rho = (M-1)(M-2)/(N-1)(N-2)$: the two other outputs of the candidate's input must be distinct unused values other than t. A free candidate $x\|\langle 3\rangle_2$ is accepted with exactly this probability. In the implementation below, the simulator draws a dummy value $\zeta \in \mathcal{S}$ and accepts if $t + \zeta_{j_b} \in \mathrm{Leg}(\zeta)$.

4.2 Simulator description

The simulator keeps a partial permutation P_b for each class b, with domain $\textsc{Dom}_b$ and range $\textsc{Ran}_b$. It samples Γ lazily and keeps a flag $\textsc{Failure}$. Inputs of all $w+1$ calls of an x are assigned together, so $\textsc{Dom}_b$ consists of the inputs $\iota_j(x)$ with $c(j) = b$ of the assigned x, together with the assigned free inputs. Legal sets $\mathrm{Leg}(\cdot)$ are taken with respect to the ranges $R_b := \textsc{Ran}_b$.

Before the formal definition, we describe the simulator at a high level. It extends the natural simulator for the sum of permutations, used for instance in [BDD+26, Kim26], with two additions: one anchor per input in place of one output per input, and a coin for free inputs. A complete description in pseudocode is in Algorithms 1 and 2:

1. *Forward query at a call of a new input x.* Look up the construction value $y = \Gamma(x)$, choose an anchor uniformly among the legal ones, and assign the outputs of *all* calls of x from it. This keeps every permutation injective and makes the construction value of x correct.

2. *Forward query at a free input.* No construction value constrains it, so answer with a uniform unused output, as a lazily sampled permutation would.

3. *Inverse query at a new output t of π_b.* Make up to m trials. In each trial, draw a uniform unused input of π_b as a candidate preimage. If it is a call of some input x, accept it exactly when giving it the output t is legal for the value $\Gamma(x)$; then assign the outputs of all calls of x. If it is a free input, toss a coin with exactly the acceptance probability that a family candidate has, and accept on heads. If no trial accepts, give up and answer $\perp$ from then on.

The rejection of illegal family candidates is unavoidable: an input whose value y forbids the output t cannot be the preimage of t. The coin makes free candidates face the same odds, which keeps the answer distribution of an inverse query close to that of a real permutation (see *Purpose of the coin* below).

Definition 11. The simulator answers the two kinds of primitive query as follows:

- **Forward query $S_b(u)$.**
  - If $u \in \textsc{Dom}_b$, return $P_b(u)$.
  - If u is a free input, sample $v \leftarrow \mathrm{Unif}(G \setminus \textsc{Ran}_b)$, set $P_b(u) \leftarrow v$, and return v.
  - Otherwise $u = \iota_j(x)$ for a unique unassigned x. Let $y \leftarrow \Gamma(x)$. If $y \notin \mathcal{S}$ or $\mathrm{Leg}(y) = \varnothing$, set $\textsc{Failure}$ and return $\perp$. Otherwise sample $v_0 \leftarrow \mathrm{Unif}(\mathrm{Leg}(y))$, assign $P_{c(j')}(\iota_{j'}(x)) \leftarrow \lambda_{j'}^{-1}(v_0 + y_{j'})$ for all j', and return $P_b(u)$.

- **Inverse query $S_b^{-1}(t)$.** If $t \in \textsc{Ran}_b$, return $P_b^{-1}(t)$. Otherwise make up to m *trials*. A trial samples a candidate $u \leftarrow \mathrm{Unif}(G \setminus \textsc{Dom}_b)$.
  - If $u = \iota_j(x)$ is a family input, let $y \leftarrow \Gamma(x)$. If $y \notin \mathcal{S}$, set $\textsc{Failure}$ and return $\perp$. Put $v_0 := \lambda_j(t) - y_j$. The candidate is *legal* if $v_0 \in \mathrm{Leg}(y)$. A legal candidate is accepted: assign all $w+1$ outputs of x as above, and return u.

Algorithm 1 Forward query $S_b(u)$ of $\textsc{Sim}_m$.

if $\textsc{Failure}$ **then return** $\perp$
if $u \in \textsc{Dom}_b$ **then return** $P_b(u)$
if $u \in \mathrm{Free}_b$ **then**
    $v \leftarrow \mathrm{Unif}(G \setminus \textsc{Ran}_b)$; $P_b(u) \leftarrow v$
    return v
$(x, j) \leftarrow \iota^{-1}_{c^{-1}(b)}(u)$; $y \leftarrow \Gamma(x)$
if $y \notin \mathcal{S}$ or $\mathrm{Leg}(y) = \varnothing$ **then**
    $\textsc{Failure} \leftarrow 1$; **return** $\perp$
$v_0 \leftarrow \mathrm{Unif}(\mathrm{Leg}(y))$
for $j' = 0, \ldots, w$ **do**
    $P_{c(j')}(\iota_{j'}(x)) \leftarrow \lambda^{-1}_{j'}(v_0 + y_{j'})$
return $P_b(u)$

– If u is a free input, accept it with probability ρ_b (Lemma 4, computed from the current range sizes), independently of everything else. If accepted, set $P_b(u) \leftarrow t$ and return u.

If all m trials reject, set $\textsc{Failure}$ and return $\perp$.

Once $\textsc{Failure}$ is set, every later primitive query is answered by $\perp$. The distinguisher's access to Γ is unaffected.

The simulator reads Γ at most once per forward query and at most m times per inverse query, once for each family candidate; free answers and dummy values never read Γ. So against a distinguisher with q primitive queries it makes at most mq queries to Γ.

In Algorithms 1 and 2, $\iota^{-1}_{c^{-1}(b)}(u)$ denotes the pair (x, j) with $j \in c^{-1}(b)$ and $\iota_j(x) = u$, that is, the input x and the call j of class b at which u is evaluated. The algorithms use it only for an input u of π_b that is not free, so such a pair exists, and it is unique because each ι_j is injective and the images $\iota_j(\mathcal{D})$ with $c(j) = b$ are pairwise disjoint (Definition 3). The coin is tossed as described in *Implementing the coin* below, with a fixed call $j_b \in c^{-1}(b)$ and a dummy value ζ; by Lemma 4 it accepts with probability exactly ρ_b.

Implementing the coin. The coin needs no computation of ρ_b. Fix any call $j_b \in c^{-1}(b)$, sample a *dummy value* $\zeta \leftarrow \mathrm{Unif}(\mathcal{S})$ by rejection sampling from G^w, and accept iff $\lambda_{j_b}(t) - \zeta_{j_b} \in \mathrm{Leg}(\zeta)$. By Lemma 4 this happens with probability exactly ρ_b, whatever j_b and t are. Thus a free candidate is treated like a family candidate whose oracle value is fresh and is never seen by anyone. For balanced families there are no free inputs and the coin is never tossed. In that case, for $k = 2$ and $n = w = 1$, $\textsc{Sim}_m$ is the uniform simulator $\textsc{USim}_m$ of [Kim26, Algorithm 1]. Over $\mathbb{F}_2^d$ with $k = 1$ and identity masks, $\textsc{Sim}_2$ is the simulator of [BDD+26, Algorithms 2–3], which makes two attempts. We analyse $m = 3$. Fewer trials do not suffice for the exponent $3/4$ in a theorem covering all admissible families: already on the full-domain two-permutation instance, with $m \le 2$ trials the simulator's failures alone distinguish it at about $N^{m/(m+1)}$ queries (Section 6.2).

Running time. $\mathrm{Unif}(\mathrm{Leg}(y))$ is sampled by rejection from G. Membership in $\mathrm{Leg}(y)$ takes $w + 1$ range look-ups and at most $\binom{w+1}{2}$ comparisons. For admissible families, under the standing assumptions of Section 5.1, $|\mathrm{Leg}(y)| \ge 15N/16$ for every $y \in \mathcal{S}$ at every point of the game. (For non-admissible masks this can fail, and $\mathrm{Leg}(y)$ can even be empty (Remark 6); the simulator then answers $\perp$, and no running-time bound is claimed.) So the branch $\mathrm{Leg}(y) = \varnothing$ never occurs, and the expected number of draws per call is at most $16/15$. There are two kinds of rejection sampler:

Algorithm 2 Inverse query $S_b^{-1}(t)$ of $\textsc{Sim}_m$.

if $\textsc{Failure}$ **then return** $\perp$
if $t \in \textsc{Ran}_b$ **then return** $P_b^{-1}(t)$
for $trial = 1, \ldots, m$ **do**
  $u \leftarrow \mathrm{Unif}(G \setminus \textsc{Dom}_b)$
  if $u \in \mathrm{Free}_b$ **then**
    $\zeta \leftarrow \mathrm{Unif}(\mathcal{S})$ ▷ the coin, with a dummy value
    if $\lambda_{j_b}(t) - \zeta_{j_b} \in \mathrm{Leg}(\zeta)$ **then**
      $P_b(u) \leftarrow t$; **return** u
  else
    $(x, j) \leftarrow \iota^{-1}_{c^{-1}(b)}(u)$; $y \leftarrow \Gamma(x)$
    if $y \notin \mathcal{S}$ **then**
      $\textsc{Failure} \leftarrow 1$; **return** $\perp$
    $v_0 \leftarrow \lambda_j(t) - y_j$
    if $v_0 \in \mathrm{Leg}(y)$ **then**
      for $j' = 0, \ldots, w$ **do**
        $P_{c(j')}(\iota_{j'}(x)) \leftarrow \lambda_{j'}^{-1}(v_0 + y_{j'})$
      return u
Failure $\leftarrow 1$; **return** $\perp$

- *The anchor sampler.* A forward query runs at most one. Each of its draws is rejected with probability at most $1/16$.
- *The dummy sampler* of the coin. An inverse query runs at most three, one per trial. Each draw from G^w falls outside $\mathcal{S}$ with probability at most $|\Lambda|/N$, where $|\Lambda| \le \sum_b \binom{n_b}{2} \le (n-1)(w+1)/2$. The regime (3) is nonempty only for $N \ge 128\, nw^2$, so this probability is at most $1/128$.

Cap every sampler at $c \ge 1$ draws: a sampler fails, and the simulator answers $\perp$, once c draws in a row have been rejected. A forward query then fails with probability at most 16^{-c}, and an inverse query with probability at most $3 \cdot 128^{-c} \le 16^{-c}$. A union bound over the queries shows that, for admissible families, capping adds at most $q\,16^{-c}$ to all advantage bounds.

Purpose of the coin. The real world gives a fresh target the same preimage probability at every unassigned input of π_b, family or free. A family candidate passes a trial with probability about $\rho_b < 1$, because its oracle value may make it illegal. Accepting free candidates always would therefore favour free inputs by a factor of about $1/\rho_b$. The coin removes this bias. Section 6.1 shows that without it the simulator is distinguishable at about $N^{2/3}$ queries.

4.3 Equivalent presentations

Different parameters can describe constructions that are equivalent for both security notions. The following moves preserve both advantages, against the coin simulator.

Lemma 5. *Let* $\mathrm{F} = \mathsf{GXP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}]$, *and let* F' *be obtained from* F *by one of the following moves:*

(a) Mask absorption. *For automorphisms* $\theta_1, \ldots, \theta_k$ *of* G, *replace each* λ_j *by* $\lambda_j \theta_{c_j}^{-1}$.

(b) Change of reference. *For a call* r, *reorder the calls so that call* r *becomes call 0, keeping all other parameters attached to their calls.*

(c) Change of layout. *Replace* $(\mathcal{D}, \iota)$ *by any* $(\mathcal{D}', \iota')$ *with* $|\mathcal{D}'| = |\mathcal{D}|$ *that satisfies the disjointness condition of Definition 2.*

(d) Relabelling. *Permute the permutation indices, or the calls other than call 0.*

Then $\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{F}}(q) = \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{F}'}(q)$ *and* $\mathbf{Adv}^{\mathrm{indiff}}_{\mathsf{F},\textsc{Sim}_m}(q) = \mathbf{Adv}^{\mathrm{indiff}}_{\mathsf{F}',\textsc{Sim}_m}(q)$ *for all* m *and* q*. The moves preserve admissibility, the set of rigid pairs, the class sizes and the densities* $n_b|\mathcal{D}|/N$*, up to the relabelling of classes and calls.*

In each case the proof gives a bijection of the permutation tuples under which the constructions agree, translates the primitive interfaces at no query cost, and couples the two simulators so that their answers correspond pathwise; it is in Section D.1. The lemma concerns GXP only. The coin simulator is defined only for the difference combiner, and for a general combiner an invertible change of output coordinates need not produce an instance of GXP. The most frequently used special case is the following.

Lemma 6. *Suppose* $\lambda_j = \lambda_{j'}$ *whenever* $c(j) = c(j')$*, and write* λ_b *for the mask of class* b*. Let* F' *be the family with the same parameters and identity masks. Then* $\mathbf{Adv}^{\mathrm{indiff}}_{\mathsf{F},\textsc{Sim}_m}(q) = \mathbf{Adv}^{\mathrm{indiff}}_{\mathsf{F}',\textsc{Sim}_m}(q)$ *and* $\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{F}}(q) = \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{F}'}(q)$ *for all* m *and* q*. This applies in particular when* $n = 1$*.*

Proof. Apply Lemma 5(a) with $\theta_b := \lambda_b$. Every mask becomes $\lambda_b\lambda_b^{-1} = \mathrm{id}$. □

Remark 2. By (b), (d) and (a) of Lemma 5 (choose call *0*, relabel its permutation as permutation *1*, then absorb masks), every instance is equivalent to one in which call *0* is a call of permutation *1* and the first call of every permutation has the identity mask. Such a *normalised representative* is not unique. For example, on $G = \mathbb{Z}_7$, take one permutation with two calls and masks $(1, 2)$; this is admissible, since $1 - 2^{-1} = 1 - 4 = 4$ is a unit. Taking the second call as call *0* gives the masks $(2, 1)$, and absorbing $\theta = 2$ gives $(1, 4)$. Both $(1, 2)$ and $(1, 4)$ are normalised, and they describe equivalent instances. A canonical form would need a deterministic choice within these orbits; our theorems do not need one. Two special cases are useful:

- With one call per permutation ($n = 1$), all masks normalise to the identity. So XORP with independent permutations and arbitrary masks is equivalent to the same construction with identity masks.
- After Lemma 18, ForkEDMD$[w]$ becomes an instance with $n = 1$ and the full domain, and so it is equivalent to XORP with $w + 1$ independent permutations and identity masks.

5 Indifferentiability of the GXP family

This section proves Theorem 1. After the standing assumptions and the statement, Section 5.3 introduces the hybrid world and the plan of the proof, which occupies Sections 5.4 and 5.5. Section 5.6 derives the bounds for the instances of Section 3.4. Section 7 extends the theorem to arbitrary automorphism masks.

5.1 Standing assumptions and query types

From now on we assume

$$N \geq 2^{10}, \qquad w^2 \leq N/32, \qquad 2w \leq q \leq \frac{N}{64\, n\, w}. \tag{3}$$

The regime is nonempty only for $N \geq 128\, n\, w^2$. For $q < 2w$ the advantage is at most its value at $q = 2w$. The proof classifies the inputs at each point of the game as follows.

Definition 12. At any point of the game, an input $x \in \mathcal{D}$ is:

- *revealed* if its tuple is in the transcript;
- *construction-only* if its construction value is known but its tuple is not;
- *fresh* otherwise.

A permutation input of class b is *fresh* if it is $\iota_j(x)$ for a fresh x (a fresh family input) or an unassigned free input of class b.

Let r, h and u be the numbers of revealed, construction-only and fresh inputs $x \in \mathcal{D}$. Let f_b be the number of assigned free inputs of class b, $f := \sum_b f_b$, and

$$\bar{s} := r + h + f \le q, \qquad |R_b| = n_b r + f_b \le n\bar{s}.$$

Each query exposes at most one input or one free input, so $\bar{s} \le q$. Every fresh query has one of five types:

(1) a construction query at a fresh input;
(2) a forward query at a construction-only input;
(3) a forward query at a fresh family input;
(4) an inverse query at a target outside the range;
(5) a forward query at an unassigned free input.

Derived bounds. Put

$$\Phi := \bar{F}\bar{s} + 2w^2, \qquad \ell := N - \Phi.$$

Then (3) gives:

- $\bar{F}q \le N/32$, hence $\Phi \le 3N/32$ and $\ell \ge 29N/32$;
- $|G \setminus \textsc{Dom}_b| \ge N - nq \ge 63N/64$;
- $\phi_0 \le (wnq + w^2)/N \le 1/16$, and $1 - \rho_b \le w(nq + n)/N \le 1/4$ by Lemma 4 with $s_{\max} = n\bar{s}$.

As shown in Section 5.4, Φ bounds every set of excluded anchor values that appears in the proof.

In the proofs of Sections 5.4 and 5.5 we track orders of magnitude, and write $O(\cdot)$ for a bound whose implied constant is absolute. The statements carry explicit constants; their evaluation is collected, step by step, in Section B.1.

5.2 Main theorem

Theorem 1. *Let the family be admissible, with any class sizes and with free inputs allowed. Put $n := \max_b n_b$ and assume (3). Then the coin simulator with three trials satisfies*

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\textsc{Sim}_3}(q) \;\le\; 10\,n^{3/2}(w+1)^2\frac{q^2}{N^{3/2}} \;+\; 2w(w+1)\frac{q}{N} \;+\; 3w^2\frac{\sqrt{q}}{N}.$$

The second term is present only if there are rigid pairs. The third is present only if there are non-rigid same-class pairs. More precisely, the first coefficient is $\sqrt{K_{\mathrm{RH}}/2} + \sqrt{K_{\mathrm{HI}}/2}$, with $K_{\mathrm{RH}} \le 85\,n^3(w+1)^4$ and $K_{\mathrm{HI}} \le 228\,n^3w^3$ (Appendix B).

For fixed (k, n, w) the advantage tends to 0 when $q = o(N^{3/4})$. At $q = cN^{3/4}$ it is bounded by a constant depending on c, n and w. "Indifferentiable up to $N^{3/4}$ queries" is therefore a statement about the exponent, not a claim of negligible advantage at exactly $N^{3/4}$ queries. Unless $n = 1$, at least one of the last two terms is present, since every same-class pair is either rigid or non-rigid.

5.3 Hybrid and plan of the proof

The proof follows the architecture of Kim [Kim26]. It compares three worlds through a hybrid **Hyb** that answers every query by a simple explicit law. For inverse queries it uses the following notion.

Definition 13. Fix an inverse query at t on class b. A *legal pair* is a construction-only input x_i, with construction value z_i, together with a call $j \in c^{-1}(b)$ such that $\lambda_j(t) - z_{i,j} \in \mathrm{Leg}(z_i)$.

The hybrid answers a fresh query according to its type:

- **Type (1).** Answer $z \leftarrow \mathrm{Unif}(\mathcal{S})$.
- **Types (2) and (3).** For a forward query at x with construction value z (known, or just sampled as in type (1)), answer the tuple with anchor $v_0 \leftarrow \mathrm{Unif}(\mathrm{Leg}(z))$.
- **Type (5).** Answer $v \leftarrow \mathrm{Unif}(G \setminus \textsc{Ran}_b)$.
- **Type (4).** For an inverse query at t on class b, choose an available input with the following weights:
  - *1* for each legal pair (x_i, j);
  - *0* for the other construction-only pairs;
  - ρ_b for *every* fresh input of class b, family or free.

  If the chosen input is a fresh family input $\iota_j(x)$, complete its tuple by sampling $z \leftarrow \mathrm{Unif}(\mathcal{S} \setminus \mathrm{Forb}_{t,j})$ and the anchor $\lambda_j(t) - z_j$. A free output has no companion.

Write $e_b := n_b u + g_b$ for the number of fresh inputs of class b, where $g_b := |\mathrm{Free}_b| - f_b$ counts the unassigned free inputs, and k_{leg} for the number of legal pairs. The hybrid's normaliser for a type-(4) query is $N_{\textbf{Hyb}} := k_{\mathrm{leg}} + e_b\rho_b$. For $k = 2$ and $n = w = 1$, **Hyb** is the hybrid of [Kim26, Algorithm 2], with Kim's weight $(N - r)/N$ replaced by ρ_b, and the uniform completion replaced by the conditional law of the tuple. The proof of Theorem 1 has three steps:

1. **Real to hybrid** (Section 5.4). By the chain rule, $\mathrm{KL}(T_{\textbf{Real}} \| T_{\textbf{Hyb}})$ is the sum over queries of the expected divergence of the next response. Proposition 1 bounds this by $K_{\mathrm{RH}} q^4/N^3 + 12w^4 q/N^2$. The work lies in the real law of each response type, conditioned on what the distinguisher has seen. The anchors of construction-only inputs are hidden, and they shift the probabilities of everything else.
2. **Hybrid to ideal** (Section 5.5). First replace Γ by a uniform function into $\mathcal{S}$; this coupling costs the rigid-pair term. Then compare **Hyb** with the simulator's answers, conditioned on success. The work lies in the posterior law of the unrevealed oracle values, which the simulator's past rejections have tilted. Proposition 2 bounds the divergence by $K_{\mathrm{HI}} q^4/N^3$, including the failure probability.
3. **Assembly** (Section 5.5). Pinsker's inequality and the triangle inequality.

Remark 3 (Support conditions). The chain rule for relative entropy, used in Section 5.5, requires $\mathrm{supp}(T_{W_0}) \subseteq \mathrm{supp}(T_{W_1})$ for the two transcript laws of each comparison. Augmentation needs no check: revealing the full tuple at every primitive query cannot decrease the advantage, since the distinguisher can discard the extra coordinates (Definition 1). From **Real** to **Hyb**, the hybrid gives positive mass to every real response: real construction values lie in $\mathcal{S}$, and real anchors lie in the legal sets; a real inverse output is a legal construction-only input, a fresh family input or a free input, which **Hyb** weights by *1*, ρ_b and ρ_b, with $\rho_b \geq 3/4$; real companions lie in $\mathcal{T}_0$; and a real answer at a free input lies in $G \setminus \textsc{Ran}_b$, where **Hyb** is uniform. From **Hyb** to **Ideal′**, a hybrid inverse response is a legal input or a free input, and **Ideal′** returns a legal input on the first trial with positive probability, and a free input with probability $\rho_b/L > 0$. The other query types have identical supports in the two worlds.

5.4 From the real world to the hybrid

We first bound $\mathrm{KL}(T_{\textsc{Real}} \| T_{\textsc{Hyb}})$. The difficulty is the construction-only inputs. Their construction values are known to the distinguisher, but their anchors, and hence their raw outputs, are hidden. They occupy unknown outputs of the permutations, and this shifts the real law of every later response away from the simple law of **Hyb**.

Section 5.4 describes the hidden anchors. Section 5.4 then bound the divergence of each query type. Some of these bounds contain *coincidence counts*, which are not small for every prefix when the masks are merely admissible. Section 5.4 bounds them in expectation over the game, and Section 5.4 assembles the result.

Exposure slots and the revealing game. Some of the bounds below depend on coincidences between the construction values of pairs of inputs, and the distinguisher chooses these inputs adaptively. To average such quantities over the game, we need names for the inputs that are fixed before any adaptive choice. The exposure slots defined here provide them: an input is named by the query at which it was first exposed.

Definition 14. Number the queries $1, \dots, q$. An input $x \in \mathcal{D}$ is *exposed* at the first query at which the real world samples its raw tuple, which is one of the following:

- a construction query at an input whose tuple is not yet assigned;
- a type-(3) forward query;
- an inverse query whose answer lies in the tuple of x.

A free input is exposed by the query that first evaluates π_b there. The *exposure slot* of an input is the number ω of the query that exposes it, and $x^{(\omega)}$ denotes the input exposed at slot ω; it is undefined if the game ends earlier or the query exposes nothing.

Each query exposes at most one input or one free input, so there are at most q exposure slots. A slot is *ordinary-fresh* if its query is a construction query at a not-yet-exposed input, or a type-(3) forward query. Then $x^{(\omega)}$ is named by the query and is determined by the history before the slot. An inverse-exposure slot is not ordinary-fresh: there $x^{(\omega)}$ is part of the random response. We note two facts about eligibility:

- an input is construction-only at some time only if its exposure slot is a construction query;
- an input is the subject of a type-(2) or type-(3) forward query only if its exposure slot is a construction query (type 2) or that forward query itself (type 3), and each input is the subject of such a query at most once.

The *revealing game* is the real augmented game in which every exposure, of every kind, reveals the whole raw tuple. A deterministic distinguisher that ignores the extra information makes the same queries. So all values, and the real-game transcript, have the same law as in the real game, and the real-game prefix is a function of the revealing-game history.

Let $\mathcal{H}_\omega$ be the revealing-game history just before slot ω. It consists of the exposed inputs with their raw tuples and the distinguisher's choices; at most nq values of each permutation are fixed. Conditional on $\mathcal{H}_\omega$, the permutation values at unexposed inputs are a uniform injective completion. Hence, at an *ordinary-fresh* slot ω, for a call j and a set $S \subseteq G$ determined by $\mathcal{H}_\omega$ and the coordinates of the tuple of $x^{(\omega)}$ other than $p_j(x^{(\omega)})$,

$$\Pr\left[\lambda_j p_j(x^{(\omega)}) \in S \,\middle|\, \mathcal{H}_\omega,\ \text{the other coordinates}\right] \;\le\; \frac{|S|}{N - nq} \;\le\; \frac{32}{31} \cdot \frac{|S|}{N}. \tag{4}$$

The reason is that $p_j(x^{(\omega)})$ is uniform on G minus at most nq fixed values of its permutation, and these include the other same-class coordinates of the current tuple. Indeed, before the slot at most $q - 1$ exposures have fixed at most n values each, and the current tuple adds at most $n - 1$. The second inequality in (4) holds because $nq \le N/64$ by (3). Adaptivity enters (4) only through the choice of $x^{(\omega)}$, which $\mathcal{H}_\omega$ fixes. (4) is *false* at inverse-exposure slots, where $p_j(x^{(\omega)}) = t$ with probability about $1/n_b$, and it is never used there.

The hidden anchors. Fix an augmented prefix. Let $x_1, \dots, x_b$ be the construction-only inputs, with revealed values z_i and hidden anchors $U_i := v_0(x_i)$. Let Comp be the set of vectors $(U_i)_{i \le b}$ such that:

- each $U_i \in \mathrm{Leg}(z_i)$, with respect to the current ranges;
- the resulting raw outputs are class-wise distinct across different i.

Conditional on the transcript, U is uniform on Comp. The permutations are uniform, so every assignment of outputs at the construction-only inputs that is consistent with the transcript and with injectivity is equally likely, and these assignments correspond bijectively to Comp. Conditional on U, the permutations are independent uniform bijections from the remaining inputs onto the remaining outputs, class by class. We write $R_{b,U}$ for the ranges augmented by the outputs at the construction-only inputs, and Leg_U for the corresponding legal sets.

Exclusion count. Two construction inputs $x \neq x'$, with anchors a, a' and values z, z', interact only through class-wise distinctness. Since $v_j = v_0 + z_j$, the raw outputs are $p_j(x) = \lambda_j^{-1}(a + z_j)$ and $p_{j'}(x') = \lambda_{j'}^{-1}(a' + z'_{j'})$. So, for $c(j) = c(j')$, the two coincide iff

$$a = \alpha_{jj'}(a' + z'_{j'}) - z_j, \qquad \alpha_{jj'} := \lambda_j \lambda_{j'}^{-1} \in \mathrm{Aut}(G). \tag{5}$$

So a fixed x' excludes at most $F = \sum_b n_b^2$ values of the anchor of x. For a fixed pair (j, j'), the excluded values are distinct across different x', being affine images of distinct masked outputs of one permutation.

Given the transcript and the other anchors, the anchor of a construction-only input must avoid three kinds of value:

- at most $\sum_j |\tilde{R}_j| = \sum_b n_b |R_b| \le \bar{F}(r + f)$ masked range values;
- at most $F(b-1)$ values excluded by (5);
- at most $\binom{w+1}{2} \le w^2$ collision values.

These are at most $\bar{F}\bar{s} + w^2 \le \Phi$ values in total. Hence every conditional support of an anchor has size at least $\ell = N - \Phi \ge 29N/32$.

Construction queries. We start with the simplest query type, a construction query at a fresh input, and compute its response law in the real world.

Lemma 7. *Conditional on $U \in$ Comp, a construction query at a fresh input has response law*

$$p_U(z) = \frac{|\mathrm{Leg}_U(z)|}{\prod_b (N - |R_{b,U}|)_{n_b}} \qquad (z \in G^w).$$

In particular p_U is supported on $\mathcal{S}$, and $p_U(z) \propto |\mathrm{Leg}_U(z)|$.

Proof. The $w+1$ outputs are fresh points of their permutations. Their joint law is uniform over the tuples whose entries avoid $R_{c(j),U}$ and are class-wise distinct; there are $\prod_b (N - |R_{b,U}|)_{n_b}$ such tuples. The tuples with construction value z correspond bijectively to $\mathrm{Leg}_U(z)$, via the anchor. □

In the next lemma the multi-block structure would otherwise require a case analysis. In Fourier form the cross terms vanish.

Lemma 8. *Let $R_0, \dots, R_w$ be any sets, not necessarily distinct, with $\tilde{R}_j = \lambda_j(R_j)$ and $s_j = |R_j| \le s_{\max}$, where $(w+1)s_{\max} \le N/4$ and $w^2 \le N/32$. Let $p(z) \propto |\mathrm{Leg}(z)|$ on $\mathcal{S}$, with $|\mathrm{Leg}|$ as in (2). Then*

$$\chi^2\big(p \,\big\|\, \mathrm{Unif}(\mathcal{S})\big) \le \frac{12}{N^3} \sum_{0 \le j < j' \le w} s_j^2 s_{j'} + \frac{12\, w^5 s_{\max}^4}{N^4} + \frac{12\, w^4}{N^2} \le \frac{C_{\mathrm{con}}(w)\, s_{\max}^3}{N^3} + \frac{12 w^4}{N^2},$$

where $C_{\mathrm{con}}(w) \le 6w(w+1)+3w^4$. If no same-class pair is non-rigid, then $e \equiv 0$ in (2) and the term $12w^4/N^2$ is absent.

Proof. Since $p(z) = |\operatorname{Leg}(z)|/\sum_{z' \in \mathcal{S}} |\operatorname{Leg}(z')|$,

$$\chi^2(p \| \mathrm{Unif}(\mathcal{S})) = \frac{\mathrm{Var}_{z \leftarrow \mathrm{Unif}(\mathcal{S})} |\operatorname{Leg}(z)|}{\left(\mathbb{E}_{z \leftarrow \mathrm{Unif}(\mathcal{S})} [|\operatorname{Leg}(z)|]\right)^2}.$$

For the denominator, if $z \in \mathcal{S}$ then $|\operatorname{Leg}(z)| \ge N - \sum_j s_j - w^2 \ge (1 - \frac{1}{4} - \frac{1}{32})N$, whose square exceeds $N^2/2$.

Expanding $|\operatorname{Leg}(z)|$. For the numerator, let $\Upsilon(z) := |\bigcup_j (\tilde{R}_j - z_j)|$. By the Bonferroni inequalities,

$$\Upsilon(z) = \sum_j s_j - \Sigma_2(z) + \Sigma_3(z),$$

$$\Sigma_2(z) := \sum_{j<j'} c_{jj'}(z_{j'} - z_j), \qquad 0 \le \Sigma_3(z) \le \sum_{j<j'<j''} c_{jj'j''}(z),$$

where $c_{jj'}(d) := |\tilde{R}_j \cap (\tilde{R}_{j'} - d)|$ and $c_{jj'j''}$ is the triple intersection count. These formulas define $\Upsilon, \Sigma_2, \Sigma_3, e$ on all of G^w. Introduce the auxiliary function $\operatorname{Lg}(z) := N - \sum_j s_j + \Sigma_2(z) - \Sigma_3(z) - e(z)$ on G^w. By (2), $\operatorname{Lg} = |\operatorname{Leg}|$ on $\mathcal{S}$. (Off $\mathcal{S}$, Lg is not $|\operatorname{Leg}| = 0$; only its values on $\mathcal{S}$ matter.) With $0 \le e \le w^2$,

$$\operatorname{Var} \operatorname{Lg} \le 3 \operatorname{Var} \Sigma_2 + 3\mathbb{E}[\Sigma_3^2] + 3w^4.$$

We bound $\operatorname{Var} \Sigma_2$ for z uniform on G^w, and then pass to $\mathcal{S}$.

Bounding $\operatorname{Var} \Sigma_2$. First, the cross terms vanish. For $z \leftarrow \mathrm{Unif}(G^w)$ with $z_0 = 0$, the variables $X_{jj'} := c_{jj'}(z_{j'} - z_j)$ for distinct unordered pairs are pairwise uncorrelated. Disjoint pairs give independent variables. Suppose the pairs share exactly one index, say $\{j, j'\}$ and $\{j, j''\}$, and at least one of j', j'' is nonzero, say j'. Conditional on z_j, the two variables are independent, being functions of $z_{j'}$ and of $z_{j''}$. Since $z_{j'} - z_j$ is uniform given z_j, $\mathbb{E}[X_{jj'} \mid z_j] = s_j s_{j'}/N$ does not depend on z_j, and so the covariance is 0. (If $j = 0$ the two variables are independent outright.) Hence $\operatorname{Var} \Sigma_2 = \sum_{j<j'} \operatorname{Var} X_{jj'}$.

For one pair we use Parseval. Let $A, B \subseteq G$ with $|A| = a$, $|B| = a'$, and $f(d) := |A \cap (B - d)| = \sum_{g \in G} \mathbf{1}_A(g)\mathbf{1}_B(g+d)$. We use the normalised transform $\hat{h}(\psi) = \mathbb{E}_g[h(g)\overline{\mathrm{e}(\psi, g)}]$ of Section 2.3. Substituting $u = g + d$, so that $\overline{\mathrm{e}(\psi, d)} = \overline{\mathrm{e}(\psi, u)}\,\mathrm{e}(\psi, g)$,

$$\hat{f}(\psi) = \frac{1}{N} \sum_g \mathbf{1}_A(g)\mathrm{e}(\psi, g) \sum_u \mathbf{1}_B(u)\overline{\mathrm{e}(\psi, u)} = \frac{1}{N} \cdot N\overline{\hat{\mathbf{1}}_A(\psi)} \cdot N\hat{\mathbf{1}}_B(\psi) = N\,\overline{\hat{\mathbf{1}}_A(\psi)}\,\hat{\mathbf{1}}_B(\psi).$$

The mean of f is $\hat{f}(0)$, so by Parseval the variance of $f(d)$ for uniform d is the sum over nonzero ψ:

$$\mathrm{Var}_d f(d) = \sum_{\psi \ne 0} |\hat{f}(\psi)|^2 = N^2 \sum_{\psi \ne 0} |\hat{\mathbf{1}}_A(\psi)|^2 |\hat{\mathbf{1}}_B(\psi)|^2 \le N^2 \Big(\frac{a}{N}\Big)^2 \frac{a'}{N} = \frac{a^2 a'}{N}.$$

Here we bounded $|\hat{\mathbf{1}}_A(\psi)| \le \mathbb{E}_g[\mathbf{1}_A(g)] = a/N$ for each ψ, and used Parseval once more for $\sum_\psi |\hat{\mathbf{1}}_B(\psi)|^2 = \mathbb{E}_g[\mathbf{1}_B(g)^2] = a'/N$. Applying this with $A = \tilde{R}_j$, $B = \tilde{R}_{j'}$ and $d = z_{j'} - z_j$, which is uniform on G (if $j = 0$ it is $z_{j'}$, which is uniform), gives $\operatorname{Var} X_{jj'} \le s_j^2 s_{j'}/N$, and so $\operatorname{Var} \Sigma_2 \le N^{-1} \sum_{j<j'} s_j^2 s_{j'}$. For a subgroup $A = B$ this is tight up to a constant. The attack of [Kim26, Sec. 4] and Section 6.2 fills a set of this kind, a subgroup or a coset progression; it detects the resulting structure through the simulator's posterior rather than through this variance.

Bounding $\mathbb{E}[\Sigma_3^2]$. Pointwise, $0 \le \Sigma_3 \le \Sigma_2 \le \binom{w+1}{2} s_{\max}$. Also $(z_{j'} - z_j, z_{j''} - z_j)$ is uniform on G^2, so $\mathbb{E}[c_{jj'j''}] = s_j s_{j'} s_{j''}/N^2$. Hence

$$\mathbb{E}[\Sigma_3^2] \le \max \Sigma_3 \cdot \mathbb{E}[\Sigma_3] \le \binom{w+1}{2} s_{\max} \cdot \binom{w+1}{3} \frac{s_{\max}^3}{N^2} \le \frac{w^5 s_{\max}^4}{N^2}.$$

It remains to pass from G^w to $\mathcal{S}$. For any X on G^w, $\mathrm{Var}_{\mathcal{S}} X \le (|G^w|/|\mathcal{S}|)\,\mathrm{Var}_{G^w} X \le 2\,\mathrm{Var}_{G^w} X$. So

$$\mathrm{Var}_{\mathcal{S}} |\,\mathrm{Leg}\,| \le 2\big(3N^{-1} \textstyle\sum s_j^2 s_{j'} + 3w^5 s_{\max}^4/N^2 + 3w^4\big),$$

and dividing by $N^2/2$ gives the first bound. For the second, $\sum_{j<j'} s_j^2 s_{j'} \le \binom{w+1}{2} s_{\max}^3$, and $12w^5 s_{\max}^4/N^4 \le 3w^4 s_{\max}^3/N^3$ because $s_{\max} \le N/4(w+1)$. If there are no non-rigid same-class pairs then $e \equiv 0$. □

For $w = 1$ the lemma contains [Kim26, Sec. 3.5], up to the constant: there $\Sigma_3 \equiv 0$, $e \equiv 0$ and $\sum s_j^2 s_{j'} = s_0^2 s_1$.

Forward queries. There are three forward-query situations. At a construction-only input (type 2), only its anchor is unknown. At a fresh family input (type 3), the construction value comes first, and that is Lemma 8. At a free input (type 5), a single output is sampled. In each case, the hidden anchors of the other construction-only inputs make some outputs slightly less likely.

Before stating the lemma we describe the mechanism. In the real world the anchor of x is uniform among the legal anchors whose raw outputs do not collide with the hidden raw outputs of the other construction-only inputs. Each hidden anchor is close to uniform, so a fixed anchor a of x is excluded with probability about Fh_x/N, and this probability is nearly the same for every a. The chi-square distance from the uniform law on $\mathrm{Leg}(z_0)$ is governed by the variance of this exclusion probability over a. The variance is small, except for a contribution from anchors a at which several of the affine maps involved coincide; the count of such coincidences is $\varkappa(a)$ below.

Lemma 9. *Let the family be admissible, and let x be an input whose construction value z_0 is known: either x is construction-only, or x is fresh and z_0 is the value just sampled in a type-(3) query. Let $E_0 := \mathrm{Leg}(z_0)$ with respect to the current ranges, and let $x_1, \dots, x_{h_x}$ be the construction-only inputs other than x, so $h_x \le \bar{s}$. Then*

$$\chi^2\big(\mathcal{L}(v_0(x) \mid \mathit{transcript}, z_0) \,\big\|\, \mathrm{Unif}(E_0)\big) \le \frac{4h_x^2 F^2 \Phi}{N^3} + \frac{6.73}{N^2}\,\mathbb{E}_{a \leftarrow \mathrm{Unif}(G)}\big[\varkappa(a)^2\big],$$

where $\varkappa(a)$ is the coincidence count of (6) below. If the masks are constant on classes, then $\varkappa \equiv 0$.

Proof. The law of the anchor is $\mathrm{Unif}(E_0)$ reweighted by the probability $1 - p_a$ that the anchor a avoids the hidden outputs, so its distance from $\mathrm{Unif}(E_0)$ is governed by the variance of p_a.

Step 1: deriving P. Write $X = (U_i)_{i \le h_x}$ for the hidden anchors of $x_1, \dots, x_{h_x}$, and Comp_{-x} for the set of their compatible values, defined as in Section 5.4 but without x. The permutations are uniform, so, conditional on the transcript and on z_0, every assignment of raw outputs to $x, x_1, \dots, x_{h_x}$ that is consistent with the transcript and with injectivity is equally likely. These assignments correspond bijectively to the pairs (a, X) with

$$a \in E_0, \qquad X \in \mathrm{Comp}_{-x}, \qquad a \notin \mathrm{Occ}(X),$$

where a is the anchor of x, and $a \notin \mathrm{Occ}(X)$ means that no raw output of x equals a raw output, in the same class, of some x_i. So $(v_0(x), X)$ is uniform on this set. Summing over X, the law of $a := v_0(x)$ is

$$P(a) = \frac{\mathbf{1}[a \in E_0](1 - p_a)}{\sum_{a' \in E_0}(1 - p_{a'})}, \qquad p_a := \Pr_{X \leftarrow \mathrm{Unif}(\mathrm{Comp}_{-x})}[a \in \mathrm{Occ}(X)].$$

Put $\bar{\pi} := \mathbb{E}_{a \leftarrow \mathrm{Unif}(E_0)}[1 - p_a]$. Then $|E_0|P(a) = (1 - p_a)/\bar{\pi}$ on E_0, and

$$\chi^2(P \| \mathrm{Unif}(E_0)) = \mathbb{E}_{a \leftarrow \mathrm{Unif}(E_0)}\left[\left(\frac{1 - p_a}{\bar{\pi}} - 1\right)^2\right] = \frac{\mathrm{Var}_{\mathrm{Unif}(E_0)}[p_a]}{\bar{\pi}^2}.$$

Step 2: decomposing $\mathrm{Occ}(X)$. By (5) with the roles of the two inputs as here, the raw output $p_j(x)$ equals $p_{j'}(x_i)$, for $c(j) = c(j')$, exactly when $a = \alpha_{jj'}(U_i + z_{i,j'}) - z_{0,j}$, that is, when $U_i = \alpha_{jj'}^{-1}(a + z_{0,j}) - z_{i,j'}$. Hence $a \in \mathrm{Occ}(X)$ iff $U_i \in S_i(a)$ for some i, where

$$S_i(a) := \{g_{i,f}(a) : f \in \mathcal{P}\}, \qquad g_{i,f}(a) := \alpha_f^{-1}(a + z_{0,j}) - z_{i,j'} = \alpha_f^{-1}(a) + c_{i,f}$$

for $f = (j, j')$ in the set $\mathcal{P}$ of ordered same-class pairs, and $c_{i,f} := \alpha_f^{-1}(z_{0,j}) - z_{i,j'}$. Each $g_{i,f}$ is an affine bijection of G. Different pairs f may give the same map; keep one representative of each distinct map in a set $\mathcal{P}_i \subseteq \mathcal{P}$, so $|\mathcal{P}_i| \le F$ and $|S_i(a)| \le F$. Two distinct maps $g_{i,f}, g_{i,f'}$ can still agree at particular points. Put $\psi_{ff'} := \alpha_f^{-1} - \alpha_{f'}^{-1} \in \mathrm{End}(G)$, so that $g_{i,f}(a) = g_{i,f'}(a)$ iff $\psi_{ff'}(a) = c_{i,f'} - c_{i,f}$.

Definition 15. For $i \le h_x$ and distinct $f, f' \in \mathcal{P}_i$, the *coincidence set* is

$$C_{i,ff'} := \{a : \psi_{ff'}(a) = c_{i,f'} - c_{i,f}\},$$

the set of anchors a at which $g_{i,f}$ and $g_{i,f'}$ agree. The *coincidence count* $\kappa_i(a)$ is the number of unordered pairs $\{f, f'\} \subseteq \mathcal{P}_i$ with $a \in C_{i,ff'}$, and $\kappa(a) := \sum_i \kappa_i(a)$.

Each coincidence set is empty or a coset of $\ker \psi_{ff'}$. Index the pairs $(i, \{f, f'\})$ by θ. Expanding the square of $\kappa(a) = \sum_\theta \mathbf{1}[a \in C_\theta]$ and averaging over a gives

$$\mathbb{E}_{a \leftarrow \mathrm{Unif}(G)}\left[\kappa(a)^2\right] = \frac{1}{N} \sum_{\theta, \theta'} |C_\theta \cap C_{\theta'}|. \tag{6}$$

If the masks are constant on classes, every α_f is the identity, so all the maps $g_{i,f}$ for fixed i are translations. Distinct translations never agree, so $\kappa \equiv 0$.

Step 3: marginals of hidden anchors. The constraints that define Comp_{-x} are of two kinds: unary ($U_i \in \mathrm{Leg}(z_i)$) and pairwise (the raw outputs of x_i and $x_{i'}$ are distinct in each class). So, conditional on the other anchors X_{-i}, the anchor U_i is uniform on the set $D_i \subseteq \mathrm{Leg}(z_i)$ of values compatible with X_{-i}. By Section 5.4, $|D_i| \ge \ell$. Consequently

$$\Pr[U_i \in S \mid X_{-i}] \le \frac{|S|}{\ell} \qquad \text{for every } S \subseteq G. \tag{7}$$

In particular $p_a \le \sum_i \Pr[U_i \in S_i(a)] \le F h_x/\ell$. Under (3) this is at most a small constant, so $\bar{\pi} = 1 - \mathbb{E}[p_a]$ is bounded below by a constant (Section B.1 (1)).

We also need the one-dimensional marginals $\mu_i(a) := \Pr[U_i = a]$. Averaging (7) with $S = \{a\}$ gives $\mu_i(a) \le 1/\ell$. For a lower bound, let $a \in \mathrm{Leg}(z_i)$. Then $a \in D_i$ unless one of the other $h_x - 1$ anchors takes one of the at most F values that would make an output of x_i with anchor a collide with it; by (7) this has probability at most $F(h_x - 1)/\ell$. Since $|D_i| \le N$,

$$\frac{1 - F(h_x - 1)/\ell}{N} \le \mu_i(a) \le \frac{1}{\ell} \qquad (a \in \mathrm{Leg}(z_i)),$$

and $\mu_i(a) = 0$ for the at most Φ values $a \notin \mathrm{Leg}(z_i)$. Both deviations from $1/N$ on $\mathrm{Leg}(z_i)$ are at most $\Phi/\ell N$: the upper one is $1/\ell - 1/N = \Phi/\ell N$, since $\ell = N - \Phi$, and the lower one is $F(h_x - 1)/\ell N \le \Phi/\ell N$. So the variance of $\mu_i(a)$ for uniform $a \in G$ satisfies

$$\mathrm{Var}_a[\mu_i(a)] \le \mathbb{E}_a\big[(\mu_i(a) - \tfrac{1}{N})^2\big] \le \frac{1}{N}\Big(\Phi \cdot \frac{1}{N^2} + N \cdot \frac{\Phi^2}{\ell^2 N^2}\Big) = \frac{\Phi}{N^3}\Big(1 + \frac{\Phi N}{\ell^2}\Big) = O\Big(\frac{\Phi}{N^3}\Big),$$

where the first term counts the illegal values and the second the legal ones. The last step uses that $\Phi N/\ell^2$ is bounded by a constant, because $\Phi \le 3N/32$ and $\ell \ge 29N/32$ (Section B.1 (2)).

Step 4: approximating p_a by $A(a)$. For each fixed X, Bonferroni's inequalities give

$$\mathbf{1}[a \in \mathrm{Occ}(X)] = \sum_i \mathbf{1}[U_i \in S_i(a)] - X_a, \qquad 0 \le X_a \le \sum_{i<i'} \mathbf{1}[U_i \in S_i(a),\ U_{i'} \in S_{i'}(a)].$$

For $i \ne i'$, conditioning first on $X_{-i'}$ and then using (7) for U_i gives $\Pr[U_i \in S, U_{i'} \in S'] \le |S|\,|S'|/\ell^2$. So the expected overcount satisfies

$$\mathbb{E}[X_a] \le D_{\mathrm{f}} := \tbinom{h_x}{2}(F/\ell)^2.$$

For the first sum, $\Pr[U_i \in S_i(a)] = \sum_{u \in S_i(a)} \mu_i(u)$. We replace it by the simpler quantity

$$A(a) := \sum_i \sum_{f \in \mathcal{P}_i} \mu_i(g_{i,f}(a)),$$

which counts each $u \in S_i(a)$ with its multiplicity $\mathrm{mult}_i(u) := \#\{f \in \mathcal{P}_i : g_{i,f}(a) = u\} \ge 1$. The excess $\sum_u (\mathrm{mult}_i(u) - 1)\mu_i(u)$ is at most $\frac{1}{\ell}\sum_u \binom{\mathrm{mult}_i(u)}{2} = \kappa_i(a)/\ell$, because $m - 1 \le \binom{m}{2}$ for every integer $m \ge 1$ and $\sum_u \binom{\mathrm{mult}_i(u)}{2}$ is the number of coinciding pairs. Taking expectations over X in the Bonferroni identity therefore gives

$$p_a = A(a) - E_a, \qquad 0 \le E_a \le \frac{\kappa(a)}{\ell} + D_{\mathrm{f}}.$$

Step 5: bounding $\mathrm{Var}_a[p_a]$. We bound the two parts separately:

- A is a sum of at most Fh_x functions of the form $\mu_i \circ g_{i,f}$. Since $g_{i,f}$ is a bijection, $g_{i,f}(a)$ is uniform when a is, so each function has variance $O(\Phi/N^3)$ by Step 3. By Cauchy–Schwarz, the variance of a sum of M terms is at most M^2 times the largest variance, so $\mathrm{Var}_a[A] = O((Fh_x)^2\Phi/N^3)$.
- $\mathrm{Var}_a[E_a] \le \mathbb{E}[E_a^2] \le 2\mathbb{E}[\kappa^2]/\ell^2 + 2D_{\mathrm{f}}^2$, by $(u+v)^2 \le 2u^2 + 2v^2$.
- $D_{\mathrm{f}}^2 \le h_x^4 F^4/4\ell^4 = O(h_x^2 F^2 \Phi/N^3)$. Here we write $h_x^4F^4 = h_x^2F^2 \cdot (h_xF)^2$ and use $(h_xF)^2 \le (N/32)\Phi$, which holds because $h_xF \le N/32$ and $h_xF \le \Phi$.

With $\mathrm{Var}(u - v) \le 2\,\mathrm{Var}\,u + 2\,\mathrm{Var}\,v$, these give

$$\mathrm{Var}_{\mathrm{Unif}(G)}[p_a] \le O\Big(\frac{h_x^2F^2\Phi}{N^3}\Big) + \frac{4\,\mathbb{E}[\kappa^2]}{\ell^2}.$$

Step 6: restricting to E_0. Step 1 needs the variance under $\mathrm{Unif}(E_0)$ rather than $\mathrm{Unif}(G)$. For any function h and constant c, $\mathrm{Var}_{\mathrm{Unif}(E_0)} h \le \mathbb{E}_{\mathrm{Unif}(E_0)}[(h-c)^2] \le \frac{N}{|E_0|}\mathbb{E}_{\mathrm{Unif}(G)}[(h-c)^2]$; taking c to be the mean under $\mathrm{Unif}(G)$ shows that restricting to E_0 costs at most a factor $N/|E_0| \le N/\ell$. Dividing by $\bar{\pi}^2$, which is bounded below by a constant (Step 3),

$$\chi^2 \le \frac{N}{\ell\,\bar{\pi}^2}\,\mathrm{Var}_{\mathrm{Unif}(G)}[p_a] \le O\Big(\frac{h_x^2F^2\Phi}{N^3}\Big) + \frac{4N}{\ell^3\bar{\pi}^2}\,\mathbb{E}[\kappa^2].$$

The explicit constants 4 and 6.73 of the statement are evaluated in Section B.1 (3). □

Lemma 10. *Let the family be admissible, and let $u^\circ \in \mathrm{Free}_b$ be unassigned. Let P be the real law of $\pi_b(u^\circ)$ given the transcript. Put $\varepsilon := 0$ if every same-class pair is rigid, and $\varepsilon := 1$ otherwise. Then*

$$\chi^2\big(P \,\big\|\, \mathrm{Unif}(G \setminus R_b)\big) \le \frac{4n^2h^2\Phi + 2\varepsilon n^4h^2}{N^3} \le \frac{4n^2\bar{F}\bar{s}^3 + (8n^2w^2 + 2\varepsilon n^4)\bar{s}^2}{N^3}.$$

Proof. Conditional on $U \in \mathrm{Comp}$, π_b restricted to its unassigned inputs is a uniform bijection onto $G \setminus (R_b \cup \mathrm{Occ}_b(U))$. Here

$$\mathrm{Occ}_b(U) := \{\lambda_a^{-1}(U_i + z_{i,a}) : i \le h,\ c(a) = b\}$$

consists of exactly n_bh elements outside R_b. So $\pi_b(u^\circ)$ is uniform on a set whose size $N - |R_b| - n_bh$ does not depend on U. Hence $P(o) = (1 - p_o)/(N - |R_b| - n_bh)$ for $o \notin R_b$, with $p_o := \Pr_U[o \in \mathrm{Occ}_b(U)]$, and $\chi^2 = \mathrm{Var}_{\mathrm{Unif}(G\setminus R_b)}[p_o]/(1-\bar{p})^2$.

Now $o \in \mathrm{Occ}_b(U)$ iff $U_i \in S_i(o) := \{g_{i,a}(o) : c(a) = b\}$ for some i, where $g_{i,a}(o) := \lambda_a(o) - z_{i,a}$ is an affine bijection and $|S_i(o)| \le n$. This is the structure of Lemma 9, with $(a, S_i(a), Fh_x, E_0)$ replaced by $(o, S_i(o), nh, G \setminus R_b)$, and its proof applies line by line. The supports have size at least ℓ, the marginals satisfy the same variance bound, the overcount is $\binom{h}{2}(n/\ell)^2$, and $|G \setminus R_b| \ge N - n\bar{s} \ge \ell$.

The coincidence count is now bounded pointwise. The equation $g_{i,a}(o) = g_{i,a'}(o)$ reads $(\lambda_a - \lambda_{a'})(o) = z_{i,a} - z_{i,a'}$, and $\lambda_a - \lambda_{a'} = -\lambda_a\mu_{aa'}\lambda_{a'}$ is 0 or an automorphism by admissibility. If it is 0 (a rigid pair), structural goodness of z_i gives $z_{i,a} \ne z_{i,a'}$ and there is no coincidence. Otherwise exactly one o coincides. Hence $\kappa_i \le \binom{n}{2}$ pointwise and $\sum_o \kappa_i(o) \le \binom{n}{2}$, so $\mathbb{E}_o[\kappa^2] \le h^2n^4/4N$, and $\kappa \equiv 0$ if every same-class pair is rigid.

Steps 5 and 6 of the proof of Lemma 9 then give $\chi^2 \le (4n^2h^2\Phi + 2\varepsilon n^4h^2)/N^3$ (Section B.1 (4)). The second form uses $h \le \bar{s}$ and $\Phi = \bar{F}\bar{s} + 2w^2$. □

Inverse queries. Fix an inverse query at a target $t \notin R_b$ on class b. For a construction-only x_i and a call $j \in c^{-1}(b)$, put $a_{ij} := \lambda_j(t) - z_{i,j}$, the anchor that would make $p_j(x_i) = t$. Recall that the hybrid's law K of the inverse output gives:

- weight $1/N_{\mathrm{HYB}}$ to each legal pair;
- weight 0 to the other construction-only pairs;
- weight ρ_b/N_{HYB} to each of the e_b fresh inputs of class b.

Here $N_{\mathrm{HYB}} = k_{\mathrm{leg}} + e_b\rho_b$.

Lemma 11. *Let R be the real law of the inverse output. Then*

$$\chi^2(R\|K) \le \frac{26\,nh\,\Phi^2}{N^3} \le \frac{52\,n\bar{F}^2\bar{s}^3 + 208\,nw^4\bar{s}}{N^3}.$$

Proof. The idea is as follows. In the real world the preimage of t is a construction-only input exactly when its hidden anchor takes one specific value, which happens with probability close to $1/N$; otherwise it is uniform over the fresh inputs. The hybrid uses exactly $1/N_{\mathrm{HYB}}$ and ρ_b/N_{HYB} for these two kinds of input. So both laws give each legal construction-only pair a mass close to $1/N$, and both are constant on the fresh inputs; we bound the two discrepancies separately.

Step 1: deriving R. Let $U \leftarrow \mathrm{Unif}(\mathrm{Comp})$ be the hidden anchors (Section 5.4). The preimage of t is $\iota_j(x_i)$ iff $p_j(x_i) = t$, that is, iff $U_i = a_{ij}$. For fixed i at most one j can satisfy this, since two would make two raw outputs of x_i in class b equal. And at most one i can satisfy it, since t has only one preimage. If no pair (i, j) matches, then conditional on U the permutation π_b, restricted to its unassigned inputs, is a

uniform bijection onto the unoccupied outputs, and t is one of them. So the preimage is then uniform over all e_b fresh inputs of class b, family and free alike, irrespective of the call index. Hence

$$R(\iota_j(x_i)) = r_{ij} := \Pr[U_i = a_{ij}], \qquad R(v) = \frac{1 - \sum_{i,j} r_{ij}}{e_b} \quad \text{for every fresh input } v.$$

Step 2: bounding $N - N_{\text{Hyb}}$. Every input of class b is assigned, an input of a construction-only x_i, or fresh, so $N = |\textsc{Dom}_b| + n_b h + e_b$ with $|\textsc{Dom}_b| = n_b r + f_b$. Since $N_{\text{Hyb}} = k_{\text{leg}} + e_b \rho_b \geq e_b \rho_b$,

$$N - N_{\text{Hyb}} \leq n_b(r + h) + f_b + (1 - \rho_b)e_b \leq n\bar{s} + wn\bar{s} + wn \leq \Phi,$$

where we used $n_b(r + h) + f_b \leq n\bar{s}$, and $(1 - \rho_b)e_b \leq w(n\bar{s} + n)$ from Lemma 4 with $s_{\max} = n\bar{s}$ and $e_b \leq N$; the last step is $(w + 1)n\bar{s} + wn \leq \bar{F}\bar{s} + 2w^2$. Also $e_b \geq N - n\bar{s} \geq 3N/4$.

Step 3: construction-only inputs. If (i, j) is illegal, then $U_i \in \text{Leg}(z_i)$ forces $U_i \neq a_{ij}$, so $r_{ij} = 0 = K(\iota_j(x_i))$ and the pair contributes nothing. Let (i, j) be legal. Conditional on U_{-i}, the anchor U_i is uniform on a set $D_i \subseteq \text{Leg}(z_i)$ with $\ell \leq |D_i| \leq N$, and $a_{ij} \in D_i$ unless one of the other anchors takes one of F specific values, which has probability at most $F(h - 1)/\ell$. So, exactly as for the marginals in Lemma 9,

$$\frac{1 - F(h-1)/\ell}{N} \leq r_{ij} \leq \frac{1}{\ell}, \qquad \left| r_{ij} - \frac{1}{N} \right| \leq \frac{\Phi}{\ell N}.$$

By Step 2, also $|1/N_{\text{Hyb}} - 1/N| = (N - N_{\text{Hyb}})/N N_{\text{Hyb}} \leq \Phi/N(N - \Phi)$. So $|R - K| \leq \Phi/\ell N + \Phi/N(N - \Phi) = O(\Phi/N^2)$ at every legal pair. On its support, K is at least $\rho_b/N_{\text{Hyb}} \geq \rho_b/N \geq 3/4N$. A pair therefore contributes $O((\Phi/N^2)^2 \cdot N) = O(\Phi^2/N^3)$ to $\chi^2(R\|K) = \sum_v (R(v) - K(v))^2/K(v)$, and the $k_{\text{leg}} \leq nh$ legal pairs contribute

$$O\left(\frac{nh\Phi^2}{N^3}\right).$$

Step 4: fresh inputs. Both laws are constant on the fresh inputs: R by Step 1, and K because the hybrid gives the same weight ρ_b to family and free inputs. Let $m_R := 1 - \sum r_{ij}$ and $m_K := 1 - k_{\text{leg}}/N_{\text{Hyb}}$ be their total masses there. They differ by the total discrepancy on the legal pairs, so by Step 3

$$|m_R - m_K| \leq k_{\text{leg}} \max_{\text{legal pairs}} |R - K| =: \Delta, \qquad \Delta = O\left(\frac{nh\Phi}{N^2}\right).$$

Each of the e_b fresh inputs has $R - K = (m_R - m_K)/e_b$ and $K \geq 3/4N$, so together they contribute at most

$$e_b \left(\frac{\Delta}{e_b}\right)^2 \frac{4N}{3} = \frac{4N}{3e_b}\Delta^2 = O(\Delta^2) = O\left(\frac{n^2h^2\Phi^2}{N^4}\right) = O\left(\frac{nh\Phi^2}{N^3}\right),$$

using $e_b \geq 3N/4$ and $nh \leq N$. Adding the two contributions, the total is $O(nh\Phi^2/N^3)$; the constant 26 of the statement is evaluated in Section B.1 (5). The second form of the bound follows from $\Phi^2 \leq 2\bar{F}^2\bar{s}^2 + 8w^4$ and $h \leq \bar{s}$. □

If the inverse output is a fresh family input, the rest of its tuple is revealed as well.

Definition 16. If an inverse query on class b returns a fresh family input $\iota_j(x)$, its *companion tuple* is the tuple $(p_{j'}(x))_{j' \neq j}$ of the remaining raw outputs of x. A free output has no companion tuple.

The real law of the companion tuple differs from the hybrid's $\text{Unif}(\mathcal{S} \setminus \text{Forb}_{t,j})$ because the tuple must also avoid the hidden outputs of the construction-only inputs.

Lemma 12. *Let the family be admissible. Conditional on the inverse output being a fresh family input $\iota_j(x)$, let R^{comp} be the real law of the companion tuple $(p_{j'}(x))_{j'\neq j}$, and K^{comp} the hybrid's law. Let $\ell' := \ell - n$, and let $Y_{j'j''}$ be the coincidence counts defined in the proof; they are functions of the construction-only values only. Then*

$$\chi^2(R^{\mathrm{comp}}\|K^{\mathrm{comp}}) \le \frac{36\,w^2n^2h^2\Phi}{N^3} + \frac{32}{\ell'^2}\sum_{j'<j'',\ j',j''\neq j} \mathrm{Var}_{v\leftarrow\mathrm{Unif}(G^w)}\left[Y_{j'j''}\right].$$

Proof. The structure of the proof is the same as for Lemma 9. The real companion tuple is uniform among the hybrid's tuples that avoid the hidden outputs of the construction-only inputs, so the chi-square distance is again the variance of an exclusion probability, now over the w coordinates of the tuple. Each coordinate on its own is easy to control. The difficulty is the overlap of two coordinates that are excluded by the *same* hidden anchor, which is counted by the coincidence counts $Y_{j'j''}$.

Step 1: conditioning on the inverse output. By Step 1 of the proof of Lemma 11, the inverse output is a fresh input exactly when no anchor U_i equals one of the values $a_{ij'}$, $j' \in c^{-1}(b)$, and in that case it is uniform over the e_b fresh inputs, whatever U is. So, conditional on the inverse output being $\iota_j(x)$, the anchors U are uniform on $\mathrm{Comp}\cap\mathcal{A}$, where

$$\mathcal{A} := \{U_i \neq a_{ij'} \text{ for all } i \text{ and all } j' \in c^{-1}(b)\}.$$

The event $\mathcal{A}$ adds at most n forbidden values for each U_i. Hence, conditional on U_{-i}, U_i is uniform on a set D_i that excludes at most $\Phi + n$ values, so $|D_i| \ge \ell' \ge 7N/8$, and $\Pr[U_i \in S \mid U_{-i}] \le |S|/\ell'$ for every S.

For a class b', the hidden outputs of the construction-only inputs form the set

$$\mathrm{Occ}_{b'}(U) := \{\lambda_a^{-1}(U_i + z_{i,a}) : i \le h,\ c(a) = b'\}.$$

It has exactly $n_{b'}h$ distinct elements, all outside $R_{b'}$, and, for $b' = b$, all different from t on $\mathcal{A}$.

Step 2: deriving R^{comp}. The hybrid samples $z \leftarrow \mathrm{Unif}(\mathcal{S}\setminus \mathrm{Forb}_{t,j})$ and completes the tuple. By the bijection in the proof of Lemma 4, the completed tuple $(p_{j'})_{j'\neq j}$ is uniform on the set $\mathcal{T}_0$ of class-wise injective tuples that avoid the ranges and are consistent with $p_j = t$; so $K^{\mathrm{comp}} = \mathrm{Unif}(\mathcal{T}_0)$. In Real, conditional on U, the permutations are uniform bijections on what remains. So the remaining outputs of x are uniform over the set $\mathcal{T}_U \subseteq \mathcal{T}_0$ of tuples that also avoid $\mathrm{Occ}(U)$. Its size,

$$|\mathcal{T}_U| = \prod_{b'}(N - |R_{b'}| - n_{b'}h - \mathbf{1}[b' = b])_{n_{b'}-\mathbf{1}[b'=b]},$$

does not depend on U. Averaging over U therefore gives $R^{\mathrm{comp}}(v) = (1-p_v)/|\mathcal{T}_U|$, with $p_v := \Pr_U[v \text{ hits } \mathrm{Occ}(U)]$. Since R^{comp} sums to 1, $|\mathcal{T}_U| = |\mathcal{T}_0|(1-\bar p)$ with $\bar p := \mathbb{E}_{v\leftarrow K^{\mathrm{comp}}}[p_v]$. As in Step 1 of Lemma 9,

$$\chi^2(R^{\mathrm{comp}}\|K^{\mathrm{comp}}) = \mathbb{E}_{v\leftarrow K^{\mathrm{comp}}}\left[\left(\frac{1-p_v}{1-\bar p} - 1\right)^2\right] = \frac{\mathrm{Var}_{K^{\mathrm{comp}}}[p_v]}{(1-\bar p)^2}.$$

Step 3: bounding $\mathrm{Var}_v[p_v^{(j')}]$. For $v = (p_{j'})_{j'\neq j}$ and $j' \neq j$, let $H_{j'}$ be the event that the coordinate $p_{j'}$ is a hidden output. It is a union over the anchors,

$$H_{j'} := \{p_{j'} \in \mathrm{Occ}_{c(j')}(U)\} = \bigcup_i H_{j',i},$$

$$H_{j',i} := \{U_i \in S_i^{j'}\},$$

$$S_i^{j'} := \{\lambda_a(p_{j'}) - z_{i,a} : c(a) = c(j')\},$$

because $p_{j'} = \lambda_a^{-1}(U_i + z_{i,a})$ iff $U_i = \lambda_a(p_{j'}) - z_{i,a}$. So $|S_i^{j'}| \le n_{c(j')} \le n$ (two of the affine expressions may coincide), and $p_v^{(j')} := \Pr_U[H_{j'}] \le nh/\ell'$ by Step 1.

Under K^{comp}, the coordinate $p_{j'}$ is uniform on the $M_{j'} := N - |R_{c(j')}| - \mathbf{1}[c(j') = b]$ unused outputs of its class: this follows from the description of $\mathcal{T}_0$, which treats all unused outputs symmetrically. For each fixed U, exactly $n_{c(j')}h$ of these outputs lie in $\mathrm{Occ}_{c(j')}(U)$. Exchanging the order of the two averages, $\mathbb{E}_v[p_v^{(j')}] = n_{c(j')}h/M_{j'}$ exactly. Combined with the pointwise bound $p_v^{(j')} \le n_{c(j')}h/\ell'$, this gives

$$\mathrm{Var}_v[p_v^{(j')}] \le \mathbb{E}_v[p_v^{(j')}]\Big(\frac{n_{c(j')}h}{\ell'} - \frac{n_{c(j')}h}{M_{j'}}\Big) = \frac{(n_{c(j')}h)^2(M_{j'} - \ell')}{\ell' M_{j'}^2} = O\Big(\frac{n^2h^2\Phi}{N^3}\Big),$$

since $M_{j'} - \ell' \le \Phi + n \le 2\Phi$, and ℓ' and $M_{j'}$ are at least a constant times N (Section B.1 (6)).

Step 4: bounding ov_v. The tuple v hits $\mathrm{Occ}(U)$ iff some $H_{j'}$ occurs. Bonferroni's inequalities, averaged over U, give

$$p_v = \sum_{j' \ne j} p_v^{(j')} - \mathrm{ov}_v, \qquad 0 \le \mathrm{ov}_v \le \sum_{j' < j''} \Pr[H_{j'} \cap H_{j''}].$$

We split each intersection according to the anchors responsible, $H_{j'} \cap H_{j''} \subseteq \bigcup_{i,i'}(H_{j',i} \cap H_{j'',i'})$, and distinguish two cases:

- *Different anchors $i \ne i'$.* Conditioning on $U_{-i'}$ and using Step 1 twice, $\Pr[H_{j',i} \cap H_{j'',i'}] \le (n/\ell')^2$. Summing over the ordered pairs (i, i') gives at most $(nh/\ell')^2$.
- *The same anchor i.* Here $H_{j',i} \cap H_{j'',i} = \{U_i \in S_i^{j'} \cap S_i^{j''}\}$, which has probability at most $|S_i^{j'} \cap S_i^{j''}|/\ell'$. An element of the intersection is a value $\lambda_a(p_{j'}) - z_{i,a} = \lambda_{a'}(p_{j''}) - z_{i,a'}$, so the size of the intersection is at most

  $$m_i^{j'j''}(v) := \#\big\{(a, a') : c(a) = c(j'),\ c(a') = c(j''),\ a \ne a',\ \varphi_{aa'}(v) = z_{i,a} - z_{i,a'}\big\},$$

  where $\varphi_{aa'}(v) := \lambda_a(p_{j'}) - \lambda_{a'}(p_{j''})$. Pairs with $a = a'$ can be omitted. For them the equation reads $\lambda_a(p_{j'}) = \lambda_a(p_{j''})$, so $p_{j'} = p_{j''}$, which is impossible on $\mathcal{T}_0$ because then $c(j') = c(j'')$ and the tuple is class-wise injective. Off $\mathcal{T}_0$, m is only an auxiliary function whose second moment we use.

Summing over the $\binom{w}{2}$ pairs $j' < j''$ gives $0 \le \mathrm{ov}_v \le C_v/\ell' + D$, where

$$C_v := \sum_{j' < j''} Y_{j'j''}, \qquad Y_{j'j''} := \sum_i m_i^{j'j''}(v), \qquad D := \tbinom{w}{2}(nh/\ell')^2.$$

Step 5: bounding $\mathbb{E}[C_v^2]$. We first compute under $v \leftarrow \mathrm{Unif}(G^w)$, with independent uniform coordinates, and then pass to $\mathcal{T}_0$. Write $Y_{j'j''} = \sum_\theta \mathbf{1}[\varphi_{aa'}(v) = \delta_\theta]$ as a sum over the *triples* $\theta = (i, a, a')$, of which there are at most hn^2, with $\delta_\theta := z_{i,a} - z_{i,a'}$. Each $\varphi_{aa'}(v)$ is uniform on G, since $p_{j'}$ is uniform and λ_a is an automorphism, so $\mathbb{E}[Y_{j'j''}] \le hn^2/N$.

Two of the variables Y are uncorrelated if their index pairs are disjoint, since they are then independent. They are also uncorrelated if the index pairs share exactly one index, say j'. Conditional on $p_{j'}$, each indicator in $Y_{j'j''}$ has probability exactly $1/N$, because $p_{j''}$ is uniform and $\lambda_{a'}$ is an automorphism. So $\mathbb{E}[Y_{j'j''} \mid p_{j'}]$ is a constant, and given $p_{j'}$ the two variables are functions of the independent coordinates $p_{j''}$ and $p_{j'''}$. Hence $\mathrm{Var}\, C_v = \sum_{j' < j''} \mathrm{Var}\, Y_{j'j''}$ and

$$\mathbb{E}_{G^w}[C_v^2] = \mathrm{Var}\, C_v + (\mathbb{E}[C_v])^2 \le \sum_{j' < j''} \mathrm{Var}_{G^w}[Y_{j'j''}] + \tbinom{w}{2}^2 \frac{h^2n^4}{N^2}.$$

Finally $\mathbb{E}_{\mathcal{T}_0}[C_v^2] \le (N^w/|\mathcal{T}_0|)\,\mathbb{E}_{G^w}[C_v^2] \le 2\,\mathbb{E}_{G^w}[C_v^2]$, since $C_v^2 \ge 0$ and $|\mathcal{T}_0| = \rho_b|\mathcal{S}| \ge N^w/2$ (Section B.1 (7)).

Step 6: bounding $\mathrm{Var}_v[p_v]$. By $\mathrm{Var}(u-v) \le 2\,\mathrm{Var}\,u + 2\,\mathrm{Var}\,v$, Cauchy–Schwarz over the w coordinates (which gives $\mathrm{Var}\sum_{j'} p^{(j')} \le w\sum_{j'} \mathrm{Var}\,p^{(j')}$), and $\mathrm{Var}[\mathrm{ov}_v] \le \mathbb{E}[\mathrm{ov}_v^2] \le 2\mathbb{E}[C_v^2]/\ell'^2 + 2D^2$,

$$\mathrm{Var}_v[p_v] \le O\Big(\frac{w^2n^2h^2\Phi}{N^3}\Big) + \frac{8}{\ell'^2}\Big(\sum_{j'<j''} \mathrm{Var}[Y_{j'j''}] + \tbinom{w}{2}^2\frac{h^2n^4}{N^2}\Big) + 4\tbinom{w}{2}^2\frac{n^4h^4}{\ell'^4}.$$

With $\ell' \ge 7N/8$, the last two terms are $O(w^4h^2n^4/N^4)$ and $O(w^4n^4h^4/N^4)$ respectively. Both are $O(w^2n^2h^2\Phi/N^3)$, because $w^2n^2/N\Phi \le n^2/2N \le 1$ and $w^2n^2h^2/N\Phi \le wnh/N \le 1/64$. So

$$\mathrm{Var}_v[p_v] \le O\Big(\frac{w^2n^2h^2\Phi}{N^3}\Big) + \frac{8}{\ell'^2}\sum \mathrm{Var}[Y_{j'j''}].$$

Finally $\bar p \le \sum_{j'} n_{c(j')}h/M_{j'} \le 1/2$, so $(1-\bar p)^{-2} \le 4$ and $\chi^2 \le 4\,\mathrm{Var}_v[p_v]$. This gives the statement, with the factor $32/\ell'^2$ in front of the variance sum; the constant 36 is evaluated in Section B.1 (8). □

For the next subsection we need a pointwise upper bound on these variances. Fix the pair $\{j', j''\}$. For two triples $\theta = (i, a, a')$ and $\theta' = (i', \bar a, \bar a')$ of this pair, let

$$\Psi_{\theta\theta'}(v) := (\varphi_{aa'}(v), \varphi_{\bar a\bar a'}(v)) \in G^2.$$

Both coordinates depend only on $(p_{j'}, p_{j''})$, so we regard $\Psi_{\theta\theta'}$ as a homomorphism $G^2 \to G^2$, and write $\ker\Psi_{\theta\theta'} \le G^2$ for its kernel. Its nonempty fibres are cosets of the kernel, so they have size $|\ker\Psi_{\theta\theta'}|$. For $v \leftarrow \mathrm{Unif}(G^w)$ the pair $(p_{j'}, p_{j''})$ is uniform on G^2, and each indicator $\mathbf{1}[\varphi_{aa'}(v) = \delta_\theta]$ has mean $1/N$. Hence

$$\mathrm{Cov}\big(\mathbf{1}[\varphi_{aa'}(v) = \delta_\theta], \mathbf{1}[\varphi_{\bar a\bar a'}(v) = \delta_{\theta'}]\big) = \frac{|\ker\Psi_{\theta\theta'}|}{N^2}\mathbf{1}\big[(\delta_\theta, \delta_{\theta'}) \in \mathrm{Im}\,\Psi_{\theta\theta'}\big] - \frac{1}{N^2}. \tag{8}$$

Writing $\mathrm{Var}\,Y_{j'j''}$ as the sum of the variances of the indicators, each at most $1/N$, plus the sum of the covariances over $\theta \ne \theta'$, and dropping the negative terms $-1/N^2$, gives $\mathrm{Var}_{G^w}[Y_{j'j''}] \le V_{j'j''}$. Here the *envelope* is

$$V_{j'j''} := \sum_\theta \frac{1}{N} + \sum_{\theta\ne\theta'} \frac{|\ker\Psi_{\theta\theta'}|}{N^2}\mathbf{1}\big[(\delta_\theta, \delta_{\theta'}) \in \mathrm{Im}\,\Psi_{\theta\theta'}\big]. \tag{9}$$

All its terms are nonnegative, so $V_{j'j''}$ can only grow when triples are added. The variance itself need not have this property, which is why we work with the envelope.

Coincidences in expectation over the game. The coincidence terms of Lemmas 9 and 12 can be of order 1 for a particular prefix. Take, for example, an admissible family on $\mathbb{F}_{16}^d$ whose masks multiply different coordinates by different scalars. There, two construction-only inputs with $z_0 = z_1 = 0$ already produce coincidence sets of size $N/16$. Such prefixes are rare. The construction values are outputs of random permutations, so a coincidence set is large only when a random value falls into a subgroup. The two lemmas below make this precise. They must charge *fixed* labels, because adaptivity can select large coincidences (Remark 4).

Lemma 13. *Let the family be admissible and fix a deterministic distinguisher. Take slot labels $\omega \ne \omega'$, a role assignment ($x := x^{(\omega)}$ and $x_i := x^{(\omega')}$, or the reverse), and distinct $f, f' \in \mathcal{P}$. Let $C_{\omega,\omega',f,f'}$ be the coincidence set $C_{i,ff'}$ of the proof of Lemma 9, computed from the construction values in these roles, if both slots occur and are ordinary-fresh; otherwise let it be $\varnothing$. Then, in the real world,*

$$\mathbb{E}[|C_{\omega,\omega',f,f'}|] \le \frac{32}{31} \qquad \textit{for all fixed } \omega, \omega', f, f' \textit{ and roles,}$$

and, summing over all type-(2) and type-(3) queries of the game,

$$\mathbb{E}\Big[\sum_{\text{type-(2),(3) queries}} \mathbb{E}_{a\leftarrow\mathrm{Unif}(G)}\big[\kappa_x(a)^2\big]\Big] \le \frac{32}{31}\cdot\frac{q\,(q\binom{F}{2})^2}{N} \le \frac{0.26\,q^3F^4}{N}. \tag{10}$$

Proof. A coincidence set is the preimage $\psi^{-1}(c)$ of a constant c under an endomorphism ψ. It has size $|\ker\psi|$ if $c\in\operatorname{Im}\psi$, and is empty otherwise. The constant c involves a raw output that is still nearly uniform when it is sampled, so c lands in the subgroup $\operatorname{Im}\psi$ with probability about $|\operatorname{Im}\psi|/N$. Since $|\ker\psi|\,|\operatorname{Im}\psi| = N$, the expected size is about *1*, whatever ψ is. The work is in choosing, case by case, a raw output that is fresh at the moment it is sampled, so that (4) applies.

Step 1: decomposing c. Write $f=(j,j')$, $f'=(\tilde j,\tilde j')$ and $\psi := \psi_{ff'}$. With $c_{i,f} = \alpha_f^{-1}(z_{0,j}) - z_{i,j'}$ as in the proof of Lemma 9, the constant $c := c_{i,f'} - c_{i,f}$ splits as $c = u + u'$ with:

- $u := \alpha_{f'}^{-1}(z_{0,\tilde j}) - \alpha_f^{-1}(z_{0,j})$ depends only on the tuple of x;
- $u' := z_{i,j'} - z_{i,\tilde j'} = \lambda_{j'}p_{j'}(x_i) - \lambda_{\tilde j'}p_{\tilde j'}(x_i)$ depends only on the tuple of x_i; the anchor of x_i cancels in the difference.

The coincidence set is $C=\psi^{-1}(c)$, so $|C| = |\ker\psi|\cdot\mathbf{1}[c\in\operatorname{Im}\psi]$, and $|\ker\psi|\,|\operatorname{Im}\psi| = N$. (If $\psi = 0$ the two maps are the same and the set is empty after deduplication. Counting it as $\ker\psi = G$ with constant *0* only enlarges the quantity, and the same argument covers it.) Since $f\neq f'$, either $j\neq\tilde j$ or $j'\neq\tilde j'$. It therefore suffices to show

$$\Pr\big[c\in\operatorname{Im}\psi \text{ and both slots are eligible}\big] \le \frac{32}{31}\cdot\frac{|\operatorname{Im}\psi|}{N}. \tag{11}$$

Step 2: eligibility indicators. Write the event that both slots occur and are ordinary-fresh as $E_{\text{early}}\cap E_{\text{late}}$, one event for each slot. Only the *later* event is bounded by *1*. The earlier one is determined at the earlier slot, and we keep it, so that (4) is only ever applied on the event that its slot is ordinary-fresh. When the randomness we use is sampled at the earlier slot,

$$\mathbb{E}\big[\mathbf{1}_{E_{\text{early}}}\mathbf{1}_{E_{\text{late}}}|\ker\psi|\,\mathbf{1}[c\in\operatorname{Im}\psi]\big] \le |\ker\psi|\,\mathbb{E}\big[\mathbf{1}_{E_{\text{early}}}\Pr(c\in\operatorname{Im}\psi \mid \mathcal{H}_{\text{early}})\big].$$

When it is sampled at the later slot, we use that the ordinary-fresh status of that slot is determined before its tuple is sampled, so (4) applies there on E_{late}.

Step 3: bounding $\Pr[c\in\operatorname{Im}\psi]$. Let $\omega_{\text{late}} := \max(\omega,\omega')$. In each of the four cases below we name one raw output p, with mask λ, such that $c\in\operatorname{Im}\psi$ is equivalent to λp lying in a set S of size $|\operatorname{Im}\psi|$ that is fixed by the history at the slot of p and by the other coordinates of the tuple of p. (The set S is the image of a coset of $\operatorname{Im}\psi$ under an automorphism.) Then (4) gives (11). The cases, and the choice of p in each, are the following:

(i) $x_i = x^{(\omega_{\text{late}})}$ and $j'\neq\tilde j'$. Take $p := p_{j'}(x_i)$. Here u is fixed by $\mathcal{H}_{\omega_{\text{late}}}$, and $c\in\operatorname{Im}\psi$ iff $\lambda_{j'}p\in\operatorname{Im}\psi - u + \lambda_{\tilde j'}p_{\tilde j'}(x_i)$, a coset of $\operatorname{Im}\psi$ determined by $\mathcal{H}_{\omega_{\text{late}}}$ and by another coordinate of the tuple of x_i.

(ii) $x_i = x^{(\omega_{\text{late}})}$ and $j'=\tilde j'$. Then $u'=0$, so $c=u$ depends only on the tuple of x, which was sampled at the earlier slot; and $j\neq\tilde j$. If $\tilde j\neq 0$, take $p := p_{\tilde j}(x)$. Since $z_{0,\tilde j} = \lambda_{\tilde j}p_{\tilde j}(x) - \lambda_0 p_0(x)$, the constant u is $\alpha_{f'}^{-1}\lambda_{\tilde j}\,p$ plus terms in the other coordinates of the tuple of x, and $\alpha_{f'}^{-1}\lambda_{\tilde j}$ is an automorphism. If $\tilde j = 0$, then $j\neq 0$, and we take $p := p_j(x)$ instead.

(iii) $x = x^{(\omega_{\text{late}})}$ and $j\neq\tilde j$. Take $p := p_{\tilde j}(x)$, or $p_j(x)$ if $\tilde j = 0$, as in (ii). Now u' is fixed by $\mathcal{H}_{\omega_{\text{late}}}$, because x_i was exposed earlier.

(iv) $x = x^{(\omega_{\text{late}})}$ and $j = \tilde{\jmath}$, so $j' \neq \tilde{\jmath}'$. Then $u = (\alpha_{f'}^{-1} - \alpha_f^{-1})(z_{0,j}) = -\psi(z_{0,j})$ lies in the subgroup $\operatorname{Im}\psi$, so $c \in \operatorname{Im}\psi$ iff $u' \in \operatorname{Im}\psi$. This depends only on the tuple of x_i, which was sampled at the earlier slot. Take $p := p_{j'}(x_i)$, as in (i).

In every case $\Pr[c \in \operatorname{Im}\psi] \leq \frac{32}{31}|\operatorname{Im}\psi|/N$ on the relevant eligibility event. Hence

$$\mathbb{E}[|C|] \leq \frac{32}{31} \cdot \frac{|\ker\psi|\,|\operatorname{Im}\psi|}{N} = \frac{32}{31}.$$

Step 4: summing over the game. Consider a type-(2) or type-(3) query at x, and let $x_1, \ldots, x_{h_x}$ be the construction-only inputs at that time. The exposure slots of x and of every x_i are ordinary-fresh: x was exposed by a construction query (type 2) or by this forward query (type 3), and a construction-only input was exposed by a construction query. So, by (6), $\mathbb{E}_a[\kappa_x(a)^2]$ is a sub-sum, with nonnegative terms, of

$$\frac{1}{N}\sum_{\theta,\theta'} |C_\theta \cap C_{\theta'}|,$$

where now θ and θ' range over all labels (slot of a partner, pair $\{f, f'\}$), for the slot of x fixed. Each slot is the subject of a type-(2) or type-(3) query at most once, so the sum over these queries is at most the sum of this expression over all q possible slots of x. For fixed labels, $\mathbb{E}[|C_\theta \cap C_{\theta'}|] \leq \mathbb{E}[|C_\theta|] \leq \frac{32}{31}$. There are at most $q\binom{F}{2}$ labels θ, hence at most $(q\binom{F}{2})^2$ pairs. Therefore

$$\mathbb{E}\Big[\sum_{\text{type-(2),(3) queries}} \mathbb{E}_a[\kappa_x(a)^2]\Big] \leq q \cdot \frac{(q\binom{F}{2})^2}{N} \cdot \frac{32}{31},$$

which is at most $0.26\, q^3F^4/N$ (Section B.1 (9)). □

Remark 4. The per-pair statement concerns fixed labels only. Suppose the distinguisher construction-queries three inputs R, X_1, X_2, and then forward-queries whichever of X_1, X_2 has a nonempty coincidence set with R. For a map pair with $|\ker\psi_{ff'}| = N/16$ (an $\mathbb{F}_{16}^4$ family), the *selected* pair has expected coincidence size about $1.998 > 32/31$. The slot sum charges the labels (X_1, R) and (X_2, R) separately, each at most $32/31$, and so dominates the selected pair; only (10) is used. Likewise, (4) must not be applied at inverse-exposure slots. The eligibility condition excludes them at no cost, because neither a forward-queried input of type (2) or (3) nor a construction-only partner is ever exposed by an inverse query.

Lemma 14. *Let the family be admissible and fix a deterministic distinguisher. For each unordered pair $\{j_1, j_2\}$ of call indices, let $V_{j_1 j_2}$ be the envelope (9) computed with the triples of* all *inputs exposed by construction queries during the game. Here the slots are fixed labels, the terms are nonnegative, and undefined slots contribute nothing. Let $W := \sum_{j_1<j_2} V_{j_1 j_2}$. Then at every inverse query with output index j, on every complete execution,*

$$\sum_{j'<j'',\, j',j''\neq j} \operatorname{Var}_{G^w}[Y_{j'j''}] \leq W.$$

Also

$$\mathbb{E}[V_{j_1 j_2}] \leq \frac{2qn^2}{N} + \frac{32}{31}\frac{(qn^2)^2}{N^2} \leq \frac{1.01\,q^2n^3}{N}, \qquad \mathbb{E}[W] \leq \tbinom{w+1}{2}\frac{1.01\,q^2n^3}{N}.$$

Proof. The first claim holds pointwise, and the second follows by bounding each term of (9) in expectation.

Pointwise inequality. At an inverse query, the variances $\mathrm{Var}[Y_{j'j''}]$ of Lemma 12 are computed from the triples of the construction-only inputs at that time, for the pairs $\{j', j''\}$ with $j', j'' \neq j$. Each is at most the envelope of the same triples, by (9). The construction-only inputs at that time are among the inputs exposed by construction queries during the whole game, and the pairs $\{j', j''\}$ are among all pairs. Since the envelope only grows when triples are added, each of these envelopes is at most $V_{j'j''}$, and their sum is at most W. Nothing here conditions on j, on the time of the query, or on the decision to make it. So the inequality holds on every complete execution, and by the tower property the expected contribution of any realised inverse response is at most $\mathbb{E}[W]$.

Bounding $\mathbb{E}[V_{j_1 j_2}]$. Fix $\{j_1, j_2\}$, and index the triples by construction-exposure slots, $\theta = (\omega, a, a')$ with $c(a) = c(j_1)$, $c(a') = c(j_2)$ and $a \neq a'$. These slots are ordinary-fresh, and whether a slot is a construction query at an unassigned input is decided before its tuple is sampled. So (4) applies at each of them, with the eligibility indicator bounded by *1*. We bound the expectations of the terms of (9) by type:

- *Diagonal terms.* There are at most qn^2 triples, each contributing $1/N$.
- *Swapped pairs* $\theta' = (\omega, a', a)$: the same slot and the same index set, which exist only when $c(j_1) = c(j_2)$. For such a pair the term equals $\Pr_v[\Psi_{\theta\theta'}(v) = (\delta_\theta, \delta_{\theta'})]$, which is at most $\Pr_v[\varphi_{aa'}(v) = \delta_\theta] = 1/N$. There are at most qn^2 of them.
- *All other pairs* $\theta = (\omega, a, a')$, $\theta' = (\omega', \tilde{a}, \tilde{a}')$, where the slots differ or the index sets differ. In both of the following sub-cases we find one raw output p, with mask λ, such that one of $\delta_\theta, \delta_{\theta'}$ equals $\pm\lambda p$ plus a quantity fixed by the conditioning, and the other is fixed by the conditioning:
  - *Different slots*, say $\omega < \omega'$. Take $p := p_{\tilde{a}}(x^{(\omega')})$ and condition on $\mathcal{H}_{\omega'}$ and the other coordinates of the tuple of $x^{(\omega')}$. Then δ_θ is fixed by $\mathcal{H}_{\omega'}$, and $\delta_{\theta'} = \lambda_{\tilde{a}} p - \lambda_{\tilde{a}'} p_{\tilde{a}'}(x^{(\omega')})$, whose second term is another coordinate because $\tilde{a} \neq \tilde{a}'$.
  - *The same slot with* $\{a, a'\} \neq \{\tilde{a}, \tilde{a}'\}$. Choose $a^\circ \in \{a, a'\} \setminus \{\tilde{a}, \tilde{a}'\}$ (exchanging θ and θ' if necessary), and take $p := p_{a^\circ}(x^{(\omega)})$. Then δ_θ contains $\pm\lambda_{a^\circ} p$, while $\delta_{\theta'}$ does not involve p.

  Given the conditioned value of one coordinate, the set of values of the other that put $(\delta_\theta, \delta_{\theta'})$ in $\mathrm{Im}\,\Psi_{\theta\theta'}$ is empty or a coset of size $|\mathrm{Im}\,\Psi_{\theta\theta'}|/N = N/|\ker\Psi_{\theta\theta'}|$. This uses that each coordinate projection of $\Psi_{\theta\theta'}$ is onto G, so that every fibre of the projection of $\mathrm{Im}\,\Psi_{\theta\theta'}$ has the same size. By (4), $\pm\lambda p$ falls into this coset with probability at most $\frac{32}{31}/|\ker\Psi_{\theta\theta'}|$. So the expected term is at most $\frac{32}{31}/N^2$, and there are at most $(qn^2)^2$ such pairs.

Adding the three types gives $\mathbb{E}[V_{j_1 j_2}] \leq 2qn^2/N + \frac{32}{31}(qn^2)^2/N^2$, which is at most $1.01\,q^2n^3/N$ (Section B.1 (10)). □

The real-to-hybrid bound. We now combine the per-query bounds of this subsection into a bound on the relative entropy between the real world and the hybrid.

Proposition 1. *Let the family be admissible and assume* (3)*. Put*

$$C_{\mathrm{RH}} := \max\bigl(n^3 C_{\mathrm{con}}(w) + 4\bar{F}^3,\ 52n\bar{F}^2 + 36w^2n^2\bar{F}\bigr),$$
$$C'_{\mathrm{RH}} := \max(8\bar{F}^2w^2, 72w^4n^2) + 22w(w+1)n^3 + 1.75\bar{F}^4,$$
$$C''_{\mathrm{RH}} := 208nw^4.$$

For every deterministic distinguisher, the sum over the q responses of the expected conditional divergence from Real *to* Hyb *is at most*

$$\frac{C_{\mathrm{RH}}q^4 + C'_{\mathrm{RH}}q^3 + C''_{\mathrm{RH}}q^2}{N^3} + \frac{12w^4q}{N^2},$$

where the expectation is over the real prefix. The last term is present only if some same-class pair is non-rigid. In terms of $w+1$: $C_{\mathrm{RH}} \le 31n^3(w+1)^4$, $C'_{\mathrm{RH}} \le 78n^3(w+1)^4 + 2n^4(w+1)^4$ and $C''_{\mathrm{RH}} \le 208n(w+1)^4$ (Appendix B).

Proof. We bound the divergence of each query type at a fixed prefix, and then sum these bounds over the game.

Per-prefix bounds. We take one bound per query type, each using $\mathrm{KL} \le \chi^2$ and, for types (3) and (4), the chain rule for their two parts:

- *Type (1).* Apply Lemma 8 with $R_j = R_{c(j),U}$ and $s_{\max} = n\bar{s}$; then $(w+1)n\bar{s} \le \bar{F}q \le N/32$. KL is jointly convex, so the divergence of the mixture over $U \in \mathrm{Comp}$ is at most the average divergence. This gives $n^3 C_{\mathrm{con}} \bar{s}^3/N^3 + 12w^4/N^2$.
- *Type (2).* Lemma 9 with $h_x = h - 1 \le \bar{s}$ gives $(4\bar{F}^3\bar{s}^3 + 8\bar{F}^2w^2\bar{s}^2)/N^3$, plus the coincidence term $6.73\,\mathbb{E}_a[\kappa_x(a)^2]/N^2$.
- *Type (3).* The value by Lemma 8, then the anchor by Lemma 9 with $h_x = h$: the sum of the two previous bounds.
- *Type (4).* The output by Lemma 11, then the companion by Lemma 12. With $\Phi^2 \le 2\bar{F}^2\bar{s}^2 + 8w^4$ and $h \le \bar{s}$, this gives
$$\frac{(52n\bar{F}^2 + 36w^2n^2\bar{F})\bar{s}^3 + 72w^4n^2\bar{s}^2 + 208nw^4\bar{s}}{N^3},$$
plus the companion variance term $(32/\ell'^2)\sum \mathrm{Var}[Y_{j'j''}]$. A free output has no companion.
- *Type (5).* By Lemma 10: $(4n^2\bar{F}\bar{s}^3 + (8n^2w^2 + 2n^4)\bar{s}^2)/N^3$. This is dominated coefficient by coefficient by type (2): $4n^2\bar{F} \le 4\bar{F}^3$, and $8n^2w^2 + 2n^4 \le 8\bar{F}^2w^2 = 8n^2(w+1)^2w^2$, which holds because $n^2 \le (w+1)^2 \le 4w^2(w^2+2w)$.

Summing over the game. Bound the per-prefix terms, other than the two coincidence terms, by their maximum over the types, with $\bar{s} \le q$. Their sum over the q queries is at most
$$\frac{q\left(C_{\mathrm{RH}}q^3 + \max(8\bar{F}^2w^2, 72w^4n^2)q^2 + C''_{\mathrm{RH}}q\right)}{N^3} + \frac{12w^4q}{N^2}.$$
The two coincidence terms are charged over the whole game:

- By Lemma 13, the forward coincidence terms $6.73\,\mathbb{E}_a[\kappa_x(a)^2]/N^2$ sum in expectation to at most $1.75\,q^3\bar{F}^4/N^3$ (Section B.1 (11)).
- By Lemma 14, at each of the at most q inverse queries the companion term is at most $(32/\ell'^2)W$ on every complete execution. By the tower property and the bound on $\mathbb{E}[W]$, these terms sum in expectation to at most $22\,w(w+1)n^3q^3/N^3$ (Section B.1 (11)).

Adding the three contributions proves the proposition. The adaptive choice of query types, query times and output indices does not affect the bound, because every charged quantity is bounded pointwise by a fixed-label sum over the whole game, whose expectation is what the two lemmas bound. The evaluations in terms of m are in Appendix B. □

5.5 From the hybrid to the ideal world

In the ideal world the responses are produced by the simulator from the oracle Γ. A forward response at a fresh input reveals $\Gamma(x)$ and samples the anchor exactly as **Hyb** does. The difference lies in inverse queries. The simulator accepts a family candidate only if its oracle value makes it legal, and the rejected candidates are not shown to the distinguisher. A table with many illegal values makes rejections, and hence the observed history, more likely. So, conditional on the history, the unrevealed oracle values are *tilted* slightly towards illegal values. We now bound the effect of that tilt.

Removing structurally bad values. Let Ideal′ be Ideal with Γ replaced by a uniform function $\mathcal{D} \to \mathcal{S}$. Couple the two worlds by resampling only the structurally bad entries. The worlds agree unless some entry read by the distinguisher or by the simulator is bad. Free answers and dummy values never read Γ, and dummy values are drawn from $\mathcal{S}$ directly. Each query reads at most $1 + m$ entries, so

$$\mathrm{TV}(\textsc{Ideal}, \textsc{Ideal}') \le (1+m)q \cdot \frac{|\Lambda|}{N} \le \frac{4\binom{w+1}{2}q}{N} \qquad (m = 3). \tag{12}$$

In Ideal′ the simulator never fails structurally. All statements below are in Ideal′. For families without rigid pairs, Ideal′ = Ideal and (12) is 0.

Trials and the posterior of the table. Fix an inverse query at t on class b and a table $\Gamma = \Xi$ with values in $\mathcal{S}$. Let $L := |G \setminus \textsc{Dom}_b|$ be the number of available inputs and g the number of available free inputs. An available family input $\iota_j(x)$ is *illegal* if $\lambda_j(t) - \Xi(x)_j \notin \mathrm{Leg}(\Xi(x))$, that is, if $\Xi(x) \in \mathrm{Forb}_{t,j}$. Let $I(\Xi)$ be the number of illegal available inputs; an unassigned x contributes $I(x; \Xi(x)) \in \{0, \dots, n_b\}$ of them, one for each illegal call $j \in c^{-1}(b)$.

Definition 17. The *acceptance weight* $a(v)$ of an available input v is 1 for a legal family input, 0 for an illegal one, and ρ_b for a free input.

A trial then rejects with probability

$$R(\Xi) := \frac{I(\Xi) + g(1-\rho_b)}{L}.$$

For the call c in question, put $\vartheta_c(I) := L^{-1}(1 + R + R^2)$, where $R = (I + g(1-\rho_b))/L$. The data L, g and ρ_b of a call are fixed by the prefix.

For a prefix τ without failure, index its inverse calls by $c \in \mathcal{I}_\tau$, and let $I_c(\Xi)$ be the number of illegal inputs at call c, computed from the state before that call. Let μ_τ be the law that fixes the revealed entries of the table and makes the unrevealed ones independent and uniform on $\mathcal{S}$.

Lemma 15. *The following hold for every prefix τ without failure.*

(a) For fixed $\Gamma = \Xi$, a call with three trials returns the available input v with probability $a(v)\,\vartheta_c(I(\Xi))$, and fails with probability $R(\Xi)^3$. Also $0 \le \log \vartheta_c(I+1) - \log \vartheta_c(I) \le 3/L$.

(b) The conditional law of the table given τ is

$$\pi_\tau(\Xi) = \frac{W_\tau(\Xi)\,\mu_\tau(\Xi)}{\mathbb{E}_{\mu_\tau}[W_\tau]}, \qquad W_\tau(\Xi) := \prod_{c \in \mathcal{I}_\tau} \vartheta_c(I_c(\Xi)).$$

(c) Let $\pi_\tau^\vartheta \propto \pi_\tau\,\vartheta_c(I(\cdot))$ be the posterior reweighted by the current call. Conditional on success, the current call returns v with probability proportional to $\mathbb{E}_{\pi_\tau^\vartheta}[a(v)]$. This is:

- ρ_b *for a free input;*
- $\mathbf{1}[\mathit{legal}]$ *for a construction-only input;*
- $\theta_j := \Pr_{\pi_\tau^\vartheta}[\Gamma(x) \notin \mathrm{Forb}_{t,j}]$ *for* $v = \iota_j(x)$ *with* x *unrevealed, which does not depend on* x.

Conditional on success and on $v = \iota_j(x)$ *with* x *unrevealed, the law of* Γ *is* π_τ^ϑ *conditioned on* $\Gamma(x) \in \mathcal{S} \setminus \mathrm{Forb}_{t,j}$.

The proof is in Appendix C. By (c), the simulator's success-conditioned answer has the same form as the hybrid's. It puts weight *1* on legal construction-only pairs and ρ_b on free inputs, exactly like **Hyb**, and weight θ_j on fresh family inputs, where **Hyb** puts ρ_b. Since θ_j is the tilted probability of the event whose untilted probability is ρ_b, it remains to show that the tilt is small.

The bias of one unrevealed value. Each rejection by the simulator changes the posterior of the unrevealed oracle values. The next lemma bounds the effect on one value.

Lemma 16. *Fix an unrevealed* x *and all other entries of the table. Assume* (3). *The conditional law* P *of* $\Gamma(x)$, *optionally reweighted by the current factor* $\vartheta_c(I)$, *satisfies*

$$\frac{e^{-A}}{|\mathcal{S}|} \le P(z) \le \frac{e^{A}}{|\mathcal{S}|}, \qquad A := \frac{4nq}{N} \le \frac{1}{16}.$$

For $V = \mathcal{S}$, *or any* $V \subseteq \mathcal{S}$ *with* $|V| \ge (1 - \phi_{\mathcal{S}})|\mathcal{S}|$,

$$\chi^2\big(\mathrm{Unif}(V) \,\big\|\, P(\cdot \mid V)\big) \le \frac{44\, n^3 w\, q^3}{N^3}.$$

The bounds persist after averaging over the other entries under their conditional law.

Proof. The law P is $\mathrm{Unif}(\mathcal{S})$ reweighted by e^{β} for a bounded function β, and Lemma 1 turns bounds on β into the statement.

Step 1: deriving $\beta(z)$. Fix the other entries of the table. At each call c, the number of illegal inputs splits as $I_c = I_c^{-x} + I_c(x; z)$: here I_c^{-x} counts the illegal inputs of the other entries and is fixed, and $I_c(x; z) \in \{0, \dots, n\}$ counts the illegal inputs of x at call c when $\Gamma(x) = z$ (it is *0* if x was not available at call c). The relevant forbidden sets are built from ranges of size at most nq, so they have density at most $\phi_{\mathcal{S}}$ in $\mathcal{S}$. Let $\mathcal{J}$ be the set of calls in the product of Lemma 15(b): it is $\mathcal{I}_\tau$, together with the current call in the reweighted case, so $|\mathcal{J}| \le q$. Since μ_τ is uniform on $\mathcal{S}$ in the coordinate x, the conditional law of $z = \Gamma(x)$ is proportional to

$$\prod_{c \in \mathcal{J}} \vartheta_c\big(I_c^{-x} + I_c(x; z)\big).$$

Write each factor as a telescoping product from I_c^{-x} up to $I_c^{-x} + I_c(x; z)$:

$$\vartheta_c\big(I_c^{-x} + I_c(x; z)\big) = \vartheta_c(I_c^{-x}) \prod_{r=1}^{I_c(x;z)} \frac{\vartheta_c(I_c^{-x} + r)}{\vartheta_c(I_c^{-x} + r - 1)} = \vartheta_c(I_c^{-x}) \exp\Big(\sum_{r=1}^{n} \delta_{c,r} \mathbf{1}[I_c(x; z) \ge r]\Big),$$

with $\delta_{c,r} := \log\big(\vartheta_c(I_c^{-x} + r)/\vartheta_c(I_c^{-x} + r - 1)\big)$. The first factor does not depend on z. So $P(z) \propto e^{\beta(z)}$ on $\mathcal{S}$, where

$$\beta(z) = \sum_{c \in \mathcal{J}} \sum_{r=1}^{n} \delta_{c,r} \mathbf{1}[I_c(x; z) \ge r].$$

By Lemma 15(a), $0 \le \delta_{c,r} \le 3/L_c \le 4/N$, since $L_c \ge 63N/64$.

Step 2: bounding β and $\mathbb{E}[\beta^2]$. The increments $\delta_{c,r}$ depend on r and on the other entries, so β is *not* linear in $I_c(x;z)$. We use only the two-sided bound

$$0 \le \beta(z) \le \frac{4}{N} S(z), \qquad S(z) := \sum_{c \in \mathcal{J}} I_c(x;z) \le n|\mathcal{J}| \le nq,$$

which gives $0 \le \beta \le A$. For the second moment we use $\mathbb{E}[S^2] \le \max S \cdot \mathbb{E}[S]$. Under $z \leftarrow \mathrm{Unif}(\mathcal{S})$, each call contributes $\mathbb{E}[I_c(x;z)] = \sum_{j'' \in c^{-1}(b_c)} \Pr[z \in \mathrm{Forb}_{t_c,j''}] \le n\phi_{\mathcal{S}}$, so $\mathbb{E}[S] \le qn\phi_{\mathcal{S}}$ and

$$\mathbb{E}_{z \leftarrow \mathrm{Unif}(\mathcal{S})}[S^2] \le nq \cdot qn\phi_{\mathcal{S}} = n^2 q^2 \phi_{\mathcal{S}}.$$

With $\phi_{\mathcal{S}} \le \frac{32}{31}\phi_0$ and $\phi_0 \le (wnq + w^2)/N$ (Section 4.1, with $s_{\max} = nq$),

$$\mathbb{E}[\beta^2] \le \frac{16}{N^2}\mathbb{E}[S^2] \le \frac{16 n^2 q^2}{N^2} \cdot \frac{32}{31} \cdot \frac{wnq + w^2}{N} = O\Big(\frac{n^3 w q^3}{N^3}\Big),$$

using $w^2 \le wnq$, which follows from $q \ge 2w$. The three factors of n arise as follows: one from $S \le nq$, one from $\mathbb{E}[I_c(x;z)] \le n\phi_{\mathcal{S}}$, and one from the range sizes inside ϕ_0.

Step 3: bounding $\chi^2(\mathrm{Unif}(V) \| P(\cdot \mid V))$. The law $P(\cdot \mid V)$ is $\mathrm{Unif}(V)$ reweighted by e^{β}. So Lemma 1, with $U = \mathrm{Unif}(V)$, $g = \beta|_V$ and $0 \le g \le A$, gives

$$\chi^2(\mathrm{Unif}(V) \| P(\cdot|V)) \le e^{3A}\mathbb{E}_{\mathrm{Unif}(V)}[\beta^2] \le e^{3/16}(1-\phi_{\mathcal{S}})^{-1}\mathbb{E}_{\mathrm{Unif}(\mathcal{S})}[\beta^2] = O(n^3 w q^3/N^3).$$

The middle step uses $\beta^2 \ge 0$ and $|V| \ge (1-\phi_{\mathcal{S}})|\mathcal{S}|$ to pass from $\mathrm{Unif}(V)$ to $\mathrm{Unif}(\mathcal{S})$. The constants 34 and 44 are evaluated in Section B.1 (12).

Finally, when the other entries are averaged over their conditional law, the law of $\Gamma(x)$ given $\Gamma(x) \in V$ is a mixture of the conditioned laws above. By the convexity statement of Lemma 1, the chi-square bound survives this averaging. The density bounds $e^{\pm A}$ are the ratio bound of Lemma 1, and a mixture of laws that satisfy them also satisfies them. □

Inverse responses and failure. With the bias bound in hand, we compare the inverse responses of the simulator with those of the hybrid, and bound the probability of failure.

Lemma 17. *Fix a prefix without failure and an inverse query at t on class b, and assume (3). Let P be the law of the inverse output in* Ideal′ *conditioned on success, and P^{comp} the conditional law of the companion tuple. Then*

$$\chi^2(K\|P) \le \frac{56\, n^3 w q^3}{N^3}, \qquad \mathrm{KL}(K^{\mathrm{comp}}\|P^{\mathrm{comp}}) \le \frac{44\, n^3 w q^3}{N^3},$$

so by the chain rule the augmented inverse response has divergence at most $100\, n^3 w q^3/N^3$.

Proof. By Lemma 15(c), the ideal law gives a construction-only pair (x_i, j), a fresh family input $\iota_j(x)$ and a free input v the probabilities

$$P(\iota_j(x_i)) = \frac{\mathbf{1}[(i,j)\ \text{legal}]}{Z}, \qquad P(\iota_j(x)) = \frac{\theta_j}{Z}, \qquad P(v) = \frac{\rho_b}{Z},$$

with $Z := k_{\mathrm{leg}} + u\sum_{j \in c^{-1}(b)} \theta_j + g\rho_b$, while the hybrid's normaliser is $N_{\mathrm{Hyb}} = k_{\mathrm{leg}} + n_b u \rho_b + g\rho_b$.

Comparing θ_j with ρ_b. By Lemma 16, applied to $B = \mathcal{S} \cap \mathrm{Forb}_{t,j}$, whose density in $\mathcal{S}$ is $1 - \rho_b \le \phi_{\mathcal{S}} \le \frac{32}{31}\phi_0$,

$$|\delta_j| \le (e^A - 1)(1 - \rho_b) \le 2A\phi_{\mathcal{S}} \le \frac{9\,nq\,\phi_0}{N} =: \delta, \qquad \delta_j := \theta_j - \rho_b.$$

Bounding $\chi^2(K\|P)$. Both laws are normalised weights, $K(v) = w_K(v)/N_{\text{Hyb}}$ and $P(v) = w_P(v)/Z$, and the two normalisers differ by $Z - N_{\text{Hyb}} = u\sum_j \delta_j$. We compute $\chi^2(K\|P) = \sum_v (K(v)-P(v))^2/P(v)$ by kind of input:

- *Legal construction-only inputs and free inputs.* Here the two weights agree, $w_K(v) = w_P(v) =: w$, so $K(v) - P(v) = w(Z - N_{\text{Hyb}})/N_{\text{Hyb}}Z$, and the contribution of v is $w(Z - N_{\text{Hyb}})^2/N_{\text{Hyb}}^2 Z$. Summing the weights, these inputs contribute $(k_{\text{leg}} + g\rho_b)(Z - N_{\text{Hyb}})^2/N_{\text{Hyb}}^2 Z$.
- *Fresh family inputs* $\iota_j(x)$. Here $K = \rho_b/N_{\text{Hyb}}$ and $P = \theta_j/Z$ with $\theta_j = \rho_b + \delta_j$, so

$$K - P = \frac{\rho_b Z - \theta_j N_{\text{Hyb}}}{N_{\text{Hyb}}Z} = \frac{\rho_b(Z - N_{\text{Hyb}}) - \delta_j N_{\text{Hyb}}}{N_{\text{Hyb}}Z} = \frac{\rho_b u\sum_{j'}\delta_{j'} - \delta_j N_{\text{Hyb}}}{N_{\text{Hyb}}Z}.$$

  Dividing its square by $P = \theta_j/Z$ and multiplying by the u unrevealed inputs for each j gives the second sum below.

Altogether,

$$\chi^2(K\|P) = (k_{\text{leg}} + g\rho_b)\frac{(Z - N_{\text{Hyb}})^2}{N_{\text{Hyb}}^2 Z} + u\sum_{j\in c^{-1}(b)} \frac{\left(\rho_b u\sum_{j'}\delta_{j'} - \delta_j N_{\text{Hyb}}\right)^2}{N_{\text{Hyb}}^2 Z\,\theta_j}.$$

Under (3), δ is at most a small constant. Hence N_{Hyb}, Z and θ_j are bounded below by constant multiples of N, N and 1 respectively; in particular $N_{\text{Hyb}} \ge N - \Phi$ (Section B.1 (13)). Also $|Z - N_{\text{Hyb}}| \le n_b u\delta$, and $|\rho_b u\sum\delta_{j'} - \delta_j N_{\text{Hyb}}| \le n_b u\delta + N\delta \le 2N\delta$. In the first sum, $k_{\text{leg}} \le n_b h$, and the three counts $n_b h$, g and $n_b u$ add up to at most N, so the first sum is at most

$$\frac{n_b h + g}{N}\left(\frac{n_b u}{N}\right)^2 \frac{N^3\delta^2}{N_{\text{Hyb}}^2 Z} = O(\delta^2).$$

In the second sum there are $n_b u \le N$ terms, each with numerator at most $(2N\delta)^2$ and denominator at least a constant times N^3, so it is $O(\delta^2)$ as well. Hence $\chi^2(K\|P) = O(\delta^2)$. Since $\phi_0 \le 2wnq/N$ for $q \ge 2w$, we have $\delta \le 18wn^2q^2/N^2$, and therefore

$$\chi^2(K\|P) = O\left(\frac{w^2n^4q^4}{N^4}\right) = O\left(\frac{n^3wq^3}{N^3}\right),$$

using $wnq/N \le 1/64$. The constant 56 of the statement is evaluated in Section B.1 (13).

Companion tuple. If the output is a construction-only input, the tuple is the same fixed value in both worlds, and a free output has no companion. If it is $\iota_j(x)$ with x unrevealed, then the table law given this response is π_τ^ϑ conditioned on $\Gamma(x) \in V := \mathcal{S} \setminus \mathrm{Forb}_{t,j}$ in Ideal′ (Lemma 15(c)), and $\mathrm{Unif}(V)$ in Hyb. The tuple is a bijective function of $\Gamma(x)$ (proof of Lemma 4). Lemma 16 with this V, averaged over the other entries, gives $\mathrm{KL} \le \chi^2 \le 44\,n^3wq^3/N^3$. □

Proposition 2. *At every attainable prefix without failure, and under (3), consider the next response of* $\textsc{Ideal}'_3$ *conditioned on success. Its divergence from* $\textsc{Hyb}$ *is at most* $100\,n^3wq^3/N^3$*, for every query type. An inverse call fails with probability*

$$p_{\text{fail}} \le \left(\frac{nh + 2n + e^{1/16}\phi_{\mathcal{S}}N}{L}\right)^3 \le \left(\frac{4nwq}{N}\right)^3.$$

Consequently $\Pr_{\textsc{Ideal}'_3}[\textsc{Failure}] \le 64\,n^3w^3q^4/N^3$.

Proof. The divergence of each query type is bounded as follows:

- types (1) and (3): Lemma 16 for the construction value, since given the value both worlds sample v_0 uniformly on $\mathrm{Leg}(z)$;
- types (2) and (5): identical laws (for type (5), both worlds answer $\mathrm{Unif}(G \setminus \textsc{Ran}_b)$);
- type (4): Lemma 17.

$\textsc{Bounding}$ p_{fail}. The three trials sample candidates u_1, u_2, u_3 independently and uniformly from the L available inputs, independently of Γ, and the coins are fresh. Let $\mathcal{H}_{k-1}$ be the history of the call before trial k: the candidates so far, the oracle values they revealed, and the fact that each was rejected. Then

$$p_{\text{fail}} \le \prod_{k=1}^{3} \sup_{\mathcal{H}_{k-1}} \Pr[\text{trial } k \text{ rejects} \mid \mathcal{H}_{k-1}].$$

Conditional on τ and $\mathcal{H}_{k-1}$, the table has law proportional to $\pi_\tau(\Xi)\prod_{i<k}\mathbf{1}[\Xi(x_i) = \Gamma(x_i)]$ over the family candidates x_i read so far. A rejected free candidate constrains no entry. So for an unrevealed x not yet read, $\Gamma(x)$ has the law of Lemma 16 conditional on all other entries, and $\Pr[\iota_j(x) \text{ illegal} \mid \mathcal{H}_{k-1}] \le e^A\phi_{\mathcal{S}}$. Since u_k is uniform and independent of $\mathcal{H}_{k-1}$, the conditional rejection probability of trial k is at most

$$\frac{1}{L}\Big(nh + (k-1)n + \sum_{\text{other inputs}} \Pr[\text{reject}]\Big) \le \frac{nh + (k-1)n + e^A\phi_{\mathcal{S}}N}{L},$$

where the three terms arise as follows:

- nh bounds the inputs of construction-only x, legal or not;
- $(k-1)n$ counts *all* inputs of each family candidate already read, since one of them is known to be illegal and the others may be;
- each other unrevealed family input rejects with probability at most $e^A\phi_{\mathcal{S}}$, and each free input with probability $1 - \rho_b \le \phi_{\mathcal{S}}$, and these inputs number at most N.

With $k - 1 \le 2$ this gives the first inequality. For the second, we use:

- $nh \le nq \le wnq$;
- $2n \le wnq$, since $wq \ge 2$;
- $e^A\phi_{\mathcal{S}}N \le 1.65\,wnq$, from $A \le 1/16$, $\phi_{\mathcal{S}} \le \frac{32}{31}\phi_0$ and $w^2 \le wnq/2$;
- $L \ge 63N/64$, so that the per-trial bound is at most $4wnq/N$ (Section B.1 (14)).

Finally, take a union bound over the at most q inverse calls. □

Proof of the main theorem.

Proof of Theorem 1. By (12), $\mathrm{TV}(\textsc{Ideal}_3, \textsc{Ideal}'_3) \le 2w(w+1)q/N$, and this is 0 without rigid pairs. For the real world and the hybrid, by the chain rule and Proposition 1, with $q \ge 2w$ in the form $q^3 \le q^4/2w$ and $q^2 \le q^4/4w^2$,

$$\mathrm{KL}(T_{\textsc{Real}} \| T_{\textsc{Hyb}}) \le \frac{K_{\mathrm{RH}}q^4}{N^3} + \frac{12w^4q}{N^2}, \qquad K_{\mathrm{RH}} := C_{\mathrm{RH}} + \frac{C'_{\mathrm{RH}}}{2w} + \frac{C''_{\mathrm{RH}}}{4w^2} \le 85\,n^3(w+1)^4$$

(Appendix B); the last term is absent when every same-class pair is rigid. For the hybrid and the ideal world, we use Lemma 2, Proposition 2, and $\log \frac{1}{1-p} \le 2p$ for $p \le \frac{1}{2}$, applied to $p_{\text{fail}} \le (4nwq/N)^3 \le 16^{-3}$. Each of the at most q responses contributes at most $100\,n^3wq^3/N^3 + 2(4nwq/N)^3$, so

$$\mathrm{KL}(T_{\textsc{Hyb}} \| T_{\textsc{Ideal}'_3}) \le \frac{K_{\mathrm{HI}} q^4}{N^3}, \qquad K_{\mathrm{HI}} := 100n^3w + 128n^3w^3 \le 228\,n^3w^3.$$

Finally, Pinsker's inequality and the triangle inequality give

$$\mathbf{Adv} \le \left(\sqrt{\tfrac{K_{\mathrm{RH}}}{2}} + \sqrt{\tfrac{K_{\mathrm{HI}}}{2}}\right)\frac{q^2}{N^{3/2}} + \sqrt{6}\,w^2\frac{\sqrt{q}}{N} + 2w(w+1)\frac{q}{N}.$$

Here $\sqrt{6} \le 3$, and the first coefficient is at most $10\,n^{3/2}(w+1)^2$ (Appendix B). □

5.6 Bounds for the instances

Direct instances. All instances below are admissible, so Theorem 1 applies directly. We state them with their constants. Write $\textsc{Sim}_3$ for the coin simulator with three trials; for balanced instances it is the plain uniform simulator.

Corollary 1. *Let $w \ge 1$, and let π be a uniform permutation of G. Fix injections $\iota_0, \dots, \iota_w : \mathcal{D} \to G$ with pairwise disjoint images, and put*

$$\mathrm{XORP}[w](x) := \big(\pi(\iota_1 x) - \pi(\iota_0 x),\ \dots,\ \pi(\iota_w x) - \pi(\iota_0 x)\big).$$

If $N \ge 2^{10}$, $w^2 \le N/32$ and $2w \le q \le N/64w(w+1)$, then

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{XORP}[w],\textsc{Sim}_3}(q) \le 10\,(w+1)^{7/2}\frac{q^2}{N^{3/2}} + 2w(w+1)\frac{q}{N}.$$

Proof. Theorem 1 with $k = 1$, $n = w+1$ and identity masks. All pairs are rigid, so the coupling term $2w(w+1)q/N$ is present and there is no $\sqrt{q}$ term. □

Over $G = \mathbb{F}_2^d$ with the layout $\iota_j(x) = x \| \langle j \rangle_s$, $s = \lceil \log_2(w+1) \rceil < d$, the tags $j \ge w+1$ are free inputs. When $w+1$ is a power of two the family is balanced, and this is the setting of [BDD+26, Sec. 4], where the construction is called XORP_{w+1}. For other w the corollary covers this layout with the unused tags free. [BDD+26, eq. (8)] defines XORP at every width, but its indifferentiability proof assumes that the number of calls is a power of two. The two-trial simulator of [BDD+26] is not covered.

Corollary 2. *Let $G = \mathbb{F}_2^d$, $w \ge 1$ with $s := \lceil \log_2(w+1) \rceil < d$, and let $\sigma \in \mathrm{GL}(d,2)$ be w-admissible ($I + \sigma^r$ invertible for $1 \le r \le w$). If $N \ge 2^{10}$, $w^2 \le N/32$ and $2w \le q \le N/64w(w+1)$, then LXoP$[\sigma, w]$ of [Din25, BGGJ26] satisfies*

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{LXoP},\textsc{Sim}_3}(q) \le 10\,(w+1)^{7/2}\frac{q^2}{N^{3/2}} + 3w^2\frac{\sqrt{q}}{N}.$$

Proof. By Section 3.4, LXoP is the family with $k = 1$, $n = w+1$, $\lambda_j = (-\sigma)^j$ and the tag layout $\iota_j(x) = \langle j \rangle_s \| x$, up to an invertible linear map M on the output: LXoP $= M \circ \mathrm{F}$. The simulator for LXoP runs $\textsc{Sim}_3$ with the oracle $M^{-1} \circ \Gamma$, as with the map T in Lemma 18, so the advantage is unchanged. The unused tags are free inputs. By w-admissibility every same-class pair is non-rigid, so the coupling term is absent and the $\sqrt{q}$ term is present. □

If σ is not w-admissible, the family can still be admissible: this happens when every $I + \sigma^r$ with $r \le w$ is either invertible or 0. A rigid pair exists exactly when $\sigma^r = I$ for some $1 \le r \le w$, and then the coupling term $2w(w+1)q/N$ appears. If some $I + \sigma^r$ is singular but nonzero, the family is not admissible and is not covered by this corollary; Theorem 3 still applies to it, under its hypotheses.

Corollary 3. *Over any finite abelian group of order $N \ge 2^{10}$ and in the regime (3), the following bounds hold:*

(i) $\pi_1 + \pi_2$*:* $\mathbf{Adv} \le 40\, q^2/N^{3/2}$*, which is [Kim26, Theorem 1] up to constants, now over every finite abelian group rather than $\mathbb{F}_2^n$;*

(ii) the domain-separated sum $\pi(0\|x) \oplus \pi(1\|x)$ over $\mathbb{F}_2^d$: $\mathbf{Adv} \le 114\, q^2/N^{3/2} + 4q/N$*;*

(iii) XORP with $w+1$ independent permutations and arbitrary automorphism masks: $\mathbf{Adv} \le 10\,(w+1)^2 q^2/N^{3/2}$*;*

(iv) every family whose masks are multiplications by nonzero scalars, over a vector space over a finite field: such families are admissible, with the bound of Theorem 1.

Proof. Substitute (k, n, w) into Theorem 1: $(2,1,1)$ gives $10 \cdot 4 = 40$; $(1,2,1)$ gives $10 \cdot 2^{3/2} \cdot 4 \le 114$, plus the coupling term $4q/N$; and $n = 1$ gives $10(w+1)^2$ with no lower-order terms. Field-scalar masks are admissible, because the difference of two scalars is 0 or invertible. □

Fork constructions. The fork constructions of Dutta et al. [DGL22] reduce to the family in three exact steps, none of which loses queries or advantage:

- a composition bijection with an interface wrapper (Lemma 18);
- the removal of a disposable pre-permutation (Lemma 19);
- the restriction of the construction domain (Lemma 20).

Lemma 18. *Let* $\mathsf{Fork} := \mathsf{Fork}[G; \beta; \boldsymbol{\gamma}_0; \boldsymbol{\gamma}_1]$ *be a fork (Definition 7), with $w+1$ independent uniform permutations $\pi_0, \dots, \pi_w$ of G, and write $v_0 := \pi_0(x)$ and $v_j := \pi_j(\beta_j v_0)$. Let* F *be the family instance with $k = w+1$, $n = 1$, $\mathcal{D} = G$, $\iota_j = \mathrm{id}$, class map $c(j) = j$ and masks*

$$\lambda_0 := \gamma_{01}, \qquad \lambda_j := \gamma_{01}\gamma_{0j}^{-1}\gamma_{1j} \quad (1 \le j \le w),$$

and let $T \in \mathrm{Aut}(G^w)$ be the blockwise map $T(y)_j := \gamma_{0j}\gamma_{01}^{-1} y_j$. Then:

- $\mathsf{Fork}^{\pi} = T \circ \mathrm{F}^{\varphi(\pi)}$ *for a bijection φ of $\mathrm{Perm}(G)^{w+1}$;*
- *for every simulator S for* F *in the augmented interface there is a simulator S^{Fork} for* Fork*, in the ordinary interface, such that every q-query distinguisher against $(\mathsf{Fork}, S^{\mathsf{Fork}})$ is matched by a q-query augmented distinguisher against (F, S) with the same advantage;*
- *if $S = \textsc{Sim}_m$, then S^{Fork} is the natural uniform simulator of* Fork *with m trials.*

Proof. We first change variables. Define $\varphi(\pi_0, \dots, \pi_w) := (\pi_0, \pi'_1, \dots, \pi'_w)$ with $\pi'_j := \pi_j \circ \beta_j \circ \pi_0$. Since $\beta_j \circ \pi_0$ is a permutation, φ is a bijection, with inverse $\pi_j = \pi'_j \circ \pi_0^{-1} \circ \beta_j^{-1}$. So $\varphi(\pi)$ is uniform when π is. By construction $v_0 = \pi_0(x) = p_0(x)$ and $v_j = \pi'_j(x) = p_j(x)$, so

$$T\big(\mathrm{F}^{\varphi(\pi)}(x)\big)_j = \gamma_{0j}\gamma_{01}^{-1}\big(\gamma_{01}\gamma_{0j}^{-1}\gamma_{1j}v_j - \gamma_{01}v_0\big) = \gamma_{1j}v_j - \gamma_{0j}v_0 = \mathsf{Fork}^{\pi}(x)_j.$$

The order of composition matters: $\lambda_j = \gamma_{1j}$ would be wrong unless the γ's commute.

Next we translate the queries. A family primitive query at x, on any class, reveals the tuple $(x, v_0, \dots, v_w)$, from which $\pi_j(\beta_j v_0) = v_j$ and $\pi_j^{-1}(v_j) = \beta_j v_0$ can be read off. Each fork query is answered by one family query, as listed in Table 3. Since $n = 1$ and $\mathcal{D} = G$, every element of G is a valid input of every class.

fork query	family query (augmented)	answer
$\pi_0(u)$	$\pi'_0(u)$	v_0 from the tuple of u
$\pi_0^{-1}(v)$	$\pi_0'^{-1}(v)$	the returned x
$\pi_j(u)$, $j \geq 1$	$\pi_0'^{-1}(\beta_j^{-1}u)$	v_j from the tuple of the returned x
$\pi_j^{-1}(v)$, $j \geq 1$	$\pi_j'^{-1}(v)$	$\beta_j v_0$ from the tuple of the returned x

Table 3: Translation of fork queries into augmented family queries in the proof of Lemma 18.

In the real world the translated answers are exactly the fork's answers, because $\pi_j = \pi'_j \circ \pi_0^{-1} \circ \beta_j^{-1}$ and $\pi_j^{-1} = \beta_j \circ \pi_0 \circ \pi_j'^{-1}$.

The simulator S^{Fork}, given oracle access to a random function $\Gamma' : G \to G^w$, runs S with the oracle $T^{-1} \circ \Gamma'$ and applies the translation, passing on any $\perp$. The translation costs one query per fork query only because the family interface is augmented. In the ordinary interface the third row would cost two queries, $\pi_0'^{-1}$ and then π'_j.

Given a fork distinguisher, the family distinguisher runs it, translating primitive queries by the table and answering a construction query x by $T(\mathbf{F}(x))$. In the real world it presents $(\mathsf{Fork}^\pi, \pi^\pm)$ for uniform π. In the ideal world it presents $(T \circ \Gamma,$ translated $S^\Gamma)$. Here $T \circ \Gamma$ is a uniform random function, since T is a bijection of G^w, and this is the law of $(\Gamma', S^{\mathsf{Fork},\Gamma'})$. So the advantages coincide.

Finally, for $S = \textsc{Sim}_m$, the simulator S^{Fork} answers queries as follows:

- $\pi_0(u)$: read $\Gamma'(u)$ and sample v_0 uniformly among the values consistent with the ranges;
- $\pi_j(u)$: sample a uniform unassigned x up to m times, accepting iff $v_0 = \beta_j^{-1}u$ is consistent with $\Gamma'(x)$ and the ranges;
- inverse queries likewise.

This is the uniform simulator one would write down for Fork directly. □

Lemma 19. *Let C' be a construction that uses only independent uniform permutations $\pi'_1, \ldots, \pi'_m$ of G. Let C be the construction with $m+1$ independent uniform permutations, $C^{\pi_0,\ldots,\pi_m} := C'^{\pi_1 \circ \pi_0, \ldots, \pi_m \circ \pi_0}$, with primitive interfaces $\pi_0^\pm, \ldots, \pi_m^\pm$. For every simulator S for C' there is a simulator S^C for C such that every q-query distinguisher against (C, S^C) is matched by a q-query distinguisher against (C', S) with the same advantage, in the ordinary interfaces.*

Proof. The map $(\pi_0, \ldots, \pi_m) \mapsto (\pi_0, \pi_1 \circ \pi_0, \ldots, \pi_m \circ \pi_0)$ is a bijection of $\mathrm{Perm}(G)^{m+1}$. So π_0 and $(\pi'_i := \pi_i \circ \pi_0)_i$ are independent and uniform. S^C samples its own uniform π_0 lazily, runs S, and translates:

- $\pi_0^\pm$ by its own table;
- $\pi_i(u)$ by $\pi'_i(\pi_0^{-1}(u))$: one query to S after a look-up in its own table, extended lazily if u is new;
- $\pi_i^{-1}(v)$ by $\pi_0(x)$, where $x = \pi_i'^{-1}(v)$ is one query to S.

If S returns $\perp$, so does S^C.

Given a distinguisher against C, the distinguisher against C' samples its own π_0 and applies the same translation to its oracles. Construction queries pass through unchanged, since $C(x) = C'(x)$. In the real world this presents $(C^\pi, \pi^\pm)$ for $\pi_i := \pi'_i \circ \pi_0^{-1}$, which are independent and uniform together with π_0. In the ideal world it presents $(\Gamma, S^{C,\Gamma})$. □

Lemma 20. *Let $C : \mathcal{D} \to G^w$ be a construction and $\mathcal{X} \subseteq \mathcal{D}$, and let $C|_\mathcal{X}$ be the same construction queried only on $\mathcal{X}$. For every simulator S for C there is a simulator $S_\mathcal{X}$ for $C|_\mathcal{X}$ such that every q-query distinguisher against $(C|_\mathcal{X}, S_\mathcal{X})$ is matched by a q-query distinguisher against (C, S) with the same advantage.*

Proof. $S_{\mathcal{X}}$, with oracle access to $\Gamma' : \mathcal{X} \to G^w$, runs S with the oracle Γ that equals Γ' on $\mathcal{X}$ and is an independent uniform value, sampled lazily, elsewhere. Since $\Gamma|_{\mathcal{X}}$ and $\Gamma|_{\mathcal{D}\setminus\mathcal{X}}$ are independent and uniform for a uniform $\Gamma : \mathcal{D} \to G^w$, the ideal worlds agree. The real worlds agree trivially, and so do the query counts. □

Together, the three lemmas give the security of every construction in the class GXP^+ of Definition 8, in terms of its base.

Definition 18. Let $C \in \mathsf{GXP}^+$ with base $\mathbf{F}$, and let $m \geq 1$. The *wrapped simulator* S_C of C is obtained from $\textsc{Sim}_m$ for $\mathbf{F}$ by applying the simulator maps $S \mapsto S^{\mathsf{Fork}}$, $S \mapsto S^C$ and $S \mapsto S_{\mathcal{X}}$ of Lemmas 18 to 20 along the rules that produce C from $\mathbf{F}$.

Proposition 3. *Let $C \in \mathsf{GXP}^+$ with base $\mathbf{F}$, let $m \geq 1$, and let S_C be the wrapped simulator of C (Definition 18). For every q*

$$\mathbf{Adv}^{\mathrm{indiff}}_{C,S_C}(q) \leq \mathbf{Adv}^{\mathrm{indiff}}_{\mathbf{F},\textsc{Sim}_m}(q) \qquad \textit{and} \qquad \mathbf{Adv}^{\mathrm{prf}}_{C}(q) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathbf{F}}(q),$$

where on the right the indifferentiability advantage is taken in the augmented interface of Definition 1. In particular, for $m = 3$, Theorem 1 (if $\mathbf{F}$ is admissible) and Theorem 3 (in general) bound the indifferentiability of C, with the parameters of $\mathbf{F}$.

Proof. We induct on the number of rules used to produce C; each step preserves the number of queries, and the rules are handled as follows:

- (R1): $C = \mathbf{F}$ and $S_C = \textsc{Sim}_m$. An ordinary distinguisher is an augmented one that ignores part of the answers (Definition 1), so the ordinary advantage is at most the augmented one. The pseudorandomness claim is an equality.
- (R2): Lemma 18 matches every q-query distinguisher against the fork and its wrapped simulator with a q-query augmented distinguisher against $(\mathbf{F}, \textsc{Sim}_m)$ of the same advantage. For pseudorandomness, $\mathsf{Fork}^{\pi} = T \circ \mathbf{F}^{\varphi(\pi)}$ with $\varphi(\pi)$ uniform and T a fixed bijection, so a distinguisher against Fork is one against $\mathbf{F}$ that applies T to the answers.
- (R3): Lemma 19 matches every q-query distinguisher against C and S^C with a q-query distinguisher against C' and $S_{C'}$ of the same advantage, in the ordinary interfaces; the induction hypothesis bounds the latter. For pseudorandomness, C and C' have the same construction law, since $\pi_1 \circ \pi_0, \ldots, \pi_m \circ \pi_0$ are independent and uniform.
- (R4): Lemma 20 does the same for the restriction, and a pseudorandomness distinguisher against C restricted to $\mathcal{X}$ is one against C. □

Corollary 4. *Let $G = \mathbb{F}_2^d$ with $N = 2^d \geq 2^{10}$, let $w \geq 1$ with $w^2 \leq N/32$, and let $2w \leq q \leq N/64w$. Wrap $\textsc{Sim}_3$ through Lemmas 18 to 20 as indicated below. Each of the following constructions is in GXP^+ and satisfies*

$$\mathbf{Adv}^{\mathrm{indiff}}(q) \leq 10\,(w+1)^2\,\frac{q^2}{N^{3/2}},$$

with no lower-order term:

(i) ForkEDMD$[w]$ *[DGL22, Sec. 3.5]: $x \in \{0,1\}^{d-s}$, $X = \pi_0(x\|0^s)$, with output $(\pi_i(X) \oplus X)_{i=1}^w$ from $w+1$ independent public permutations, for any $0 \leq s \leq d$. Apply Lemma 18 with $\beta_j = \gamma_{0j} = \gamma_{1j} = \mathrm{id}$, then Lemma 20 with $\mathcal{X} = \{x\|0^s\}$.*

(ii) EDMD *[MN17], $\pi_1(\pi_0(x)) \oplus \pi_0(x)$ (for EDMD itself the reduction to $\pi_1 + \pi_2$ is [BDD+26, proof of Theorem 3, App. A.2]), and more generally $\gamma_1\pi_1(\beta\pi_0(x)) - \gamma_0\pi_0(x)$ with $\beta, \gamma_0, \gamma_1 \in \mathrm{Aut}(G)$, over any finite abelian group. Apply Lemma 18 with $w = 1$; the coefficient is 40.*

(iii) ForkCENC$[w]$ *[DGL22, Sec. 3.4]:* $x \in \{0,1\}^d$, $X = \pi_0(x)$, $Y_i = \pi_i(X)$, *with output* $(Y_1 \oplus Y_{i+1})_{i=1}^{w}$ *from* $w+2$ *independent public permutations. Lemma 19, with* $\pi'_i := \pi_i \circ \pi_0$, *turns it into XORP with the* $w+1$ *independent permutations* $\pi'_1, \ldots, \pi'_{w+1}$ *and identity masks.*

Proof. Each construction is in GXP^+, so Proposition 3 applies with its base. In each case the base has $k = w+1$, $n = 1$, $\mathcal{D} = G$, and a single call per class. So Theorem 1 applies, there are no same-class pairs, and $n^{3/2}(w+1)^2 = (w+1)^2$. The reductions preserve the query count and the advantage. Lemma 18 produces an augmented family distinguisher, which is what Theorem 1 bounds; an ordinary family distinguisher, as in ForkCENC through Lemma 19, is dominated by an augmented one. [DGL22] requires $w \geq 2$ and $s \geq 1$ for ForkEDMD. The corollary also covers $w = 1$ (EDMD), $s = 0$ (the full domain) and $s = d$ (a one-element domain). □

ForkSTHCENC and ForkEDM-CTR of [DGL22] are not covered. The first truncates and leaves one block unsummed. The second feeds $x\|0^s$, multiplied by field elements, into the bottom permutations, which is an EDM rather than an EDMD shape.

6 Lower bounds for the coin simulator

The upper bound of Theorem 1 rests on two features of the simulator: the acceptance coin, and the three trials of an inverse query. This section shows that neither the coin nor the exponent $3/4$ can be improved. Section 6.1 shows that a simulator without the coin is distinguishable at about $N^{2/3}$ queries on families with free inputs. Section 6.2 gives a matching attack on the coin simulator itself, with $O(N^{3/4})$ primitive queries, for every family over every finite abelian group with a construction domain of density bounded below, and for the fork constructions.

6.1 The coin is necessary

The *naive* simulator $\textsc{Sim}_m^{\text{naive}}$ is $\textsc{Sim}_m$ with every free candidate accepted. It answers the free inputs as an independent lazy permutation, which is the natural first choice. The real world gives the preimage of a fresh target the same probability at every unassigned input, family or free. The naive simulator instead favours free inputs, because family candidates are sometimes rejected. The bias per query is about $w|R|/N \approx wnq/N$. Over $q/2$ inverse queries it gives a signal of order q^2/N against a noise of order $\sqrt{q}$. So the bias is detectable at $q \approx N^{2/3}$, well inside the range where Theorem 1 gives a bound of order $q^2/N^{3/2}$ for the coin simulator.

Proposition 4. *Let* $k = 1$, *so that* $n = w+1$ *and all calls belong to one class. Let the family be admissible, with free density* $\gamma := |\text{Free}|/N$. *Let* $N \geq 2^{10}$ *and* $w^2 \leq N/32$. *Let* q *be even, with*

$$4w \leq q \leq N/64nw, \qquad \gamma N \geq q, \qquad (1-\gamma)N \geq 2nq.$$

There is a distinguisher making $q/2$ *forward and* $q/2$ *inverse* augmented *primitive queries (Definition 1), and no construction query, with*

$$\mathbf{Adv}^{\text{indiff}}_{\mathrm{F},\textsc{Sim}_3^{\text{naive}}} \geq 1 - 2\exp\left(-2^{-13}\gamma^2(1-\gamma)^2w^2n^2\frac{q^3}{N^2}\right) - 2w(w+1)\frac{q}{N},$$

where the last term is present only if some pair is rigid. For fixed $\gamma \in (0,1)$ *and* w, *the advantage tends to 1 once* $q \gg N^{2/3}$.

Proof. The distinguisher works as follows:

1. Make forward queries at $\iota_0(x)$ for $q/2$ distinct inputs x.
2. Make $I := q/2$ inverse queries, at targets drawn uniformly from $G \setminus \textsc{Ran}$.
3. Before inverse query k, compute from the transcript the fraction p_k of fresh inputs that are free. The transcript determines the assigned inputs exactly, because every tuple and every free answer is revealed.
4. Let $X_k := \mathbf{1}[\text{the answer is free}]$ and $T := \sum_k (X_k - p_k)$. Output "ideal" if some answer is $\perp$ or $T \geq I\beta/2$, where
$$\beta := 0.0586\,\gamma(1-\gamma)\,wnq/N.$$

Real world. There are no construction-only inputs, so given the transcript the preimage of a fresh target is uniform over the unassigned inputs. Thus $\mathbb{E}[X_k - p_k \mid \text{past}] = 0$. By Azuma's inequality, with increments in $[-1,1]$, $\Pr[T \geq I\beta/2] \leq \exp(-I\beta^2/8)$.

Ideal world. By (12) it suffices to work in $\textsc{Ideal}'$, at a cost of $2w(w+1)q/N$, which is 0 without rigid pairs. The proofs of Lemmas 15 and 16 apply verbatim with acceptance weight 1 for free inputs, that is with $R = I(\Xi)/L$. Write u for the number of unrevealed inputs x, g for the number of available free inputs, and $\eta := e^{-A}(1-\rho)$. We note the following facts:

- By the lower density bound of Lemma 16, $\theta_j = 1 - \Pr_{\pi_\tau^\vartheta}[\Gamma(x) \in \mathrm{Forb}_{t,j}] \leq 1 - \eta$ for every j.
- By Lemma 15(c), on success,
$$\Pr[X_k = 1 \mid \text{past}] = \frac{g}{g + u\sum_j \theta_j} \geq \frac{g}{g + nu(1-\eta)}.$$
- Since $p_k = g/(g+nu)$, the right-hand side is $p_k/(1-(1-p_k)\eta) \geq p_k + p_k(1-p_k)\eta$.
- By Lemma 4, with $|\mathcal{S}| \geq (N-1)_w$ and $|R| \leq nq$ (so $w|R| \leq N$),
$$\rho \leq \left(\frac{N - |R| - 1}{N-1}\right)^w \leq e^{-w|R|/N} \leq 1 - \frac{w|R|}{2N}.$$
As $|R| \geq nq/2$ after the forward phase and $e^{-A} \geq e^{-1/16}$, this gives $\eta \geq e^{-A} w|R|/2N \geq 0.2348\,wnq/N$ (Section B.2 (1)).
- At every inverse query, at least $\gamma N - q/2 \geq \gamma N/2$ free inputs and at least $(1-\gamma)N - nq \geq (1-\gamma)N/2$ family inputs are fresh. So $p_k \geq \gamma/2$ and $1 - p_k \geq (1-\gamma)/2$.

Hence, conditional on the past and on the success of the call, the drift is at least $p_k(1-p_k)\eta \geq \gamma(1-\gamma)\eta/4 \geq \beta$ (Section B.2 (2)).

Handling failures. Define X'_k as follows: $X'_k := X_k$ if call k and all earlier calls succeed, and otherwise X'_k is an independent Bernoulli variable with mean $p_k + \beta$. That mean is at most 1, because $\beta \leq (1-\gamma)/2 \leq 1 - p_k$. Put $T' := \sum_k (X'_k - p_k)$. Then $\mathbb{E}[X'_k - p_k \mid \text{past}] \geq \beta$ for every k, and $T' = T$ on the event of no failure. Hence $T' \geq I\beta + M$, where $M := \sum_k (X'_k - \mathbb{E}[X'_k \mid \text{past}])$ is a martingale with increments in $[-1,1]$. So
$$\Pr_{\textsc{Ideal}'}[\text{no } \perp,\ T < I\beta/2] \leq \Pr[T' < I\beta/2] \leq \exp(-I\beta^2/8).$$
Subtracting the two bounds, the advantage is at least $1 - 2\exp(-I\beta^2/8)$ minus the coupling cost, and $I\beta^2/8 \geq 2^{-13}\gamma^2(1-\gamma)^2w^2n^2q^3/N^2$ (Section B.2 (3)). □

Remark 5. Proposition 4 counts augmented queries: a forward query reveals the whole tuple, and an inverse answer reveals its tuple. Against the ordinary interfaces, simulate each augmented query by the ordinary query followed by at most $w = n - 1$ forward completion queries; the real and simulated tuples then agree with the augmented responses, and failures are seen at once. So the attack costs at most $Q := nq$ ordinary primitive queries and no construction query. In terms of Q, the exponent of Proposition 4 reads $2^{-13}\gamma^2(1-\gamma)^2\frac{w^2}{n}\frac{Q^3}{N^2}$, the coupling term reads $2w(w+1)Q/nN$, and the side conditions hold with $q = Q/n$. For fixed width the exponent $2/3$ is unchanged.

6.2 Matching attack on the coin simulator

We show that Theorem 1 cannot be improved in the exponent for the coin simulator. For every family over every finite abelian group, with arbitrary automorphism masks, fixed width and class sizes and a construction domain of density bounded below, a distinguisher with $O(N^{3/4})$ primitive queries and no construction queries tells $\textsc{Sim}_3$ from the real permutations with advantage $3/4$ (Theorem 2 and Corollary 5). The same holds for the fork constructions against their wrapped simulators (Corollary 6).

The distinguisher and the structure of the proof are Kim's [Kim26, Sec. 4 and App. E], for $\pi_1+\pi_2$ over $\mathbb{F}_2^d$. In outline: Phase 1 fills a structured set of outputs of one call; Phase 2 makes inverse queries in which a candidate of the selected call is rejected whenever its oracle value lies in a fixed test event E; and Phase 3 tests, on fresh inverse answers, whether the hidden values have become biased towards E. In the real world there is no bias, while the simulator's rejections in Phase 2 tilt the posterior of the oracle towards E. The new ingredients are:

- *coset progressions* in place of subgroups, which exist in every finite abelian group at every scale (Lemma 31); Phase 1 fills a progression twice as long as the target sets, so that one fixed test event is still forbidden at every Phase-2 call;
- two *separating shifts*, which keep the target and test sets apart when the two attacked calls use the same permutation;
- masks, several calls per permutation through the threshold form of the tilt, free inputs and the acceptance coin, and the random number of useful test iterations;
- converses of the reductions of Section 5.6, which carry the attack to the fork constructions.

Local notation. In Section 6.2, r denotes the size of the target progression P (not the number of revealed inputs), P and I denote progressions (not the partial permutations P_b of the simulator), T is the test statistic, and m is the number of trials of the simulator.

Setting and result. The attack fills a set of outputs that is nearly closed under addition, at the scale $N^{3/4}$. In $\mathbb{F}_2^d$, Kim uses a subgroup; in a general group we use a coset progression, which exists at every scale (Definition 26 and Lemma 31 in Section A.2). Fix a coset progression (H, g) of length L, and write

$$P := P([0, L)), \qquad I := P([0, 2L-1)), \qquad r := |P| = L|H|.$$

So P is the target progression, and I is the doubled progression that contains all sums of two elements of P; Phase 1 fills the preimage A_1 of I defined below. The cosets $H + ig$ for $i \in [0, 2L-1)$ need not be distinct, so $r \le |I| \le (2L-1)|H| < 2r$. We use three properties of these sets:

(P1) $u + P \subseteq I$ for every $u \in P$, since $(H + ig) + (H + i'g) = H + (i+i')g$ and $i + i' \le 2L - 2$;
(P2) $P + P \subseteq I$, which is (P1) again;
(P3) $|I - I| \le (4L-3)|H| < 4r$, since $I - I = P((-2L+1, 2L-1))$.

When $L = 1$, $P = I = H$ is a subgroup of order r, and the construction below reduces to Kim's construction with a subgroup.

Sets of the attack. Fix two distinct calls $j \neq j'$, with classes $b := c(j)$ and $b' := c(j')$; these may be equal. Every family with $w \geq 1$ has such a pair. For a construction value y put $D(y) := y_{j'} - y_j$, so that $D(\mathrm{F}(x)) = v_{j'}(x) - v_j(x)$. With the progression fixed above, let $c_1, c_2 \in G$. Put

$$A_1 := \lambda_{j'}^{-1}(I), \qquad A_2 := \lambda_j^{-1}(c_1 + P), \qquad A_3 := \lambda_j^{-1}(c_2 + P), \qquad C^{\cup} := \lambda_{j'}^{-1}(c_2 - c_1 + I).$$

A_2 and A_3 have size r; A_1 and $C^{\cup}$ have size $|I| < 2r$.

Definition 19. A pair of shifts $(c_1, c_2) \in G^2$ is *separating* if:
(S1) $c_2 - c_1 \notin I - I$;
(S2) $|A_1 \cap A_2|$, $|A_1 \cap A_3|$, $|C^{\cup} \cap A_2|$ and $|C^{\cup} \cap A_3|$ are all at most $r/16$.

By (S1), $A_1 \cap C^{\cup} = \varnothing$, and also $A_2 \cap A_3 = \varnothing$, since $P - P \subseteq I - I$. In each case the two sets are preimages under one mask of disjoint translates.

Lemma 21. *If c_1, c_2 are independent and uniform on G, they are separating with probability at least $1 - 132r/N$.*

Proof. By (P3), (S1) fails with probability $|I - I|/N < 4r/N$. For (S2), take $\mathbb{E}[|A_1 \cap A_2|] = \sum_{g' \in A_1} \Pr[\lambda_j(g') - c_1 \in P] = |I|r/N < 2r^2/N$. The other three overlaps are bounded similarly:

- $\mathbb{E}[|A_1 \cap A_3|] < 2r^2/N$;
- for the overlap of $C^{\cup}$ with A_2, conditioning on c_1,

$$\begin{aligned}\mathbb{E}[|C^{\cup} \cap A_2|] &= \sum_{g' \in G} \Pr[\lambda_j(g') - c_1 \in P] \cdot \Pr[\lambda_{j'}(g') - c_2 + c_1 \in I \mid c_1] \\ &= N \cdot \frac{r}{N} \cdot \frac{|I|}{N} < 2r^2/N;\end{aligned}$$

- $\mathbb{E}[|C^{\cup} \cap A_3|] < 2r^2/N$, conditioning on c_2 instead.

By Markov's inequality, each overlap exceeds $r/16$ with probability less than $32r/N$. A union bound gives $4r/N + 4 \cdot 32r/N$. □

Fix a separating pair. The test event is

$$E := \{z \in G^w : D(z) \in -c_1 + P\}.$$

Let $\mathrm{dens} := n_b|\mathcal{D}|/N \in (0, 1]$ be the fraction of inputs of π_b used by the construction.

Regime. We assume the following:
(R1) $N \geq 2^{16}$ and $1000\, n\, w^2 \leq N$;
(R2) $600\, n^3 w\, r \ln N \leq N$;
(R3) $30\, n_b\, r \leq \mathrm{dens} N$.

Theorem 2. *Let the masks be arbitrary automorphisms of G. Assume (R1)–(R3), a coset progression with $|P| = r$, a separating pair (c_1, c_2), and $m \geq 1$. The distinguisher $\mathcal{A}$ of Section 6.2 makes fewer than $4r(w+1)$ primitive queries and no construction queries, and satisfies*

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\mathrm{Sim}_m}(\mathcal{A}) \geq 1 - \frac{187\, n_b^2\, N^3}{\mathrm{dens}^2 r^4} - \frac{251\, n\, r}{N} - \frac{4m\, r\, |\Lambda|}{N} - e^{-r\mathrm{dens}^2/17 n_b^2} - N^{-2}.$$

Corollary 5. *Let G be any finite abelian group of order N. Put $\xi := \sqrt{n_b/\mathrm{dens}}\, N^{3/4}$ and $R := \lceil 6.3\xi \rceil$, and take a coset progression with $R \le r < 2R$ (Lemma 31). Suppose (R1)–(R3) hold for this r, and $251nr/N + 4mr|\Lambda|/N + e^{-r\mathrm{dens}^2/17n_b^2} + N^{-2} \le 1/8$. Then $\mathcal{A}$ makes at most*

$$4r(w+1) \le 51\,(w+1)\sqrt{n_b/\mathrm{dens}}\; N^{3/4}$$

queries and has advantage at least $3/4$. This holds for every family, admissible or not, and every pair of distinct calls.

Proof. Since $r \ge 6.3\xi$ and $\xi^4 = n_b^2 N^3/\mathrm{dens}^2$, the first error term of Theorem 2 is below $1/8$; the others add up to at most $1/8$ by hypothesis. Since $r < 2R$ and $\xi \ge N^{3/4} \ge 2^{12}$ by (R1), the query count is at most $51(w+1)\xi$ (Section B.2 (4)). A separating pair exists by Lemma 21, since $132r/N < 1$ by (R2). □

Scope. Since $\mathrm{dens} = n_b|\mathcal{D}|/N$, the attack scale is $\xi = \sqrt{N/|\mathcal{D}|}\, N^{3/4}$, and (R3) reads $|\mathcal{D}| \ge 30r$. So the corollary gives an attack with $O_w(N^{3/4})$ queries only when the construction domain has density bounded below; for a family with $|\mathcal{D}| = 1$, say, (R3) fails for every r. Fix n, w, m and $\mathrm{dens}_0 > 0$. For families with these parameters and $\mathrm{dens} \ge \mathrm{dens}_0$, the conditions of Corollary 5 hold for all large N. Together with Theorems 1 and 3, which are stated for $\textsc{Sim}_3$, this shows that for every such family over every finite abelian group, admissible or not, the coin simulator $\textsc{Sim}_3$ gives indifferentiability up to $\Theta(N^{3/4})$ queries and no further. The standard tagged instances of XORP and LXoP have dens bounded below for fixed width, and the full-domain instances behind the forks have dens $= 1$. For other trial counts see the remark on the number of trials in Section 6.2. With the explicit constants this requires very large N: roughly $N \ge 2^{75}$ even for $w = n = 1$.

The distinguisher. $\mathcal{A}$ uses the ordinary primitive interfaces and keeps the ranges $\textsc{Ran}_{b''}$ it has seen. *Completing* an answer u means: if $u = \iota_{j''}(x)$ is a family input, forward-query $\iota_{j'''}(x)$ for every $j''' \ne j''$; a free input needs nothing. The distinguisher runs in three phases:

Phase 1. For each $a \in A_1$ with $a \notin \textsc{Ran}_{b'}$: $u \leftarrow \pi_{b'}^{-1}(a)$, and complete u.

Phase 2. For each $t \in A_2$ with $t \notin \textsc{Ran}_b$: $u \leftarrow \pi_b^{-1}(t)$, and complete u.

Phase 3. For $k = 1, \dots, r$, with t_k the k-th element of A_3: skip if $t_k \in \textsc{Ran}_b$. Otherwise $u \leftarrow \pi_b^{-1}(t_k)$.

- If $u = \iota_j(x_k)$, the iteration is useful (Definition 20). Put

$$C_k^* := \lambda_{j'}^{-1}\big(\lambda_j(t_k) - c_1 + P\big), \qquad \pi_k := \frac{|C_k^* \setminus \textsc{Ran}_{b'}|}{N - |\textsc{Ran}_{b'}|},$$

  with $\textsc{Ran}_{b'}$ taken after this answer. Query $p \leftarrow \pi_{b'}(\iota_{j'}(x_k))$ *first*, and add $T_k := \mathbf{1}[p \in C_k^*] - \pi_k$ to T.
- In every case, complete u.

Output 1 if some answer was $\perp$, or if $T \ge \theta := 0.105\,\mathrm{dens} r^3/n_b N^2$; otherwise output 0.

Definition 20. A Phase-3 iteration k of $\mathcal{A}$ is *useful* if its inverse query returns the input $\iota_j(x_k)$ of the attacked call j. Only useful iterations make a test query and add a term T_k to T.

There are at most $|I| + 2r < 4r$ iterations. Each makes one inverse query and at most w forward queries. It assigns the outputs of at most one input or one free input, and $\mathcal{A}$ completes every answer at once. Hence the following hold:

- $\mathcal{A}$ knows every range exactly;
- there are never construction-only inputs;
- every range satisfies $|\textsc{Ran}_{b''}| \le s_{\max} := 4rn$.

Test event and alignment. Every test set has size r. Since $\lambda_j(t_k) - c_2 \in P$, (P2) gives $C_k^* \subseteq C^\cup$. If $p_j(x) = t_k$, then

$$D(\mathrm{F}(x)) \in -c_1 + P \iff \lambda_{j'}(p_{j'}(x)) = \lambda_j(t_k) + D \in \lambda_j(t_k) - c_1 + P \iff p_{j'}(x) \in C_k^*.$$

If $\lambda_j(t) \in c_1 + P$ and $I \subseteq \lambda_{j'}(\textsc{Ran}_{b'})$, consider a candidate $\iota_j(x)$ with value $z \in E$ at target t. Its call-j' masked output is $\lambda_j(t) + D(z) \in (\lambda_j(t) - c_1) + P \subseteq I$, by (P1). So the candidate is illegal at target t. That is, $E \subseteq \mathrm{Forb}_{t,j}$. The doubled length of I is needed for this: the target offsets $\lambda_j(t) - c_1$ range over all of P, and one fixed event E must be forbidden at each of them. For a subgroup ($L = 1$) this is automatic, since $u + H = H$.

Overview of the analysis. We first explain heuristically why the statistic T separates the two worlds at about $N^{3/4}$ queries, and then list the steps of the proof.

In the real world, each test outcome $\mathbf{1}[p \in C_k^*]$ is the indicator that a fresh output of a real permutation lands in a set of size about r. Its conditional mean is exactly π_k, which the distinguisher subtracts. So T is a sum of martingale differences, with mean 0 and variance about $r \cdot r/N$, and it rarely exceeds a few multiples of $r/\sqrt{N}$.

In the ideal world, at every Phase-2 call, a candidate $\iota_j(x)$ whose oracle value lies in E is illegal, so the simulator rejects it and draws again. The distinguisher does not see the rejections. It sees which input each call returned, and a call returns a given input with probability a/V, where a is the acceptance weight of that input and V is the total acceptance weight of all available inputs. A table in which many unrevealed values lie in E has fewer legal candidates at each Phase-2 call, hence a smaller V, and so makes the observed answers more likely. By Bayes' rule this tilts the posterior of every unrevealed value towards E. Each Phase-2 call multiplies the posterior weight of a value in E by about $1 + 1/N$, so after about r calls a value in E has posterior probability about $\pi(1 + r/N)$ instead of $\pi \approx r/N$. The Phase-3 answers are inputs whose values were unrevealed during Phase 2, and the test reads off whether the value lies in E. So each useful iteration has a drift of about $\pi \cdot r/N \approx r^2/N^2$. With about $\mathrm{dens}\, r/n_b$ useful iterations the total drift is about $\mathrm{dens}\, r^3/n_b N^2$. This exceeds the noise $r/\sqrt{N}$ once $r \gg \sqrt{n_b/\mathrm{dens}}\, N^{3/4}$.

The proof makes this rigorous in six steps:

- Section D.2 shows that T has mean 0 in Real and bounds its variance.
- Lemma 33 replaces $\textsc{Sim}_m$ by the simulator with unlimited trials and a structurally good oracle. This only helps the distinguisher, because it outputs 1 on $\perp$.
- The posterior of the oracle is a product over the inverse calls (Section D.2). On a high-probability event every factor is close to N (Lemma 34), and the tilt of a single value takes the threshold form of Lemma 35.
- Lemma 38 computes the first-order drift and controls the nonlinear remainder with Lemma 36. Lemma 37 bounds the probabilities of the hits that could spoil the alignment.
- Lemma 39 shows that, with high probability, few targets are skipped and enough iterations are useful.
- Section D.2 combines these with Chebyshev's inequality.

The full proof of Theorem 2, which follows the steps listed above, is in Section D.2.

Tightness for the forks. Converses of the three reductions of Section 5.6 carry the attack to the fork constructions, against the same wrapped simulators. They are proved in Section D.2.

Corollary 6. *Let $G = \mathbb{F}_2^d$ and $w \geq 1$, and let $m \geq 1$. Consider ForkEDMD$[w]$ (any s) and ForkCENC$[w]$, as in Corollary 4, with the simulators obtained by wrapping $\textsc{Sim}_m$ as there. Put $r := 2^{\lceil 3d/4 \rceil + a}$ with $a \geq 0$*

least such that $r \geq 6.3\,N^{3/4}$. *Suppose that*

$$N \geq 2^{16}, \quad 1000\,w^2 \leq N, \quad 600\,w\,r \ln N \leq N, \quad 134\,r/N + e^{-r/17} + N^{-2} \leq 1/8.$$

Then there is a distinguisher with no construction queries and advantage at least $3/4$. *It makes at most* $3r(w+1) \leq 38\,(w+1)N^{3/4}$ *primitive queries against ForkEDMD, and at most* $3r(w+2) \leq 38\,(w+2)N^{3/4}$ *against ForkCENC. For EDMD and its generalisation Corollary 4(ii) over any finite abelian group, the same holds with* $R := \lceil 6.3\,N^{3/4} \rceil$, *a coset progression with* $R \leq r < 2R$, *the term* $251r/N$, *and at most* $8r \leq 102\,N^{3/4}$ *queries.*

The proof, in Section D.2, translates the distinguisher of Theorem 2 through the converses of the three reductions. Together with Corollary 4, which is stated for $\textsc{Sim}_3$, the wrapped three-trial coin simulator is therefore indifferentiable for these forks up to $\Theta(N^{3/4})$ queries and no further. The attack works for every m, but the matching upper bound is proved only for $m = 3$; for $m \leq 2$ the wrapped EDMD simulator is already distinguishable earlier (see the remark on the number of trials in Section 6.2); for wider forks this would need a short extension of that remark. In the parameters a gap remains. The upper bound $10(w+1)^2q^2/N^{3/2}$ allows about $N^{3/4}/\sqrt{20}(w+1)$ queries, and the attack needs about $38(w+1)N^{3/4}$.

Remarks. We make two remarks, on the scope of the attack and on the number of trials:

- **Coverage.** Every family with $w \geq 1$ has two distinct calls, and every finite abelian group has coset progressions at every scale (Lemma 31). So Theorem 2 applies to every family over every finite abelian group, admissible or not, subject to its regime; in particular the domain density dens must be bounded below for an $O_w(N^{3/4})$ attack. By Corollary 6 it also applies to the fork corollaries. So every indifferentiability bound of this paper, all of which are stated for $\textsc{Sim}_3$, is tight in the exponent for its simulator, at fixed width and class sizes and density bounded below: Corollary 5 for Theorem 1 and its direct instances (Corollaries 1 to 3), and Corollary 6 for Corollary 4. When $b \neq b'$, condition (S2) is not needed: the sets in (S2) lie in the ranges of different permutations. Requiring it anyway is harmless, since Lemma 21 already accounts for it.

- **Number of trials.** The attack works against $\textsc{Sim}_m$ for every m. For $m \leq 2$, however, the coin simulator is distinguishable much earlier by its failures alone, already on one simple instance: the full-domain instance with two independent permutations, $\mathsf{GXP}[G; G; (1,2); \textsc{id}; (-1,1)]$, which is $\pi_1+\pi_2$. Fix m and Q with $(m+1)Q < N$, and put $a := 1-mQ/(N-Q) \in (0,1]$. The distinguisher makes Q inverse queries at fresh targets on the first permutation, and completes each answer by one forward query on the second.

  Consider iteration i, after i successful iterations, so that both ranges have size i. Condition on the simulator's full internal history: all candidates sampled so far, including those already rejected in the current call, and all oracle entries read. At trial r of iteration i, at most $mi + (r-1) \leq mQ$ inputs have been sampled, out of at least $N - i \geq N - Q$ available ones. So the trial picks an input that has never been sampled with conditional probability at least a. The oracle entry of such an input has never been read, so it is uniform and independent of the history. With the target fixed, the candidate is illegal exactly when its output on the second permutation lands in that permutation's range, which has size i; this has probability i/N. Hence each trial rejects with conditional probability at least ai/N. Multiplying these sequential conditional bounds over the m trials, iteration i fails with conditional probability at least $(ai/N)^m$, whatever happened before. The real world never fails, so the $2Q$ queries

distinguish with advantage at least

$$1 - \prod_{i<Q}\Big(1 - \Big(\frac{ai}{N}\Big)^m\Big) \;\geq\; 1 - \exp\Big(-a^m \sum_{i<Q}\Big(\frac{i}{N}\Big)^m\Big).$$

For fixed m and $Q = CN^{m/(m+1)}$, we have $a \to 1$, and the exponent tends to $C^{m+1}/(m+1)$, so the advantage is bounded below by a positive constant. This happens at about $N^{1/2}$ queries for $m = 1$, $N^{2/3}$ for $m = 2$ and $N^{3/4}$ for $m = 3$. The last matches Theorem 1, and explains the choice $m = 3$ there. The calculation is for fixed m; it does not give a statement that is uniform in a growing m.

The same failure attack applies to the wrapped EDMD simulator of Corollary 4(ii). Each inverse query together with its completion is one augmented family query, and Lemma 41 with $w = 1$ turns it into at most two fork queries, so at most $2Q$ fork queries suffice. For $m \geq 4$ no upper bound is proved here; the attack of Section 6.2 shows that $N^{3/4}$ cannot be exceeded.

7 Arbitrary masks

This section removes the admissibility hypothesis from Theorem 1. For admissible families nothing changes, and Theorem 1 with its proof in Sections 5.4 and 5.5 remains the sharper statement. For arbitrary automorphism masks we prove the same exponent $3/4$ for the same simulator $\textsc{Sim}_3$, with additional lower-order terms (Theorem 3). The proof reuses the architecture of Section 5. Every lemma of Sections 5.4 and 5.5 is either used verbatim, with a stated reason why it does not depend on admissibility, or replaced by a version proved here.

The simulator (Definition 11), the legal sets Leg, the forbidden sets Forb and the weights ρ_b are already defined for arbitrary masks: (1) is imposed for every non-rigid same-class pair. Nothing in the simulator changes. The *global* hybrid of Section 5.3, however, is not a well-defined process for arbitrary masks. Its type-(3) value is drawn from $\mathrm{Unif}(\mathcal{S})$, and $\mathcal{S}$ can contain values y with $\mathrm{Leg}(y) = \varnothing$ even when all ranges are empty (Remark 6). We therefore never use it as a process; we define the comparison kernels directly (Section 7.3).

7.1 Activity of same-class pairs

For a same-class pair $p = \{a, a'\}$ with $a < a'$, write $\mu_p := \mu_{aa'} = \lambda_a^{-1} - \lambda_{a'}^{-1} \in \mathrm{End}(G)$ and $\kappa_p := |\ker \mu_p|$.

Definition 21. A same-class pair p is *rigid* if $\kappa_p = N$ ($\mu_p = 0$), *regular* if $\kappa_p = 1$ ($\mu_p \in \mathrm{Aut}(G)$), and *singular* otherwise.

The family is admissible iff it has no singular pair. Put

$$b_0 := \tbinom{w+1}{2}, \qquad \tau_0 := \frac{N}{64\, b_0}.$$

A singular pair is *small* if $\kappa_p \leq \tau_0$ and *large* otherwise. Let $\mathcal{P}_S, \mathcal{P}_L$ be the sets of small and large singular pairs, $B_S := |\mathcal{P}_S|$, $B_L := |\mathcal{P}_L|$ (so $B_S + B_L \leq b_0$), $\kappa_S := \max_{p\in\mathcal{P}_S} \kappa_p$ ($\kappa_S := 1$ if $\mathcal{P}_S = \varnothing$). Let B_R be the number of regular pairs, and put $b_1 := B_R + B_S \leq b_0$, the number of regular and small singular pairs. Finally, let

$$\beta_L := \sum_{p\in\mathcal{P}_L} \frac{1}{\kappa_p} \leq \frac{B_L}{\tau_0} = \frac{64\, b_0 B_L}{N}.$$

Definition 22. For $y \in G^w$, with $y_0 := 0$, put $c_p(y) := \lambda_{a'}^{-1}(y_{a'}) - \lambda_a^{-1}(y_a)$. The pair p is *active at* y if $c_p(y) \in \mathrm{Im}\,\mu_p$.

In this notation (1) for p reads $\mu_p(v_0) \neq c_p(y)$.

Lemma 22. *Let $p = \{a, a'\}$ be a same-class pair.*

(a) If p is active at y, exactly κ_p anchors v_0 (a coset of $\ker \mu_p$) violate (1) for p; otherwise none does.

(b) For a tuple with anchor v_0, raw outputs $p_j = \lambda_j^{-1}(v_0 + y_j)$ and value y, $c_p(y) = p_{a'} - p_a + \mu_p(v_0)$. Hence p is active at y iff $p_{a'} - p_a \in \operatorname{Im} \mu_p$, a condition on the raw outputs alone.

(c) For $y \leftarrow \mathrm{Unif}(G^w)$, $c_p(y)$ is uniform on G, so $\Pr[p \text{ active}] = 1/\kappa_p$. For distinct pairs $p \neq p'$, $(c_p(y), c_{p'}(y))$ is uniform on G^2, so their activities are independent.

(d) For i with value z_i and raw outputs $p_\cdot$, and for $a \neq a'$ in one class, the set $C := \{o \in G : (\lambda_a - \lambda_{a'})(o) = z_{i,a} - z_{i,a'}\}$ has $|C| = \kappa_p \, \mathbf{1}[p \text{ active at } z_i]$.

(e) At an ordinary-fresh slot ω (Section 5.4), $\Pr[p \text{ active at the value of } x^{(\omega)} \mid \mathcal{H}_\omega] \leq \frac{32}{31} \cdot \frac{1}{\kappa_p}$.

Proof. (a) The solutions of $\mu_p(v_0) = c$ form a coset of $\ker \mu_p$ if $c \in \operatorname{Im} \mu_p$, and are empty otherwise.

(b) $\lambda_{a'}^{-1}(y_{a'}) = \lambda_{a'}^{-1}(v_{a'} - v_0) = p_{a'} - \lambda_{a'}^{-1}(v_0)$, and likewise for a; for $a = 0$ use $y_0 = 0 = p_0 - \lambda_0^{-1}(v_0)$. Subtracting gives $c_p(y) = p_{a'} - p_a + (\lambda_a^{-1} - \lambda_{a'}^{-1})(v_0)$, and the last term lies in $\operatorname{Im} \mu_p$.

(c) c_p is a homomorphism, and $y_{a'}$ (with $a' \geq 1$) enters only through the automorphism $\lambda_{a'}^{-1}$, so c_p is onto. For $p \neq p'$ we show that $y \mapsto (c_p(y), c_{p'}(y))$ is onto. Some index $s \geq 1$ lies in exactly one of the two pairs: if $p \setminus p' = \{0\}$, then $p = \{0, a'\}$ with $a' \in p'$, and the other element of p' is nonzero and not in p. Say $s \in p \setminus p'$, the other case being symmetric, and let $s' \geq 1$ be an element of p'. Given targets (u, u'), set all coordinates other than y_s, $y_{s'}$ to 0, choose $y_{s'}$ with $c_{p'}(y) = u'$ (possible since $y_{s'}$ enters $c_{p'}$ through an automorphism and y_s does not enter $c_{p'}$), and then choose y_s with $c_p(y) = u$.

(d) Expanding the masked outputs,

$$z_{i,a} - z_{i,a'} = \lambda_a p_a - \lambda_{a'} p_{a'} = (\lambda_a - \lambda_{a'})(p_a) + \lambda_{a'}(p_a - p_{a'}).$$

So $C \neq \varnothing$ iff $p_a - p_{a'} \in \lambda_{a'}^{-1} \operatorname{Im}(\lambda_a - \lambda_{a'}) = \operatorname{Im}(\lambda_{a'}^{-1}\lambda_a - I)$. Now $\mu_p = (I - \lambda_{a'}^{-1}\lambda_a)\lambda_a^{-1}$, so $\operatorname{Im} \mu_p = \operatorname{Im}(I - \lambda_{a'}^{-1}\lambda_a)$ and $\ker \mu_p = \lambda_a \ker(I - \lambda_{a'}^{-1}\lambda_a)$. Since $\lambda_a - \lambda_{a'} = \lambda_{a'}(\lambda_{a'}^{-1}\lambda_a - I)$, also $|\ker(\lambda_a - \lambda_{a'})| = \kappa_p$, and by (b) $C \neq \varnothing$ iff p is active at z_i. No commutativity of masks is used.

(e) By (b) the event is $\lambda_{a'} p_{a'}(x^{(\omega)}) \in \lambda_{a'}\big(p_a(x^{(\omega)}) + \operatorname{Im} \mu_p\big)$, a set of size N/κ_p fixed by $\mathcal{H}_\omega$ and the coordinate p_a. Apply (4). □

Definition 23. A value $y \in G^w$ is *bad* if some rigid pair or some large singular pair is active at y. Let $\Lambda_\star \subseteq G^w$ be the set of bad values, and $\mathcal{S}_\star := G^w \setminus \Lambda_\star$.

Definition 24. The *singular excess* of $y \in G^w$ is $e_S(y) := \sum_{p \in \mathcal{P}_S} \kappa_p \, \mathbf{1}[p \text{ active at } y]$.

Then $\mathcal{S}_\star \subseteq \mathcal{S}$, and $e_S(y) \leq B_S \tau_0 \leq N/64$. By Lemma 22(a), for $y \in \mathcal{S}_\star$ the collision count of (2) satisfies $e(y) \leq B_R + e_S(y)$: each regular pair forbids one anchor, each active small singular pair forbids κ_p anchors, and rigid pairs and inactive pairs forbid none. By Lemma 22(c), $|\mathcal{S} \setminus \mathcal{S}_\star| \leq \beta_L N^w$.

7.2 The theorem

In addition to (3) we assume, when $B_L > 0$,

$$b_0^2 \leq N/2^{12}, \qquad \text{so that} \qquad \beta_L \leq 64 b_0^2/N \leq 1/64. \tag{13}$$

Theorem 3. *Let the masks be arbitrary automorphisms, with any class sizes and free inputs allowed. Assume* (3) *and* (13)*. Then*

$$\begin{aligned}\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\textsc{Sim}_3}(q) \;\le\; & C_\star \frac{q^2}{N^{3/2}} + 2w(w+1)\frac{q}{N} + 3.6\, b_1 \frac{\sqrt{q}}{N} \\ & + 1.8\, \frac{\sqrt{B_S \kappa_S q}}{N} + 1.6\, B_S \sqrt{n\kappa_S}\, \frac{q}{N^{3/2}} + 9.6\, \beta_L q + 22.6\, \beta_L \sqrt{q},\end{aligned}$$

with

$$\begin{aligned} C_\star &:= \sqrt{K^\star_{\mathrm{RH}}/2} + \sqrt{K^\star_{\mathrm{HI}}/2}, \\ K^\star_{\mathrm{RH}} &:= 1.6\, K_{\mathrm{RH}} + \frac{5.8F^2B_S^2 + 5.8n^2B_S^2 + 3.5n^2B_S + 11.8w^2n^2B_S + 0.05nB_S^2}{2w}, \\ K^\star_{\mathrm{HI}} &:= 111.2\, n^3 w + 128\, n^3 w^3, \end{aligned}$$

where $K_{\mathrm{RH}} \le 85n^3(w+1)^4$ is the constant of Theorem 1. The term in b_1 vanishes without regular and small singular pairs (in particular when $n = 1$), the other terms in B_S, κ_S without small singular pairs, those in β_L without large ones, and the term $2w(w+1)q/N$ without rigid pairs. Using $\beta_L \le 64 b_0 B_L/N$, the last two terms are at most $615\, b_0 B_L q/N + 1447\, b_0 B_L \sqrt{q}/N$.

For fixed (k, n, w) the terms other than the first are of order q/N, $\sqrt{q}/N$, $\sqrt{q\kappa_S}/N$ or $q\sqrt{\kappa_S}/N^{3/2}$. Since $\kappa_S \le N$, each of these tends to 0 for $q = o(N^{3/4})$, so the exponent is $3/4$ for every choice of masks. For admissible families $B_S = B_L = 0$ and Theorem 1 gives better constants. Lower bounds, and what is and is not tight, are collected in Section 7.7. In short, the query exponent $3/4$ is matched for $\textsc{Sim}_m$ on positive-density families, but the advantage formula of Theorem 3 is *not* matched term by term.

7.3 Plan of the proof

A response *reveals* a value if it is a type-(1) answer, the value sampled in a type-(3) query, or the value of a fresh family input returned by a type-(4) query (through its companion). A transcript is *bad* once a revealed value lies in $\Lambda_\star$. In $\textsc{Real}$ values activating a rigid pair never occur, but values activating a large singular pair do, with small probability. We remove them by total variation, through the following worlds:

- $\textsc{Hyb}_\star$, defined directly by its next-response kernel $K_\star(\tau)$ at every prefix τ whose revealed values lie in $\mathcal{S}_\star$:
  - type (1): $z \leftarrow \mathrm{Unif}(\mathcal{S}_\star)$;
  - types (2) and (3): the anchor $v_0 \leftarrow \mathrm{Unif}(\mathrm{Leg}(z))$, where for type (3) the value z is first drawn from $\mathrm{Unif}(\mathcal{S}_\star)$; this is well defined because $|\mathrm{Leg}(z)| \ge \ell_\star > 0$ for $z \in \mathcal{S}_\star$ ((16) below);
  - type (5): $\mathrm{Unif}(G \setminus \textsc{Ran}_b)$;
  - type (4): weights 1 for legal construction-only pairs, ρ_b for fresh free inputs and $\rho''_j := |\mathcal{S}_\star \setminus \mathrm{Forb}_{t,j}|/|\mathcal{S}|$ for fresh family inputs $\iota_j(x)$; a fresh family output is completed with $z \leftarrow \mathrm{Unif}(\mathcal{S}_\star \setminus \mathrm{Forb}_{t,j})$ and anchor $\lambda_j(t) - z_j$.

  For type (4) we also use the *old* inverse kernel K of Section 5.3 (weights 1, ρ_b, ρ_b, completion from $\mathrm{Unif}(\mathcal{S} \setminus \mathrm{Forb}_{t,j})$). It is well defined: its tuple set $\mathcal{T}_0$ is nonempty and its weights are positive. $K_\star(\tau)$ is exactly K conditioned on its revealed value not lying in $\Lambda_\star$, with $\rho''_j = \rho_b(1 - \beta^K_j)$ and $\beta^K_j := |(\mathcal{S} \setminus \mathrm{Forb}_{t,j}) \cap \Lambda_\star|/|\mathcal{S} \setminus \mathrm{Forb}_{t,j}|$;
- $\textsc{Real}^\perp$ and $\textsc{Hyb}^\perp$, the stopped real world and the stopped hybrid of Definition 25 below;

- $\textsc{Ideal}_\star$: $\textsc{Ideal}_3$ with Γ replaced by a uniform function $\mathcal{D} \to \mathcal{S}_\star$; the simulator is unchanged, and its dummy values are still drawn from $\mathcal{S}$.

Definition 25. The *stopped real world* $\textsc{Real}^\perp$ runs as $\textsc{Real}$, but at the first bad response it answers $\perp$, and $\perp$ to every later query. At a prefix τ that is not bad, let β_τ be the probability, under the real conditional law given τ, that the next response is bad, if τ has positive probability under $\textsc{Real}$, and $\beta_\tau := 0$ otherwise. The *stopped hybrid* $\textsc{Hyb}^\perp$ answers at such a prefix from $\beta_\tau \delta_\perp + (1-\beta_\tau) K_\star(\tau)$, and answers $\perp$ thereafter.

Since β_τ is a function of the prefix, $\textsc{Hyb}^\perp$ is a well-defined process. The real-side KL below only involves real-supported prefixes, and the bound $\beta_\tau \le \beta_{\max}$ of Lemma 24 holds on all prefixes. Then, by the triangle inequality and Pinsker's inequality,

$$\begin{aligned}\mathbf{Adv} \le \Pr_{\textsc{Real}}[\mathrm{bad}] + \sqrt{\tfrac{1}{2}\,\mathrm{KL}(T_{\textsc{Real}^\perp} \| T_{\textsc{Hyb}^\perp})} + \Pr_{\textsc{Hyb}^\perp}[\perp] \\ + \sqrt{\tfrac{1}{2}\,\mathrm{KL}(T_{\textsc{Hyb}_\star} \| T_{\textsc{Ideal}_\star})} + \mathrm{TV}(T_{\textsc{Ideal}_\star}, T_{\textsc{Ideal}_3}),\end{aligned} \tag{14}$$

where we used $\mathrm{TV}(T_{\textsc{Real}}, T_{\textsc{Real}^\perp}) \le \Pr_{\textsc{Real}}[\mathrm{bad}]$ and $\mathrm{TV}(T_{\textsc{Hyb}^\perp}, T_{\textsc{Hyb}_\star}) \le \Pr_{\textsc{Hyb}^\perp}[\perp]$. The coupling behind the last inequality runs both worlds with identical answers until $\textsc{Hyb}^\perp$ outputs $\perp$. $\textsc{Hyb}_\star$ only visits prefixes whose revealed values lie in $\mathcal{S}_\star$, so $K_\star$ is always defined along its executions.

Lemma 23. *At a real-supported prefix τ that is neither bad nor ended, let R_G be the real law of the next response conditioned on not being bad, and $K_G := K_\star(\tau)$. Then the next-response laws $\tilde R, \tilde K$ of $\textsc{Real}^\perp, \textsc{Hyb}^\perp$ satisfy* $\mathrm{KL}(\tilde R \| \tilde K) = (1-\beta_\tau)\,\mathrm{KL}(R_G \| K_G) \le \chi^2(R_G \| K_G)$.

Proof. $\tilde R = \beta_\tau \delta_\perp + (1-\beta_\tau) R_G$ and $\tilde K = \beta_\tau \delta_\perp + (1-\beta_\tau) K_G$, so the $\perp$ atom contributes $\beta_\tau \log 1 = 0$. Finally $\mathrm{KL} \le \chi^2$. □

Remark 6. Over $\mathbb{F}_2^{18}$ take one permutation, $w = 3$, and masks equal to the identity and the three transvections of $\mathrm{GL}(2,2)$ on the first two coordinates, extended by the identity. There are no rigid pairs, so $\mathcal{S} = G^3$. At $y = 0$ the four raw outputs are $\lambda_j^{-1}(v_0)$. Every vector of $\mathbb{F}_2^2$ is fixed by some transvection, so the identity output collides with another output for every anchor, and $\mathrm{Leg}(0) = \varnothing$. With $q = 6$ both (3) and (13) hold. The value 0 activates a large singular pair, so it lies in $\Lambda_\star$, and $\textsc{Hyb}_\star$ and $\textsc{Ideal}_\star$ never use it. The simulator itself returns $\perp$ on such values (Definition 11). The same example shows that the bound $|\mathrm{Leg}(y)| \ge 15N/16$ for all $y \in \mathcal{S}$, in the running-time paragraph of Section 4.2, holds only for admissible families. For arbitrary masks a running-time claim needs a separate capping argument; Theorem 3 is a statement about advantage only.

Lemma 24. *Assume (3) and (13). At every prefix that is not bad, $\beta_\tau \le \beta_{\max} := 2.8\,\beta_L \le 0.044$, and $\beta_\tau = 0$ for queries of types (2) and (5). Under the old laws* $\mathrm{Unif}(\mathcal{S})$ *(types (1) and (3)) and K (type (4)), the corresponding probabilities are at most $1.38\,\beta_L$; in particular $\beta_j^K \le 1.38\,\beta_L$. Consequently* $\Pr_{\textsc{Real}}[\mathrm{bad}] \le 2.8\,q\beta_L$ *and* $\Pr_{\textsc{Hyb}^\perp}[\perp] \le 2.8\,q\beta_L$.

Proof. For types (1) and (3), by Lemma 7, $p_U(z) = |\mathrm{Leg}_U(z)|/\prod_b (N - |R_{b,U}|)_{n_b}$. Here $|R_{b,U}| \le n\bar s \le nq$ and $(M)_{n_b} \ge (M-n)^{n_b}$, so $p_U(z) \le N^{-w}(1-x)^{-(w+1)}$ with $x := (nq+n)/N$. Since $(w+1)(nq+n) = \bar F(q+1)$ and $\bar F q \le N/32$, $(1-x)^{-(w+1)} \le 1.05$ (Section B.2 (5)). Values activating a rigid pair have $|\mathrm{Leg}_U| = 0$. So the probability of a bad value is at most $1.05 \cdot |\mathcal{S} \cap \Lambda_\star|/N^w \le 1.05\beta_L$, by Lemma 22(c).

For type (4), only a fresh family output $\iota_j(x)$ reveals a value. By the proof of Lemma 12, conditional on this output and on U, the companion is uniform on $\mathcal{T}_U \subseteq \mathcal{T}_0$ with $|\mathcal{T}_U| = (1-\bar p)|\mathcal{T}_0| \ge |\mathcal{T}_0|/2$.

Under the bijection of Lemma 4, the bad tuples of $\mathcal{T}_0$ correspond to values in $(\mathcal{S} \setminus \mathrm{Forb}_{t,j}) \cap \Lambda_\star$, of which there are at most $\beta_L N^w$. With $|\mathcal{T}_0| = \rho_b|\mathcal{S}| \geq \frac{3}{4} \cdot \frac{31}{32} N^w$, the conditional probability is at most $2\beta_L N^w/|\mathcal{T}_0| \leq 2.8\beta_L$ (Section B.2 (5)). Types (2) and (5) reveal no new value. Under the old laws the same counts give $|\mathcal{S} \cap \Lambda_\star|/|\mathcal{S}| \leq \frac{32}{31}\beta_L$ for types (1) and (3), and $\beta_j^K \leq \beta_L N^w/|\mathcal{T}_0| \leq 1.38\beta_L$. The last two claims follow by a union bound over the at most q steps, since $\textsc{Real}^\perp$ and $\textsc{Hyb}^\perp$ only continue on prefixes that are not bad. Finally $\beta_L \leq 1/64$ by (13), so $2.8\beta_L \leq 0.044$. □

The coupling term is handled exactly as in Section 5.5: each query reads at most $1 + m = 4$ entries of Γ (dummy values read none), and each read entry lies in $\Lambda_\star$ with probability at most $|\Lambda|/N + \beta_L$. Hence

$$\mathrm{TV}(T_{\textsc{Ideal}_\star}, T_{\textsc{Ideal}_3}) \leq 4q\Big(\frac{|\Lambda|}{N} + \beta_L\Big) \leq 2w(w+1)\frac{q}{N} + 4\beta_L q. \tag{15}$$

7.4 Charges for singular pairs

Index the inputs exposed by construction queries by their slots ω (Section 5.4); these slots are ordinary-fresh. For such a slot let $A_{\omega,p}$ be the event that $p \in \mathcal{P}_S$ is active at the value of $x^{(\omega)}$, and $\varepsilon_\omega := \sum_{p \in \mathcal{P}_S} \kappa_p \mathbf{1}[A_{\omega,p}]$ its singular excess. Define the game-wide quantities

$$A_{\mathrm{tot}} := \sum_\omega \sum_{p \in \mathcal{P}_S} \mathbf{1}[A_{\omega,p}], \qquad \Xi := \sum_\omega \varepsilon_\omega, \qquad \Xi_2 := \sum_\omega \varepsilon_\omega^2,$$

computed over the whole real game, with undefined slots contributing 0. All three are nonnegative and only grow as the game proceeds. In particular, at any time the construction-only inputs $x_1, \ldots, x_h$ satisfy $\#\{i : \varepsilon_i > 0\} \leq A_{\mathrm{tot}}$, $\sum_i \varepsilon_i \leq \Xi$ and $\sum_i \varepsilon_i^2 \leq \Xi_2$, where ε_i denotes the excess of x_i. Since $\textsc{Real}^\perp$ follows $\textsc{Real}$ until it stops, these bounds also hold in $\textsc{Real}^\perp$, with the same real-game quantities.

Lemma 25. *For every deterministic distinguisher, in* $\textsc{Real}$,

$$\mathbb{E}[\Xi] \leq 1.04\, q B_S, \qquad \mathbb{E}[\Xi_2] \leq 1.04\, q B_S^2 \kappa_S,$$
$$\mathbb{E}[A_{\mathrm{tot}}\Xi] \leq 1.05\, q^2 B_S^2, \qquad \mathbb{E}[\Xi^2] \leq 1.07\, q^2 B_S^2 + 1.04\, q B_S^2 \kappa_S.$$

Proof. By Lemma 22(e), $\Pr[A_{\omega,p}] \leq c/\kappa_p$ with $c := 32/31$, and for slots $\omega < \omega'$

$$\Pr[A_{\omega,p} \cap A_{\omega',p'}] = \mathbb{E}\big[\mathbf{1}_{A_{\omega,p}} \Pr(A_{\omega',p'} \mid \mathcal{H}_{\omega'})\big] \leq \frac{c^2}{\kappa_p \kappa_{p'}},$$

because $A_{\omega,p}$ is determined by $\mathcal{H}_{\omega'}$. Whether a slot is a construction exposure is decided before its tuple is sampled, so the eligibility indicator causes no selection. On one slot, $\Pr[A_{\omega,p} \cap A_{\omega,p'}] \leq c/\max(\kappa_p, \kappa_{p'})$. Since there are at most q slots, we obtain the following:

- $\mathbb{E}[\Xi] \leq \sum_\omega \sum_p \kappa_p\, c/\kappa_p = c\, q B_S$.
- The Cauchy–Schwarz inequality over at most B_S terms gives $\varepsilon_\omega^2 \leq B_S \sum_p \kappa_p^2 \mathbf{1}[A_{\omega,p}]$, so $\mathbb{E}[\Xi_2] \leq cqB_S \sum_p \kappa_p \leq c\, q B_S^2 \kappa_S$.
- $\mathbb{E}[A_{\mathrm{tot}}\Xi] = \sum_{(\omega,p),(\omega',p')} \kappa_{p'} \Pr[A_{\omega,p} \cap A_{\omega',p'}]$. For $\omega \neq \omega'$ the terms are at most $c^2/\kappa_p \leq c^2/2$ (as $\kappa_p \geq 2$), so they total at most $c^2 q^2 B_S^2/2$. For $\omega = \omega'$ the terms are at most $\kappa_{p'} \cdot c/\kappa_{p'} = c$, totalling at most cqB_S^2. With $q \geq 2$ the sum is at most $(c^2/2 + c/2)q^2 B_S^2$.
- $\mathbb{E}[\Xi^2] = \sum \kappa_p \kappa_{p'} \Pr[A_{\omega,p} \cap A_{\omega',p'}]$. For $\omega \neq \omega'$ the total is at most $c^2 q^2 B_S^2$. For $\omega = \omega'$ the terms are at most $c \min(\kappa_p, \kappa_{p'}) \leq c\kappa_S$, totalling at most $cqB_S^2\kappa_S$.

With $c = 32/31$ these are the stated bounds (Section B.2 (6)). □

7.5 Real to hybrid

Unchanged ingredients. We use the following verbatim, for arbitrary automorphism masks:

- (4), Lemma 7, (5) and the description of the hidden anchors in Section 5.4, except for the count of collision values, which is treated below. These use only that the masks are automorphisms ($\alpha_{jj'} \in \mathrm{Aut}(G)$) and injectivity.
- Lemma 4, its bijection $\mathcal{S} \setminus \mathrm{Forb}_{t,j} \leftrightarrow \mathcal{T}_0$, and $1 - \rho_b \le w(s_{\max} + n)/N$: the proof counts tuples and never uses (1).
- Lemma 11, Steps 1–2, and Lemma 12, Steps 1–2, 4 and 5. Step 5 needs only that $\varphi_{aa'}(v) = \lambda_a p_{j'} - \lambda_{a'} p_{j''}$, for distinct coordinates $j' \neq j''$, is uniform, which holds because $\lambda_a, \lambda_{a'}$ are automorphisms.
- Lemma 13 and Remark 4. The four cases of its proof use only automorphism masks, (4) and $f \neq f'$; $\psi_{ff'}$ may be singular, as it already can be for admissible families.
- Lemma 14 and (8). Each coordinate projection of $\Psi_{\theta\theta'}$ is onto, because each $\varphi_{aa'}$ contains an automorphism image of a coordinate.

Support sizes. At a prefix that is not bad, every construction-only value z_i lies in $\mathcal{S}_\star$. Given the transcript and U_{-i}, the anchor U_i must avoid at most $\bar{F}(r+f) + F(h-1) + b_0 \le \Phi$ values (masked ranges, exclusions (5), regular collisions), plus at most ε_i singular collision values (Lemma 22(a)). So U_i is uniform on a set D_i with

$$|D_i| \ge N - \Phi - \varepsilon_i \ge \ell_\star := \ell - \frac{N}{64} \ge \frac{57}{64} N, \tag{16}$$

and on $\mathcal{A}$ (Lemma 12, Step 1), $|D_i| \ge N - \Phi - n - \varepsilon_i \ge \ell'_\star := \ell' - N/64 \ge \frac{55}{64} N$. Likewise $|\mathrm{Leg}(z_0)| \ge \ell_\star$ for every $z_0 \in \mathcal{S}_\star$. Throughout, $\Pr[U_i \in S \mid U_{-i}] \le |S|/\ell_\star$ replaces (7).

Lemma 26. *At a prefix that is not bad, the next type-(1) answer, and the value part of a type-(3) answer, satisfy*

$$\chi^2(R_G \| \mathrm{Unif}(\mathcal{S}_\star)) \le \frac{C_{\mathrm{con}}(w) s_{\max}^3}{N^3} + \frac{25.6\, b_1^2 + 6.4\, B_S \kappa_S}{N^2},$$

with $s_{\max} = n\bar{s}$ and C_{con} as in Lemma 8.

Proof. Conditional on U, $R_{G,U}(z) \propto |\mathrm{Leg}_U(z)|$ on $\mathcal{S}_\star$. R_G is a mixture over U (with weights reweighted by the conditioning). Since $\chi^2(\cdot \,\|\, \mathrm{Unif}(\mathcal{S}_\star))$ is convex in its first argument and the second argument is the same for every component, it suffices to bound each $\chi^2(R_{G,U} \| \mathrm{Unif}(\mathcal{S}_\star)) = \mathrm{Var}_{\mathcal{S}_\star} |\mathrm{Leg}_U| / (\mathbb{E}_{\mathcal{S}_\star}[|\mathrm{Leg}_U|])^2$, as in Lemma 8.

For the denominator, on $\mathcal{S}_\star$ we have $|\mathrm{Leg}_U| \ge N - \sum_j s_j - b_0 - e_S \ge c_0 N$, where c_0 is the absolute constant of Section B.2 (7). For the numerator, define $\mathrm{Lg} = N - \sum_j s_j + \Sigma_2 - \Sigma_3 - \bar{e}$ on G^w, where $\bar{e} := e$ on $\mathcal{S}_\star$ and $\bar{e} := 0$ elsewhere. Then $\mathrm{Lg} = |\mathrm{Leg}_U|$ on $\mathcal{S}_\star$ and $0 \le \bar{e} \le B_R + e_S$. For any X, $\mathrm{Var}_{\mathcal{S}_\star} X \le (N^w/|\mathcal{S}_\star|) \mathbb{E}_{G^w}[(X - \mathbb{E}_{G^w}[X])^2]$, and $N^w/|\mathcal{S}_\star| \le \frac{64}{61}$, because $|\mathcal{S}| \ge \frac{31}{32} N^w$ and $|\mathcal{S} \setminus \mathcal{S}_\star| \le \beta_L N^w \le N^w/64$. So

$$\mathrm{Var}_{\mathcal{S}_\star} \mathrm{Lg} \le \tfrac{64}{61} \cdot 3\big(\mathrm{Var}\, \Sigma_2 + \mathbb{E}[\Sigma_3^2] + \mathbb{E}_{G^w}[(B_R + e_S)^2]\big).$$

The first two terms are bounded exactly as in Lemma 8. For the third, Lemma 22(c) gives $\mathbb{E}[e_S] = B_S$ and $\mathbb{E}[e_S^2] = \sum_p \kappa_p + \sum_{p \neq p'} 1 \le B_S \kappa_S + B_S^2$. So

$$\mathbb{E}[(B_R + e_S)^2] \le B_R^2 + 2 B_R B_S + B_S \kappa_S + B_S^2 = b_1^2 + B_S \kappa_S \le 4 b_1^2 + B_S \kappa_S.$$

Dividing by $(c_0 N)^2$ gives the prefactor $c_1 := \frac{64}{61} \cdot 3/c_0^2$. It is below the coefficient of the range terms in Lemma 8, so these terms are at most those of Lemma 8, and the collision term is at most $c_1(4b_1^2 + B_S\kappa_S)/N^2$, which is the stated term (Section B.2 (7)). □

Lemma 27. *At a prefix that is not bad:*

(a) for a type-(2) query, or the anchor part of a type-(3) query, with h_x, $\kappa(a)$ as in Lemma 9,

$$\chi^2 \le \frac{5.6\, h_x^2 F^2 \Phi}{N^3} + \frac{6.1}{N^2}\mathbb{E}_a[\kappa(a)^2] + \frac{5.5\, F^2 A_{\mathrm{tot}}\Xi}{N^3};$$

(b) for a type-(5) query,

$$\chi^2 \le \frac{5.6\, n^2 h^2 \Phi}{N^3} + \frac{1.6\, n^4 h^2}{N^3} + \frac{5.5\, n^2 A_{\mathrm{tot}}\Xi + 3.1\, n^2 h\, \Xi}{N^3}.$$

Proof. (a) Follow the proof of Lemma 9. Steps 1, 2 and 4 are unchanged, with $E_0 = \mathrm{Leg}(z_0)$, $|E_0| \ge \ell_\star$, and $\ell_\star$ in place of ℓ in Step 4. In Step 3, which bounds $\mathrm{Var}_a[\mu_i(a)]$, we have $\mu_i(a) \le 1/|D_i| \le 1/(N - \Phi - \varepsilon_i)$ and $\mu_i(a) \ge (1 - F(h_x - 1)/\ell_\star)/N$ on $\mathrm{Leg}(z_i)$. So $|\mu_i(a) - 1/N| \le (\Phi + \varepsilon_i)/\ell_\star N$ there, and $\mu_i = 0$ on the at most $\Phi + \varepsilon_i$ values outside $\mathrm{Leg}(z_i)$. Hence

$$\mathrm{Var}_a[\mu_i(a)] \le \frac{\Phi + \varepsilon_i}{N^3}\Big(1 + \frac{(\Phi + \varepsilon_i)N}{\ell_\star^2}\Big) \le c_3\, \frac{\Phi + \varepsilon_i}{N^3},$$

where the absolute constant c_3 bounds the bracket. Also $p_a \le F h_x/\ell_\star$ is at most a small absolute constant, which bounds $\bar\pi$ from below (Section B.2 (8)).

In Step 5, which bounds $\mathrm{Var}[p_a]$, we have $A = \sum_i \sum_{f \in \mathcal{P}_i} \mu_i \circ g_{i,f}$, and each $g_{i,f}$ is a bijection, so the standard deviation of each term is $\sqrt{\mathrm{Var}\,\mu_i}$. Using $\mathrm{Var}(\sum X_k) \le (\sum \sqrt{\mathrm{Var}\, X_k})^2$, $\sqrt{\Phi + \varepsilon_i} \le \sqrt{\Phi} + \sqrt{\varepsilon_i}$ and $(\sum_i \sqrt{\varepsilon_i})^2 \le \#\{i : \varepsilon_i > 0\} \sum_i \varepsilon_i \le A_{\mathrm{tot}}\Xi$,

$$\mathrm{Var}_a[A] \le \frac{c_3 F^2}{N^3}\Big(\sum_i (\sqrt{\Phi} + \sqrt{\varepsilon_i})\Big)^2 \le \frac{2c_3 F^2}{N^3}\big(h_x^2\Phi + A_{\mathrm{tot}}\Xi\big).$$

The remaining parts of Step 5 hold with $\ell_\star$ in place of ℓ: $E_a \le \kappa(a)/\ell_\star + D_{\mathrm{f}}$ with $D_{\mathrm{f}} = \binom{h_x}{2}(F/\ell_\star)^2$, and $D_{\mathrm{f}}^2 \le c_D\, h_x^2 F^2 \Phi/N^3$ with $c_D := (N/\ell_\star)^4/128$. With $\mathrm{Var}_a[E_a] \le 2\mathbb{E}[\kappa^2]/\ell_\star^2 + 2D_{\mathrm{f}}^2$ and $\mathrm{Var}(u - v) \le 2\,\mathrm{Var}\, u + 2\,\mathrm{Var}\, v$,

$$\mathrm{Var}[p_a] \le \frac{(4c_3 + 4c_D)\, h_x^2 F^2 \Phi + 4c_3\, F^2 A_{\mathrm{tot}}\Xi}{N^3} + \frac{4\mathbb{E}[\kappa^2]}{\ell_\star^2}.$$

In Step 6 the bound is multiplied by $(N/\ell_\star)/\bar\pi^2$, which gives the constants 5.6, 5.5 and 6.1 of (a) (Section B.2 (9)).

(b) Follow the proof of Lemma 10 with the same changes, (Fh_x, F) replaced by (nh, n). The marginal part gives $5.6n^2h^2\Phi/N^3 + 5.5n^2 A_{\mathrm{tot}}\Xi/N^3$. For the coincidence count, the pointwise bound of Lemma 10 is replaced as follows. Index by $\theta = (i, \{a, a'\})$ with $a \ne a'$ in class b; there are at most $hn^2/2$ labels. By Lemma 22(d), $|C_\theta| = \kappa_p \mathbf{1}[p \text{ active at } z_i]$. On a prefix that is not bad, rigid and large singular pairs are inactive at every z_i, regular pairs give $|C_\theta| \le 1$, and small singular pairs give ε_i in total. So $\sum_\theta |C_\theta| \le hn^2/2 + \Xi$, and

$$\mathbb{E}_o[\kappa^2] = \frac{1}{N}\sum_{\theta,\theta'} |C_\theta \cap C_{\theta'}| \le \frac{1}{N}\Big(\frac{hn^2}{2}\Big)\Big(\frac{hn^2}{2} + \Xi\Big).$$

With the coefficient $6.1/N^2$ of (a), this gives the terms in n^4h^2 and $n^2h\,\Xi$ of (b) (Section B.2 (9)). □

Lemma 28. *Let R, K be laws on a finite set, and $a, b \ge 0$ functions with $R_G := aR$ and $K_G := bK$ probability laws, $a, b \in [1 - \beta, (1 - \beta)^{-1}]$, and $\beta \le 0.044$. Then $\chi^2(R_G \| K_G) \le 1.26\, \chi^2(R\|K) + 11.5 \max|a - b|^2$, and $\max|a - b| \le 2.1\beta$.*

Proof. $R_G - K_G = a(R-K) + K(a-b)$ and, by Young's inequality with parameter $\eta = 0.1001$, $(x+y)^2 \le (1+\eta)x^2 + (1+1/\eta)y^2$. Divide by $K_G = bK \ge (1-\beta)K$, and use $a^2 \le (1-\beta)^{-2}$. The two coefficients are $(1+\eta)(1-\beta)^{-3} \le 1.26$ and $(1+1/\eta)(1-\beta)^{-1} \le 11.5$. Finally $|a-b| \le (1-\beta)^{-1} - (1-\beta) \le 2.1\beta$ (Section B.2 (10)). □

Lemma 29. *At a prefix that is not bad, the next type-(4) response satisfies, with $K_{\mathrm{inv}} := 39.1$,*

$$\mathrm{KL}(R_G \| K_G) \le \frac{K_{\mathrm{inv}} nh\Phi^2}{N^3} + \frac{4.3\, n\Xi_2}{N^3} + \frac{5.7\, n^2\Xi^2}{N^4} + 1.03\left(\frac{36 w^2 n^2 h^2 \Phi}{N^3} + \frac{43.4}{N^2} W\right) + \frac{11.3\, w^2 n^2 q\, \Xi}{N^3} + 400\, \beta_L^2,$$

with W as in Lemma 14.

Proof. By the chain rule, $\mathrm{KL}(R_G\|K_G) \le \chi^2(R_G^{\mathrm{out}}\|K_G^{\mathrm{out}}) + \mathbb{E}[\chi^2(R_G^{\mathrm{comp}}\|K_G^{\mathrm{comp}})]$.

Unconditioned output. Follow Lemma 11. Steps 1–2 are unchanged. In Step 3, at a legal pair (i,j), $|r_{ij} - 1/N| \le (\Phi + \varepsilon_i)/\ell_\star N$ and $|1/N_{\mathrm{HYB}} - 1/N| \le \Phi/N(N-\Phi)$, so $|R-K| \le (c_1\Phi + c_2\varepsilon_i)/N^2$ with $c_2 := N/\ell_\star$ and $c_1 := c_2 + N/(N-\Phi)$. With $(x+y)^2 \le 2x^2 + 2y^2$, $K \ge 3/4N$ and at most n legal pairs per i, the legal pairs contribute at most $\frac{8}{3}(c_1^2 nh\Phi^2 + c_2^2 n\Xi_2)/N^3$. In Step 4, $\Delta \le (c_1 nh\Phi + c_2 n\Xi)/N^2$, and the fresh inputs contribute at most $\frac{16}{9}\Delta^2 \le \frac{32}{9}(c_1^2 n^2 h^2 \Phi^2 + c_2^2 n^2 \Xi^2)/N^4$, where $n^2h^2\Phi^2/N^4 \le nh\Phi^2/N^3$ since $nh \le N$. So

$$\chi^2(R^{\mathrm{out}}\|K^{\mathrm{out}}) \le c_{\mathrm{o}} \frac{nh\Phi^2}{N^3} + \frac{8}{3}c_2^2 \frac{n\Xi_2}{N^3} + \frac{32}{9}c_2^2 \frac{n^2\Xi^2}{N^4}, \qquad c_{\mathrm{o}} := \left(\frac{8}{3} + \frac{32}{9}\right)c_1^2.$$

Conditioned output. $R_G^{\mathrm{out}}(v) = R^{\mathrm{out}}(v) a_v$ with $a_v = (1-\beta_v^R)/(1-\beta_\tau)$, where β_v^R is the real conditional probability that the companion of v is bad (zero unless v is a fresh family input). Likewise $K_G^{\mathrm{out}}(v) = K^{\mathrm{out}}(v) b_v$ with $b_v = (1-\beta_v^K)/(1-\beta^K)$, where $\beta_v^K := \beta_j^K$ if $v = \iota_j(x)$ is a fresh family input and $\beta_v^K := 0$ for free and construction-only outputs, and β^K is the total bad probability of the old kernel K; this is the output law of $K_\star(\tau)$. By Lemma 24 all these probabilities are at most $\beta_{\max} \le 0.044$. Lemma 28 gives

$$\chi^2(R_G^{\mathrm{out}}\|K_G^{\mathrm{out}}) \le 1.26\, \chi^2(R^{\mathrm{out}}\|K^{\mathrm{out}}) + 11.5\, (2.1\beta_{\max})^2,$$

with $\beta_{\max} = 2.8\beta_L$. Evaluating c_1 and c_2 gives the constants $K_{\mathrm{inv}} = 39.1$, 4.3, 5.7 and 400 of the statement (Section B.2 (11)).

Companion. Given output $\iota_j(x)$ and U, the real companion is uniform on $\mathcal{T}_U$, and $|\mathcal{T}_U|$ does not depend on U (Lemma 12, Step 2). Conditioning on not being bad restricts the companion to $\mathcal{T}_0' := \{v \in \mathcal{T}_0 : \text{value} \notin \Lambda_\star\}$. The joint law of (U,v) is proportional to $P(U)\mathbf{1}[v \in \mathcal{T}_U \cap \mathcal{T}_0']$, so $R_G^{\mathrm{comp}}(v) \propto 1 - p_v$ on $\mathcal{T}_0'$, while $K_G^{\mathrm{comp}} = \mathrm{Unif}(\mathcal{T}_0')$. Hence $\chi^2 = \mathrm{Var}_{\mathrm{Unif}(\mathcal{T}_0')}[p_v]/(1-\bar p')^2$, and

$$\mathrm{Var}_{\mathrm{Unif}(\mathcal{T}_0')}[p_v] \le \frac{|\mathcal{T}_0|}{|\mathcal{T}_0'|}\mathrm{Var}_{\mathrm{Unif}(\mathcal{T}_0)}[p_v], \qquad \frac{|\mathcal{T}_0|}{|\mathcal{T}_0'|} = \frac{1}{1-\beta_j^K} \le 1.03,$$

by Lemma 24. Also $\bar p \le \sum_{j'} n_{c(j')} h/M_{j'}$ is at most a small absolute constant, since $wnh \le N/64$ and $M_{j'}$ is close to N (Lemma 12, Step 3), so $\bar p' \le 1.03\,\bar p \le 1/2$ and $(1-\bar p')^{-2} \le 4$ (Section B.2 (12)). It

remains to bound $\mathrm{Var}_{\mathrm{Unif}(\mathcal{T}_0)}[p_v]$ as in Lemma 12. Steps 4–5 are unchanged, with $\ell'_\star$ in place of ℓ'. In Step 3, $p_v^{(j')} \le \sum_i n/|D_i|$ with $|D_i| \ge \ell' - \varepsilon_i$, and $1/(\ell'-\varepsilon) - 1/\ell' \le \varepsilon/\ell'\ell'_\star$. So

$$\max p^{(j')} - \mathbb{E}[p^{(j')}] \le \frac{n_{c(j')} h(M_{j'} - \ell')}{\ell' M_{j'}} + \frac{n\Xi}{\ell'\ell'_\star},$$

and multiplying by $\mathbb{E}[p^{(j')}] = n_{c(j')} h/M_{j'}$ adds at most $n^2 h\Xi/M\ell'\ell'_\star \le c_4\, n^2 q\Xi/N^3$ to the variance of one coordinate, with $c_4 := \max_{j'} N^3/M_{j'}\ell'\ell'_\star$. Step 6 (Cauchy–Schwarz over w coordinates, $\mathrm{Var}(u-v) \le 2\,\mathrm{Var}\,u + 2\,\mathrm{Var}\,v$, $\chi^2 \le 4\,\mathrm{Var}$) multiplies this by $8w^2$, and with the factor 1.03 it becomes at most $11.3\, w^2 n^2 q\Xi/N^3$. The other constants of Lemma 12, re-evaluated with $\ell'_\star$ in place of ℓ', give the coefficients 36 and 43.4 of the statement (Section B.2 (12)). □

Proposition 5. *Assume* (3) *and* (13). *For every deterministic distinguisher,*

$$\mathrm{KL}(T_{\textsc{Real}^\perp} \| T_{\textsc{Hyb}^\perp}) \le \frac{K^\star_{\mathrm{RH}} q^4}{N^3} + \frac{(25.6\, b_1^2 + 6.4\, B_S\kappa_S)\, q}{N^2} + \frac{4.6\, nB_S^2\kappa_S\, q^2}{N^3} + 400\, q\beta_L^2.$$

Proof. By the chain rule and Lemma 23, the left side is the expected sum, over surviving real-supported prefixes, of $(1-\beta_\tau)\,\mathrm{KL}(R_G \| K_\star(\tau))$. We apply the per-type bounds of Lemmas 26, 27 and 29 (with $\mathrm{KL} \le \chi^2$), *retaining the factor* $1-\beta_\tau$ for quantities that depend on the current response. For types (2) and (5) nothing is revealed, so $R_G = R$, $\beta_\tau = 0$, and $K_\star$ is the corresponding old law, which is well defined at such prefixes. Type (3) splits into a value part and an anchor part by the chain rule. The anchor bound of Lemma 27(a) contains the coincidence charge $f(z) := \mathbb{E}_a[\kappa(a)^2]$, a nonnegative function of the value z just sampled. Since $(1-\beta_\tau)R_G(dz) = \mathbf{1}[z \in \mathcal{S}_\star]\, R(dz)$ for the value marginal, $(1-\beta_\tau)\mathbb{E}_{R_G}[f] \le \mathbb{E}_R[f \mid \tau]$. Averaging over the stopped real prefix, and coupling $\textsc{Real}^\perp$ with $\textsc{Real}$, the total is at most the full-real-game expectation of the same charge, so the slot sum of Lemma 13 applies with no loss from conditioning. The same holds for every other nonnegative response-dependent charge. Prefix-measurable bounds may discard the factor directly, and the game-wide quantities A_{tot}, Ξ, Ξ_2 and W dominate the local charges along the coupled full real execution.

Old terms. Each per-step bound of Lemmas 26, 27 and 29, without its new terms, is at most 1.6 times the corresponding bound in the proof of Proposition 1 (Section B.2 (13)). The game sums of Proposition 1, including Lemmas 13 and 14, therefore apply and give at most $1.6\,(C_{\mathrm{RH}}q^4 + C'_{\mathrm{RH}}q^3 + C''_{\mathrm{RH}}q^2)/N^3$.

New terms. Each step uses whole-game quantities, and there are at most q steps of each type, with $h \le q$. Summing the new terms of Lemmas 26, 27 and 29 over the steps and inserting the bounds of Lemma 25 gives the following contributions:

- construction values: $q(25.6\, b_1^2 + 6.4 B_S\kappa_S)/N^2$;
- types (2)/(3): a constant multiple of $F^2 B_S^2 q^3/N^3$, from $\mathbb{E}[A_{\mathrm{tot}}\Xi]$;
- type (5): constant multiples of $n^2 B_S^2 q^3/N^3$ and $n^2 B_S q^3/N^3$, from $\mathbb{E}[A_{\mathrm{tot}}\Xi]$ and $\mathbb{E}[\Xi]$;
- type (4): constant multiples of $nB_S^2\kappa_S q^2/N^3$, from $\mathbb{E}[\Xi_2]$; of $nB_S^2 q^3/N^3$ and $nB_S^2\kappa_S q^2/N^3$, from $\mathbb{E}[\Xi^2]$ and $n/N \le 1/128$; and of $w^2 n^2 B_S q^3/N^3$, from $\mathbb{E}[\Xi]$; together with $400 q\beta_L^2$.

The q^3 coefficients are at most those in the numerator of $K^\star_{\mathrm{RH}}$, and the κ_S-terms add up to at most $4.6\, nB_S^2\kappa_S q^2/N^3$ (Section B.2 (13)). With $q^3 \le q^4/2w$ the q^3 terms join $K^\star_{\mathrm{RH}}$. The old q^3, q^2 terms are absorbed as in the proof of Theorem 1, giving $1.6\,K_{\mathrm{RH}}$. □

7.6 Hybrid to ideal

Proposition 6. *Under* (3) *and* (13), $\mathrm{KL}(T_{\mathrm{HYB}_\star} \| T_{\mathrm{IDEAL}_\star}) \le K^\star_{\mathrm{HI}} q^4/N^3 + 143\, q\beta_L^2$.

Proof. Everything in Section 5.5 is stated for a table with values in a fixed set, here $\mathcal{S}_\star$. We check the inputs:

- *Forbidden density.* The second alternative (F2) of Section 4.1 is a union of at most b_0 sets, each of relative density exactly $1/N$ in G^w, for every pair, regular or singular. For a pair $\{a, a'\}$ and target call j, the equation $\mu_{aa'}(\lambda_j t - y_j) = c_{aa'}(y)$ contains a free coordinate with an automorphism coefficient: some y_s, $s \notin \{0, j\}$, if $\{a, a'\} \neq \{0, j\}$, and y_j itself if $\{a, a'\} = \{0, j\}$, whose coefficient is then $\lambda_j^{-1} + \mu_{0j} = \lambda_0^{-1}$. Outside $\mathcal{S}$ the full set $\mathrm{Forb}_{t,j}$ also contains the rigid hyperplanes; rigid and non-rigid pair constraints together number at most b_0, and only $\mathcal{S}_\star \cap \mathrm{Forb}_{t,j}$ is used below. Hence $|\mathrm{Forb}_{t,j}| \le \phi_0 N^w$ as in Section 4.1, and $\phi_{\mathcal{S}_\star} := \phi_0 N^w/|\mathcal{S}_\star| \le \frac{64}{61}\phi_0$.
- *Lemma 15* is generic (it counts illegal inputs).
- *Lemma 16* holds with $\mathcal{S}_\star$ in place of $\mathcal{S}$ and the same constants A, *34*, *44*. Its Step 2 uses $\mathbb{E}[S^2] \le \max S \cdot \mathbb{E}[S]$, valid even if all forbidden sets coincide, as they may for κ_p targets in one coset of a singular kernel. The constants *34* and *44* survive the replacement of $\frac{32}{31}$ by $\frac{64}{61}$ (Section B.2 (14)).
- *Failure.* As in Proposition 2, with $e^{1/16}\phi_{\mathcal{S}_\star} N \le 1.68\, wnq$; the per-trial bound $4nwq/N$ is unchanged.
- *Inverse output.* In $\mathrm{IDEAL}_\star$, the untilted acceptance probability of a fresh family input is $\rho'_j := |\mathcal{S}_\star \setminus \mathrm{Forb}_{t,j}|/|\mathcal{S}_\star|$, and Lemma 16 gives $|\theta_j - \rho'_j| \le 2A\phi_{\mathcal{S}_\star} \le \delta$ with δ of Lemma 17. $\mathrm{HYB}_\star$ uses $\rho''_j = \rho'_j|\mathcal{S}_\star|/|\mathcal{S}|$, so $|\theta_j - \rho''_j| \le \delta + \frac{32}{31}\beta_L =: \delta_\star$. The proof of Lemma 17 goes through with ρ_b replaced by $\rho''_j \le 1$ in the fresh-family weights. Free and legal construction-only weights agree. The lower bounds on N_{HYB}, θ_j and Z in Lemma 17 weaken only by the extra terms $1.38\beta_L N$ and $\frac{32}{31}\beta_L$, and the two sums together are still at most $12\,\delta_\star^2$. Splitting $\delta_\star^2$ by Young's inequality and using $11\delta^2 \le 56\, n^3 wq^3/N^3$ gives (Section B.2 (14))
$$\chi^2 \le 12\delta_\star^2 \le 67.2\,\frac{n^3 wq^3}{N^3} + 143\beta_L^2.$$
The companion is Lemma 16 with $V = \mathcal{S}_\star \setminus \mathrm{Forb}_{t,j}$: $44n^3wq^3/N^3$.

Summing as in the proof of Theorem 1, each response contributes at most $111.2n^3wq^3/N^3 + 2(4nwq/N)^3 + 143\beta_L^2$. □

7.7 Lower bounds and the meaning of tightness

Proposition 7. *Suppose some same-class pair (j, j') has $\kappa := |\ker \mu_{jj'}| \ge 2$, and let $N \ge 16$ and $q \le \min(|\mathcal{D}|, N/2)$. There is a non-adaptive distinguisher making q construction queries and no primitive query whose advantage against every simulator satisfies*
$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\mathrm{SIM}}(q) \ge \max\left\{\frac{q}{4N}\,\mathbf{1}\left[q \le \frac{\kappa^2}{2N}\right],\ \frac{q\kappa}{24N^2 + q\kappa}\right\}.$$

Proof. Let $H := \mathrm{Im}(\lambda_j - \lambda_{j'})$, $h := N/\kappa = |H|$ and $Z(x) := \mathrm{F}(x)_j - \mathrm{F}(x)_{j'}$ (with $\mathrm{F}(x)_0 := 0$). Query distinct $x_1, \ldots, x_q$ and let $S := \#\{i : Z(x_i) \in H\}$. Since $Z = \lambda_j(p_j - p_{j'}) + (\lambda_j - \lambda_{j'})(p_{j'})$, $Z \in H$

iff $p_j - p_{j'} \in H' := \lambda_j^{-1}(H)$, a subgroup of order h. In the ideal world $S \sim \mathrm{Bin}(q, 1/\kappa)$. In $\textsc{Real}$ the $2q$ raw outputs at calls j, j' form a uniform injection. Each indicator has mean $p = (h-1)/(N-1)$, so $\Delta := q(1/\kappa - p) = q(\kappa-1)/\kappa(N-1) \geq q/2N$. Counting ordered pairs with difference in $H' \setminus \{0\}$ that avoid two given values gives the exact covariance $2(h-1)(N-h)/(N-1)^2(N-2)(N-3) \leq 4/\kappa N^2$ for $N \geq 12$, so $\mathrm{Var}_{\textsc{Real}}\, S \leq 2q/\kappa$. We use two tests on S:

- *Rare-event test.* Output "ideal" iff $S \geq 1$. Then $\Pr_{\mathrm{ideal}}[S \geq 1] \geq q/\kappa - q^2/2\kappa^2$ and $\Pr_{\textsc{Real}}[S \geq 1] \leq qp$, so the advantage is at least $q/2N - q^2/2\kappa^2 \geq q/4N$ when $q \leq \kappa^2/2N$.
- *Second-moment test.* For any statistic f, $\mathrm{TV}(P, Q) \geq \Delta_f^2/(2(\mathrm{Var}_P f + \mathrm{Var}_Q f) + \Delta_f^2)$, by Cauchy–Schwarz with $|dP - dQ| \leq dP + dQ$. With $f = S$ this gives $q\kappa/(24N^2 + q\kappa)$. □

The hypotheses $N \geq 16$ and $q \leq \min(|\mathcal{D}|, N/2)$ are part of the statement. Section 8.1 discusses the same statistic from the side of pseudorandomness.

Remark 7. Theorem 2 and Corollary 5 hold for arbitrary automorphism masks (the proof in Section D.2 does not use admissibility). Hence, at fixed width and class sizes and positive construction-domain density, the exponent $3/4$ of Theorem 3 for $\textsc{Sim}_3$ is optimal, for every family, admissible or not. The same attack limits security against each fixed $\textsc{Sim}_m$ to $O(N^{3/4})$ queries; without rigid pairs this limit is uniform in m, including unlimited retries (below). No upper security theorem for general m is asserted. As for retry counts, the attack's lower bound contains a term $4mr|\Lambda|/N$ from the coupling to a structurally good oracle, and we distinguish two cases:

- Without rigid pairs, $\Lambda = \varnothing$. The bound is then uniform in m, and the analysis of the unlimited-retry simulator also covers $m = \infty$. An example is the non-admissible width-one family over $\mathbb{F}_2^d$ with masks I and $I + J$, where J has rank one and $J^2 = 0$ ($\kappa = N/2$), on a domain of density $1/2$.
- With rigid pairs, the bound holds for each fixed m but is not uniform in arbitrarily large m. That would need a sharper coupling, and none is claimed.

This tightness concerns the family $\{\textsc{Sim}_m\}$ and the wrapped simulators already treated. It is not a claim about every simulator that might be called coin-based or uniform; such a class would have to be defined before a broader claim is made.

Status of the results. For arbitrary masks the status is as follows:

- *Established:* the upper exponent $3/4$ for all masks (Theorem 3); the sharper formula of Theorem 1 for admissible masks; a matching query exponent for $\textsc{Sim}_3$ on positive-density families, and an $O(N^{3/4})$ query limit for each fixed $\textsc{Sim}_m$, uniform in m when there are no rigid pairs (Corollary 5 and Remark 7).
- *Not established:* a matching advantage formula uniformly in q, κ and w. In particular, $\sqrt{q\kappa_S}/N$ is not matched by Proposition 7, whose second bound is of order $q\kappa/N^2$. The linear bound $q/4N$ matches the large-pair term only when $q \leq \kappa^2/2N$. A pair just above τ_0 need not satisfy this throughout the regime, although it does at the $N^{3/4}$ scale for fixed width and large N (for example when $\kappa = N/2$). Also not established: a uniform all-retry attack for families with rigid pairs, and any construction-wide impossibility result.

7.8 Proof of Theorem 3

Insert Lemma 24, (15) and Propositions 5 and 6 into (14), and use $\sqrt{\sum x_k} \leq \sum \sqrt{x_k}$. The linear terms are $2.8q\beta_L + 2.8q\beta_L + 4q\beta_L = 9.6q\beta_L$ and $2w(w+1)q/N$. The terms $K_{\mathrm{RH}}^\star q^4/N^3$ and $K_{\mathrm{HI}}^\star q^4/N^3$ give $C_\star q^2/N^{3/2}$. Halving the remaining terms of Propositions 5 and 6 and taking square roots gives the terms of Theorem 3 in b_1, in B_S and κ_S, and in $\beta_L\sqrt{q}$ (Section B.2 (15)). The support conditions are those of Remark 3,

masks	PRF upper bound	PRF lower bound	source
rigid-free	$C_w q/N^{3/2}$, $q \le c_w N$	—	Theorem 4, [BGGJ26]
admissible, with a rigid pair	$O_w(q/N + q^2/N^{3/2})$	$q/2N$	Corollary 7
non-admissible, $\kappa \ge 2$	exponent $3/4$ (Theorem 3)	Proposition 7	Section 7

Table 4: Pseudorandomness bounds by type of mask. The first two rows are explained in the paragraphs on rigid-free families and rigid pairs at the start of Section 8, and the last row in Section 8.1.

applied to non-bad responses, together with the common $\perp$ atom of Lemma 23. The weights ρ''_j are positive because $\beta_j^K \le 0.022$. □

8 Pseudorandomness

A pseudorandomness distinguisher is an indifferentiability distinguisher that never calls the primitives, so $\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \le \mathbf{Adv}^{\mathrm{indiff}}_{\mathrm{F},\mathrm{SIM}}(q)$ (Section 2.1). The indifferentiability theorems therefore already settle the pseudorandomness of every family up to about $N^{3/4}$ queries. This section records that consequence, compares it with the Fourier bound for rigid-free families proved in Appendix E, and discusses what changes for non-admissible masks.

Corollary 7. *Under the hypotheses of Theorem 1, for every classical distinguisher making at most q construction queries,*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \le 10\, n^{3/2}(w+1)^2 \frac{q^2}{N^{3/2}} + 2w(w+1)\frac{q}{N} + 3w^2\frac{\sqrt{q}}{N},$$

where the second term is present only if some same-class pair is rigid, and the third only if some same-class pair is non-rigid. For arbitrary automorphism masks, the bound of Theorem 3 holds in the same way under its hypotheses.

Proof. Restrict Theorem 1 or Theorem 3 to distinguishers without primitive queries. □

Rigid pairs. If some same-class pair (j, j') is rigid, then $\mathrm{F}(x)_j \ne \mathrm{F}(x)_{j'}$ always in $\textsc{Real}$ (with $\mathrm{F}(x)_0 := 0$). Testing this on $q \le \min(|\mathcal{D}|, N)$ distinct inputs gives advantage $1-(1-1/N)^q \ge q/2N$. So Corollary 7 is of the right order, $\Theta_w(q/N)$, for $q \lesssim N^{1/2}$. For XORP with one permutation, sharper bounds of order q/N are known [BN18b]; for the whole family no such bound is proved here.

Rigid-free families. For rigid-free families (admissible with no rigid pairs, Definition 5; a non-admissible family without rigid pairs is *not* rigid-free), the Fourier method of [BGGJ26] improves the exponent from $3/4$ to essentially 1: $\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \le C_w q/N^{3/2}$ for $q \le c_w N$, with explicit constants and over every finite abelian group (Theorem 4). The proof is in Appendix E. Table 4 summarises the pseudorandomness bounds by type of mask. Prior pseudorandomness results for named instances are [Ebe17, BGGJ26] for $\pi_1 + \pi_2$, [Din25, BGGJ26] for LXoP, and [Iwa06, BN18b, BDD+26] for XORP.

8.1 Non-admissible masks

Recall from Definition 21 that a same-class pair $p = \{j, j'\}$ is *singular* if $\mu_{jj'}$ is neither 0 nor an automorphism, with kernel size $\kappa_p := |\ker \mu_{jj'}| = |\ker(\lambda_j - \lambda_{j'})|$; rigid pairs have $\kappa_p = N$ and regular pairs $\kappa_p = 1$. Theorem 4 does not apply to families with a singular pair, while Theorem 3 gives the indifferentiability

exponent $3/4$ for them. For pseudorandomness beyond $N^{3/4}$ queries, the masks matter: security degrades with $\varkappa$, interpolating between the regular and the rigid case.

Proposition 8. *Let (j, j') be a same-class pair, $\varkappa := \varkappa_{jj'}$, $H := \mathrm{Im}(\lambda_j - \lambda_{j'})$, so $|H| = N/\varkappa$, and $Z(x) := \mathrm{F}(x)_j - \mathrm{F}(x)_{j'}$ with $\mathrm{F}(x)_0 := 0$. In* Real*, for every x and $z \in G$,*

$$\Pr[Z(x) = z] = \frac{N - \varkappa \mathbf{1}[z \in H]}{N(N-1)},$$

while $Z(x)$ is uniform in the ideal world. Consequently the law ν of $Z(x)$ satisfies $\mathrm{TV}(\nu, \mathrm{Unif}(G)) = \frac{\varkappa-1}{\varkappa(N-1)}$ and $\chi^2(\nu \| \mathrm{Unif}(G)) = \frac{\varkappa-1}{(N-1)^2}$, and $\Pr[Z(x) \in H] = \frac{1}{\varkappa} - \frac{\varkappa-1}{\varkappa(N-1)}$.

Proof. $Z(x) = \lambda_j(p_j) - \lambda_{j'}(p_{j'})$, where $(p_j, p_{j'})$ is a uniform pair of distinct elements. Given $p_{j'}$, the value $Z = z$ forces $p_j = \lambda_j^{-1}(z + \lambda_{j'}(p_{j'}))$, and this differs from $p_{j'}$ iff $z \neq (\lambda_j - \lambda_{j'})(p_{j'})$. Exactly $\varkappa \mathbf{1}[z \in H]$ values of $p_{j'}$ violate this. The remaining statements are direct summations over H and its complement. □

For $\varkappa = N$ this is the rigid case, $Z \neq 0$. For $\varkappa = 1$, Z is exactly uniform. In between, Z avoids the subgroup H slightly: each query shifts the expected number of hits in H by about $1/N$, against a standard deviation of about $\varkappa^{-1/2}$. Counting these hits gives the construction-only lower bounds of Proposition 7, which hold against every simulator and hence also for pseudorandomness. Their first bound shows that a pair with $\varkappa \gtrsim \sqrt{qN}$ is no better than a rigid pair. For example, over $\mathbb{F}_2^d$ take $w = 1$, one permutation and $\lambda_1 = I + J$ with $J \neq 0$, $J^2 = 0$ of rank one (a transvection; this is the non-admissible $\mathrm{LXoP}[\sigma, 1]$ with $\sigma = I + J$). Then $\varkappa = N/2$, and the advantage is $\Omega(q/N)$ for $q \leq N/8$.

The second bound of Proposition 7 is the square of the natural order $\sqrt{q\varkappa}/N$, which is what one expects from a normal approximation of the hit count: shift about q/N against standard deviation about $\sqrt{q/\varkappa}$. Proving $\mathbf{Adv} \geq c \min(q, \sqrt{q\varkappa})/N$ in the intermediate range needs a Berry–Esseen-type estimate for the hit count in Real; the exact covariance in the proof of Proposition 7 is the main input. Already for $\varkappa = 2$, $\sqrt{q}/N$ exceeds the rigid-free order $q/N^{3/2}$ for every $q \leq N$. So, if that estimate holds, a single singular pair with the smallest possible kernel defeats the bound of Theorem 4.

Upper bounds beyond $N^{3/4}$. Level one of the Fourier expansion can be handled for all masks.

Lemma 30. *For every family, with Ω_t as in Appendix E and ν_1 the law of $\mathrm{F}(x)$ in* Real*, every classical distinguisher with at most q construction queries satisfies*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \leq \min\Big\{2q_* \,\mathrm{TV}\big(\nu_1, \mathrm{Unif}(G^w)\big),\ \tfrac{1}{2}\sqrt{q_*\Omega_1}\Big\} + \frac{1}{2}\Big(\sum_{t=2}^{q_*} \binom{q_*}{t} \Omega_t\Big)^{1/2},$$

$$\mathrm{TV}\big(\nu_1, \mathrm{Unif}(G^w)\big) \leq \sum_b \binom{n_b}{2} \frac{1}{N}.$$

Proof. As in Lemma 46, reduce to a fixed set Q of q_* inputs and write $\mathbf{Adv} = |\langle a, \xi - 1\rangle|$. Split the Fourier expansion into level one and the rest, as in [BGGJ26, Lemma 1] with $r = 1$. The rest is bounded by Cauchy–Schwarz and $\|a - \mathbb{E}[a]\|_2 \leq \frac{1}{2}$. The level-one part is $\sum_{x \in Q} \mathbb{E}_{\mathrm{Unif}}[(\Pi_x a)(\nu_1/\mathrm{Unif} - 1)(y_x)]$ with $|\Pi_x a| \leq 1$, hence at most $2q_* \,\mathrm{TV}(\nu_1, \mathrm{Unif})$; alternatively, Cauchy–Schwarz gives $\frac{1}{2}\sqrt{q_*\Omega_1}$. For the last claim, independent uniform raw outputs make $\mathrm{F}(x)$ exactly uniform, and a uniform injection of n_b values is within $\binom{n_b}{2}/N$ of independent uniform values. □

So the first term is $O_w(q/N)$ for every family, by the total-variation route. Let $\kappa_{\max}$ be the largest same-class kernel size. If $\Omega_1 = O_w(\kappa_{\max}/N^2)$, the first term would be $O_w(\min\{q, \sqrt{q\kappa_{\max}}\}/N)$, matching Proposition 7 in shape. This is shown only for the support-two part of Ω_1, which is exactly $\sum_{(j,j')}(\kappa_{jj'} - 1)/(N-1)^2$ by Proposition 8; the larger supports still need to be bounded for singular pairs. What is missing is control of the levels $t \geq 2$.

Remark 8. In the merging recursion of Appendix E, a merge inside one block at a singular pair loses injectivity (Lemma 47). The child map then has fibres of size at most $\kappa_{\max}$, because $(C_j - C_i)\beta_j$ determines β_j up to $\ker \mu^*_{ji}$. Every other step is unchanged, which suggests

$$\Omega_t \leq \kappa_{\max}^{\lceil t/2 \rceil} \cdot 2^{(w+1)t+d+w+1}\left(\frac{k_{\max}}{N - k_{\max}}\right)^{t+d} \qquad (t \geq 2).$$

If the display holds, the series in Lemma 30 converges for $q \leq c_w N/\kappa_{\max}^{1/3}$, and its levels $t \geq 2$ contribute $O_w(q\sqrt{\kappa_{\max}}/N^{3/2})$, which is dominated by the first term. Hence we conjecture that, for every family and $q \leq c_w N/\kappa_{\max}^{1/3}$,

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) = O_w\left(\frac{\min\{q, \sqrt{q\kappa_{\max}}\}}{N}\right) \quad (\kappa_{\max} \geq 2), \qquad \mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) = O_w\left(\frac{q}{N^{3/2}}\right) \quad (\kappa_{\max} = 1).$$

In particular $\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) = O_w(q/N)$ for every choice of masks when $q \leq c_w N^{2/3}$. By Proposition 7, this worst case is attained by non-admissible families (for example $\kappa = N/2$ above), as well as by rigid ones. Nothing in this remark is proved.

9 Conclusions

We introduced the family LCP of linear combinations of masked permutation outputs, and studied its subfamily GXP, in which every output block is the difference of two masked calls. GXP contains XORP with one or several permutations, LXoP and the sum of two permutations. The fork constructions are not members, but they lie in the class GXP$^+$ of constructions that reduce to GXP exactly, and every result for an instance carries over to the members of GXP$^+$ built on it (Proposition 3).

For indifferentiability, every member of GXP, with arbitrary automorphism masks (with the width condition (13) when some mask difference has a large kernel), is indifferentiable from a random function up to about $N^{3/4}$ queries, with the coin simulator $\textsc{Sim}_3$ (Theorem 3); for admissible families, Theorem 1 gives a sharper bound. The acceptance coin cannot be omitted in a theorem that covers free inputs: without it, the simulator is distinguishable at about $N^{2/3}$ queries on single-permutation families with a constant fraction of free inputs (Proposition 4). For fixed width and class sizes, the exponent $3/4$ is optimal for $\textsc{Sim}_3$, for every family with domain density bounded below and for the fork constructions (Corollaries 5 and 6 and Remark 7).

For pseudorandomness, every rigid-free member of GXP has advantage $O_w(q/N^{3/2})$ for $q \leq c_w N$, over every finite abelian group (Theorem 4). Rigid-freeness is the exact condition under which the Fourier-analytic recursion of [BGGJ26] applies. Families with a rigid pair are limited to about q/N, and non-admissible masks can also be strictly worse than $q/N^{3/2}$ (Section 8.1). So pseudorandomness is governed by rigid-freeness, a strengthening of admissibility, while neither rigid pairs nor non-admissible masks affect the indifferentiability exponent.

9.1 Open problems

We list the questions that this work leaves open:

- **Parameters of the matching attack.** Section 6.2 shows that the exponent $3/4$ is tight for $\textsc{Sim}_3$, but not the dependence on the parameters. For the forks, the attack needs about $170(w+1)^2$ times more queries than the upper bound allows. It is open whether the factor $n^{3/2}(w+1)^2$ of Theorem 1 reflects the construction or the analysis. Families whose construction domain has vanishing density are not covered by the attack. Its explicit constants also need very large N, roughly $N \geq 2^{75}$.

- **More trials and other simulators.** With $m \leq 2$ trials the coin simulator is already distinguishable at about $N^{m/(m+1)}$ queries on the full-domain two-permutation instance (Section 6.2), and for $m \geq 4$ no upper bound is proved. The two-trial simulator of [BDD+26] is not analysed for the general family. Kim's Gyroscope simulator [Kim26] reaches about $N^{4/5}$ queries for $\pi_1 + \pi_2$; whether it extends to the family is open.

- **Lower-order terms for non-admissible masks.** Theorem 3 adds a term of order $\sqrt{q\kappa_S}/N$ for small singular kernels, while the construction-only lower bound of Proposition 7 is only of order $q\kappa/N^2$ in that range. Closing the gap needs a normal approximation of the hit count. The extra hypothesis (13) for large singular kernels, a running-time bound for $\textsc{Sim}_3$ with arbitrary masks (Remark 6), and a matching attack uniform in the number of trials for families with rigid pairs are also open.

- **Pseudorandomness for non-admissible masks.** Remark 8 conjectures that the advantage is $O_w(\min\{q, \sqrt{q\kappa_{\max}}\}/N)$. For the level-one part of this bound (Lemma 30), only the $O_w(q/N)$ form is proved. The $\sqrt{q\kappa_{\max}}$ form also needs a bound on the larger supports of Ω_1.

- **Pseudorandomness with rigid pairs.** Families with a rigid pair admit an $\Omega(q/N)$ attack. An $O_w(q/N)$ upper bound for the whole family, extending the known bounds for XORP [BN18b], is not proved here.

- **Sharper constants.** The constants of Theorem 4 grow exponentially in w because of crude pattern counting. The constant 10 of Theorem 1 has not been optimised either.

- **Quantum queries.** The method of [BGGJ26] also yields quantum bounds through Fourier interpolation. We have not pursued them for the general family.

AI Disclosure. During the preparation of this manuscript, large language models, especially Claude Opus 5.5, were used in several sections to work out technical details and numerical derivations as well as to draft the body text. Other models like Claude Fable 5.1 and GPT-6 Astra were used to check the technical content and identify gaps in previous versions of the proofs.

References

[AEG+26] Elena Andreeva, Maria Eichlseder, Simon Gerhalter, Marcel Nageler, Oliver Christoph Popa, and Andreas Weninger. Forkcipher-based committing security from explicit collision security assumptions. Cryptology ePrint Archive, Paper 2026/2110, 2026. *(Cited on p. 11.)*

[BBB+26] Carsten Baum, Ward Beullens, Lennart Braun, Cyprien Delpech de Saint Guilhem, Michael Klooß, Christian Majenz, Shibam Mukherjee, Emmanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl. Shorter, tighter, FAESTer: Optimizations and improved (QROM) analysis for VOLE-in-the-head signatures. Cryptology ePrint Archive, Paper 2026/164, 2026. *(Cited on p. 5.)*

[BCEJ23] Ritam Bhaumik, Benoît Cogliati, Jordan Ethan, and Ashwin Jha. On quantum secure compressing pseudorandom functions. In *Advances in Cryptology – ASIACRYPT 2023*, volume 14440 of *Lecture Notes in Computer Science*, pages 34–66. Springer, 2023. *(Cited on p. 15.)*

[BCG+12] Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knezevic, Lars R. Knudsen, Gregor Leander, Ventzislav Nikov, Christof Paar, Christian Rechberger, Peter Rombouts, Søren S. Thomsen, and Tolga Yalçın. PRINCE – a low-latency block cipher for pervasive computing applications. In *Advances in Cryptology – ASIACRYPT 2012*, Lecture Notes in Computer Science, pages 208–225. Springer, 2012. *(Cited on p. 4.)*

[BDD+26] Ritam Bhaumik, Nilanjan Datta, Avijit Dutta, Ashwin Jha, Sougata Mandal, Bart Mennink, Hrithik Nandi, and Yaobin Shen. How to build a short-input random oracle from public random permutations. In *Advances in Cryptology – EUROCRYPT 2026, Part VI*, volume 16546 of *Lecture Notes in Computer Science*, pages 394–423. Springer, 2026. Full version: Cryptology ePrint Archive, Paper 2026/336; section and theorem numbers refer to it. *(Cited on pp. 4, 5, 8, 10, 15, 18, 19, 22, 23, 48, 51, 70, 73.)*

[BGGJ26] Ritam Bhaumik, Chun Guo, Xiaoning Guo, and Ashwin Jha. Indistinguishability of sum of permutations: A Fourier analytic route to classical and quantum security. Cryptology ePrint Archive, Paper 2026/2230, 2026. *(Cited on pp. 4, 7, 8, 9, 10, 14, 18, 48, 70, 71, 72, 73, 109, 111, 113, 115.)*

[BI99] Mihir Bellare and Russell 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. *(Cited on p. 4.)*

[BJK+16] Christof Beierle, Jérémy Jean, Stefan Kölbl, Gregor Leander, Amir Moradi, Thomas Peyrin, Yu Sasaki, Pascal Sasdrich, and Siang Meng Sim. The SKINNY family of block ciphers and its low-latency variant MANTIS. In *Advances in Cryptology – CRYPTO 2016*, Lecture Notes in Computer Science, pages 123–153. Springer, 2016. *(Cited on p. 4.)*

[BKL+07] Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: An ultra-lightweight block cipher. In *Cryptographic Hardware and Embedded Systems – CHES 2007*, Lecture Notes in Computer Science, pages 450–466. Springer, 2007. *(Cited on p. 4.)*

[BKR98] Mihir Bellare, Ted Krovetz, and Phillip Rogaway. Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible. In *Advances in Cryptology – EUROCRYPT '98*, volume 1403 of *Lecture Notes in Computer Science*, pages 266–280. Springer, 1998. *(Cited on pp. 4, 18.)*

[BL16] Karthikeyan Bhargavan and Gaëtan Leurent. On the practical (in-)security of 64-bit block ciphers. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS 2016*, pages 456–467. ACM, 2016. *(Cited on p. 4.)*

[BN18a] Srimanta Bhattacharya and Mridul Nandi. Full indifferentiable security of the xor of two or more random permutations using the χ^2 method. In *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 387–412. Springer, 2018. *(Cited on p. 5.)*

[BN18b] Srimanta Bhattacharya and Mridul Nandi. Revisiting variable output length XOR pseudorandom function. *IACR Transactions on Symmetric Cryptology*, 2018(1):314–335, 2018. *(Cited on pp. 11, 70, 73.)*

[BPP+17] Subhadeep Banik, Sumit Kumar Pandey, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, and Yosuke Todo. GIFT: A small present. In *Cryptographic Hardware and Embedded Systems – CHES 2017*, Lecture Notes in Computer Science, pages 321–345. Springer, 2017. *(Cited on p. 4.)*

[CDN+23] Benoît Cogliati, Avijit Dutta, Mridul Nandi, Jacques Patarin, and Abishanka Saha. Proof of mirror theory for a wide range of $\xi_{\max}$. In *Advances in Cryptology – EUROCRYPT 2023, Part IV*, volume 14007 of *Lecture Notes in Computer Science*, pages 470–501. Springer, 2023. *(Cited on p. 5.)*

[CHJ26] Wonseok Choi, Minki Hhan, and Junyoung Jang. Quantum security of XOR of permutations via Fourier analysis. Cryptology ePrint Archive, Paper 2026/2241, 2026. *(Cited on p. 10.)*

[CLL19] Wonseok Choi, Byeonghak Lee, and Jooyoung Lee. Indifferentiability of truncated random permutations. In *Advances in Cryptology – ASIACRYPT 2019, Part I*, volume 11921 of *Lecture Notes in Computer Science*, pages 175–195. Springer, 2019. *(Cited on p. 10.)*

[CLM19] Yu Long Chen, Eran Lambooij, and Bart Mennink. How to build pseudorandom functions from public random permutations. In *Advances in Cryptology – CRYPTO 2019*, Lecture Notes in Computer Science, pages 266–293. Springer, 2019. *(Cited on p. 10.)*

[CLP15] Benoit Cogliati, Rodolphe Lampe, and Jacques Patarin. The indistinguishability of the XOR of k permutations. In *Fast Software Encryption, FSE 2014*, volume 8540 of *Lecture Notes in Computer Science*, pages 285–302. Springer, 2015. *(Cited on p. 5.)*

[CMP21] Yu Long Chen, Bart Mennink, and Bart Preneel. Categorization of faulty nonce misuse resistant message authentication. In *Advances in Cryptology – ASIACRYPT 2021*, volume 13092 of *Lecture Notes in Computer Science*, pages 520–550. Springer, 2021. *(Cited on p. 15.)*

[CS16] Benoît Cogliati and Yannick Seurin. EWCDM: An efficient, beyond-birthday secure, nonce-misuse resistant MAC. In *Advances in Cryptology – CRYPTO 2016, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 121–149. Springer, 2016. *(Cited on p. 4.)*

[CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, second edition, 2006. *(Cited on p. 84.)*

[DDM+26] Nilanjan Datta, Avijit Dutta, Sougata Mandal, Hrithik Nandi, and Amlan Sinha. Post-quantum security of keyed sum of permutations and its siblings. Cryptology ePrint Archive, Paper 2026/426, 2026. *(Cited on p. 11.)*

[DDMN25] Nilanjan Datta, Avijit Dutta, Sougata Mandal, and Hrithik Nandi. Sequential indifferentiability of STH and EDM. Cryptology ePrint Archive, Paper 2025/1518, 2025. *(Cited on p. 5.)*

[DF04] David S. Dummit and Richard M. Foote. *Abstract Algebra*. Wiley, third edition, 2004. *(Cited on p. 86.)*

[DGL22] Avijit Dutta, Jian Guo, and Eik List. Forking sums of permutations for optimally secure and highly efficient PRFs. Cryptology ePrint Archive, Paper 2022/1609, 2022. *(Cited on pp. 4, 5, 10, 15, 18, 19, 49, 51, 52, 106.)*

[DHT17] Wei Dai, Viet Tung Hoang, and Stefano Tessaro. Information-theoretic indistinguishability via the chi-squared method. In *Advances in Cryptology – CRYPTO 2017*, volume 10403 of *Lecture Notes in Computer Science*, pages 497–523. Springer, 2017. *(Cited on pp. 5, 11.)*

[Din24] Itai Dinur. Tight indistinguishability bounds for the XOR of independent random permutations by Fourier analysis. In *Advances in Cryptology – EUROCRYPT 2024, Part I*, volume 14651 of *Lecture Notes in Computer Science*, pages 33–62. Springer, 2024. *(Cited on pp. 5, 10.)*

[Din25] Itai Dinur. Combining outputs of a random permutation: New constructions and tight security bounds by Fourier analysis. In *Advances in Cryptology – EUROCRYPT 2025, Part I*, volume 15601 of *Lecture Notes in Computer Science*, pages 244–273. Springer, 2025. *(Cited on pp. 4, 5, 9, 10, 15, 18, 48, 70, 109.)*

[DL25] Avijit Dutta and Eik List. Forking sums of permutations for highly secure and efficient PRFs. *Designs, Codes and Cryptography*, 93(10):4627–4658, 2025. *(Cited on pp. 5, 10.)*

[DNS22] Avijit Dutta, Mridul Nandi, and Abishanka Saha. Proof of mirror theory for $\xi_{\max} = 2$. *IEEE Transactions on Information Theory*, 68(9):6218–6232, 2022. *(Cited on p. 5.)*

[Ebe17] Sean Eberhard. More on additive triples of bijections. arXiv:1704.02407, 2017. *(Cited on pp. 5, 8, 9, 10, 14, 70, 109.)*

[FR26] Thibauld Feneuil and Matthieu Rivain. On the security of MPC-in-the-head signatures with correlated GGM trees. Cryptology ePrint Archive, Paper 2026/615, 2026. *(Cited on p. 5.)*

[GBJ+23] Aldo Gunsing, Ritam Bhaumik, Ashwin Jha, Bart Mennink, and Yaobin Shen. Revisiting the indifferentiability of the sum of permutations. In *Advances in Cryptology – CRYPTO 2023*, volume 14083 of *Lecture Notes in Computer Science*, pages 628–660. Springer, 2023. *(Cited on p. 5.)*

[Gun22] Aldo Gunsing. Block-cipher-based tree hashing. In *Advances in Cryptology – CRYPTO 2022*, volume 13510 of *Lecture Notes in Computer Science*, pages 205–233. Springer, 2022. *(Cited on p. 5.)*

[IMV16] Tetsu Iwata, Bart Mennink, and Damian Vizár. CENC is optimally secure. Cryptology ePrint Archive, Paper 2016/1087, 2016. *(Cited on p. 5.)*

[Iwa06] Tetsu Iwata. New blockcipher modes of operation with beyond the birthday bound security. In *Fast Software Encryption, FSE 2006*, volume 4047 of *Lecture Notes in Computer Science*, pages 310–327. Springer, 2006. *(Cited on pp. 4, 5, 15, 18, 70.)*

[Kim26] Sunyeop Kim. Indifferentiability of the sum of two permutations: Tight $3n/4$ security for the uniform simulator and a new simulator for $4n/5$ security. Cryptology ePrint Archive, Paper 2026/2092, 2026. *(Cited on pp. 4, 5, 7, 8, 10, 22, 23, 27, 30, 31, 49, 54, 73, 101, 103.)*

[Lee17] Jooyoung Lee. Indifferentiability of the sum of random permutations toward optimal security. *IEEE Transactions on Information Theory*, 63(6):4050–4054, 2017. *(Cited on p. 5.)*

[Luc00] Stefan Lucks. The sum of PRPs is a secure PRF. In *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 470–484. Springer, 2000. *(Cited on p. 4.)*

[Men25] Bart Mennink. Minimized PRFs from public permutations. *IACR Transactions on Symmetric Cryptology*, 2025(3):230–260, 2025. *(Cited on p. 11.)*

[MN17] Bart Mennink and Samuel Neves. Encrypted Davies-Meyer and its dual: Towards optimal security using mirror theory. In *Advances in Cryptology – CRYPTO 2017*, volume 10403 of *Lecture Notes in Computer Science*, pages 556–583. Springer, 2017. *(Cited on pp. 5, 51.)*

[MP15] Bart Mennink and Bart Preneel. On the XOR of multiple random permutations. In *Applied Cryptography and Network Security, ACNS 2015*, volume 9092 of *Lecture Notes in Computer Science*, pages 619–634. Springer, 2015. *(Cited on p. 5.)*

[MPN10] Avradip Mandal, Jacques Patarin, and Valérie Nachef. Indifferentiability beyond the birthday bound for the xor of two public random permutations. In *Progress in Cryptology – INDOCRYPT 2010*, volume 6498 of *Lecture Notes in Computer Science*, pages 69–81. Springer, 2010. *(Cited on p. 5.)*

[MRH04] Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *Theory of Cryptography Conference, TCC 2004*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2004. *(Cited on pp. 4, 5, 12.)*

[Pat08] Jacques Patarin. A proof of security in $O(2^n)$ for the xor of two random permutations. In *Information Theoretic Security, ICITS 2008*, volume 5155 of *Lecture Notes in Computer Science*, pages 232–248. Springer, 2008. *(Cited on p. 5.)*

[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. *(Cited on p. 5.)*

[RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indifferentiability framework. In *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 487–506. Springer, 2011. *(Cited on p. 5.)*

[TGL25] YanJin Tan, JunTao Gao, and XueLian Li. Post-quantum security of the sum of Even-Mansour. Cryptology ePrint Archive, Paper 2025/2274, 2025. *(Cited on p. 11.)*

Index of notation

The symbols below are listed by the part of the paper that introduces them, with the place of their definition. Symbols marked *local* have a different meaning, or no meaning, outside the indicated part.

Symbol	*Meaning*	*Defined in*
General		
G, N	a finite abelian group, written additively, and its order	Section 2, p. 11
$\widehat{G}$, $\mathrm{e}(\psi, g)$	the character group and the pairing	Section 2.3, p. 13
$\mathrm{Aut}(G)$, $\mathrm{End}(G)$	automorphisms and endomorphisms of G	Section 2.3, p. 13
$(a)_t$	falling factorial $a(a-1)\cdots(a-t+1)$	Section 2, p. 11
$\textsc{Real}$, $\textsc{Ideal}$, Γ	the two worlds, and the random function of the ideal world	Section 2.1, p. 12
q	the total number of queries of a distinguisher	Section 2.1, p. 12
Family (Section 3)		
$\mathsf{LCP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}; \mathbf{M}]$	linear combination of masked permutation outputs, with combiner $\mathbf{M}$	Definition 2, p. 15
$\mathsf{GXP}[G; \mathcal{D}; \mathbf{c}; \boldsymbol{\iota}; \boldsymbol{\lambda}]$, Δ_{w+1}	the subfamily with the difference combiner Δ_{w+1}	Section 3.2, p. 16
w; calls $j = 0, \ldots, w$	number of output blocks; the $w+1$ calls	Definition 3, p. 16
k, $c(j) = c_j$	number of permutations; the class map	Definition 3, p. 16
n_b, n	class sizes and $n := \max_b n_b$	Definition 3, p. 16
$\mathcal{D}$, ι_j	the construction domain and the input injections	Definition 3, p. 16
λ_j	the masks, automorphisms of G	Definition 3, p. 16
$p_j(x)$, $v_j(x)$	raw and masked outputs, $v_j = \lambda_j(p_j)$	Definition 3, p. 16
y, $y_0 := 0$; v_0	a construction value, with $v_j = v_0 + y_j$; the anchor	Definition 4, p. 16
Free_b	the free (unused) inputs of permutation b	Definition 4, p. 16
F, $\bar{F}$	$\sum_b n_b^2$ and $(w+1)n$	Section 3.2, p. 16
$\mu_{jj'}$	$\lambda_j^{-1} - \lambda_{j'}^{-1}$ for a same-class pair	(1), p. 17
Λ; $\mathcal{S}$	the set of rigid pairs; the structurally good values	Definitions 5 and 6, p. 17
Simulator (Section 4)		
R_b, $\tilde{R}_j$	ranges, and masked ranges $\lambda_j(R_{c(j)})$	Section 4.1, p. 19
$\mathrm{Leg}(y)$, $e(y)$	legal anchors; the collision correction	Definition 9, (2), p. 20
$\mathrm{Forb}_{t,j}$	values for which call j cannot have output t	Definition 10, p. 20
$s_{\max}$, ϕ_0, $\phi_{\mathcal{S}}$	a bound on the ranges; forbidden densities	Section 4.1, p. 19
ρ_b	the index-free weight, the coin probability	Lemma 4, p. 21
$\textsc{Sim}_m$; m	the coin simulator; its number of trials	Definition 11, p. 22
P_b, $\textsc{Dom}_b$, $\textsc{Ran}_b$, $\textsc{Failure}$	the simulator's tables and failure flag	Section 4.2, p. 22
ζ	the dummy value used to toss the coin	Section 4.2, p. 22
Indifferentiability proof (Section 5)		
r, h, u	numbers of revealed, construction-only and fresh inputs	Section 5.1, p. 25
f_b, f, $\bar{s}$	assigned free inputs; $\bar{s} := r + h + f$	Section 5.1, p. 25
Φ, ℓ	$\bar{F}\bar{s} + 2w^2$, a bound on excluded anchors; $\ell := N - \Phi$	Section 5.1, p. 25
$\textsc{Hyb}$; e_b, g_b, k_{leg}, $N_{\textsc{Hyb}}$	the hybrid; fresh inputs, fresh free inputs, legal pairs, normaliser	Section 5.3, p. 27

Symbol	*Meaning*	*Defined in*
U, Comp	hidden anchors of the construction-only inputs, and their compatible values	Section 5.4, p. 29
ω, $x^{(\omega)}$, $\mathcal{H}_\omega$	an exposure slot, its input, the history before it	Definition 14, p. 28
$\kappa(a)$; $V_{j'j''}$, W	coincidence counts; envelopes	Definition 15, (9), pp. 32 and 38
Ideal′	the ideal world with Γ uniform into $\mathcal{S}$	Section 5.5, p. 43
ϑ_c, I_c, R	trial factor, number of illegal inputs, rejection probability	Section 5.5, p. 43
μ_τ, π_τ, W_τ; θ_j	prior, posterior and weight of the table; tilted legality probability	Lemma 15, p. 43
K_{RH}, K_{HI}	the constants of the two comparisons	Theorem 1, p. 26
Arbitrary masks (Section 7)		
p, μ_p, κ_p	a same-class pair $\{a, a'\}$, its mask difference $\mu_{aa'}$, and $\lvert\ker \mu_p\rvert$	Section 7.1, p. 59
rigid, regular, singular	$\kappa_p = N$, $\kappa_p = 1$, and otherwise	Definition 21, p. 59
$c_p(y)$; active	the collision constant of p; p is active at y if $c_p(y) \in \operatorname{Im} \mu_p$	Definition 22, p. 59
b_0, τ_0	$\binom{w+1}{2}$; the threshold $N/64b_0$ between small and large singular pairs	Section 7.1, p. 59
$\mathcal{P}_S$, $\mathcal{P}_L$; B_S, B_L; κ_S	small and large singular pairs; their numbers; the largest small kernel (1 if none)	Section 7.1, p. 59
B_R; b_1	the number of regular pairs; $B_R + B_S$	Section 7.1, p. 59
β_L	$\sum_{p\in\mathcal{P}_L} 1/\kappa_p$, the density of values activating a large pair	Section 7.1, p. 59
$\Lambda_\star$, $\mathcal{S}_\star$; $e_S(y)$	bad values (a rigid or large singular pair active) and their complement; the singular excess	Definitions 23 and 24, p. 60
Hyb$_\star$, Real$^\perp$, Hyb$^\perp$, Ideal$_\star$; β_τ	the comparison worlds; the probability of a bad next response	Section 7.3, Definition 25, pp. 61 and 62
$A_{\omega,p}$, ε_ω	$p \in \mathcal{P}_S$ active at the value of slot ω; the excess of that value	Section 7.4, p. 63
A_{tot}, Ξ, Ξ_2	game-wide counts of small-pair activations, and the sums of ε_ω and ε_ω^2	Section 7.4, p. 63
Pseudorandomness (Appendix E)		
$\hat{\mu}_r(\beta)$, Θ_k	injection coefficients; the fixed-support moment	Section 2.3, p. 13
$T(\alpha)$, C_j	pullback of an output character; coefficients of the block relation	(20), p. 111
$\hat{\xi}_t$, Ω_t	level-t coefficients and their moment	Lemma 45, p. 111
q_*	$\min(q, \lvert\mathcal{D}\rvert)$	Appendix E, p. 109
A_w, K_w, C_w, c_w	the constants of the theorem	Proposition 9 and Theorem 4, pp. 113 and 114
Matching attack (Section 6.2; local)		
(H, g), L; P, I	a coset progression and its length; the target and doubled progressions	Definition 26, p. 88; Section 6.2, p. 54
r	$\lvert P\rvert$, the size of the target progression	Section 6.2, p. 54
c_1, c_2; A_1, A_2, A_3, $C^\cup$	separating shifts; target sets and the union of the test sets	Definition 19, p. 55
$D(y)$, E; dens	$y_{j'} - y_j$ and the test event; the domain density $n_b\lvert\mathcal{D}\rvert/N$	Section 6.2, p. 54
C_k^*, π_k, T, θ	test sets, baselines, the statistic and its threshold	Section 6.2, p. 56
ξ	the attack scale $\sqrt{n_b/\mathrm{dens}}\, N^{3/4}$	Corollary 5, p. 56

Glossary of terms

The terms below are introduced in this paper, or used here in a specific sense. Each entry points to the definition.

Term	*Meaning*	*Defined in*
acceptance coin	the simulator accepts a free candidate with probability ρ_b	Definition 11, p. 22
acceptance weight	1, 0 or ρ_b for a legal, illegal or free input	Definition 17, p. 43
active pair	a same-class pair whose collision constant lies in Im μ_p	Definition 22, p. 59
admissible	every same-class $\mu_{jj'}$ is 0 or an automorphism	Definition 5, p. 17
anchor	the masked output v_0, which fixes all outputs of an input	Definition 4, p. 16
augmented interface	primitive queries at a call return the whole tuple	Definition 1, p. 12
bad value	a value at which a rigid or large singular pair is active	Definition 23, p. 60
balanced	a family without free inputs	Definition 4, p. 16
base	the GXP instance that a construction in GXP$^+$ reduces to	Definition 8, p. 19
calls; tuple	the permutation inputs of one construction input; their outputs	Definitions 1 and 3, pp. 12 and 16
class; class map	the calls that use one permutation; the map c	Definition 3, p. 16
coincidence count	the number of coinciding affine images at an anchor	Definition 15, p. 32
combiner	the matrix $\mathbf{M}$ applied to the masked outputs	Definition 2, p. 15
companion tuple	the other outputs of an input returned by an inverse query	Definition 16, p. 35
coset progression	$H + \{0, g, \ldots, (L-1)g\}$ for a subgroup H	Definition 26, p. 88
exposure slot	the query that first samples the tuple of an input	Definition 14, p. 28
forbidden values	values for which a call cannot have a given output	Definition 10, p. 20
fork	a construction feeding one output into the other calls	Definition 7, p. 19
free input	a permutation input never used by the construction	Definition 4, p. 16
fresh, construction-only, revealed	the three kinds of construction input in the game	Definition 12, p. 25
index-free weight	the legality probability ρ_b, the same for every call	Lemma 4, p. 21
legal anchor; legal set	an anchor consistent with the ranges; the set $\mathrm{Leg}(y)$	Definition 9, p. 20
legal pair	a construction-only input and a call that may have the target	Definition 13, p. 27
raw, masked outputs	$p_j(x)$ and $v_j(x) = \lambda_j(p_j(x))$	Definition 3, p. 16
reducible	reducible to GXP by the moves defining GXP$^+$	Definition 8, p. 19
rigid pair; rigid-free	a same-class pair with equal masks; admissible with none	Definition 5, p. 17
same-class pair	two calls that use the same permutation	Section 3.3, p. 17
separating shifts	shifts keeping the target and test sets of the attack apart	Definition 19, p. 55
singular, regular pair	$1 < \kappa_p < N$; $\kappa_p = 1$	Definition 21, p. 59
singular excess	$\sum_p \kappa_p$ over small singular pairs active at y	Definition 24, p. 60
stopped hybrid	a world halted at the first bad response	Definition 25, p. 62
structurally good	a value with $y_j \neq y_{j'}$ for every rigid pair	Definition 6, p. 17
trials	the candidates drawn by one inverse query of $\textsc{Sim}_m$	Definition 11, p. 22
useful test	an iteration of the attack whose inverse query hits the progression	Definition 20, p. 56
wrapped simulator	the simulator for a construction in GXP$^+$, built from $\textsc{Sim}_m$	Definition 18, p. 51

Appendix

A Mathematical background

This appendix recalls the standard material that Section 2 assumes: probability and information theory in Section A.1, finite abelian groups in Section A.2, and Fourier analysis on finite abelian groups in Section A.3. None of it is new. It is included so that the paper is self-contained, and it also contains the proofs of the lemmas stated in Section 2. We assume familiarity with basic probability (random variables, expectation, independence), with complex numbers, and with the definition of a group.

A.1 Probability and information theory

Finite distributions. All distributions in this paper are defined on finite sets, so every probability below is a finite sum and no measure theory is needed. A *distribution* on a finite set Ω is a function $P : \Omega \to [0,1]$ with $\sum_{x\in\Omega} P(x) = 1$. Its *support* is $\mathrm{supp}(P) := \{x : P(x) > 0\}$, and the probability of an event $A \subseteq \Omega$ is $P(A) := \sum_{x\in A} P(x)$. We write $X \leftarrow P$ for a random variable X with distribution P, and say that X is *sampled from* P. For a function $f : \Omega \to \mathbb{R}$ we write

$$\mathbb{E}_P[f] := \sum_{x\in\Omega} P(x)f(x), \qquad \mathrm{Var}_P(f) := \mathbb{E}_P\big[(f - \mathbb{E}_P[f])^2\big] = \mathbb{E}_P[f^2] - (\mathbb{E}_P[f])^2.$$

If $X, Y \leftarrow P$ are independent, then expanding the square gives

$$\mathrm{Var}_P(f) = \tfrac{1}{2}\mathbb{E}\big[(f(X) - f(Y))^2\big]. \tag{17}$$

For an event A with $P(A) > 0$, the *conditional distribution* is $P(x \mid A) := P(x)/P(A)$ for $x \in A$. A *mixture* of distributions $P_1, \ldots, P_m$ with weights $w_i \geq 0$, $\sum_i w_i = 1$, is the distribution $\sum_i w_i P_i$: first pick i with probability w_i, then sample from P_i.

A random sequence $Z = (Z_1, \ldots, Z_q)$ can be generated one entry at a time: first Z_1, then Z_2 given Z_1, and so on. Accordingly, its distribution P is the product of the distributions of the successive entries, each given the earlier ones. For every $z = (z_1, \ldots, z_q)$,

$$P(z) = \prod_{i=1}^{q} P(z_i \mid z_1, \ldots, z_{i-1}), \tag{18}$$

where, for a prefix τ of positive probability, $P(\cdot \mid \tau)$ is the conditional distribution of the next entry given that the sequence starts with τ. We write $P^{(i)}$ for the distribution of the prefix $(Z_1, \ldots, Z_i)$. We also use the word *law* for the distribution of a random variable, as in "the law of the next response".

In a security game, a distinguisher interacts with a system of oracles, and its *transcript* is the sequence of its queries and of the responses it receives. A deterministic distinguisher chooses each query as a function of the responses it has received so far. So, for a fixed distinguisher, the transcript is determined by the sequence of responses, and the distribution of the transcript has the form (18): the i-th factor is the distribution of the system's response to the i-th query, given the earlier interaction. This is what lets us analyse a game one query at a time. Two systems that answer every query with the same conditional distribution produce the same transcript distribution against every distinguisher.

Total variation. The *total variation* between distributions P, Q on Ω is

$$\mathrm{TV}(P,Q) := \max_{A \subseteq \Omega} |P(A) - Q(A)| = \frac{1}{2} \sum_{x \in \Omega} |P(x) - Q(x)|.$$

To see that the two expressions agree, note that $P(A) - Q(A) = \sum_{x \in A}(P(x) - Q(x))$ is largest for $A = \{x : P(x) > Q(x)\}$. Since $\sum_x (P(x) - Q(x)) = 0$, the positive and the negative parts of $P - Q$ have the same total, so this maximum is half of $\sum_x |P(x) - Q(x)|$; the same holds for $Q(A) - P(A)$. The second expression is half the ℓ_1 distance, so TV is a metric and satisfies the triangle inequality.

Total variation measures how well two distributions can be told apart from a single sample. A test is an event A: it says "P" when the sample lies in A, and the difference $|P(A) - Q(A)|$ of the probabilities with which the two distributions trigger it is the advantage of the test. Total variation is the best advantage of any test, and this is how it enters cryptography.

Fact 1. *Let T_0, T_1 be the transcript laws of a deterministic distinguisher $\mathcal{A}$ in two worlds. Then* $|\Pr[\mathcal{A}^{(0)} = 1] - \Pr[\mathcal{A}^{(1)} = 1]| \le \mathrm{TV}(T_0, T_1)$.

Proof. The output of $\mathcal{A}$ is a function of its transcript, so $\mathcal{A}$ outputs 1 exactly when the transcript lies in some fixed set E. The difference of the two probabilities is $T_0(E) - T_1(E)$, which is at most $\mathrm{TV}(T_0, T_1)$ in absolute value by the definition of TV. □

Relative entropy and chi-square divergence. The *relative entropy* (or Kullback–Leibler divergence) and the *chi-square divergence* of P from Q are

$$\mathrm{KL}(P\|Q) := \sum_x P(x) \log \frac{P(x)}{Q(x)}, \qquad \chi^2(P\|Q) := \sum_x \frac{(P(x) - Q(x))^2}{Q(x)}.$$

Terms with $P(x) = Q(x) = 0$ are 0, and both divergences are $+\infty$ if $P(x) > 0 = Q(x)$ for some x. So they are finite exactly when $\mathrm{supp}(P) \subseteq \mathrm{supp}(Q)$, which we call the *support condition*. Neither is symmetric in P and Q, and neither is a metric. Both measure how surprising samples from P look to someone who expects Q. The relative entropy is the expected value, under P, of the log-likelihood ratio $\log(P/Q)$, that is, the average evidence that one sample from P provides against Q. The chi-square divergence measures the same discrepancy quadratically, and it is often easier to compute. Their use is that they behave well along sequences (Fact 6), while total variation does not, and that they control total variation (Fact 5).

Under the support condition, let $L := P/Q$ be the likelihood ratio on $\mathrm{supp}(Q)$. Then $\mathbb{E}_Q[L] = 1$ and

$$\chi^2(P\|Q) = \sum_x \frac{P(x)^2}{Q(x)} - 1 = \mathbb{E}_Q[(L-1)^2] = \mathrm{Var}_Q(L), \tag{19}$$

so the chi-square divergence is the variance of the likelihood ratio. This form is the one used in Lemma 1.

Fact 2. $0 \le \mathrm{KL}(P\|Q) \le \chi^2(P\|Q)$.

Proof. Both bounds use $\log t \le t - 1$ for $t > 0$. Assume the support condition, since otherwise $\chi^2 = \infty$. For the upper bound, $P(x) \log \frac{P(x)}{Q(x)} \le P(x)\left(\frac{P(x)}{Q(x)} - 1\right)$ for every x in $\mathrm{supp}(P)$; summing gives $\mathrm{KL}(P\|Q) \le \sum_x P(x)^2/Q(x) - 1 = \chi^2(P\|Q)$ by (19). For the lower bound, $-\mathrm{KL}(P\|Q) = \sum_x P(x) \log \frac{Q(x)}{P(x)} \le \sum_{x \in \mathrm{supp}(P)} (Q(x) - P(x)) \le 0$. □

Convexity and data processing. Both divergences have the form

$$D_\phi(P\|Q) = \sum_x Q(x)\,\phi\Big(\frac{P(x)}{Q(x)}\Big), \qquad \phi(t) = t\log t \ \text{(KL)}, \qquad \phi(t) = (t-1)^2 \ (\chi^2),$$

with the conventions above, for a convex function ϕ on $[0,\infty)$ with $\phi(1) = 0$. Two properties follow from convexity alone.

Fact 3. KL *and* χ^2 *are jointly convex: for distributions* P_i, Q_i *and weights* $w_i \geq 0$ *with* $\sum_i w_i = 1$,

$$D_\phi\Big(\sum_i w_i P_i \,\Big\|\, \sum_i w_i Q_i\Big) \leq \sum_i w_i\, D_\phi(P_i\|Q_i).$$

In particular, each is convex in its second argument for a fixed first argument.

Proof. It suffices to treat one point x. Write $p_i = P_i(x)$, $q_i = Q_i(x)$, $p = \sum_i w_i p_i$ and $q = \sum_i w_i q_i$, and assume all $q_i > 0$ (otherwise the right-hand side is infinite, or the corresponding terms vanish on both sides). Since $\frac{p}{q} = \sum_i \frac{w_i q_i}{q} \cdot \frac{p_i}{q_i}$ is an average of the ratios p_i/q_i with weights $w_i q_i/q$, Jensen's inequality for the convex ϕ gives $q\,\phi(p/q) \leq \sum_i w_i q_i\,\phi(p_i/q_i)$. □

Fact 4. *For any map* $h : \Omega \to \Omega'$, *let* P_h *and* Q_h *be the laws of* $h(X)$ *when* $X \leftarrow P$ *and* $X \leftarrow Q$. *Then* $D_\phi(P_h\|Q_h) \leq D_\phi(P\|Q)$.

Proof. Fix $o \in \Omega'$ and the fibre $F := h^{-1}(o)$, with $Q(F) > 0$ and $Q(x) > 0$ on F (otherwise the right-hand side is infinite, or the fibre contributes 0 to both sides). The ratio $P(F)/Q(F)$ is the average of the ratios $P(x)/Q(x)$, $x \in F$, with weights $Q(x)/Q(F)$. By Jensen's inequality, $Q(F)\,\phi(P(F)/Q(F)) \leq \sum_{x\in F} Q(x)\,\phi(P(x)/Q(x))$. Sum over o. □

Processing a sample, in particular forgetting part of it, can only make two distributions harder to tell apart. The same holds for total variation, by its first expression. Finally, a bound on the relative entropy gives a bound on the total variation. This is Pinsker's inequality.

Fact 5 (see [CT06, Lemma 11.6.1])**.** $\mathrm{TV}(P,Q) \leq \sqrt{\mathrm{KL}(P\|Q)/2}$.

Proof. Let $A := \{x : P(x) > Q(x)\}$, $p := P(A)$ and $q := Q(A)$, so that $\mathrm{TV}(P,Q) = p - q$. Applying Fact 4 to the indicator of A reduces the claim to the two-point distributions with probabilities $(p, 1-p)$ and $(q, 1-q)$:

$$\mathrm{KL}(P\|Q) \geq p\log\frac{p}{q} + (1-p)\log\frac{1-p}{1-q} =: k(q).$$

We show $k(q) \geq 2(p-q)^2$ for $0 < q < 1$; the cases $q \in \{0,1\}$ follow by letting q tend to the endpoint, or because the divergence is infinite. The function $g(q) := k(q) - 2(p-q)^2$ satisfies $g(p) = 0$ and

$$g'(q) = -\frac{p}{q} + \frac{1-p}{1-q} + 4(p-q) = (q-p)\Big(\frac{1}{q(1-q)} - 4\Big).$$

Since $q(1-q) \leq \frac{1}{4}$, the bracket is nonnegative, so g decreases for $q < p$ and increases for $q > p$. Hence $g \geq g(p) = 0$. □

Chain rule. Relative entropy splits along a sequence into the expected divergences of the next entry, and this is the main reason we work with it.

Fact 6. *Let P, Q be laws of sequences $(Z_1, \dots, Z_q)$ with $\operatorname{supp}(P) \subseteq \operatorname{supp}(Q)$. Then*

$$\mathrm{KL}(P\|Q) = \sum_{i=1}^{q} \mathbb{E}_{\tau \leftarrow P^{(i-1)}} \big[\mathrm{KL}\big(P(\cdot \mid \tau) \,\|\, Q(\cdot \mid \tau)\big)\big].$$

Proof. For $z \in \operatorname{supp}(P)$, all the conditional probabilities in (18) are positive under both laws, and taking logarithms turns the products into sums:

$$\log \frac{P(z)}{Q(z)} = \sum_{i=1}^{q} \log \frac{P(z_i \mid \tau_i)}{Q(z_i \mid \tau_i)}, \qquad \tau_i := (z_1, \dots, z_{i-1}).$$

Take the expectation under P. In the i-th term, first average z_i over $P(\cdot \mid \tau_i)$ with the prefix fixed, which gives $\mathrm{KL}(P(\cdot \mid \tau_i)\|Q(\cdot \mid \tau_i))$; then average the prefix over $P^{(i-1)}$. □

By Facts 1 and 5, a bound $\mathrm{KL}(T_0\|T_1) \le \varepsilon$ on two transcript laws bounds the advantage by $\sqrt{\varepsilon/2}$. The chain rule reduces $\mathrm{KL}(T_0\|T_1)$ to one query at a time, and on each query Fact 2 lets us compute a chi-square divergence instead. When the comparison goes through an intermediate world, the triangle inequality for TV adds the two contributions. This is the structure of the proofs in Sections 5 and 7. The support condition is checked separately (Remark 3).

Proofs of the lemmas of Section 2.2.

Proof of Lemma 1. Write $Z := \mathbb{E}_Q[e^g]$, so that $P(x) = Q(x)e^{g(x)}/Z$, and let m and M be the minimum and the maximum of g on $\operatorname{supp}(Q)$, so that $d = M - m$. Then P is a distribution with the same support as Q. We first bound the ratio. On $\operatorname{supp}(Q)$, $P(x)/Q(x) = e^{g(x)}/Z$. The normaliser Z is an average of values in $[e^m, e^M]$, so it lies in that interval too, and $e^{g(x)}/Z$ lies in $[e^{m-M}, e^{M-m}] = [e^{-d}, e^d]$.

Next we bound $\chi^2(P\|Q)$. By (19), $\chi^2(P\|Q) = \operatorname{Var}_Q(P/Q) = \operatorname{Var}_Q(e^g)/Z^2$. Let $X, Y \leftarrow Q$ be independent. For real $a, b \le M$, the mean value theorem gives $|e^a - e^b| \le e^M|a-b|$, because the derivative of e^t is at most e^M between a and b. With (17),

$$\operatorname{Var}_Q(e^g) = \tfrac{1}{2}\mathbb{E}\big[(e^{g(X)} - e^{g(Y)})^2\big] \le \tfrac{1}{2}e^{2M}\,\mathbb{E}\big[(g(X) - g(Y))^2\big] = e^{2M}\operatorname{Var}_Q(g).$$

Dividing by $Z^2 \ge e^{2m}$ gives $\chi^2(P\|Q) \le e^{2(M-m)}\operatorname{Var}_Q(g) = e^{2d}\operatorname{Var}_Q(g)$.

For the consequence, U is uniform on V and $U_g(x) = U(x)e^{g(x)}/\mathbb{E}_U[e^g]$ with $0 \le g \le A$. Reweighting U_g by e^{-g} gives back U: indeed $U_g(x)e^{-g(x)} = U(x)/\mathbb{E}_U[e^g]$ is proportional to $U(x)$. So U is obtained from $Q := U_g$ by the first part with the function $-g$, whose oscillation is at most A:

$$\chi^2(U\|U_g) \le e^{2A}\operatorname{Var}_{U_g}(g) \le e^{2A}\,\mathbb{E}_{U_g}[g^2] \le e^{3A}\,\mathbb{E}_U[g^2].$$

The second step uses that a variance is at most the second moment. The last step uses $U_g/U \le e^A$, which is the ratio bound of the first part applied to U_g and U. Finally, for mixtures, Fact 3 shows that $Q \mapsto \chi^2(U\|Q)$ is convex. So if every Q_i satisfies $\chi^2(U\|Q_i) \le B$, then $\chi^2(U\|\sum_i w_iQ_i) \le \sum_i w_iB = B$. □

Proof of Lemma 2. Q^+ is the law of a Q-sample conditioned on not being $\perp$, and $Q(x) = (1-p)\,Q^+(x)$ for $x \in \Omega$. Since P puts no mass on $\perp$, only $x \in \Omega$ contribute to $\mathrm{KL}(P\|Q)$, and for each of them

$$\log \frac{P(x)}{Q(x)} = \log \frac{P(x)}{Q^+(x)} + \log \frac{1}{1-p}.$$

Multiply by $P(x)$ and sum over $x \in \Omega$; since $\sum_x P(x) = 1$, the constant contributes exactly $\log \frac{1}{1-p}$. If the support condition fails, both sides are infinite. □

A.2 Finite abelian groups

Our permutations act on a finite abelian group, and the masks are maps of this group that respect its addition. This subsection recalls the few notions we need, with the examples that matter for cryptography.

Groups. A *finite abelian group* is a finite set G with an addition that is associative and commutative, has a neutral element 0, and gives every g an inverse $-g$. Its *order* is $N := |G|$. The examples to keep in mind are the following:

- $\mathbb{Z}_N = \{0, \dots, N-1\}$ with addition modulo N. This is the natural setting for arithmetisation-oriented designs.
- $\mathbb{F}_2^d$, the bit strings of length d with bitwise XOR, of order 2^d; more generally $\mathbb{F}^d$ for a finite field $\mathbb{F}$. This is the setting of block ciphers, where $N = 2^d$ and addition is XOR.
- Products $G_1 \times G_2$ with coordinatewise addition, and in particular the powers G^r, whose elements are tuples $y = (y_1, \dots, y_r)$. Construction values of our families lie in G^w.

A *subgroup* is a subset H containing 0 and closed under addition and negation. For example, the even residues form a subgroup of $\mathbb{Z}_N$ when N is even, and any linear subspace is a subgroup of $\mathbb{F}_2^d$. The *cosets* of H are its translates $g + H = \{g + h : h \in H\}$. Two cosets are either equal or disjoint, since $g + H$ and $g' + H$ meet exactly when $g - g' \in H$, and then they coincide. So the cosets partition G into $N/|H|$ sets of size $|H|$. In particular $|H|$ divides N (Lagrange's theorem). Applied to the subgroup $\{0, g, 2g, \dots\}$ generated by g, this shows $Ng = g + \cdots + g = 0$ for every $g \in G$. The cosets of H themselves form a group, the *quotient* G/H, with the addition $(g + H) + (g' + H) := (g + g') + H$; it has order $N/|H|$.

Homomorphisms. A *homomorphism* $\varphi : G \to G'$ between abelian groups is a map that respects addition: $\varphi(a + b) = \varphi(a) + \varphi(b)$ for all a, b. It follows that $\varphi(0) = 0$ and $\varphi(-a) = -\varphi(a)$. A bijective homomorphism is an *isomorphism*, and two groups are *isomorphic*, written $G \cong G'$, if there is an isomorphism between them; isomorphic groups are the same group with the elements renamed. Every finite abelian group is isomorphic to a product $\mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_s}$ of cyclic groups [DF04, Sec. 5.2]. We use this decomposition only to construct characters (Fact 11); none of our statements depends on it.

The *kernel* $\ker \varphi := \{a : \varphi(a) = 0\}$ and the *image* $\operatorname{Im} \varphi := \varphi(G)$ of a homomorphism are subgroups of G and of G'. For $c \in G'$, the *fibre* of φ over c is the set $\varphi^{-1}(c) := \{v \in G : \varphi(v) = c\}$ of solutions of the equation $\varphi(v) = c$. The fibres over the elements of G' partition G, and the fibre over c is nonempty exactly when $c \in \operatorname{Im} \varphi$. The next fact describes the fibres. It is the reason that kernel sizes appear in our bounds: a same-class pair of calls with mask difference μ excludes, for a given construction value, either no anchor or a whole fibre of μ, which has $|\ker \mu|$ elements (Section 7.1).

Fact 7. *Let $\varphi : G \to G'$ be a homomorphism and $c \in G'$. The fibre $\varphi^{-1}(c)$ is empty if $c \notin \operatorname{Im} \varphi$, and is a coset of $\ker \varphi$, with $|\ker \varphi|$ elements, if $c \in \operatorname{Im} \varphi$. Consequently $N = |\ker \varphi| \cdot |\operatorname{Im} \varphi|$, and if $v \leftarrow \mathrm{Unif}(G)$ then $\varphi(v)$ is uniform on $\operatorname{Im} \varphi$.*

Proof. Suppose $\varphi(v_0) = c$ for some v_0. Since φ respects addition, $\varphi(v) = c$ holds iff $\varphi(v - v_0) = 0$, that is, iff $v - v_0 \in \ker\varphi$. So the fibre over c is the coset $v_0 + \ker\varphi$, which has $|\ker\varphi|$ elements. The $|\operatorname{Im}\varphi|$ nonempty fibres partition G and all have this size, which gives $N = |\ker\varphi| \cdot |\operatorname{Im}\varphi|$. A uniform v lands in each nonempty fibre with the same probability $|\ker\varphi|/N$, so $\varphi(v)$ is uniform on $\operatorname{Im}\varphi$. □

Endomorphisms and automorphisms. An *endomorphism* of G is a homomorphism from G to itself, and $\operatorname{End}(G)$ is the set of endomorphisms. Endomorphisms can be added and composed:

- the sum $(\varphi + \varphi')(g) := \varphi(g) + \varphi'(g)$ is again an endomorphism, because G is abelian;
- the product is composition, $\varphi\varphi' := \varphi \circ \varphi'$, that is, first φ' and then φ.

These operations obey the usual rules of arithmetic, such as $\varphi(\varphi' + \varphi'') = \varphi\varphi' + \varphi\varphi''$, except that the product need not be commutative; in algebraic language, $\operatorname{End}(G)$ is a *ring*. For $\mathbb{F}_2^d$ this is the familiar arithmetic of $d \times d$ matrices. An *automorphism* is a bijective endomorphism, that is, an isomorphism from G to itself. Its inverse is again an endomorphism, and the automorphisms form a group $\operatorname{Aut}(G)$ under composition. In our families the masks are automorphisms, but differences of masks, such as $\lambda_j^{-1} - \lambda_{j'}^{-1}$, are in general only endomorphisms.

Fact 8. *For $\varphi \in \operatorname{End}(G)$ the following are equivalent: φ is injective;* $\ker\varphi = \{0\}$*; φ is surjective;* $\varphi \in \operatorname{Aut}(G)$.

Proof. By Fact 7, all nonempty fibres have $|\ker\varphi|$ elements. So φ is injective, meaning that every fibre has at most one element, iff $\ker\varphi = \{0\}$. By $N = |\ker\varphi| \cdot |\operatorname{Im}\varphi|$, $\ker\varphi = \{0\}$ iff $\operatorname{Im}\varphi = G$, that is, iff φ is surjective. A map that is both injective and surjective is a bijection. □

So an endomorphism of a finite abelian group is either an automorphism or has a nontrivial kernel. This is the trichotomy behind rigid, regular and singular pairs (Section 7.1): a difference of masks is 0 (kernel of size N), an automorphism (kernel of size 1), or neither.

Example 2. The two standard groups show all three cases:

- On $\mathbb{Z}_N$, every endomorphism is multiplication by some $a \in \mathbb{Z}_N$, namely $a = \varphi(1)$, since $\varphi(g) = \varphi(1 + \cdots + 1) = g\varphi(1)$. It is an automorphism iff $\gcd(a, N) = 1$. For N even, the masks 1 and -1 give the difference $1^{-1} - (-1)^{-1} = 2$, whose kernel $\{0, N/2\}$ has size 2.
- On $\mathbb{F}_2^d$, the endomorphisms are the $d \times d$ matrices over $\mathbb{F}_2$, acting on column vectors, and the automorphisms are the invertible ones. A difference of two invertible matrices can be singular without being 0: for $J \neq 0$ with $J^2 = 0$, both I and $I + J$ are invertible (each is its own inverse), and $I^{-1} - (I + J)^{-1} = J$.

Maps between powers. Our constructions combine several group elements linearly, and the natural tool for that is a matrix whose entries are endomorphisms. A $w' \times m$ matrix $\mathbf{M} = (M_{ij})$ with entries in $\operatorname{End}(G)$ defines a homomorphism $G^m \to G^{w'}$ by

$$(\mathbf{M}y)_i := \sum_{j=1}^{m} M_{ij}(y_j),$$

exactly as a real matrix acts on a column vector. Composing two such maps corresponds to multiplying the matrices, where the entries are multiplied by composition, in order. Integer entries act by repeated addition, so the integer matrix $(-1\ 1)$ maps (y_1, y_2) to $y_2 - y_1$. Blockwise maps, such as the map T of Lemma 18, are the diagonal matrices. This is how the combiner of Definition 2 acts.

The sets that our simulator must avoid are described by linear equations in several variables, that is, in a tuple $y \in G^w$, such as $\varphi_1(y_1) + \varphi_2(y_2) = c$. We call the solution set of one such equation an *affine hyperplane*, by analogy with linear algebra. When one coefficient is invertible, a hyperplane contains exactly a $1/N$ fraction of all tuples.

Fact 9. *Let $\varphi_1, \dots, \varphi_w \in \mathrm{End}(G)$ and $c \in G$, and suppose that some φ_s is an automorphism. If $y \leftarrow \mathrm{Unif}(G^w)$, then $\Pr[\sum_j \varphi_j(y_j) = c] = 1/N$.*

Proof. Condition on the coordinates y_j with $j \neq s$. The equation becomes $\varphi_s(y_s) = c'$ for a fixed c'. Since φ_s is a bijection and y_s is uniform and independent of the other coordinates, $\varphi_s(y_s)$ is uniform, and it equals c' with probability $1/N$. □

A simpler kind of set will also appear. If a set $C \subseteq G^w$ is described by a condition on a single difference $y_{j'} - y_j$, say $C = \{y : y_{j'} - y_j \in B\}$ for a set $B \subseteq G$, then for uniform y the difference $y_{j'} - y_j$ is uniform on G, so C contains a $|B|/N$ fraction of all tuples.

Coset progressions. The matching attack of Section 6.2 needs, at every scale R, a set of about R group elements that is nearly closed under addition. In $\mathbb{F}_2^d$ a subgroup of the right size does the job. In $\mathbb{Z}_p$ with p prime there are no nontrivial subgroups, and an interval $\{0, 1, \dots, L-1\}$ takes their place: the sum of two elements of the interval lies in the interval of twice the length. Coset progressions combine the two.

Definition 26. A *coset progression* of length $L \geq 1$ is a pair (H, g) of a subgroup $H \leq G$ and an element $g \in G$ such that the cosets $H, H + g, \dots, H + (L-1)g$ are distinct. For a set J of integers we write $P(J) := \bigcup_{i \in J} (H + ig)$.

So $P([0, L))$ is a union of L distinct cosets of H lined up along g, and it has $L|H|$ elements. A subgroup is the case $L = 1$, and an interval in $\mathbb{Z}_p$ is the case $H = \{0\}$, $g = 1$. Since $(H + ig) + (H + i'g) = H + (i + i')g$, sums of elements of $P([0, L))$ lie in $P([0, 2L-1))$, which has fewer than twice as many elements.

Lemma 31. *For every integer $1 \leq R \leq N$ there is a coset progression (H, g) of some length L with $R \leq L|H| < 2R$.*

Proof. We climb a chain of subgroups in steps of prime index and stop where the size first reaches R: either that subgroup is already small enough, or the previous subgroup, repeated along an element of the next one, gives the right size. A finite abelian group has a chain of subgroups $0 = H_0 < H_1 < \cdots < H_s = G$ in which every index $p_i := |H_i|/|H_{i-1}|$ is prime: given $H_{i-1} \neq G$, choose H_i so that H_i/H_{i-1} is a subgroup of prime order of the quotient G/H_{i-1}, which exists by Cauchy's theorem. Let i be least with $|H_i| \geq R$.

If $|H_i| < 2R$, take $H := H_i$ and $L := 1$ (for $R = 1$ this is $H = H_0 = \{0\}$). Otherwise $|H_i| \geq 2R$, so $i \geq 1$. Put $H := H_{i-1}$ and $h := |H|$; then $h < R$, by the minimality of i. Take $g \in H_i \setminus H_{i-1}$. The coset of g generates the quotient H_i/H_{i-1}, which is cyclic of order p_i, so the cosets $H + ig$, $0 \leq i < p_i$, are distinct. Put $L := \lceil R/h \rceil$. Then $R \leq hL < R + h < 2R \leq hp_i$, so $L < p_i$, the cosets $H, \dots, H + (L-1)g$ are distinct, and $L|H| = hL \in [R, 2R)$. □

In an elementary abelian 2-group every progression found this way is a subgroup ($p_i = 2$ forces $L = 1$). In $\mathbb{Z}_p$ with p prime it is an interval ($H = \{0\}$, $g = 1$).

A.3 Fourier analysis on finite abelian groups

Fourier analysis expands a function on G in a basis of *characters*, which are the functions that turn addition into multiplication. A sum or a uniform random choice over G behaves simply in this basis, and our injection coefficients and moments are computed there.

Characters. Let $\mathbb{T} := \{z \in \mathbb{C} : |z| = 1\}$ be the unit circle. A *character* of G is a homomorphism $\psi : G \to \mathbb{T}$, that is, $\psi(a+b) = \psi(a)\psi(b)$ for all a, b. The *trivial* or *principal* character is the constant 1.

Fact 10. *Every character satisfies $\psi(0) = 1$ and $\psi(-g) = \overline{\psi(g)}$, and its values are N-th roots of unity.*

Proof. $\psi(0) = \psi(0+0) = \psi(0)^2$ and $\psi(0) \neq 0$, so $\psi(0) = 1$. Then $\psi(g)\psi(-g) = \psi(0) = 1$, and for $|z| = 1$ the inverse of z is $\overline{z}$. Finally $\psi(g)^N = \psi(Ng) = \psi(0) = 1$, by Lagrange's theorem. □

The characters form a group $\widehat{G}$, the *dual group*, under pointwise multiplication $(\psi\psi')(g) := \psi(g)\psi'(g)$. Following the convention of Section 2.3, we write this group additively: $\psi + \psi'$ denotes the pointwise product, 0 denotes the principal character and $-\psi = \overline{\psi}$. A character thus plays two roles: it is an element of the group $\widehat{G}$, which we add to other characters, and it is a function on G, which we evaluate. To keep the two roles apart, we evaluate through the map

$$\mathrm{e} : \widehat{G} \times G \to \mathbb{T}, \qquad \mathrm{e}(\psi, g) := \psi(g),$$

which takes one element of each group and returns a complex number of modulus 1. Such a map is called a *pairing* of $\widehat{G}$ with G. It is symmetric in spirit: for fixed ψ it is the character ψ, and for fixed g it is a character of $\widehat{G}$. In this notation $\mathrm{e}(\psi + \psi', g) = \mathrm{e}(\psi, g)\mathrm{e}(\psi', g)$, $\mathrm{e}(\psi, g + g') = \mathrm{e}(\psi, g)\mathrm{e}(\psi, g')$ and $\overline{\mathrm{e}(\psi, g)} = \mathrm{e}(-\psi, g) = \mathrm{e}(\psi, -g)$.

Example 3. For the standard groups the characters are explicit:

- On $\mathbb{Z}_N$, the characters are $\psi_a(g) = e^{2\pi i\, ag/N}$ for $a \in \mathbb{Z}_N$: a character is determined by $\psi(1)$, which must be an N-th root of unity. So $\widehat{G} \cong \mathbb{Z}_N$, with $\mathrm{e}(a, g) = e^{2\pi i\, ag/N}$.
- On $\mathbb{F}_2^d$, the characters are $\psi_a(g) = (-1)^{a \cdot g}$ for $a \in \mathbb{F}_2^d$, where $a \cdot g$ is the inner product modulo 2. So $\widehat{G} \cong \mathbb{F}_2^d$.
- On a product $G_1 \times G_2$, every character satisfies $\psi(g_1, g_2) = \psi(g_1, 0)\psi(0, g_2)$, so the characters are the products $\psi_1(g_1)\psi_2(g_2)$ of characters of the factors. So $\widehat{G_1 \times G_2} = \widehat{G}_1 \times \widehat{G}_2$. For G^r this gives the characters $y \mapsto \mathrm{e}(\beta, y) = \prod_i \mathrm{e}(\beta_i, y_i)$ with $\beta \in \widehat{G}^r$, as in Section 2.3.

Fact 11. *$|\widehat{G}| = N$, and for every $g \neq 0$ there is a character ψ with $\psi(g) \neq 1$.*

Proof. By the cyclic decomposition, we may take $G = \mathbb{Z}_{n_1} \times \cdots \times \mathbb{Z}_{n_s}$. By Example 3, $\widehat{G}$ is the product of the duals of the factors, which have $n_1, \ldots, n_s$ elements; so $|\widehat{G}| = N$. If $g \neq 0$, some coordinate $g_i \in \mathbb{Z}_{n_i}$ is nonzero, and the character $e^{2\pi i\, g_i/n_i}$ of that coordinate is not 1 at g. □

Orthogonality. The basic computation of Fourier analysis is that a nonprincipal character sums to zero.

Fact 12. *For $\psi \in \widehat{G}$ and $g \in G$,*

$$\sum_{g' \in G} \mathrm{e}(\psi, g') = \begin{cases} N, & \psi = 0, \\ 0, & \psi \neq 0, \end{cases} \qquad \sum_{\psi' \in \widehat{G}} \mathrm{e}(\psi', g) = \begin{cases} N, & g = 0, \\ 0, & g \neq 0. \end{cases}$$

Proof. If $\psi \neq 0$, pick h with $\mathrm{e}(\psi, h) \neq 1$. Shifting the summation variable by h does not change the sum $S := \sum_{g'} \mathrm{e}(\psi, g')$, so $S = \sum_{g'} \mathrm{e}(\psi, g' + h) = \mathrm{e}(\psi, h)S$, and therefore $S = 0$. The second identity is proved in the same way, with a character ψ' such that $\mathrm{e}(\psi', g) \neq 1$, which exists by Fact 11. The cases $\psi = 0$ and $g = 0$ count N terms equal to 1, using $|\widehat{G}| = N$. □

Fourier basis. For functions $f, h : G^r \to \mathbb{C}$ we use the inner product $\langle f, h\rangle := \mathbb{E}_y[f(y)\overline{h(y)}]$ with $y \leftarrow \mathrm{Unif}(G^r)$, and $\|f\|_2^2 := \langle f, f\rangle = \mathbb{E}[|f|^2]$. This is the space $L^2(\mathrm{Unif}(G^r))$ of Section 2.3.

Fact 13. *The characters $y \mapsto \mathrm{e}(\beta, y)$, $\beta \in \widehat{G^r}$, form an orthonormal basis of $L^2(\mathrm{Unif}(G^r))$. Consequently every f satisfies*

$$f = \sum_{\beta \in \widehat{G^r}} \hat{f}(\beta)\,\mathrm{e}(\beta, \cdot), \qquad \hat{f}(\beta) := \langle f, \mathrm{e}(\beta, \cdot)\rangle = \mathbb{E}_y\big[f(y)\overline{\mathrm{e}(\beta, y)}\big], \qquad \mathbb{E}[|f|^2] = \sum_{\beta} |\hat{f}(\beta)|^2.$$

In addition, $\hat{f}(0) = \mathbb{E}[f]$.

Proof. For $\beta, \beta' \in \widehat{G^r}$, $\mathrm{e}(\beta, y)\overline{\mathrm{e}(\beta', y)} = \mathrm{e}(\beta - \beta', y)$, so

$$\langle \mathrm{e}(\beta, \cdot), \mathrm{e}(\beta', \cdot)\rangle = \prod_{i=1}^{r} \frac{1}{N} \sum_{y_i \in G} \mathrm{e}(\beta_i - \beta_i', y_i),$$

which is 1 if $\beta = \beta'$ and 0 otherwise, by Fact 12. So the N^r characters are orthonormal, hence linearly independent, and since the space has dimension N^r they form a basis. The expansion and Parseval's identity hold for every orthonormal basis. Finally $\mathrm{e}(0, y) = 1$, so $\hat{f}(0) = \mathbb{E}[f]$. □

The Fourier coefficients of a probability density are expectations of characters. This is how injection coefficients and the moments of Appendix E arise.

Fact 14. *Let P be a distribution on G^r with density $f(y) := N^r P(y)$ relative to $\mathrm{Unif}(G^r)$. Then $\hat{f}(\beta) = \mathbb{E}_{Y \leftarrow P}[\overline{\mathrm{e}(\beta, Y)}]$, $\hat{f}(0) = 1$, and*

$$\chi^2\big(P \,\|\, \mathrm{Unif}(G^r)\big) = \sum_{\beta \neq 0} |\hat{f}(\beta)|^2.$$

Proof. $\hat{f}(\beta) = N^{-r} \sum_y N^r P(y)\overline{\mathrm{e}(\beta, y)} = \mathbb{E}_{Y \leftarrow P}[\overline{\mathrm{e}(\beta, Y)}]$, which is 1 at $\beta = 0$. By the definition of χ^2, $\chi^2(P\|\mathrm{Unif}) = \sum_y N^{-r}(f(y) - 1)^2 = \mathbb{E}[|f - 1|^2]$. The function $f - 1$ has the same coefficients as f, except at $\beta = 0$, where its coefficient is 0. Apply Parseval's identity. □

So a distribution is close to uniform exactly when its nonprincipal Fourier coefficients are small in ℓ_2. For a uniform injection $Y : [r] \to G$, the coefficients are the $\hat{\mu}_r(\beta)$ of Section 2.3.

Annihilators and dual endomorphisms. The *annihilator* of a subgroup $H \leq G$ is $H^\perp := \{\psi \in \widehat{G} : \mathrm{e}(\psi, h) = 1 \text{ for all } h \in H\}$, a subgroup of $\widehat{G}$. For $\varphi \in \mathrm{End}(G)$ and $\psi \in \widehat{G}$, the composition $\psi \circ \varphi$ is again a character; the *dual* φ^* maps ψ to it, that is, $\mathrm{e}(\varphi^*\psi, g) = \mathrm{e}(\psi, \varphi g)$.

Fact 15. *For $\varphi, \varphi' \in \mathrm{End}(G)$ and a subgroup $H \leq G$:*

(a) $\varphi^* \in \mathrm{End}(\widehat{G})$, $(\varphi + \varphi')^* = \varphi^* + \varphi'^*$, and $(\varphi\varphi')^* = \varphi'^*\varphi^*$;
(b) $|H^\perp| = N/|H|$;
(c) $\ker \varphi^* = (\mathrm{Im}\,\varphi)^\perp$, so $|\ker \varphi^*| = |\ker \varphi|$, and φ is an automorphism iff φ^* is.

Proof. (a) All three identities are checked by pairing with an arbitrary g. First,

$$\mathrm{e}(\varphi^*(\psi + \psi'), g) = \mathrm{e}(\psi + \psi', \varphi g) = \mathrm{e}(\psi, \varphi g)\mathrm{e}(\psi', \varphi g) = \mathrm{e}(\varphi^*\psi + \varphi^*\psi', g).$$

Second, $e(\psi, \varphi g + \varphi' g) = e(\psi, \varphi g) e(\psi, \varphi' g)$, which is the pairing of $\varphi^*\psi + \varphi'^*\psi$ with g. Third, $e(\psi, \varphi(\varphi' g)) = e(\varphi^*\psi, \varphi' g) = e(\varphi'^*\varphi^*\psi, g)$; the order is reversed because φ' is applied first.

(b) Restricted to H, a character is a character of H; it is principal on H exactly when $\psi \in H^\perp$. By Fact 12 applied to H, the average $\frac{1}{|H|}\sum_{h\in H} e(\psi, h)$ is therefore *1* for $\psi \in H^\perp$ and *0* otherwise. Summing over ψ and exchanging the sums,

$$|H^\perp| = \frac{1}{|H|}\sum_{h\in H}\sum_{\psi\in\widehat{G}} e(\psi, h) = \frac{N}{|H|},$$

because by Fact 12 only $h = 0$ contributes, with N.

(c) $\varphi^*\psi = 0$ means $e(\psi, \varphi g) = 1$ for all g, that is, $\psi \in (\operatorname{Im}\varphi)^\perp$. By (b) and Fact 7, $|\ker\varphi^*| = N/|\operatorname{Im}\varphi| = |\ker\varphi|$. By Fact 8, applied to G and to $\widehat{G}$, both φ and φ^* are automorphisms exactly when this common kernel size is *1*. □

On $\mathbb{Z}_N$, the dual of multiplication by a is multiplication by a, since $e(a', ag) = e^{2\pi i a'ag/N} = e(aa', g)$. On $\mathbb{F}_2^d$, the dual of a matrix A is its transpose, since $(-1)^{a\cdot Ag} = (-1)^{(A^\intercal a)\cdot g}$.

Proof of Lemma 3. A uniform injection $Y : [r] \to G$ is a uniformly random tuple $(Y_1, \ldots, Y_r)$ of distinct elements of G; there are $(N)_r$ such tuples. We first work out the smallest case by hand.

Example 4. For $\psi \neq 0$ and $\beta = (\psi, -\psi)$, $\overline{e(\beta, y)} = \overline{e(\psi, y_1)}\, e(\psi, y_2)$. Summing over all pairs and subtracting the pairs with $y_1 = y_2$,

$$\sum_{y_1\neq y_2} \overline{e(\psi, y_1)}\, e(\psi, y_2) = \Big|\sum_{y\in G} e(\psi, y)\Big|^2 - \sum_{y\in G} |e(\psi, y)|^2 = 0 - N,$$

by Fact 12. Dividing by the $(N)_2 = N(N-1)$ injective pairs gives $\hat{\mu}_2(\beta) = -1/(N-1)$. For two independent uniform values the coefficient would be *0*; the excluded diagonal $y_1 = y_2$ creates a small negative correlation. Merging (Lemma 3(d)) gives the same value: $\hat{\mu}_2(\beta) = -\frac{1}{N-1}\hat{\mu}_2(0,0) = -\frac{1}{N-1}$.

Proof of Lemma 3. (a) Since $e(0, g) = 1$, a coordinate with label *0* contributes the factor *1* to $\overline{e(\beta, Y)}$, so $\hat{\mu}_r(\beta) = \mathbb{E}[\overline{e(\beta|_J, Y|_J)}]$. It remains to see that the restriction $Y|_J$ is a uniform injection $J \to G$. Every injective tuple indexed by J extends to the same number $(N-k)_{r-k}$ of injective tuples indexed by $[r]$, namely the ways of filling the other $r-k$ positions with distinct unused values. So every injective J-tuple is equally likely, and the expectation is $\hat{\mu}_k(\beta|_J)$.

(b) A map that sends injective tuples bijectively to injective tuples preserves the uniform law on them, and two such symmetries of the uniform injection give the three claims:

- The map $y \mapsto -y$ is such a bijection. Since $e(\beta_i, -y_i) = \overline{e(\beta_i, y_i)}$, it turns $\hat{\mu}_r(\beta)$ into its complex conjugate. So $\hat{\mu}_r(\beta)$ equals its conjugate, and is real. It is at most *1* in modulus because it is an average of numbers of modulus *1*.
- For fixed $g \in G$, the map $y \mapsto (y_1 + g, \ldots, y_r + g)$ is also such a bijection. It multiplies $\overline{e(\beta, y)}$ by $\prod_i \overline{e(\beta_i, g)} = \overline{e(\sum_i \beta_i, g)}$. Hence $\hat{\mu}_r(\beta) = \overline{e(\sum_i \beta_i, g)}\,\hat{\mu}_r(\beta)$ for every g. If $\sum_i \beta_i \neq 0$, it is a nonprincipal character, so some g makes the factor different from *1*, and then $\hat{\mu}_r(\beta) = 0$.

For $r = 1$ and $\psi \neq 0$ the label sum is $\psi \neq 0$, which gives $\hat{\mu}_1(\psi) = 0$: a single uniform value is exactly uniform.

(c) By (a) we may take $r = 3$. Then $(N)_3\,\hat{\mu}_3(\beta)$ is the sum of $\overline{e(\beta, y)}$ over the injective $y \in G^3$. To sum over injective tuples only, we write their indicator in terms of the diagonals, which are easy to sum over. Inclusion–exclusion on the three possible equalities gives, for every $y \in G^3$,

$$\mathbf{1}[y_1, y_2, y_3 \text{ distinct}] = 1 - \sum_{\text{pairs } a<b} \mathbf{1}[y_a = y_b] + 2\cdot\mathbf{1}[y_1 = y_2 = y_3].$$

Indeed, the right-hand side is 1 if the entries are distinct, $1-1=0$ if exactly two of them are equal, and $1-3+2=0$ if all three are equal. We multiply by $\overline{e(\beta, y)}$ and sum over G^3, term by term:

- The constant term gives $\prod_{i=1}^{3}\sum_{y_i\in G}\overline{e(\beta_i, y_i)} = 0$, because a nonprincipal character sums to 0 over G (Fact 12).
- The term for the pair $\{1,2\}$ sums over the tuples with $y_1 = y_2$. It gives

$$\sum_{y_1, y_3}\overline{e(\beta_1+\beta_2, y_1)}\,\overline{e(\beta_3, y_3)}.$$

  The sum over y_3 vanishes since $\beta_3 \neq 0$. The other two pairs vanish in the same way, since in each case the remaining coordinate carries a nonzero label and is summed freely.
- The last term sums over the constant tuples $y_1 = y_2 = y_3 = y$ and gives

$$2\sum_{y\in G}\overline{e(\beta_1+\beta_2+\beta_3, y)} = 2N,$$

  because $\beta_1+\beta_2+\beta_3 = 0$.

Hence $(N)_3\,\hat{\mu}_3(\beta) = 2N$, which is the claim.

(d) By (a) we may take $r = k$ and $J = [k]$, so every label is nonzero. Condition on the outputs y_i for $i \neq j$. They are $k-1$ distinct points, and given them, y_j is uniform on the remaining $N-k+1$ points of G: each completion is one injective tuple, and all injective tuples are equally likely. Since $\beta_j \neq 0$, the character $\overline{e(\beta_j, \cdot)}$ sums to 0 over G (Fact 12), so its sum over the remaining points is minus its sum over the $k-1$ points y_i, $i \neq j$:

$$\mathbb{E}\left[\overline{e(\beta_j, y_j)} \,\middle|\, (y_i)_{i\neq j}\right] = -\frac{1}{N-k+1}\sum_{i\neq j}\overline{e(\beta_j, y_i)}.$$

Multiply both sides by $\prod_{i\neq j}\overline{e(\beta_i, y_i)}$, which is fixed by the conditioning, and take expectations. In the term indexed by i, the coordinate y_i now carries the factor $\overline{e(\beta_i, y_i)}\,\overline{e(\beta_j, y_i)} = \overline{e(\beta_i+\beta_j, y_i)}$, and the coordinate j carries no factor. That term is therefore $\hat{\mu}_k(\beta^{(j,i)})$, which proves the identity. If $\beta_i+\beta_j = 0$, the term has a smaller support, and (a) applies to it. □

B Evaluation of the constants

Indifferentiability. Write $\hat{w} := w+1$, so $\bar{F} = n\hat{w}$, $n \le \hat{w}$ and $\hat{w} \ge 2$.

- $C_{\mathrm{con}}(w) \le 6w(w+1) + 3w^4 \le 3\hat{w}^4$, because $(w+1)^4 \ge w^4 + 2w^2 + 2w$.
- $C_{\mathrm{RH}} \le 31\,n^3\hat{w}^4$. Indeed, $n^3 C_{\mathrm{con}} + 4\bar{F}^3 \le 3n^3\hat{w}^4 + 4n^3\hat{w}^3 \le 5n^3\hat{w}^4$. Also

$$52n\bar{F}^2 + 36w^2n^2\bar{F} = 52n^3\hat{w}^2 + 36w^2n^3\hat{w} \le n^3\hat{w}^3(52/\hat{w} + 36) \le 62n^3\hat{w}^3 \le 31n^3\hat{w}^4.$$

- $C'_{\mathrm{RH}} \le 78\,n^3\hat{w}^4 + 2\,n^4\hat{w}^4$. Indeed:
  - $\max(8\bar{F}^2w^2, 72w^4n^2) \le 72n^2\hat{w}^4 \le 72n^3\hat{w}^4$;
  - $22w\hat{w}n^3 \le 6n^3\hat{w}^4$, because $22w/\hat{w}^3 \le 11/4$ for all $w \ge 1$ (the ratio is largest at $w = 1$);
  - $1.75\bar{F}^4 \le 2n^4\hat{w}^4$.
- $C''_{\mathrm{RH}} = 208nw^4 \le 208n\hat{w}^4$.

w	A_w	K_w	C_w	c_w
1	382.1	53.14	116.05	0.02535
2	5165	195.2	588.16	0.01065
3	$4.894 \cdot 10^4$	601.2	2515.34	0.00503
4	$3.822 \cdot 10^5$	1680	9815	0.002535
5	$2.641 \cdot 10^6$	4418	$3.615 \cdot 10^4$	0.001331

Table 5: Numerical values, for $w \le 5$, of the constant A_w of Proposition 9 and the constants K_w, C_w and c_w of Theorem 4, all defined in Appendix E.

- $K_{\mathrm{RH}} = C_{\mathrm{RH}} + C'_{\mathrm{RH}}/2w + C''_{\mathrm{RH}}/4w^2 \le 85\, n^3 \hat{w}^4$, from the following three bounds, using $2w \ge \hat{w}$, $\hat{w} \ge 2$ and $n \le \hat{w}$:
  - $C_{\mathrm{RH}} \le 31 n^3 \hat{w}^4$;
  - $C'_{\mathrm{RH}}/2w \le 78 n^3 \hat{w}^3 + 2n^4 \hat{w}^3 \le 41 n^3 \hat{w}^4$;
  - $C''_{\mathrm{RH}}/4w^2 = 52 n w^2 \le 52 n \hat{w}^2 \le 13 n \hat{w}^4 \le 13 n^3 \hat{w}^4$.
- $K_{\mathrm{HI}} = 100 n^3 w + 128 n^3 w^3 \le 228\, n^3 w^3$.
- First coefficient of Theorem 1: $\sqrt{85/2} + \sqrt{114} \cdot 0.325 = 9.9893 \le 10$. Here $w^{3/2} \le 0.325 (w+1)^2$ for all real $w \ge 1$: the ratio $w^{3/2}/(w+1)^2$ is largest at $w = 3$, where it is $5.196/16 < 0.325$. The stated 10 leaves almost no slack: any increase of K_{RH} or K_{HI} in a later revision must be followed by a re-evaluation of this coefficient.

These evaluations were also checked in exact arithmetic for $w \le 200$ and $n \le w + 1$.

Pseudorandomness. Table 5 gives numerical values of the constants of Theorem 4. The bounds quoted in Theorem 4 round C_w up and c_w down. They were checked for $w \le 40$, including the finite-tail condition $u\sqrt{2eq} \le 1/2$.

B.1 Numerical steps in Sections 5.4 and 5.5

The proofs in Sections 5.4 and 5.5 track orders of magnitude. The following items supply the explicit constants in their statements. Throughout, (3) gives $\bar{F}q \le N/32$, $\Phi \le 3N/32$, $\ell \ge 29N/32$, $nq \le N/64$ and $\phi_0 \le 1/16$ (Section 5.1).

(1) *Lemma 9, Step 3.* $Fh_x \le \bar{F}q \le N/32$ and $\ell \ge 29N/32$ give $p_a \le Fh_x/\ell \le 1/29$. Hence $\bar{\pi} \ge 28/29$ and $\bar{\pi}^2 \ge 0.80$.

(2) *Lemma 9, Step 3.* $\Phi N/\ell^2 \le (3/32)/(29/32)^2 \le 0.115$, so $\mathrm{Var}_a[\mu_i(a)] \le 1.12\, \Phi/N^3$.

(3) *Lemma 9, Steps 5 and 6.* Since $(N/\ell)^4 \le (32/29)^4 \le 1.5$ and $(h_x F)^2 \le (N/32)\Phi$,

$$D_{\mathrm{f}}^2 \le \frac{h_x^4 F^4}{4\ell^4} \le \frac{1.5}{4 \cdot 32} \cdot \frac{h_x^2 F^2 \Phi}{N^3} \le 0.012\, \frac{h_x^2 F^2 \Phi}{N^3}.$$

With $\mathrm{Var}_a[A] \le 1.12\, (Fh_x)^2 \Phi/N^3$ from item (2), $\mathrm{Var}_a[E_a] \le 2\mathbb{E}[\kappa^2]/\ell^2 + 2D_{\mathrm{f}}^2$ and $\mathrm{Var}(u - v) \le 2\,\mathrm{Var}\, u + 2\,\mathrm{Var}\, v$,

$$\mathrm{Var}_{\mathrm{Unif}(G)}[p_a] \le (2 \cdot 1.12 + 4 \cdot 0.012)\, \frac{h_x^2 F^2 \Phi}{N^3} + \frac{4\mathbb{E}[\kappa^2]}{\ell^2} \le 2.3\, \frac{h_x^2 F^2 \Phi}{N^3} + \frac{4\mathbb{E}[\kappa^2]}{\ell^2}.$$

With $N/\ell \le 32/29$ and $\bar{\pi}^2 \ge 0.80$,

$$\chi^2 \le \frac{32}{29}\cdot\frac{2.3}{0.80}\cdot\frac{h_x^2F^2\Phi}{N^3} + \frac{4}{0.80}\Big(\frac{32}{29}\Big)^3\frac{\mathbb{E}[\kappa^2]}{N^2} \le 3.2\,\frac{h_x^2F^2\Phi}{N^3} + 6.7178\,\frac{\mathbb{E}[\kappa^2]}{N^2},$$

which gives the constants $4 \ge 3.2$ and $6.73 \ge 6.7178$ of the statement. The coefficient must be carried unrounded: rounding $4/\ell^2$ to $4.88/N^2$ first would give $6.731 > 6.73$.

(4) *Lemma 10.* Here $p_o \le nh/\ell \le 1/29$, so the evaluation of item (3) applies with Fh_x replaced by nh and gives $\chi^2 \le 3.2\,n^2h^2\Phi/N^3 + 6.73\,\mathbb{E}_o[\kappa^2]/N^2$. With $\mathbb{E}_o[\kappa^2] \le \varepsilon h^2n^4/4N$, the second term is at most $\frac{6.73}{4}\,\varepsilon n^4h^2/N^3 \le 2\varepsilon n^4h^2/N^3$, and $3.2 \le 4$.

(5) *Lemma 11.* With $\ell \ge 29N/32$ and $N-\Phi \ge 29N/32$, both $\Phi/\ell N$ and $\Phi/N(N-\Phi)$ are at most $4\Phi/3N^2$. So $|R-K| \le 8\Phi/3N^2$ at every legal pair, and with $K \ge 3/4N$ the legal pairs contribute at most

$$nh\Big(\frac{8\Phi}{3N^2}\Big)^2\frac{4N}{3} = \frac{256}{27}\cdot\frac{nh\Phi^2}{N^3} \le 10\,\frac{nh\Phi^2}{N^3}.$$

Then $\Delta \le nh\cdot 8\Phi/3N^2 \le 3nh\Phi/N^2$, and with $e_b \ge 3N/4$ and $nh \le N$ the fresh inputs contribute at most $\frac{16}{9}\Delta^2 = 16n^2h^2\Phi^2/N^4 \le 16\,nh\Phi^2/N^3$. The total is at most $26\,nh\Phi^2/N^3$.

(6) *Lemma 12, Step 3.* $\ell' = \ell - n \ge 7N/8$ and $M_{j'} \ge N - n\bar{s} - 1 \ge 0.98N$, so $\ell' M_{j'}^2 \ge 0.84N^3$. With $M_{j'} - \ell' \le 2\Phi$, the variance of one coordinate is at most $2\Phi(nh)^2/0.84N^3 \le 3n^2h^2\Phi/N^3$.

(7) *Lemma 12, Step 5.* $|\mathcal{T}_0| = \rho_b|\mathcal{S}| \ge \frac{3}{4}\cdot\frac{31}{32}N^w \ge N^w/2$, so $N^w/|\mathcal{T}_0| \le 2$.

(8) *Lemma 12, Step 6.* By item (6) and Cauchy–Schwarz over the w coordinates, $\operatorname{Var}\sum_{j'} p_v^{(j')} \le 3w^2n^2h^2\Phi/N^3$, which gives the term $6w^2n^2h^2\Phi/N^3$ after $\operatorname{Var}(u-v) \le 2\operatorname{Var}u + 2\operatorname{Var}v$. With $\ell' \ge 7N/8$ the last two terms are at most

$$8\cdot\frac{64}{49}\cdot\frac{w^4}{4}\cdot\frac{h^2n^4}{N^4} + 4\cdot\frac{w^4}{4}\Big(\frac{8}{7}\Big)^4\frac{n^4h^4}{N^4} \le 2.7\,\frac{w^4h^2n^4}{N^4} + 1.8\,\frac{w^4n^4h^4}{N^4},$$

and both are absorbed into $2.73\,w^2n^2h^2\Phi/N^3$, since $2.7\cdot 1 + 1.8/64 \le 2.73$ by the two ratio bounds of Step 6. So $\operatorname{Var}_v[p_v] \le 8.73\,w^2n^2h^2\Phi/N^3 + (8/\ell'^2)\sum\operatorname{Var}[Y_{j'j''}]$, and $\chi^2 \le 4\operatorname{Var}_v[p_v]$ gives the constants $4\cdot 8.73 \le 36$ and $4\cdot 8 = 32$.

(9) *Lemma 13, Step 4.* $\binom{F}{2}^2 \le F^4/4$ and $\frac{32}{31}\cdot\frac{1}{4} \le 0.26$, so $q\,(q\binom{F}{2})^2\cdot\frac{32}{31}/N \le 0.26\,q^3F^4/N$.

(10) *Lemma 14.* $2qn^2/N \le q^2n^3/N$ since $q \ge 2$, and $(qn^2)^2/N^2 = (q^2n^3/N)\cdot n/N$ with $n/N \le 0.007$. Hence $\mathbb{E}[V_{j_1j_2}] \le (1 + \frac{32}{31}\cdot 0.007)\,q^2n^3/N \le 1.01\,q^2n^3/N$.

(11) *Proposition 1.* The forward coincidence terms sum in expectation to at most $6.73\cdot 0.26\,q^3F^4/N^3 \le 1.75\,q^3\bar{F}^4/N^3$. For the companion terms, $\ell' \ge 7N/8$ gives $32/\ell'^2 \le 32\cdot\frac{64}{49}/N^2$. With $\mathbb{E}[W] \le \binom{w+1}{2}\cdot 1.01\,q^2n^3/N$ and at most q inverse queries, these terms sum in expectation to at most

$$32\cdot\frac{64}{49}\cdot\frac{1.01}{2}\cdot\frac{w(w+1)n^3q^3}{N^3} = \frac{25856}{1225}\cdot\frac{w(w+1)n^3q^3}{N^3} \le 22\,\frac{w(w+1)n^3q^3}{N^3},$$

since $25856/1225 < 21.2$.

(12) *Lemma 16.* With $w^2 \le wnq$, $\mathbb{E}[\beta^2] \le 16 \cdot \frac{32}{31} \cdot 2 \cdot n^3wq^3/N^3 \le 34\, n^3wq^3/N^3$. With $A \le 1/16$ and $\phi_S \le \frac{32}{31}\phi_0 \le \frac{32}{31} \cdot \frac{1}{16}$, the chi-square bound is at most $e^{3/16}\left(1 - \frac{32}{31} \cdot \frac{1}{16}\right)^{-1} \cdot 34\, n^3wq^3/N^3 \le 44\, n^3wq^3/N^3$.

(13) *Lemma 17.* $N_{\text{HYB}} \ge N - \Phi \ge 29N/32$. Next, $\delta = 9nq\phi_0/N \le 9/1024$, because $nq \le N/64w$ and $\phi_0 \le 1/16$. Hence $\theta_j \ge \rho_b - \delta \ge 0.74$ and $Z \ge 0.74\, e_b \ge 0.72N$. The first sum is then at most $\delta^2/((29/32)^2 \cdot 0.72) \le 1.7\,\delta^2$, and the second at most $4\delta^2/(0.82 \cdot 0.72 \cdot 0.74) \le 9.2\,\delta^2$, where $0.82 \le (29/32)^2$. So $\chi^2(K\|P) \le 11\,\delta^2$, and with $\delta \le 18wn^2q^2/N^2$ and $wnq/N \le 1/64$,

$$\chi^2(K\|P) \le 11 \cdot 18^2\, \frac{w^2n^4q^4}{N^4} = 3564\, \frac{w^2n^4q^4}{N^4} \le \frac{3564}{64} \cdot \frac{wn^3q^3}{N^3} \le 56\, \frac{n^3wq^3}{N^3}.$$

(14) *Proposition 2.* With $A \le 1/16$, $\phi_S \le \frac{32}{31}\phi_0$ and $w^2 \le wnq/2$, $e^A\phi_S N \le e^{1/16} \cdot \frac{32}{31}(wnq + w^2) \le 1.65\, wnq$. With $nh \le wnq$, $2n \le wnq$ and $L \ge 63N/64$, the per-trial bound is at most $\frac{64}{63}(1 + 1 + 1.65)\, wnq/N \le 4wnq/N$.

B.2 Numerical steps in Sections 5.6 to 8

The standing bounds are those of Section B.1; in Section 7 also $\ell_\star \ge 57N/64$, $\ell'_\star \ge 55N/64$, $\varepsilon_i \le e_S \le N/64$ and, by (13), $\beta_L \le 1/64$. The following items supply the explicit constants in the proofs of Sections 6 and 7:

(1) *Proposition 4, ideal world.* With $|R| \ge nq/2$ and $e^{-A} \ge e^{-1/16}$,

$$\eta \ge \frac{e^{-1/16}}{2} \cdot \frac{w|R|}{N} \ge \frac{e^{-1/16}}{4} \cdot \frac{wnq}{N} \ge 0.2348\, \frac{wnq}{N},$$

since $e^{-1/16}/4 \ge 0.23485$.

(2) *Proposition 4, ideal world.* With $p_k \ge \gamma/2, 1 - p_k \ge (1-\gamma)/2$ and item (1), the drift is at least

$$0.2348 \cdot \frac{\gamma(1-\gamma)}{4} \cdot \frac{wnq}{N} \ge 0.0587\, \gamma(1-\gamma)\frac{wnq}{N} \ge \beta.$$

(3) *Proposition 4, advantage.* With $I = q/2$,

$$\frac{I\beta^2}{8} = \frac{q}{2} \cdot \frac{0.0586^2}{8}\gamma^2(1-\gamma)^2w^2n^2\frac{q^2}{N^2} \ge 2^{-13}\gamma^2(1-\gamma)^2w^2n^2\frac{q^3}{N^2},$$

since $0.0586^2/16 \ge 2.14 \cdot 10^{-4} \ge 2^{-13}$.

(4) *Corollary 5.* Since $\xi^4 = n_b^2N^3/\text{dens}^2$ and $r \ge 6.3\xi$, the first error term of Theorem 2 is $187(\xi/r)^4 \le 187/6.3^4 < 0.119 < 1/8$. Next, $r < 2R \le 12.6\xi + 2 \le 12.7\xi$, because $\xi \ge 2^{12}$, so $4r(w+1) \le 50.8\,(w+1)\xi \le 51\,(w+1)\xi$.

(5) *Lemma 24.* For types (1) and (3), with $q \ge 2$, $(w+1)x = \bar{F}(q+1)/N \le \frac{3}{2}\bar{F}q/N \le \frac{3}{64} \le 0.047$ and $x \le (w+1)x/2 \le 0.024$. Hence $(1-x)^{-(w+1)} \le \exp\big((w+1)x/(1-x)\big) \le e^{0.047/0.976} \le 1.05$. For type (4), $|\mathcal{T}_0| \ge \frac{93}{128}N^w$, so $2\beta_LN^w/|\mathcal{T}_0| \le \frac{256}{93}\beta_L \le 2.8\beta_L$ and $\beta_LN^w/|\mathcal{T}_0| \le \frac{128}{93}\beta_L \le 1.38\beta_L$. Finally $2.8\beta_L \le 2.8/64 \le 0.044$.

(6) *Lemma 25.* For $c = 32/31$: $c \le 1.04$, $c^2 \le 1.07$ and $c^2/2 + c/2 \le 1.05$.

(7) *Lemma 26.* As in Lemma 8, $\sum_j s_j \le N/4$ and $b_0 \le w^2 \le N/32$; with $e_S \le N/64$, $|\mathrm{Leg}_U| \ge N(1 - \frac{1}{4} - \frac{1}{32} - \frac{1}{64}) \ge 0.703N$, so $c_0 := 0.703$ and $c_0^2 \ge 0.494$. Also $N^w/|\mathcal{S}_\star| \le 1/(\frac{31}{32} - \frac{1}{64}) = \frac{64}{61}$, and the prefactor is $c_1 = \frac{64}{61} \cdot \frac{3}{c_0^2} \le \frac{64}{61} \cdot \frac{3}{0.494} \le 6.4 \le 12$. Hence the collision term is at most $6.4(4b_1^2 + B_S\varkappa_S)/N^2 = (25.6\, b_1^2 + 6.4\, B_S\varkappa_S)/N^2$.

(8) *Lemma 27, Step 3.* $(\Phi + \varepsilon_i)N/\ell_\star^2 \le (\frac{3}{32} + \frac{1}{64})(\frac{64}{57})^2 \le 0.14$, so $c_3 := 1.14$ bounds the bracket. Also $p_a \le Fh_x/\ell_\star \le (N/32)/(57N/64) = 2/57$, so $\bar\pi \ge 55/57$.

(9) *Lemma 27, Steps 5 and 6, and part (b).* With $c_3 = 1.14$, $2c_3 = 2.28$. Next $c_D = (N/\ell_\star)^4/128 \le (64/57)^4/128 \le 0.0125$, so the coefficients in $\mathrm{Var}[p_a]$ are $4c_3 + 4c_D \le 4.56 + 0.05 = 4.61$ and $4c_3 = 4.56$. In Step 6, $(N/\ell_\star)/\bar\pi^2 \le \frac{64}{57}\left(\frac{57}{55}\right)^2 = \frac{3648}{3025} \le 1.206$, and $4.61 \cdot 1.206 \le 5.6$, $4.56 \cdot 1.206 \le 5.5$ and $4(N/\ell_\star)^2 \cdot 1.206 \le 6.1$. In (b), the coincidence count contributes $\frac{6.1}{N^3} \cdot \frac{hn^2}{2}\left(\frac{hn^2}{2} + \Xi\right) \le 1.53\,\frac{n^4h^2}{N^3} + 3.05\,\frac{n^2h\Xi}{N^3}$, within the coefficients 1.6 and 3.1 of the statement.

(10) *Lemma 28.* With $\eta = 0.1001$ and $1 - \beta \ge 0.956$: $1.1001 \cdot 0.956^{-3} \le 1.2591 \le 1.26$, $(1+1/\eta)/0.956 \le 11.4959 \le 11.5$, and $(1-\beta)^{-1} - (1-\beta) = \beta(2-\beta)/(1-\beta) \le (1.956/0.956)\,\beta \le 2.1\beta$.

(11) *Lemma 29, output.* With $\ell_\star \ge 57N/64$ and $N - \Phi \ge 29N/32$, $c_2 \le \frac{64}{57}$ and $c_1 \le \frac{64}{57} + \frac{32}{29} \le 2.23$. The legal pairs contribute at most $\frac{8}{3}c_1^2 \le 13.3$ times $nh\Phi^2/N^3$ and $\frac{8}{3}c_2^2 \le \frac{8}{3}(\frac{64}{57})^2 \le 3.4$ times $n\Xi_2/N^3$; the fresh inputs at most $\frac{32}{9}c_1^2 \le 17.7$ times $nh\Phi^2/N^3$ and $\frac{32}{9}c_2^2 \le 4.5$ times $n^2\Xi^2/N^4$. So $c_o \le 13.3 + 17.7 = 31$ and

$$\chi^2(R^{\mathrm{out}} \| K^{\mathrm{out}}) \le 31\,\frac{nh\Phi^2}{N^3} + 3.4\,\frac{n\Xi_2}{N^3} + 4.5\,\frac{n^2\Xi^2}{N^4}.$$

For the conditioned output, $2.1\beta_{\max} = 5.88\beta_L$, and $1.26 \cdot 31 \le 39.1$, $1.26 \cdot 3.4 \le 4.3$, $1.26 \cdot 4.5 \le 5.7$ and $11.5 \cdot 5.88^2 \le 400$.

(12) *Lemma 29, companion.* By Lemma 24, $1/(1 - \beta_j^K) \le 1/(1 - 1.38\beta_L) \le 1/(1 - 1.38/64) \le 1.03$. With $M_{j'} \ge 0.98N$ (Section B.1 (6)) and $wnh \le N/64$, $\bar p \le wnh/0.98N \le 1/(64 \cdot 0.98)$, so $\bar p' \le 1.03\,\bar p \le 1/2$. With $\ell' \ge 7N/8$ and $\ell'_\star \ge 55N/64$, $M_{j'}\ell'\ell'_\star \ge 0.98 \cdot \frac{7}{8} \cdot \frac{55}{64}N^3 \ge 0.736N^3$, so $c_4 \le 1.36$, and with $h \le q$ the new variance term is at most $1.36\, n^2q\Xi/N^3$. Step 6 multiplies it by $8w^2$, and $1.03 \cdot 8 \cdot 1.36 \le 11.3$. The other constants of Lemma 12, with $\ell'_\star \ge 55N/64$, become $4 \cdot 8.83 \le 36$, and $32(64/55)^2 \le 43.4$ for the coefficient of W/N^2, in place of $32 \cdot \frac{64}{49} \le 41.8$ of Proposition 1.

(13) *Proposition 5.* Old terms: the ratios to the coefficients in the proof of Proposition 1 are ≤ 1 for Lemma 26, $5.6/4$ and $6.1/6.73$ for Lemma 27, $39.1/26 \le 1.51$ for K_{inv}, 1.03 and $1.03 \cdot 43.4/41.8 \le 1.07$ for the companion, and $1.6/2$ for the free-input coincidence term (against the rounded coefficient 2 of Lemma 10); all are at most 1.6. New terms: summing over at most q steps with $h \le q$ and inserting Lemma 25, we obtain the following:

- types (2)/(3): $5.5F^2q\,\mathbb{E}[A_{\mathrm{tot}}\Xi]/N^3 \le 5.5 \cdot 1.05\,F^2B_S^2q^3/N^3 \le 5.8\,F^2B_S^2q^3/N^3$;
- type (5): $(5.5n^2q\mathbb{E}[A_{\mathrm{tot}}\Xi] + 3.1n^2q^2\mathbb{E}[\Xi])/N^3 \le (5.8n^2B_S^2 + 3.3n^2B_S)q^3/N^3$, using $5.5 \cdot 1.05 \le 5.8$ and $3.1 \cdot 1.04 \le 3.3$;

- type (4): $4.3nq\mathbb{E}[\Xi_2]/N^3 \le 4.3\cdot 1.04\, nB_S^2\kappa_S q^2/N^3 \le 4.5\, nB_S^2\kappa_S q^2/N^3$; next

$$\frac{5.7n^2q\mathbb{E}[\Xi^2]}{N^4} \le \frac{(6.1q^3+6.0q^2\kappa_S)n^2B_S^2}{N^4} \le \frac{(0.05q^3+0.05q^2\kappa_S)nB_S^2}{N^3},$$

  using $5.7\cdot 1.07 \le 6.1$, $5.7\cdot 1.04 \le 6.0$ and $n/N \le 1/128$; and $11.3w^2n^2q^2\mathbb{E}[\Xi]/N^3 \le 11.3\cdot 1.04\, w^2n^2B_Sq^3/N^3 \le 11.8\, w^2n^2B_Sq^3/N^3$.

The q^3 coefficients $5.8, 5.8, 3.3 \le 3.5, 11.8$ and 0.05 are at most those in the numerator of $K^\star_{\mathrm{RH}}$, and the κ_S-terms add up to $(4.5+0.05)\, nB_S^2\kappa_Sq^2/N^3 \le 4.6\, nB_S^2\kappa_Sq^2/N^3$.

(14) *Proposition 6.* For Lemma 16: $16\cdot\frac{64}{61}\cdot 2 \le 34$ and $e^{3/16}(1-\frac{64}{61}\cdot\frac{1}{16})^{-1}\cdot 34 \le 44$. For the inverse output, the lower bounds of Lemma 17 (Section B.1 (13)) become $N_{\mathrm{HYB}_\star} \ge N-\Phi-1.38\beta_LN \ge N(1-\frac{3}{32}-\frac{1.38}{64}) \ge 0.884N$, $\theta_j \ge 0.72$ and $Z \ge 0.70N$, and the two sums are at most $(1+4/0.72)\,\delta_\star^2/(0.884^2\cdot 0.70) \le 11.99\,\delta_\star^2 \le 12\,\delta_\star^2$. By Young's inequality, $\delta_\star^2 \le 1.1\delta^2 + 11(\frac{32}{31})^2\beta_L^2$, so $12\delta_\star^2 \le 13.2\delta^2 + 143\beta_L^2$, and $13.2\delta^2 = 1.2\cdot 11\delta^2 \le 67.2\, n^3wq^3/N^3$.

(15) *Proof of Theorem 3.* $\sqrt{25.6b_I^2q/2} \le 3.6b_I\sqrt{q}$ and $\sqrt{6.4B_S\kappa_Sq/2} \le 1.8\sqrt{B_S\kappa_Sq}$ (both divided by N); $\sqrt{4.6nB_S^2\kappa_S/2} \le 1.6B_S\sqrt{n\kappa_S}$ (times $q/N^{3/2}$); and the terms $400q\beta_L^2$ and $143q\beta_L^2$ give $(\sqrt{200}+\sqrt{71.5})\beta_L\sqrt{q} \le 22.6\beta_L\sqrt{q}$.

C Derivations for the ideal world

Throughout, $\textsc{Ideal}'$ is the ideal world with $\Gamma \leftarrow \mathrm{Unif}(\mathcal{S}^{\mathcal{D}})$ and simulator $\textsc{Sim}_3$. Prefixes are augmented and contain no failure. We think of Γ as sampled in advance; lazy sampling gives the same joint law of all responses.

Proof of Lemma 15. (a) Fix the table Ξ and the simulator state before the call. Each trial samples a candidate uniformly from the L available inputs and accepts it with probability $a(v)$: with probability 1 or 0 for a family input, according to its legality under Ξ, and with probability ρ_b for a free input, by an independent coin. The table and the state do not change during the call, and the coins are fresh, so the trials are independent and identically distributed. A single trial accepts v with probability $a(v)/L$, and rejects with probability

$$1-\frac{1}{L}\sum_v a(v) = \frac{I(\Xi)+g(1-\rho_b)}{L} = R(\Xi),$$

because the inputs with $a(v) < 1$ are the $I(\Xi)$ illegal family inputs, with $1-a=1$, and the g free inputs, with $1-a = 1-\rho_b$. The call returns v at trial $r \in \{1,2,3\}$ if the first $r-1$ trials reject and trial r accepts v, which has probability $R^{r-1}a(v)/L$. Summing over r, the call returns v with probability

$$\frac{a(v)}{L}(1+R+R^2) = a(v)\,\vartheta_c(I(\Xi)),$$

and it fails, with all three trials rejecting, with probability R^3.

For the increment, write $\vartheta_c(I) = L^{-1}(1+R+R^2)$ with $R = (I+g(1-\rho_b))/L$. Increasing I by 1 increases R by $1/L$, so

$$L\vartheta_c(I+1) - L\vartheta_c(I) = \frac{1}{L} + \left(R+\frac{1}{L}\right)^2 - R^2 = \frac{1+2R+1/L}{L} \le \frac{3}{L}.$$

The last step uses $R + 1/L \le 1$, which holds whenever $I + 1$ illegal inputs actually occur, since $R + 1/L$ is then a rejection probability. Since $L\vartheta_c \ge 1$ and $\log(1+u) \le u$,

$$0 \le \log \vartheta_c(I+1) - \log \vartheta_c(I) = \log\Big(1 + \frac{L\vartheta_c(I+1) - L\vartheta_c(I)}{L\vartheta_c(I)}\Big) \le \frac{3}{L}.$$

(b) We compute the probability of the transcript τ given $\Gamma = \Xi$. It is the product, over the queries, of the probability of the recorded response given Ξ and the state before the query, because all simulator randomness is fresh for each query. The query types contribute as follows:

- A construction query at x with answer z has likelihood $\mathbf{1}[\Xi(x) = z]$.
- A forward query at an assigned input is answered deterministically, with likelihood 1.
- A forward query at a fresh family input $\iota_j(x)$ has likelihood $\mathbf{1}[\Xi(x) = z] \cdot \mathbf{1}[v_0 \in \mathrm{Leg}(z)]/|\mathrm{Leg}(z)|$, where z and v_0 are read from the augmented response. This depends on Ξ only through the revealed value $\Xi(x)$.
- A forward query at a free input has likelihood $1/|G \setminus \textsc{Ran}_b|$, independent of Ξ.
- An inverse call c returning $\iota_{j'}(x)$ has, by (a), likelihood $\mathbf{1}[\iota_{j'}(x)$ legal under $\Xi] \cdot \vartheta_c(I_c(\Xi))$. The augmented response reveals $\Xi(x)$, so the indicator equals $\mathbf{1}[\Xi(x) =$ revealed value] times a constant.
- An inverse call returning a free input has likelihood $\rho_b\, \vartheta_c(I_c(\Xi))$.

Multiplying, the likelihood is $\mathbf{1}[\Xi$ agrees with the revealed values$] \cdot W_\tau(\Xi) \cdot c_\tau$, where c_τ does not depend on Ξ. The rejected candidates of each call, family or free, do not appear in the transcript; they have been summed out inside $\vartheta_c(I_c)$. By Bayes' rule with the prior $\mathrm{Unif}(\mathcal{S}^{\mathcal{D}})$, the posterior $\Pr[\Gamma = \Xi \mid \tau]$ is proportional to $W_\tau(\Xi)$ on the tables that agree with the revealed values, and vanishes elsewhere. Since μ_τ is the uniform law on the tables in $\mathcal{S}^{\mathcal{D}}$ that agree with the revealed values, this reads

$$\Pr[\Gamma = \Xi \mid \tau] \propto \mu_\tau(\Xi)\, W_\tau(\Xi).$$

Normalising gives the formula for π_τ.

(c) By (b), before the current call the table has law π_τ. By (a), given $\Gamma = \Xi$ the call returns v with probability $a_\Xi(v)\vartheta_c(I(\Xi))$, and succeeds with probability $1 - R(\Xi)^3$. So, conditional on success, the joint law of (Γ, v) is

$$Q^+(\Xi, v) = \frac{\pi_\tau(\Xi)\, a_\Xi(v)\, \vartheta_c(I(\Xi))}{\mathbb{E}_{\pi_\tau}[1 - R^3]}.$$

Summing over Ξ, the returned input has law

$$\Pr[v \mid \text{success}] = \frac{\mathbb{E}_{\pi_\tau}[a_\Gamma(v)\vartheta_c(I(\Gamma))]}{\mathbb{E}_{\pi_\tau}[1 - R^3]} \propto \mathbb{E}_{\pi_\tau^\vartheta}[a_\Gamma(v)],$$

since $\pi_\tau^\vartheta \propto \pi_\tau\, \vartheta_c(I(\cdot))$. We evaluate $\mathbb{E}_{\pi_\tau^\vartheta}[a_\Gamma(v)]$ for the three kinds of input:

- for a free input, $a \equiv \rho_b$;
- for a construction-only input, the value $\Gamma(x)$ is fixed by τ, so legality, and hence $a \in \{0, 1\}$, is determined;
- for $v = \iota_j(x)$ with x unrevealed, $a_\Gamma(v)$ is the indicator of $\Gamma(x) \notin \mathrm{Forb}_{t,j}$, whose expectation is θ_j.

Dividing $Q^+(\Xi, v)$ by its v-marginal gives the conditional law of Γ given success and v: it is proportional to $\pi_\tau^\vartheta(\Xi)\, a_\Xi(v)$, which for $v = \iota_j(x)$ is π_τ^ϑ conditioned on $\Gamma(x) \notin \mathrm{Forb}_{t,j}$, as claimed.

It remains to show that θ_j does not depend on the unrevealed x. It suffices that π_τ^ϑ is exchangeable in the unrevealed entries, that is, invariant under permuting them. The prior μ_τ is invariant, since it makes these entries independent and uniform. The weight $\prod_{c'} \vartheta_{c'}(I_{c'})$ is invariant as well, because each $I_{c'}(\Xi)$ is a symmetric function of the unrevealed entries:

- every unrevealed x was available at every earlier call, since assigned inputs are revealed;
- so $I_{c'}(\Xi)$ is the sum, over the unrevealed x, of $I_{c'}(x; \Xi(x))$, plus a term fixed by τ (the free inputs contribute the Ξ-independent term $g(1-\rho_b)$ to the rejection probability, and nothing to $I_{c'}$);
- and $I_{c'}(x;\cdot)$ is the same function of the value for every x, because every x has exactly one input $\iota_j(x)$ for each call j, which does not use balance.

Finally, after the response is appended to τ, the current call joins $\mathcal{I}_\tau$. Statement (b) for the longer prefix is consistent with the conditional law just derived, because the revealed value of x now fixes the legality indicator. □

D Deferred proofs for Sections 4 and 6

D.1 Proof of Lemma 5

Proof of Lemma 5. In each case we give a bijection $\pi \mapsto \pi'$ of $\mathrm{Perm}(G)^k$ with $\mathrm{F}^\pi = \mathrm{F}'^{\pi'}$ (up to a fixed bijection of the outputs or inputs that the distinguisher can apply itself), and a translation of the primitive interfaces at no query cost. This gives equality of the pseudorandomness advantages. For indifferentiability we couple the two ideal worlds so that the simulators' answers correspond pathwise, by induction over the queries and over the trials inside each query. Then every distinguisher against one pair (construction, simulator) is a distinguisher against the other, with the same advantage, and conversely.

(a) Put $\pi'_b := \theta_b \circ \pi_b$. Since $\lambda_j \pi_b = (\lambda_j \theta_b^{-1})(\theta_b \pi_b)$, the masked outputs, and hence the construction values, are the same. The interfaces are related by $\pi_b(u) = \theta_b^{-1}(\pi'_b(u))$ and $\pi_b^{-1}(t) = \pi'^{-1}_b(\theta_b(t))$. In the ideal worlds, use the same oracle table, and maintain the invariant that $P'_b = \theta_b \circ P_b$, with the same domains and with ranges $R'_b = \theta_b(R_b)$. The masked ranges $\lambda'_j(R'_b) = \lambda_j(R_b)$ are unchanged, so the legal sets agree, and so does the set $\mathcal{S}$, because the rigid pairs are unchanged. The queries are then coupled as follows:

- A fresh family forward query reads the same oracle entry. The anchor draw is coupled identically, and the raw outputs satisfy $p'_j = \lambda'^{-1}_j(v_j) = \theta_b \lambda_j^{-1}(v_j) = \theta_b(p_j)$.
- A free forward answer is uniform on $G \setminus R'_b = \theta_b(G \setminus R_b)$, so it is coupled as $v' = \theta_b(v)$.
- An inverse query at $t' = \theta_b(t)$ uses the same available inputs, so the trials are coupled identically. A family candidate has the same forced anchor, $\lambda'_j(t') - y_j = \lambda_j(t) - y_j$, and therefore the same legality. The coin uses the same call j_b and the same dummy value ζ, and $\lambda'_{j_b}(\theta_b(t)) - \zeta_{j_b} = \lambda_{j_b}(t) - \zeta_{j_b}$, so the coin decisions agree. The weight ρ_b depends only on the range sizes and on $|\mathcal{S}|$, which are the same.

This works for non-commuting automorphisms, in the order displayed.

(b) Let $L : G^w \to G^w$ be the map that takes the differences $y_j = v_j - v_0$ to the differences $y'_j := v_j - v_r$ for $j \neq r$; explicitly $y'_j = y_j - y_r$, with $y_0 = 0$. It is an invertible integer linear map, and $\mathrm{F}' = L \circ \mathrm{F}$ with the same permutations. The distinguisher applies L^{-1} to construction answers. The map L sends $\mathcal{S}$ onto the good set of F', because $y_j = y_{j'}$ iff $y'_j = y'_{j'}$, and $y_j = y_r$ iff $y'_j = 0$. The new anchor is $a' = a + y_r = v_r$, and $a \in \mathrm{Leg}(y)$ iff $a' \in \mathrm{Leg}'(L(y))$, since both describe the same tuple of masked outputs. In the ideal worlds put $\Gamma' := L \circ \Gamma$, which is again uniform, and keep the simulator states equal: the raw outputs are identical. At an inverse query for call j, the forced anchor transforms as $\lambda_j(t) - y'_j = (\lambda_j(t) - y_j) + y_r$, so legality agrees. For the coin, couple the uniform proposals of the dummy sampler through L, so that $\zeta' = L(\zeta)$; goodness and acceptance then agree.

(c) For every class b, the rule $\beta_b(\iota_j(x)) := \iota'_j(x)$ for $c_j = b$ is a bijection between the used subsets, because the images are disjoint within the class and each injection is one-to-one. The two used subsets have the same size $n_b|\mathcal{D}|$, so the complements have the same size, and we extend β_b arbitrarily to a bijection of G. It need not be a group automorphism. (If $\mathcal{D}' \neq \mathcal{D}$, first identify $\mathcal{D}'$ with $\mathcal{D}$ by a fixed bijection, which the distinguisher

applies to construction inputs and through which the oracle table is transported.) Put $\pi'_b := \pi_b \circ \beta_b^{-1}$. The construction outputs agree at every x, and the interfaces are related by $\pi'_b(u') = \pi_b(\beta_b^{-1}(u'))$ and $\pi'^{-1}_b(t) = \beta_b(\pi_b^{-1}(t))$. In the ideal worlds, use the same oracle table and maintain $P'_b(\beta_b(u)) = P_b(u)$, $\textsc{Dom}'_b = \beta_b(\textsc{Dom}_b)$ and $\textsc{Ran}'_b = \textsc{Ran}_b$, with the same failure flag. The answers then agree for the following reasons:

- Stored answers agree, up to β_b for inverse look-ups.
- The map β_b sends free inputs to free inputs, and the ranges are equal, so a free forward answer uses the same sampled output.
- A fresh family forward query at $\iota'_j(x) = \beta_b(\iota_j(x))$ concerns the same call j and input x. It reads the same oracle entry; the good set and the legal sets agree because the masks and ranges agree. The anchor draw is coupled identically.
- In an inverse query, couple each trial's candidate as $u' = \beta_b(u)$. This maps the uniform law on the available inputs onto its counterpart. A family candidate keeps its call and its input x, so it reads the same oracle entry and has the same legality; a free candidate stays free.
- The coin uses the same call j_b and the same successive proposals from G^w for the dummy sampler. The good set is the same, so the same proposals are rejected and the same ζ is accepted. The target, masks and ranges are equal, so the acceptance test $\lambda_{j_b}(t) - \zeta_{j_b} \in \mathrm{Leg}(\zeta)$ has the same outcome. The coins therefore agree pathwise; dummy values never read Γ.
- All trial decisions agree, so failures agree, both for finitely many trials and for capped samplers with equal caps.

(d) Rename every index-dependent object consistently: the class map, the masks, the injections, the order of the output blocks, and the call j_b used by the coin. The coin may use any call of its class, by the paragraph on implementing the coin after Definition 11.

Invariants. Under (a), $\lambda'^{-1}_j = \theta_b \lambda_j^{-1}$, so $\mu'_{jj'} = \theta_b \mu_{jj'}$ for a same-class pair in class b. This is 0 or an automorphism exactly when $\mu_{jj'}$ is, and $\lambda'_j = \lambda'_{j'}$ iff $\lambda_j = \lambda_{j'}$. The other moves do not change the masks or the classes, and (c) keeps $|\mathcal{D}|$. □

D.2 Proof of the matching attack

This appendix proves Theorem 2 and Corollary 6, with the notation of Section 6.2; the plan is in Section 6.2. The proof does not use admissibility, so it covers arbitrary automorphism masks. Its inputs are the following:

- the forbidden density of a fixed target: each same-class pair constraint is a set of density exactly $1/N$, for regular and singular pairs alike, as shown in the proof of Proposition 6;
- the tuple bijection and the index-free weight of Lemma 4, which hold for arbitrary masks;
- the uniform output marginals and the separating shifts;
- the posterior, dense-event, remainder, hit and concentration estimates below, which are stated for any masks.

None of these needs the inverse of a mask difference. The attack also never needs a lower bound on $|\mathrm{Leg}(y)|$ for every y, since its forward queries only complete tuples that are already assigned; so the possibly empty legal sets of Remark 6 do not affect it.

Small quantities. We first record the bounds on small quantities that the regime (R1)–(R2) gives, and that the rest of the proof uses repeatedly:

- *Forbidden density.* Put $\phi_0 := (ws_{\max} + w^2)/N$, $\phi_{\mathcal{S}} := \frac{32}{31}\phi_0$ and $\phi_\nu := 1.1\,\phi_0$. By (R1) and (R2), $\phi_0 \le \frac{4}{600n^2 \ln N} + \frac{1}{1000n} \le \frac{0.0016}{n}$, so $n\phi_\nu \le 0.0018$.

- *Ranges.* $s_{\max} = 4rn \le N/150n^2 w \ln N \le N/1500$.
- *Weight.* By Lemma 4, $1 - \rho_{b''} \le \phi_{\mathcal{S}}$ for every class, so $\rho_{b''} \ge 3/4$.
- *Density of a forbidden set under ν.* Let ν be uniform on $\mathcal{S} \setminus \mathrm{Forb}_{t,j''}$ for some target and call. Since $|\mathcal{S}| \ge (1-w^2/N)N^w$ and $|\mathcal{S} \setminus \mathrm{Forb}| \ge (1-\phi_{\mathcal{S}})|\mathcal{S}|$, every $B \subseteq G^w$ satisfies $\nu(B) \le |B|/0.999(1-\phi_{\mathcal{S}})N^w$. In particular $\nu(\mathrm{Forb}_{t',j'''}) \le \phi_\nu$ for every target t' and call j'''.
- *Output marginals under ν.* By the bijection in the proof of Lemma 4, ν corresponds to the uniform law on class-wise injective output tuples that avoid the ranges and have $p_{j''} = t$. Hence the output of any other call j''' is uniform on $G \setminus \textsc{Ran}_{c(j''')}$, with t also excluded if $c(j''') = c(j'')$.

Real world. In the real world the statistic has mean zero, because the permutations are fixed in advance and nothing is hidden from the distinguisher.

Lemma 32. *In* Real, $\mathbb{E}[T_k \mid \tau_k] = 0$ *for every useful k, where τ_k is the transcript up to the answer $\iota_j(x_k)$. Hence $\mathbb{E}_{\textsc{Real}}[T^2] \le 1.011\, r^2/N$, and $\Pr_{\textsc{Real}}[\mathcal{A} = 1] \le 1.011\, r^2/N\theta^2$.*

Proof. There are no hidden assignments. Given τ_k, the input $\iota_{j'}(x_k)$ of $\pi_{b'}$ is unassigned; only $\iota_j(x_k)$ has been assigned, to t_k. So $p_{j'}(x_k)$ is uniform on $G \setminus \textsc{Ran}_{b'}$, where $\textsc{Ran}_{b'}$ contains t_k when $b' = b$, and $\Pr[p \in C_k^* \mid \tau_k] = \pi_k$. The T_k form a martingale difference sequence with $\mathbb{E}[T_k^2 \mid \tau_k] \le \pi_k \le r/(N - s_{\max}) \le 1.011\, r/N$. There are at most r useful iterations, Real never answers $\perp$, and Chebyshev's inequality gives the last claim. □

Ideal world. The analysis of the ideal world takes the rest of this subsection: we simplify the world, bound the tilt of one hidden value, and use the tilt to bound the hits, the drift and the good event. We first replace the ideal world by a simpler one, in which the oracle takes values in $\mathcal{S}$ and the simulator never runs out of trials.

Lemma 33. *Let* $\textsc{Ideal}'_\infty$ *be the ideal world in which Γ is uniform from $\mathcal{D}$ to $\mathcal{S}$, and the simulator repeats trials until a candidate is accepted (answering $\perp$ only if no available input has positive acceptance weight). Then*

$$\Pr_{\textsc{Ideal}_m}[\mathcal{A} = 1] \ge \Pr_{\textsc{Ideal}'_\infty}[\mathcal{A} = 1] - 4mr|\Lambda|/N.$$

Proof. Forward queries of $\mathcal{A}$ only complete inputs that the simulator has already assigned, so they read no oracle value. Each of the fewer than $4r$ inverse queries reads at most m values. Resampling the structurally bad entries therefore changes the transcript with probability at most $4mr|\Lambda|/N$; this is (12).

Couple $\textsc{Ideal}'_m$ with $\textsc{Ideal}'_\infty$ through the same Γ, and the same candidates and coins in the first m trials of every inverse query. The coin is independent of Γ, so the coupling is legitimate. The two worlds agree until $\textsc{Ideal}'_m$ fails, and after that $\mathcal{A}$ outputs 1. This is the argument in the proof of [Kim26, Theorem 2]. □

From now on we work in $\textsc{Ideal}'_\infty$. A fresh inverse call c, at target t_c on class b_c, returns an available input u with probability $a_\Phi(u)/V_c(\Phi)$. Here Φ is the table, $a_\Phi(u)$ is 1 for a legal family input, 0 for an illegal one, and ρ_{b_c} for a free input, and $V_c(\Phi) := \sum_u a_\Phi(u)$.

Posterior. Let τ be a transcript without $\perp$. Either it ends between iterations, or it ends just after the answer $u = \iota_{j''}(x)$ of an inverse call c° at target t°, before u is completed; the points τ_k of Phase 3 are of the second kind. Let J be the set of its inverse calls, so $|J| < 4r$. Every returned input is revealed, except the last one in the second case. Let μ_τ be the law of tables that agree with the revealed values and are independent

uniform on $\mathcal{S}$ elsewhere, conditioned in the second case on $\Phi(x) \notin \mathrm{Forb}_{t^\circ, j''}$. (If the last answer is a free input there is nothing to condition.) Then

$$\Pr[\Gamma = \Phi \mid \tau] = \frac{\mu_\tau(\Phi) W(\Phi)}{\mathbb{E}_{\mu_\tau}[W]}, \qquad W(\Phi) := \prod_{c \in J} V_c(\Phi)^{-1}.$$

To see this, compute the probability of τ given $\Gamma = \Phi$, as in the proof of Lemma 15(b). Forward queries of $\mathcal{A}$ are look-ups, so only the inverse calls contribute. Inverse call c returns its recorded input u_c with probability $a_\Phi(u_c)/V_c(\Phi)$. The numerator $a_\Phi(u_c)$ does not depend on the unrevealed entries: for a free input it is ρ_{b_c}, for a revealed family input it is fixed by its revealed value, and for the last, uncompleted answer it is the indicator of its legality, which is accounted for by the conditioning in μ_τ. The completions are deterministic given Φ. So the likelihood is proportional to $\mathbf{1}[\Phi \text{ agrees with } \tau]\, W(\Phi)$, and Bayes' rule with the uniform prior on $\mathcal{S}^{\mathcal{D}}$ gives the formula.

An unrevealed x was available at every call. At call c it contributes $\ell_c(x; z) := n_{b_c} - I_c(x; z)$ to V_c, where

$$I_c(x; z) := \sum_{j'' \in c^{-1}(b_c)} \mathbf{1}[z \in \mathrm{Forb}_{t_c, j''}]$$

and Forb is taken with the ranges at the time of call c. Write $V_c = V_c^{(-x)} + \ell_c(x; \Phi(x))$.

Lemma 34. *For every such τ and every unrevealed x (including the last returned input in the second case), let Dense_x be the event that $V_c^{(-x)} \geq N/2$ for all $c \in J$. Then $\Pr[\mathrm{Dense}_x^c \mid \tau] \leq \eta := e^{-N/8n} \leq N^{-5}$.*

Proof. $V_c^{(-x)}$ is at least

$$\rho_{b_c} g_c + \sum_{x'} \ell_c(x'; \Phi(x')),$$

where $g_c \geq |\mathrm{Free}_{b_c}| - 4r$ counts the free inputs available at call c, and the sum is over the unrevealed $x' \neq x$ other than the last returned input (in the second case). We drop that input, whose entry is conditioned under μ_τ; all terms are nonnegative, so this only lowers the bound. There are at least $|\mathcal{D}| - 4r - 2$ remaining x'. Under μ_τ their entries are independent and uniform on $\mathcal{S}$, so these terms are independent, lie in $[0, n_{b_c}]$, and have mean at least $n_{b_c}(1 - \phi_{\mathcal{S}})$. With $\rho \geq 1 - \phi_{\mathcal{S}}$,

$$\mathbb{E}_{\mu_\tau}[V_c^{(-x)}] \geq (1 - \phi_{\mathcal{S}})\big(N - 4r(n+1) - 2n\big) \geq \tfrac{7}{8} N.$$

Hoeffding's inequality, with at most N/n_{b_c} terms of range n_{b_c}, gives

$$\mu_\tau\big(V_c^{(-x)} < N/2\big) \leq \exp\big(-2(3N/8)^2 / N n_{b_c}\big) \leq e^{-N/4n}.$$

A union bound over J gives $\mu_\tau(\mathrm{Dense}_x^c) \leq 4re^{-N/4n}$.

On compatible tables, every factor V_c is at least the weight of its returned input, hence at least $3/4$; and every factor is at most N. So $N^{-|J|} \leq W \leq (4/3)^{|J|}$, and

$$\Pr[\mathrm{Dense}_x^c \mid \tau] \leq (4N/3)^{4r} \cdot 4re^{-N/4n} \leq e^{-N/8n},$$

because $\ln(4r) + 4r \ln(4N/3) \leq 4.5\, r \ln N \leq N/8n$ by (R2). Finally $N/8n \geq 5 \ln N$ by (R1) and (R2). ∎

Tilt of one value. The rejections made by the simulator tilt the conditional law of each hidden value. The next lemma puts this tilt in a threshold form and bounds it.

Lemma 35. *Fix τ, an unrevealed x, and the other entries of the table, on* Dense_x. *Let ν be uniform on the values $z \in \mathcal{S}$ allowed for x by τ. That is all of $\mathcal{S}$, except when x is the last returned input $\iota_{j''}(x)$ in the second case, where it is $\mathcal{S} \setminus \mathrm{Forb}_{t^\circ, j''}$ (at τ_k: $\mathcal{S} \setminus \mathrm{Forb}_{t_k, j}$). The conditional law of $z = \Gamma(x)$ is $\tilde{\nu} \propto \nu e^{\Psi}$, where*

$$\Psi(z) = \sum_{c \in J} \sum_{s \geq 1} \delta_{c,s} \mathbf{1}[I_c(x; z) \geq s], \qquad \delta_{c,s} = \log \frac{V_c^{(-x)} + n_{b_c} - s + 1}{V_c^{(-x)} + n_{b_c} - s} \in \left[\frac{1}{N+1}, \frac{2}{N}\right].$$

In addition, $0 \leq \Psi \leq \omega := 8rn/N \leq 0.002$, and $e^{-\omega} \leq \tilde{\nu}/\nu \leq e^{\omega}$.

Proof. With the other entries fixed, the posterior weight of the value z of x is $\nu(z)W$, and $W = \prod_c V_c^{-1}$ with $V_c = V_c^{(-x)} + n_{b_c} - I_c(x; z)$, since the x-inputs that are legal at call c each contribute 1 to V_c. Write each factor as a telescoping product from $V_c^{(-x)} + n_{b_c}$ down to $V_c^{(-x)} + n_{b_c} - I_c(x; z)$:

$$\begin{aligned} \frac{1}{V_c^{(-x)} + n_{b_c} - I_c(x; z)} &= \frac{1}{V_c^{(-x)} + n_{b_c}} \prod_{s=1}^{I_c(x;z)} \frac{V_c^{(-x)} + n_{b_c} - s + 1}{V_c^{(-x)} + n_{b_c} - s} \\ &= \frac{\exp\left(\sum_{s \geq 1} \delta_{c,s} \mathbf{1}[I_c(x; z) \geq s]\right)}{V_c^{(-x)} + n_{b_c}}. \end{aligned}$$

The prefactor does not depend on z, so the weight of z is proportional to $\nu(z)e^{\Psi(z)}$. Each $\delta_{c,s}$ is positive: a value that is illegal at more calls lowers the normalisers V_c, and so makes the observed answers more likely, as explained in Section 6.2.

At the call that returned x (if any), the returned input is legal, so there $I_c(x; z) \leq n_{b_c} - 1$ and the sum runs over $s \leq n_{b_c} - 1$. On Dense_x, the arguments $V^{(-x)} + n_{b_c} - s$ lie in $[N/2, N]$: they are at least $V^{(-x)} \geq N/2$, and at most the number of available inputs. Since $\log \frac{y+1}{y} \in [\frac{1}{y+1}, \frac{1}{y}]$, this gives $\delta_{c,s} \in [\frac{1}{N+1}, \frac{2}{N}]$. Each call contributes at most n increments, so $\Psi \leq |J|\, n \cdot 2/N < 8rn/N$. Finally, Ψ has oscillation at most ω on the support of ν, so the ratio bound $e^{-\omega} \leq \tilde{\nu}/\nu \leq e^{\omega}$ follows from the reweighting lemma, Lemma 1. □

Lemma 36 ([Kim26, App. E.4, inequality preceding (32)])**.** *Let a, g have mean 0 under a probability law ν. Put $\omega := \operatorname{osc} g$, $Z := \mathbb{E}_\nu[e^g]$ and $R_g := e^g/Z - 1 - g$. Then*

$$|\mathbb{E}_\nu[a R_g]| \leq \tfrac{1}{2} e^{\omega}(1 + \omega)\, \omega \sqrt{\operatorname{Var}_\nu(a) \operatorname{Var}_\nu(g)}.$$

Proof. Put $h := e^g - 1 - g$. Then $R_g = (h - (Z-1) - (Z-1)g)/Z$, and since $\mathbb{E}[a] = 0$, $\mathbb{E}[a R_g] = (\mathbb{E}[ah] - (Z-1)\mathbb{E}[ag])/Z$. We note the following bounds:

- $Z \geq 1$, by Jensen's inequality;
- $|g| \leq \omega$, since g has mean 0 and oscillation ω;
- $0 \leq h \leq \frac{1}{2} g^2 e^{|g|} \leq \frac{1}{2} e^{\omega} \omega |g|$, so $\mathbb{E}[h^2]^{1/2} \leq \frac{1}{2} e^{\omega} \omega \sigma$, where $\sigma^2 := \operatorname{Var} g$, and $0 \leq Z - 1 = \mathbb{E}[h] \leq \frac{1}{2} e^{\omega} \sigma^2$.

By Cauchy–Schwarz and $\sigma \leq \omega$,

$$|\mathbb{E}[a R_g]| \leq \sqrt{\operatorname{Var} a}\left(\tfrac{1}{2} e^{\omega} \omega \sigma + \tfrac{1}{2} e^{\omega} \sigma^2 \cdot \sigma\right) \leq \tfrac{1}{2} e^{\omega}(1 + \omega) \omega \sigma \sqrt{\operatorname{Var} a}.$$ □

Hits. The tilt is small enough that a companion output assigned by the simulator lands in a fixed small set with about the same probability as a uniform output.

Lemma 37. *Let $B \subseteq G$ be a fixed set of size at most βr, with $\beta \in \{1, 2\}$. Consider any output that the simulator assigns as a companion, that is, any output other than the target, of an input returned by a fresh inverse call. Given the past, the probability that it lies in B is at most $1.04\,\beta r/N$. In particular $\rho_k := \Pr[\Gamma(x_k) \in E \mid \tau_k] \le 1.04\,r/N$.*

Proof. Given the transcript including the answer, the value of the returned input has law $\tilde{\nu}$, with $\tilde{\nu}/\nu \le e^{\omega}$ on Dense (Lemma 35). Under ν, a companion output is uniform on at least $N - s_{\max} - 1$ values (Section D.2). Hence the probability is at most $e^{\omega}\beta r/(N - s_{\max} - 1) + \eta \le 1.04\,\beta r/N$. For ρ_k, take $B = C_k^*$ and the companion $p_{j'}(x_k)$, using the alignment. □

Drift. At a useful iteration, the statistic has a positive conditional mean in the ideal world, which is the source of the distinguishing advantage.

Lemma 38. *Let k be useful, and suppose that τ_k satisfies $|J_2| \ge 7r/8$, where J_2 is the set of Phase-2 inverse calls, and $|C^{\cup} \cap \textsc{Ran}_{b'}| \le 3r/16$. Then*

$$d_k := \Pr[\Gamma(x_k) \in E \mid \tau_k] - \pi_k \ge 0.6\,r^2/N^2.$$

Proof. Let ν be uniform on $\mathcal{S} \setminus \mathrm{Forb}_{t_k,j}$. By Section D.2 and the alignment of Section 6.2, $\nu(E) = \nu(p_{j'} \in C_k^*) = \pi_k$. Since $C_k^* \subseteq C^{\cup}$ has size r,

$$\tfrac{13}{16}\,r/N \le \pi_k \le 1.011\,r/N.$$

Fix the other entries on Dense_{x_k}, and write $\tilde{\nu} \propto \nu e^{\Psi}$ (Lemma 35). Put $a := \mathbf{1}_E - \pi_k$ and $g := \Psi - \mathbb{E}_{\nu}[\Psi]$, so that $\mathbb{E}_{\nu}[a] = \mathbb{E}_{\nu}[g] = 0$, and $Z := \mathbb{E}_{\nu}[e^{g}]$. Since $\tilde{\nu} = \nu e^{g}/Z$,

$$\tilde{\nu}(E) - \pi_k = \mathbb{E}_{\nu}\Big[a\,\frac{e^{g}}{Z}\Big] = \mathbb{E}_{\nu}[a(1 + g + R_g)] = \mathbb{E}_{\nu}[a\Psi] + \mathbb{E}_{\nu}[aR_g],$$

where $R_g := e^{g}/Z - 1 - g$ is the remainder of Lemma 36, and we used $\mathbb{E}_{\nu}[a] = 0$ twice. The first term is the first-order drift, which we now bound below; the second is the nonlinear remainder. We compute the first-order term

$$\mathbb{E}_{\nu}[(\mathbf{1}_E - \pi_k)\Psi] = \sum_{c} \sum_{s} \delta_{c,s}\big[\nu(E \cap \{I_c \ge s\}) - \pi_k \nu(I_c \ge s)\big]$$

by kind of call:

- For a Phase-2 call c, the target t_c lies in A_2, so $\lambda_j(t_c) \in c_1 + P$. Phase 1 is complete, so $A_1 \subseteq \textsc{Ran}_{b'}$, that is $I \subseteq \lambda_{j'}(\textsc{Ran}_{b'})$. By the alignment, $E \subseteq \mathrm{Forb}_{t_c,j} \subseteq \{I_c \ge 1\}$, which gives the following:
  - since $\nu(E) = \pi_k$, the $s = 1$ term equals $\delta_{c,1}\pi_k(1 - \nu(I_c \ge 1))$;
  - the terms with $s \ge 2$ are at least $-\delta_{c,s}\pi_k\nu(I_c \ge s)$.

  By Lemma 35, $\delta_{c,s} \in [\frac{1}{N+1}, \frac{2}{N}]$, and $\sum_s \nu(I_c \ge s) = \mathbb{E}_{\nu}[I_c] \le n\phi_{\nu}$. So the call contributes at least $\pi_k(\frac{1}{N+1} - \frac{2}{N}n\phi_{\nu})$.
- Every other call contributes at least $-\frac{2}{N}\pi_k n\phi_{\nu}$.

With $|J_2| \ge 7r/8$, $|J_2| \le r$, fewer than $3r$ other calls ($|A_1| < 2r$ in Phase 1, at most r in Phase 3), and $n\phi_\nu \le 0.0018$:

$$\mathbb{E}_\nu[(\mathbf{1}_E - \pi_k)\Psi] \ge \pi_k \frac{r}{N}\left(\frac{7}{8} \cdot 0.9999 - 2 \cdot 0.0018 \cdot 4\right) \ge 0.86\,\pi_k \frac{r}{N}.$$

For the remainder term, apply Lemma 36 to these a and g. Since $0 \le \Psi \le \frac{2}{N}\sum_c I_c$, $\sum_c I_c \le 4rn$ and $\mathbb{E}_\nu\left[\sum_c I_c\right] \le 4rn\phi_\nu$,

$$\operatorname{Var} g \le \mathbb{E}_\nu[\Psi^2] \le \frac{4}{N^2}\, 4rn \cdot 4rn\phi_\nu = \frac{64 r^2 n^2 \phi_\nu}{N^2}.$$

Also $\operatorname{Var} a \le \pi_k$, $\omega \le 8rn/N$, and $\sqrt{\pi_k} \ge \sqrt{13r/16N}$. So the remainder is at most $41.5\, n^2\sqrt{r\phi_\nu/N} \le 0.022$ times the first-order term: with $s_{\max} = 4rn$, $r \le N/600n^3 w \ln N$ by (R2) and $N \ge 2^{16}$, $N \ge 1000nw^2$ by (R1), we get $r\phi_\nu/N = 1.1\,(4nwr^2 + w^2 r)/N^2$, and $41.5\, n^2$ times its square root is at most 0.022 even when the two parts are bounded separately. Since $0.86\,(1 - 0.022) \ge 0.84$,

$$\tilde{\nu}(E) - \pi_k \ge 0.84\,\pi_k\, r/N \ge 0.68\, r^2/N^2.$$

Averaging over the other entries, and charging Dense^c at most η, gives $d_k \ge 0.68r^2/N^2 - 2\eta \ge 0.6\, r^2/N^2$. □

Good event. The drift bound needs a few conditions on the transcript, which hold with high probability.

Lemma 39. *In* Ideal$'_\infty$, *the probability of an execution without* $\perp$ *in which one of the following fails is at most* $251\, nr/N + e^{-r\mathrm{dens}^2/17 n_b^2}$:

(G1) at most $r/8$ *Phase-2 targets are skipped;*
(G2) $|C^\cup \cap \mathrm{Ran}_{b'}| \le 3r/16$ *at the end;*
(G3) at most $r/8$ *Phase-3 targets are skipped;*
(G4) at least $0.35\,\mathrm{dens} r/n_b$ *Phase-3 iterations are useful.*

Proof. Call an output a *companion* if the simulator assigns it to an input returned by an inverse call and it is not that call's target. For (G1), a Phase-2 target $t \in A_2$ is in Ran_b before its turn only in two ways:

- It is a Phase-1 target. This needs $b = b'$, and happens for at most $|A_1 \cap A_2| \le r/16$ targets, by (S2).
- A companion of class b landed on it. Earlier Phase-2 targets are distinct from t.

Phases 1–2 have fewer than $3r$ inverse calls, hence fewer than $3rn$ companions of any one class. By the hit bound, the expected number of landings in A_2 is at most $3rn \cdot 1.04r/N$. Markov's inequality at $r/16$ gives probability at most $50\, nr/N$ that more than $r/16$ targets are hit, and otherwise at most $r/8$ are skipped. The same argument applies to (G3), since $A_2 \cap A_3 = \emptyset$ and $|A_1 \cap A_3| \le r/16$; with fewer than $4rn$ companions, the probability is at most $16 \cdot 4 \cdot 1.04\, nr/N \le 67\, nr/N$.

For (G2), note that $C^\cup$ is disjoint from A_1, and meets A_2 and A_3 in at most $r/16$ points each when $b' = b$. Since $|C^\cup| < 2r$, the hit bound with $\beta = 2$ shows that companion landings in $C^\cup$ exceed $r/16$ with probability at most $16 \cdot 4 \cdot 2.08\, nr/N \le 134\, nr/N$.

For (G4), given the transcript before it, a fresh Phase-3 inverse query returns a call-j input with probability at least $\mathbb{E}[\sum_{x'} \mathbf{1}[\Gamma(x') \notin \mathrm{Forb}_{t,j}]/V_c] \ge \frac{|\mathcal{D}|-4r}{N}(e^{-\omega}(1-\phi_{\mathcal{S}}) - \eta) \ge p_0 := 0.8\mathrm{dens}/n_b$, using $V_c \le N$, Lemma 35 and (R3). We do not condition on the future events (G3) or no failure. Instead, for the r planned Phase-3 iterations, let A_k indicate that iteration k is queried (its target is fresh and no failure has occurred), which is known before the query, and let U_k indicate that it is useful; after a failure set $A_k = U_k = 0$. Then $\mathbb{E}[U_k \mid \mathcal{F}_{k-1}] \ge p_0 A_k$, where $\mathcal{F}_{k-1}$ is the transcript before iteration k. So $S := \sum_{k \le r}(U_k - p_0 A_k)$ is a

submartingale with increments in $[-1,1]$. On a run without failure that satisfies (G3), $\sum_k A_k \ge 7r/8$, so fewer than $0.35\mathrm{dens}r/n_b$ useful iterations force $S < 0.35\mathrm{dens}r/n_b - 0.7\mathrm{dens}r/n_b = -0.35\mathrm{dens}r/n_b$. By Azuma's inequality this has probability at most $\exp(-0.35^2\mathrm{dens}^2 r/2n_b^2) \le e^{-r\mathrm{dens}^2/17n_b^2}$. Summing gives the claim, since $50 + 67 + 134 = 251$. □

Proof of Theorem 2. Index the r planned Phase-3 iterations by k, and let K be the set of useful ones. Write $T = \sum_{k\in K} d_k + M$ in $\textsc{Ideal}'_\infty$, where $M := \sum_{k\le r} M_k$, with $M_k := \mathbf{1}[\Gamma(x_k) \in E] - \rho_k$ if k is useful and $M_k := 0$ if k is skipped, is not useful, or comes after a failure. The test outcome of a useful k is revealed by its completion, before iteration $k+1$. Conditioning on τ_k centres M_k, and the tower property carries this to the transcript before iteration k. So M is a martingale with respect to the pre-iteration transcripts, with $\mathbb{E}[M_k^2 \mid \tau_k] \le \rho_k$. Hence $\mathbb{E}[M^2] \le \sum_k \mathbb{E}[\rho_k \mathbf{1}[k \in K]] \le 1.04\, r^2/N$ by Lemma 37, without conditioning on the good event. Here $T_k = \mathbf{1}[\Gamma(x_k) \in E] - \pi_k$ by the alignment. The bound on $\Pr_{\textsc{Real}}[\mathcal{A} = 1]$ is Lemma 32, and the reduction to $\textsc{Ideal}'_\infty$ is Lemma 33. On the good event every useful k satisfies the hypotheses of Lemma 38, so

$$\sum_{k\in K} d_k \ge 0.35\frac{\mathrm{dens}r}{n_b}\cdot 0.6\frac{r^2}{N^2} = 0.21\frac{\mathrm{dens}r^3}{n_b N^2} = 2\theta.$$

Suppose that the run has no $\perp$, lies in the good event, and $\mathcal{A}$ outputs 0. Then $T < \theta$ and $\sum_{k\in K} d_k \ge 2\theta$, so $M = T - \sum_k d_k < -\theta$. By Chebyshev's inequality and the second-moment bound for M, this has probability at most $\mathbb{E}[M^2]/\theta^2 \le 1.04\, r^2/N\theta^2$. In $\textsc{Real}$, $\mathcal{A}$ outputs 1 with probability at most $1.011\, r^2/N\theta^2$ by Lemma 32. Hence

$$\mathbf{Adv} \ge 1 - \frac{4mr|\Lambda|}{N} - \frac{(1.011+1.04)\, r^2}{N\theta^2} - 251\frac{nr}{N} - e^{-r\mathrm{dens}^2/17n_b^2} - N^{-2},$$

and

$$\frac{2.051\, r^2}{N\theta^2} = \frac{2.051\, n_b^2 N^3}{0.105^2\mathrm{dens}^2 r^4} \le \frac{187\, n_b^2 N^3}{\mathrm{dens}^2 r^4}.$$ □

Fork constructions. Corollary 4 bounds the fork constructions of [DGL22] against the simulator obtained by wrapping $\textsc{Sim}_3$ through Lemmas 18 to 20. Those lemmas turn a fork distinguisher into a family distinguisher, which is the direction an upper bound needs. Here we prove the converse for each of the three steps. The attack therefore transfers to the forks, against the same wrapped simulators.

Augmented form of the attack. The converses below are stated for augmented distinguishers, so we first check that the attack can be run in the augmented interface without changing its outputs.

Lemma 40. *Let $\mathcal{A}^{\mathrm{aug}}$ be $\mathcal{A}$ with every inverse query made in the augmented interface of Definition 1 and no completion queries. It reads $p = p_{j'}(x_k)$ from the revealed tuple and computes π_k, T and its output from its transcript as $\mathcal{A}$ does. Then $\mathcal{A}^{\mathrm{aug}}$ makes fewer than $4r$ queries (at most $3r$ when $L = 1$), all inverse primitive queries. Against $\textsc{Real}$ and against $\textsc{Sim}_m$, for every m, its output has the same law as the output of $\mathcal{A}$.*

Proof. $\mathcal{A}$ completes every answer at once. Its ordinary transcript and the augmented transcript of $\mathcal{A}^{\mathrm{aug}}$ are deterministic functions of each other:

- In $\textsc{Real}$, the completion queries return the permutation values, which form the tuple.
- Against $\textsc{Sim}_m$, the simulator has assigned all $w+1$ outputs of a returned family input. So the completion queries are look-ups of the stored tuple (Definition 11). They neither read Γ nor change the state, and they cannot return $\perp$. Indeed, in a run of $\mathcal{A}$ every forward query is such a look-up, so $\textsc{Failure}$ can only be set by an inverse query, whose answer is then $\perp$ itself.

family query of $\mathcal{B}$	fork queries of $\mathcal{B}^{\mathsf{Fork}}$, in order	gives
$\pi'_b(x)$, any b	$\pi_0(x)$; then $\pi_i(\beta_i v_0)$ for $1 \le i \le w$	$v_0; v_i$
$\pi'^{-1}_0(v)$	$\pi_0^{-1}(v)$; then $\pi_i(\beta_i v)$ for $1 \le i \le w$	$x; v_i$
$\pi'^{-1}_j(v)$, $j \ge 1$	$\pi_j^{-1}(v) =: u$; $\pi_0^{-1}(\beta_j^{-1}u)$; $\pi_i(\beta_i v_0)$ for $i \notin \{0, j\}$	$v_0 = \beta_j^{-1}u; x; v_i$
construction query x	construction query x	$T^{-1}(\mathsf{Fork}(x))$

Table 6: Translation of the queries of an augmented family distinguisher $\mathcal{B}$ into fork queries, in the proof of Lemma 41. The last column lists the values that the fork answers give, in order.

- An answer that is a free input reveals only that input in both interfaces.

The quantities $\textsc{Ran}_{b'}$, π_k and T are read off the transcript: $\textsc{Ran}_{b'}$ at the time of π_k consists of the outputs of earlier completed tuples and free answers, together with t_k if $b' = b$. So the outputs agree. □

Converses of the three reductions. Each of the three reductions of Section 5.6 has a converse, which turns a distinguisher against the family into one against the fork with a comparable number of queries.

Lemma 41. *In the setting of Lemma 18, let $S = \textsc{Sim}_m$ for the family F and let S^{Fork} be its wrapped simulator. For every augmented distinguisher $\mathcal{B}$ against $(\mathrm{F}, \textsc{Sim}_m)$ with q_p primitive and q_c construction queries, there is an ordinary distinguisher $\mathcal{B}^{\mathsf{Fork}}$ against $(\mathsf{Fork}, S^{\mathsf{Fork}})$ with at most $(w+1)q_p$ primitive and q_c construction queries and the same advantage.*

Proof. The distinguisher $\mathcal{B}^{\mathsf{Fork}}$ runs $\mathcal{B}$. It answers each query of $\mathcal{B}$ by the fork queries listed in Table 6, and outputs what $\mathcal{B}$ outputs. Recall $\pi'_0 = \pi_0$ and $\pi'_j = \pi_j \circ \beta_j \circ \pi_0$, and that the tuple of x is $(x, v_0, \ldots, v_w)$. Each family primitive query costs at most $w+1$ fork queries. If some fork answer in the translation of a query is $\perp$, then $\mathcal{B}$ receives $\perp$ for that query.

In the real world, let $\pi' := \varphi(\pi)$, which is uniform. Then $\pi_0 = \pi'_0$, $\pi_i(\beta_i \pi_0(x)) = \pi'_i(x)$ and $\beta_j^{-1}\pi_j^{-1}(v) = \pi_0(\pi'^{-1}_j(v))$. So the translated answers are the tuples of π', and $T^{-1} \circ \mathsf{Fork}^{\pi} = \mathrm{F}^{\pi'}$. Thus $\mathcal{B}$ sees the real augmented game of F.

In the ideal world, S^{Fork} runs $\textsc{Sim}_m$ with the oracle $\Gamma := T^{-1} \circ \Gamma'$, which is uniform. It answers each fork query by one augmented family query, by Table 3. We claim that, while $\mathcal{B}^{\mathsf{Fork}}$ serves a query Q of $\mathcal{B}$, the family queries issued to $\textsc{Sim}_m$ are one query $\hat{Q}$ followed by look-ups, and that $\hat{Q}$ has the same effect as Q. We check this for each type of Q:

- $Q = \pi'^{-1}_j(v)$ with $j \ge 1$. The fork query $\pi_j^{-1}(v)$ issues $\hat{Q} = Q$. If it returns x, all outputs of x are assigned, with $v_0 \in \textsc{Ran}_0$. The fork queries $\pi_0^{-1}(v_0)$ and $\pi_i(\beta_i v_0)$ then issue $\pi'^{-1}_0(v_0)$, which is a look-up.
- $Q = \pi'^{-1}_0(v)$. The first fork query issues $\hat{Q} = Q$; the rest issue the look-up $\pi'^{-1}_0(v)$.
- $Q = \pi'_b(x)$. The fork query $\pi_0(x)$ issues $\hat{Q} = \pi'_0(x)$. If x is assigned, $\hat{Q}$ and Q are both look-ups. Otherwise $\textsc{Sim}_m$ answers a forward query at x in the same way on every class: it reads $\Gamma(x)$, fails if $\Gamma(x) \notin \mathcal{S}$, and otherwise samples $v_0 \leftarrow \mathrm{Unif}(\mathrm{Leg}(\Gamma(x)))$ and assigns all outputs. This holds because $n = 1$, every ι_j is the identity and there are no free inputs. The remaining fork queries issue look-ups.

Look-ups read no oracle value and change no state. Hence the state of $\textsc{Sim}_m$ and the tuples passed to $\mathcal{B}$ have the same joint law as in the augmented ideal game of F played by $\mathcal{B}$. This also holds for $\perp$: $\textsc{Failure}$ can only be set by $\hat{Q}$, which then returns $\perp$, as Q would, and is absorbing. Construction queries return $T^{-1}(\Gamma'(x)) = \Gamma(x)$. So $\mathcal{B}$ sees the ideal augmented game of F. □

Lemma 42. *In the setting of Lemma 19, let S be a simulator for C' whose $\perp$ answers are absorbing, such as* $\textsc{Sim}_m$. *For every ordinary distinguisher $\mathcal{B}'$ against (C', S) there is an ordinary distinguisher against (C, S^C) with the same advantage and the same construction queries. It answers a query of $\mathcal{B}'$ to $\pi_i'^{\pm}$ by one query to $\pi_i^{\pm}$, plus one query to $\pi_0^{\pm}$ unless it already knows the value needed.*

Proof. The new distinguisher keeps the values of π_0 it has seen, and answers the primitive queries of $\mathcal{B}'$ as follows:

- it answers $\pi_i'(x)$ by $\pi_i(X)$, where $X = \pi_0(x)$, which it queries if it does not already know it;
- it answers $\pi_i'^{-1}(v)$ by first querying $\pi_i^{-1}(v) =: X$ and then returning $\pi_0^{-1}(X)$, which it queries if it does not already know it.

Construction queries pass through, since $C = C'$ on inputs. In the real world, $(\pi_0, \pi_i' := \pi_i \circ \pi_0)$ are independent and uniform, and the answers are exactly $\pi_i'^{\pm}$.

In the ideal world, S^C answers $\pi_0^{\pm}$ from its own lazily sampled table, without querying S. It answers $\pi_i(X)$ by $\pi_i'(\pi_0^{-1}(X))$ and $\pi_i^{-1}(v)$ by $\pi_0(\pi_i'^{-1}(v))$, one query to S each. In the translation above, these are exactly the queries $\pi_i'(x)$ and $\pi_i'^{-1}(v)$ of $\mathcal{B}'$, in order. So S receives exactly the queries of $\mathcal{B}'$, and the values returned to $\mathcal{B}'$ are its answers. A $\perp$ occurs in the translation of a query exactly when S has answered $\perp$ by then, which is when $\mathcal{B}'$ receives $\perp$ in its own game. □

Lemma 43. *In the setting of Lemma 20, every distinguisher against (C, S) whose construction queries lie in $\mathcal{X}$ is also a distinguisher against $(C|_{\mathcal{X}}, S_{\mathcal{X}})$, with the same queries and the same advantage. This applies in particular to distinguishers with no construction queries.*

Proof. The real worlds coincide. In the ideal world, $S_{\mathcal{X}}$ runs S with the oracle equal to Γ' on $\mathcal{X}$ and lazily uniform elsewhere. This oracle is a uniform function on $\mathcal{D}$, and construction queries in $\mathcal{X}$ read it at points of $\mathcal{X}$. □

Proof of Corollary 6. In each case the family instance has $k = w + 1$, $n = n_b = 1$, $\mathcal{D} = G$ and one call per class. So dens $= 1$, and there are no same-class pairs, hence $|\Lambda| = 0$. Take any two distinct calls; (R1)–(R3) reduce to the displayed conditions ((R3), $30r \le N$, follows from the second).

Over $\mathbb{F}_2^d$, take a subgroup of order r, so $L = 1$. With the subgroup constants (Section D.2), Theorem 2 gives advantage at least $1 - 187N^3/r^4 - 134r/N - e^{-r/17} - N^{-2} \ge 3/4$, since $187/6.3^4 < 1/8$. The attack transfers to the two forks as follows:

- *ForkEDMD.* By Lemma 40 the attack has an augmented form with at most $3r$ queries and no construction queries. Lemma 41 turns it into a fork distinguisher with at most $3r(w+1)$ queries, and Lemma 43 carries it to the restricted domain $\{x \| 0^s\}$.
- *ForkCENC.* The family is XORP with the $w + 1$ independent permutations $\pi_i' = \pi_i \circ \pi_0$, and Lemma 42 applies to the ordinary attack $\mathcal{A}$. An iteration of $\mathcal{A}$ is one inverse query followed by at most w completion queries. It costs two fork queries for the inverse query, which also give $\pi_0(x)$, and one for each completion. So it costs at most $w + 2$ fork queries, and there are at most $3r$ iterations.

In both cases $r < 12.6\, N^{3/4}$, by the minimality of a, and $3 \cdot 12.6 < 38$.

For EDMD ($w = 1$) over a general group, use Corollary 5 with $n = n_b =$ dens $= 1$ and $|\Lambda| = 0$, and Lemmas 40 and 41. That gives fewer than $4r$ augmented queries, hence fewer than $8r$ fork queries, with $r < 12.7\, N^{3/4}$. □

Subgroup case. When the progression is a subgroup ($L = 1$), $|A_1| = |C^{\cup}| = r$ and there are at most $3r$ inverse calls. The same proof then gives smaller constants: $3r(w+1)$ queries, $134nr/N$ and $3mr|\Lambda|/N$ in Theorem 2, separating probability $1 - 67r/N$, and 38 in place of 51 in Corollary 5 (for groups of order 2^d, where r can be a power of 2 within a factor 2 of 6.3ξ).

E Pseudorandomness of rigid-free families

This appendix proves Theorem 4, the bound quoted in Section 8. Pseudorandomness is the special case of indifferentiability with no primitive queries. So Theorem 1 already gives $\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \le 10n^{3/2}(w+1)^2q^2/N^{3/2}$, plus the lower-order terms. Theorem 4 improves the dependence on q to $q/N^{3/2}$, for every rigid-free family: admissible, and with no rigid pairs (Definition 5). The proof is different from that of Theorem 1. It is a Fourier-analytic moment computation in the style of Bhaumik et al. [BGGJ26], whose treatment of Dinur's LXoP [BGGJ26, Sec. 6 and App. E] builds on Dinur's recursive moment bound [Din25] and on Eberhard's merging identity [Ebe17]. The proof has four steps:

1. Reduce the advantage to the L_2 distance of the real output density from the uniform one (Lemma 46).
2. Pull each output character back to the permutation outputs. The density becomes a product of injection coefficients, one per permutation (Lemma 45), and each input contributes one linear relation among its labels.
3. Bound the second moment of these coefficients by a merging recursion. Each merge loses a label, and the linear relation recovers it. This needs exactly that every same-class difference $\lambda_j^{-1} - \lambda_{j'}^{-1}$ is invertible, that is, that the family is rigid-free (Lemma 47).
4. Sum over levels (Theorem 4).

The new ingredients relative to [BGGJ26] are the extension to every rigid-free family and to every finite abelian group, the group-independent closed form of the fixed-support moment (Lemma 44), and explicit constants.

Model. The distinguisher is classical, possibly adaptive, and makes at most q construction queries and no primitive queries. In Real the permutations are independent and ideal. Nothing here concerns quantum queries, or keyed block ciphers in place of the ideal permutations; the latter would need a separate hybrid argument.

Repeated queries are answered consistently in both worlds and can be discarded. With free inputs the domain $\mathcal{D}$ may be smaller than q, so put $q_* := \min(q, |\mathcal{D}|)$ and assume $\mathcal{D} \neq \emptyset$. A distinguisher makes at most q_* distinct queries, and may be padded with dummy queries to exactly q_*. All bounds below are proved for q_*, and hold for q because they are increasing in q.

Rigid pairs. Suppose some same-class pair (j, j') is rigid. Then $\mathrm{F}(x)_j \neq \mathrm{F}(x)_{j'}$ always in Real; for a pair involving call 0, read this as $\mathrm{F}(x)_{j'} \neq 0$. Checking this on q_* distinct inputs distinguishes with advantage $1-(1-1/N)^{q_*} \ge q_*/2N$, for $q_* \le N$. Such families therefore admit an $\Omega(q/N)$ attack when $q \le |\mathcal{D}|$, so no bound of order $q/N^{3/2}$ can hold for them. No matching $O(q/N)$ upper bound is proved here, because the recovery argument needs the absence of rigid pairs. Distinct masks are not enough either. If some same-class difference of inverse masks is a nonzero non-automorphism, the bound can fail already at one query (Section 1). From now on the family is *rigid-free* (Definition 5).

E.1 Fixed-support moment

Lemma 44. *Let G be a finite abelian group of order N and $1 \le k \le N$.*

(a) $\Theta_k = \sum_{t=0}^{k} (-1)^{k-t}\binom{k}{t} N^t/(N)_t$. *In particular, Θ_k depends only on (N,k).*

(b) $\Theta_k = \mathbb{E}[(N-S)^k]/(N)_k$, *with* $S \sim \mathrm{Gamma}(N-k+1,1)$.

(c) If $2 \le k \le N/8$, then $\Theta_k \le 2e^{-0.122k}\big(k/(N-k)\big)^{k/2}$. In particular $\Theta_k \le 2(k/(N-k))^{k/2}$, and $\Theta_k \le (k/(N-k))^{k/2}$ for $k \ge 6$; the latter is the estimate of Dinur [Din25] used in [BGGJ26] over $\mathbb{F}_2^d$.

Proof. (a) Let $\mathcal{I} \subseteq G^k$ be the set of injective tuples, and $g := N^k (N)_k^{-1} \mathbf{1}_{\mathcal{I}}$ the density of the uniform injection relative to $\mathrm{Unif}(G^k)$. Its Fourier coefficients are $\hat{g}(\beta) = \hat{\mu}_k(\beta)$, so by Parseval Θ_k is the squared norm of the part of g that lives on characters which are nonprincipal in every coordinate. To compute this part, let E_i be the operator that averages a function over coordinate i. On characters, E_i keeps those that are principal in coordinate i and kills the others. So $\Pi_{[k]} := \prod_i (1 - E_i)$ is the orthogonal projection onto the characters that are nonprincipal in every coordinate, and $\Theta_k = \|\Pi_{[k]} g\|_2^2$.

The averages of g are again injection densities. A uniform y_i avoids the other $k-1$ entries of an injective tuple with probability $(N - k + 1)/N$, so

$$E_i g(y) = \frac{N^k}{(N)_k} \cdot \frac{N - k + 1}{N} \mathbf{1}[y_{-i} \text{ injective}] = \frac{N^{k-1}}{(N)_{k-1}} \mathbf{1}[y_{-i} \text{ injective}],$$

the injection density of the remaining coordinates. Repeating this, for $T \subseteq [k]$ the function $g_T := \left(\prod_{i \notin T} E_i\right) g$ is the injection density of y_T. Expanding the product $\prod_i (1 - E_i)$ gives the identity $\Pi_{[k]} g = \sum_{T \subseteq [k]} (-1)^{k - |T|} g_T$. Since $\Pi_{[k]}$ is an orthogonal projection, $\|\Pi_{[k]} g\|^2 = \langle \Pi_{[k]} g, g \rangle$, and since g is the density of the uniform injection, $\langle h, g \rangle = \mathbb{E}_{y \leftarrow \mathrm{Unif}(\mathcal{I})}[h(y)]$ for every real h. Hence

$$\Theta_k = \sum_{T \subseteq [k]} (-1)^{k - |T|} \mathbb{E}_{y \leftarrow \mathrm{Unif}(\mathcal{I})}[g_T(y_T)] = \sum_{t=0}^{k} (-1)^{k-t} \binom{k}{t} \frac{N^t}{(N)_t},$$

because $g_T(y_T) = N^{|T|}/(N)_{|T|}$ whenever y is injective, and there are $\binom{k}{t}$ sets T of size t.

(b) For $t \le N$, $N^t/(N)_t = N^t (N - t)!/N!$, and $(N - t)! = \int_0^\infty s^{N-t} e^{-s}\, ds$. Substituting into (a) and summing the binomial expansion under the integral,

$$\begin{aligned} \Theta_k &= \frac{1}{N!} \int_0^\infty e^{-s} \sum_{t=0}^{k} \binom{k}{t} (-1)^{k-t} N^t s^{N-t}\, ds = \frac{1}{N!} \int_0^\infty s^{N-k} (N - s)^k e^{-s}\, ds \\ &= \frac{(N-k)!}{N!} \mathbb{E}\big[(N - S)^k\big], \end{aligned}$$

since $s^{N-k} e^{-s}/(N-k)!$ is the density of $\mathrm{Gamma}(N - k + 1, 1)$.

(c) Put $m := N - k + 1$, $Y := N - S$ and $\nu := \sqrt{k/m}$. Then $\nu \le 7^{-1/2}$, because $k \le N/8$. We bound $\mathbb{E}[|Y|^k]$ through exponential moments, using that S has the moment generating function $\mathbb{E}[e^{uS}] = (1 - u)^{-m}$ for $u < 1$, in three estimates:

- For $y \ge 0$, the function $y^k e^{-\nu y}$ is maximised at $y = k/\nu$, so $|y|^k \le (k/e\nu)^k e^{\nu |y|}$. Hence $\mathbb{E}[|Y|^k] \le (k/e\nu)^k (\mathbb{E}[e^{\nu Y}] + \mathbb{E}[e^{-\nu Y}])$.
- $\mathbb{E}[e^{\nu Y}] = e^{\nu N} (1 + \nu)^{-m}$. With $N = m + k - 1$ and $\nu - \log(1 + \nu) \le \nu^2/2$,

$$\log \mathbb{E}[e^{\nu Y}] = \nu(k - 1) + m(\nu - \log(1 + \nu)) \le \nu k + \frac{m\nu^2}{2} = \nu k + \frac{k}{2} \le 0.878\, k.$$

- $\mathbb{E}[e^{-\nu Y}] = e^{-\nu N} (1 - \nu)^{-m}$. With $-\nu - \log(1 - \nu) = \sum_{i \ge 2} \nu^i / i \le \nu^2/2(1 - \nu)$,

$$\log \mathbb{E}[e^{-\nu Y}] = -\nu(k - 1) - m(\nu + \log(1 - \nu)) \le \frac{m\nu^2}{2(1-\nu)} = \frac{k}{2(1-\nu)} \le 0.804\, k.$$

Both exponential moments are at most $e^{0.878k}$, so

$$\mathbb{E}[|Y|^k] \le 2\left(\frac{k}{e\nu}\right)^k e^{0.878k} = 2\left(\frac{k}{\nu}\right)^k e^{-0.122k} = 2(km)^{k/2} e^{-0.122k}.$$

Since $\Theta_k \geq 0$, (b) gives $\Theta_k \leq \mathbb{E}[|Y|^k]/(N)_k$. With $(N)_k \geq m^k$ and $m > N - k$, this is at most $2e^{-0.122k}(k/m)^{k/2} \leq 2e^{-0.122k}(k/(N-k))^{k/2}$. The last two statements follow from $e^{-0.122k} \leq 1$, and from $2e^{-0.122k} \leq 1$ for $k \geq 6$. □

For orientation only (this is not used later), let Z be a standard normal variable. As $N \to \infty$ with k fixed, $(N - S)/\sqrt{N}$ converges in law to $-Z$. So (b) gives $\Theta_k \sim \mathbb{E}[Z^k]/N^{k/2} = (k-1)!!/N^{k/2}$ for fixed even k. Since $(k-1)!! \approx \sqrt{2}\,(k/e)^{k/2}$, the ratio of Θ_k to Dinur's bound $(k/(N-k))^{k/2}$ tends, for large even k, to about $\sqrt{2}\,e^{-k/2}$. Equivalently, the bound overestimates Θ_k by a factor of about $e^{k/2}/\sqrt{2}$.

E.2 Pullback and reduction to moments

For an output character $\alpha = (\alpha_1, \ldots, \alpha_w) \in \widehat{G}^w$ put

$$T(\alpha) := (\beta_0, \ldots, \beta_w), \qquad C_j := (\lambda_j^*)^{-1},$$
$$\beta_j := \lambda_j^* \alpha_j \quad (j \geq 1), \qquad \beta_0 := -\lambda_0^*(\alpha_1 + \cdots + \alpha_w).$$

Then $\mathrm{e}(\alpha, \mathrm{F}(x)) = \prod_j \mathrm{e}(\beta_j, p_j(x))$, T is injective with $T(0) = 0$, and every $T(\alpha)$ satisfies the *block relation*

$$\sum_{j=0}^{w} C_j \beta_j = 0. \tag{20}$$

For a same-class pair, $C_i - C_j = (\lambda_i^{-1} - \lambda_j^{-1})^* = \mu_{ij}^*$, which is an automorphism of $\widehat{G}$ iff the pair is non-rigid. A nonzero α has a pullback with at least two nonzero labels, since a single nonzero label would contradict (20). For LXoP over $\mathbb{F}_2^d$ this is [BGGJ26, eqs. (24)–(27)], up to relabelling.

For t distinct inputs $x_1, \ldots, x_t$ and $\alpha = (\alpha^{(1)}, \ldots, \alpha^{(t)}) \in (\widehat{G}^w)^t$, let

$$\hat{\xi}_t(\alpha) := \mathbb{E}_{\textsc{Real}}\Big[\overline{\textstyle\prod_i \mathrm{e}(\alpha^{(i)}, \mathrm{F}(x_i))}\Big], \qquad \Omega_t := \sum_{\alpha:\, \alpha^{(i)} \neq 0 \ \forall i} |\hat{\xi}_t(\alpha)|^2.$$

So Ω_t is the contribution of the level-t coefficients on one fixed support. For a class b, write $\beta|_b \in \widehat{G}^{n_b t}$ for the labels of $T(\alpha) = (T(\alpha^{(1)}), \ldots, T(\alpha^{(t)}))$ at the calls of class b, ordered by input.

Lemma 45 (cf. [BGGJ26, eq. (25)])**.** *$\hat{\xi}_t(\alpha) = \prod_b \hat{\mu}_{n_b t}(\beta|_b)$. In particular, $\hat{\xi}_t$ and Ω_t do not depend on the choice of the t distinct inputs.*

Proof. The $n_b t$ inputs $\iota_j(x_i)$ with $c(j) = b$ are distinct, by disjointness of the images of the ι_j. So the raw outputs of π_b at these inputs form a uniform injection, independently across classes. Expand $\prod_i \mathrm{e}(\alpha^{(i)}, \mathrm{F}(x_i)) = \prod_i \prod_j \mathrm{e}(\beta_j^{(i)}, p_j(x_i))$. □

Lemma 46 (cf. [BGGJ26, Lemma 1 and Remark 1])**.** *For every classical distinguisher making at most q construction queries,*

$$\mathbf{Adv}_{\mathrm{F}}^{\mathrm{prf}} \leq \frac{1}{2}\Big(\sum_{t=1}^{q_*} \binom{q_*}{t} \Omega_t\Big)^{1/2}.$$

Proof. The distinguisher makes exactly q_* distinct queries. In both worlds, the law of the answers at distinct inputs depends only on their number; in $\textsc{Real}$ this is Lemma 45. So answering the i-th fresh query at the

i-th point of a fixed q_*-set $Q \subseteq \mathcal{D}$ does not change the transcript law. For a deterministic distinguisher, condition on each query path; randomised distinguishers average over their coins.

The acceptance probability is then a function $a : (G^w)^Q \to [0,1]$, and

$$\mathbf{Adv} = |\mathbb{E}_{\mathrm{Unif}}[a(\xi - 1)]| = |\mathbb{E}_{\mathrm{Unif}}[(a - \tfrac{1}{2})(\xi - 1)]| \le \tfrac{1}{2}\|\xi - 1\|_2,$$

where ξ is the density of the real answers relative to $\mathrm{Unif}((G^w)^Q)$. By Parseval, $\|\xi - 1\|_2^2 = \sum_{\alpha \neq 0} |\hat{\xi}(\alpha)|^2$. Group α by the set $S \subseteq Q$ of inputs with $\alpha^{(x)} \neq 0$. Each group contributes $\Omega_{|S|}$, by Lemma 45 applied to the inputs of S. □

E.3 Merging recursion

Fix t, and a set $\mathcal{S}$ of masks α with all $\alpha^{(i)} \neq 0$ whose pullbacks have a common support J_0, $|J_0| = k_0$. Then $2t \le k_0 \le (w+1)t$. Write $C(\beta) := \prod_b \hat{\mu}(\beta|_b)$, so that $\hat{\xi}_t(\alpha) = C(T(\alpha))$.

Nodes. A *node* is a set $\mathcal{S}_v \subseteq \mathcal{S}$ together with an injective map T_v on $\mathcal{S}_v$, starting from $T_\varnothing = T$, whose images have a common support J_v. A block, meaning the labels of one input, is *untouched* at v if none of its positions has been chosen as primary or secondary on the path to v. Its labels then equal those of $T(\alpha)$, so they satisfy (20).

Merge step. At a node of depth less than d that has an untouched block, choose the *primary* j: the first position of J_v in the first untouched block. This choice depends only on J_0 and the history, not on α. Let $b := c(j)$, and let $K_v \subseteq J_v$ be the set of class-b positions in J_v. There are two cases:

- If $|K_v| = 1$, the class-b factor of every coefficient at v vanishes, by Lemma 3(a),(b), and the node contributes 0.
- Otherwise, apply Lemma 3(d) inside the class-b injection. It writes the class-b factor of $C(T_v\alpha)$ as $-(N-|K_v|+1)^{-1}$ times a sum of $|K_v|-1$ terms, one for each $i \in K_v \setminus \{j\}$; the term for i is the class-b factor of $C(T_{v,i}\alpha)$. The other class factors are unchanged, so $C(T_v\alpha) = -(N-|K_v|+1)^{-1} \sum_i C(T_{v,i}\alpha)$. By Cauchy–Schwarz, $|\sum_i u_i|^2 \le (|K_v|-1)\sum_i |u_i|^2$. Summing over α:

$$\sum_{\alpha \in \mathcal{S}_v} |C(T_v\alpha)|^2 \le \frac{|K_v|-1}{(N-|K_v|+1)^2} \sum_{i \in K_v \setminus \{j\}} \sum_{e \in \{0,1\}} \sum_{\alpha \in \mathcal{S}_{v,i,e}} |C(T_{v,i}\alpha)|^2. \tag{21}$$

  Here $T_{v,i}(\alpha) := T_v(\alpha)^{(j,i)}$, and $\mathcal{S}_{v,i,e}$ splits $\mathcal{S}_v$ according to whether the merged label at i is zero ($e = 0$) or not, so each child again has a common support. The other class factors are unchanged.

Merges stay inside one permutation. This is why the extension to several permutations needs no condition on the masks of *different* classes.

Lemma 47. *If the family is rigid-free, every child map $T_{v,i}$ is injective on $\mathcal{S}_{v,i,e}$.*

Proof. Let $\beta = T_v(\alpha)$ and $\beta' = T_{v,i}(\alpha)$. All positions other than j and i agree. In the primary's block, which is untouched, (20) holds for β, and we distinguish two cases:

- If i lies in another block, then $\beta_j = -C_j^{-1} \sum_{\ell \neq j} C_\ell \beta'_\ell$, the sum running over the primary's block.
- If i lies in the same block, it is a same-class position. Then

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

  and $C_j - C_i = \mu^*_{ji}$ is invertible because the family is rigid-free.

In both cases $\beta_i = \beta'_i - \beta_j$. So β' and the known (j, i, e) determine β, and β determines α by induction. □

Over $\mathbb{F}_2^d$ with $\lambda_j = (-\sigma)^j$, the same-block condition is invertibility of $I + \sigma^{|i-j|}$, which is w-admissibility [BGGJ26, App. E.2].

Proposition 9 (cf. [BGGJ26, Lemma 5 and Proposition 7])**.** *Let the family be rigid-free with $N \geq 8(w+1)$, and put $c_* := 8(w+1)/7$.*

(a) $\Omega_1 \leq A_w/N^3$ with $A_w := (8/7)^3\big[4\binom{w+1}{3} + 2^{2w+3}(w+1)^3\big]$. If no class has three or more calls, the support-3 term is absent. If no class has two or more calls ($n = 1$), then $\Omega_1 = 0$.

(b) For $t \geq 2$ with $k_{\max} := (w+1)t \leq N/8$ and $d := \lceil t/2 \rceil$,

$$\Omega_t \leq 2^{(w+1)t+d+w+1}\Big(\frac{k_{\max}}{N - k_{\max}}\Big)^{t+d}.$$

Proof. We begin with the tree bound. Fix one support pattern J_0 with $|J_0| = k_0$, and run the recursion to depth d. Starting from the root $\mathcal{S}$, each step replaces the sum of $|C|^2$ over a node by a weighted sum over its children, by (21). After d steps we obtain a bound by a weighted sum of leaf sums. We check that the recursion can be run, and bound the weights, the leaves and the exponent:

- *Running the recursion.* A step needs an untouched block, and each step touches at most two blocks (the primary's and the one containing i). Before depth d at most $2(d-1) \leq t-1$ blocks are touched, so at least one of the t blocks is untouched. Its labels are those of $T(\alpha)$, and at least one of them is nonzero, because α is nonzero on every block; so the primary exists.
- *Weights.* A node has at most $2(|K_v| - 1) \leq 2k_0$ children (a choice of i and of e), and the factor in (21) is at most $k_0/(N-k_0)^2$, since $|K_v| \leq k_0$. So each level multiplies the bound by at most $2k_0 \cdot k_0/(N-k_0)^2$.
- *Leaves.* At a leaf, the maps T_v are injective (Lemma 47), so the masks in the leaf give distinct label vectors with a common support. Let k_b be the size of its class-b part. By Lemma 45, the leaf sum is at most $\prod_b \Theta_{k_b}$, the sum of $|C|^2$ over *all* label vectors with this support. This is 0 if some $k_b = 1$, by Lemma 3(b). Otherwise Lemma 44(c) gives at most $2^{w+1}(k_0/(N-k_0))^{\sum_b k_b/2}$: empty classes contribute $\Theta_0 = 1$, at most $w+1$ classes are nonempty, and each nonempty one contributes a factor at most $2(k_0/(N-k_0))^{k_b/2}$.
- *Exponent.* A merge sets the primary label to 0, and in the child with $e = 0$ it also makes the merged label 0. So each merge removes one or two positions from the support, and after d merges $\sum_b k_b \geq k_0 - 2d$. This is positive: for $t \geq 2$ because $k_0 \geq 2t$ and $2d \leq t+1$, and in the case $t = 1$, $d = 1$ used in (a) below because $k_0 \geq 4$ there.

Multiplying the d levels and the leaf bound, and using $k_0/(N-k_0) < 1$ to lower the exponent,

$$\sum_{\alpha \in \mathcal{S}} |\hat{\xi}_t(\alpha)|^2 \leq (2k_0)^d \Big(\frac{k_0}{(N-k_0)^2}\Big)^d 2^{w+1}\Big(\frac{k_0}{N-k_0}\Big)^{k_0/2-d} = 2^{d+w+1}\Big(\frac{k_0}{N-k_0}\Big)^{k_0/2+d}.$$

(b) Use $k_0 \geq 2t$ and $k_0 \leq k_{\max}$, increasing the base first and then decreasing the exponent, which is allowed because $k_{\max}/(N-k_{\max}) < 1$. Sum over the at most $2^{(w+1)t}$ support patterns.

(a) Take $t = 1$. By Lemma 3(b), every class meeting the support carries at least two nonzero labels. We split by the support size:

- *Support 2.* Both labels lie in one class, with $\beta_j = -\beta_i$. Then (20) gives $\mu^*_{ij}\beta_i = 0$, so $\beta_i = 0$: there are no such masks.

- *Support 3.* All three labels lie in one class $\{i, j, \ell\}$, with $\beta_\ell = -\beta_i - \beta_j$. Then (20) gives $\mu^*_{i\ell}\beta_i + \mu^*_{j\ell}\beta_j = 0$, so β_i determines the mask. By Lemma 3(c), this gives at most $N - 1$ masks of coefficient $2/(N-1)(N-2)$. That is at most $4/(N-1)(N-2)^2 \le 4(8/7)^3/N^3$ per 3-set, and there are at most $\binom{w+1}{3}$ such sets.
- *Support $k_0 \ge 4$.* The tree bound with $d = 1$ gives

$$2^{w+2}\Big(\frac{k_0}{N-k_0}\Big)^{k_0/2+1} \le 2^{w+2}\Big(\frac{8}{7}\Big)^3\frac{(w+1)^3}{N^3},$$

  using $k_0 \le w + 1 \le N/8$. There are at most 2^{w+1} patterns.

Summing the three cases gives the bound on Ω_1. □

E.4 Summing over levels

Theorem 4. *Let the family be rigid-free, over a finite abelian group of order $N \ge 8(w+1)$, built from independent ideal permutations, with $\mathcal{D} \ne \varnothing$. Class sizes may be unequal and free inputs are allowed. Put*

$$\begin{aligned} K_w &:= e\sqrt{2}\,c_*^{3/2}2^{w+1},\\ C_w &:= \tfrac{1}{2}\big(\sqrt{A_w} + 2^{(w+3)/2}K_w\big),\\ c_w &:= \min\Big\{(2K_w\sqrt{2e})^{-2/3},\ \frac{1}{8(w+1)}\Big\}. \end{aligned}$$

Then for $1 \le q \le c_w N$ and every classical, possibly adaptive, distinguisher making at most q construction queries,

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{F}}(q) \le C_w\frac{q_*}{N^{3/2}} \le C_w\frac{q}{N^{3/2}}, \qquad q_* = \min(q, |\mathcal{D}|).$$

The constants depend on w only, and not on the class structure, the masks or the group. For example, $C_1 \le 117$ and $c_1 \ge 1/40$; $C_2 \le 589$ and $c_2 \ge 1/94$; $C_3 \le 2516$ and $c_3 \ge 1/199$ (Appendix B).

Proof. By Lemma 46 it suffices to bound $\sum_{t\le q_*}\binom{q_*}{t}\Omega_t$. Since $1 \le q_* \le q \le c_w N$, we may write q for q_*. Let $u := K_w q/N^{3/2}$. For $2 \le t \le q$ we have $k_{\max} = (w+1)t \le N/8$, so $k_{\max}/(N-k_{\max}) \le c_* t/N$. With $\binom{q}{t} \le (eq/t)^t$, Proposition 9(b) gives

$$\binom{q}{t}\Omega_t \le 2^{w+1}\Big(\frac{ec_*2^{w+1}q}{N}\Big)^t\Big(\frac{2c_*t}{N}\Big)^d \le 2^{w+1}\big(u\sqrt{t}\big)^t,$$

using $2c_*t/N \le 2/7 < 1$ and $d \ge t/2$.

Consecutive terms of $(u\sqrt{t})^t$ have ratio at most $u\sqrt{e(t+1)} \le u\sqrt{2eq}$, which is at most $1/2$ for $q \le (2K_w\sqrt{2e})^{-2/3}N$. Only the finite range $2 \le t \le q$ is used; the infinite series would diverge. So

$$\sum_{t\ge 2}\binom{q}{t}\Omega_t \le 2^{w+1}\cdot 2(u\sqrt{2})^2 = 2^{w+3}u^2.$$

With Proposition 9(a) and Lemma 46, for $q \ge 1$,

$$\mathbf{Adv} \le \tfrac{1}{2}\big(A_w q/N^3 + 2^{w+3}K_w^2q^2/N^3\big)^{1/2} \le \tfrac{1}{2}\big(\sqrt{A_w} + 2^{(w+3)/2}K_w\big)\frac{q}{N^{3/2}}.$$ □

The constants grow like about $2^{3w/2}w^{3/2}$. The growth comes from two crude counts: all $2^{(w+1)t}$ support patterns, and the factor 2 of Lemma 44(c) at small per-class supports. No attempt has been made to optimise them.

Corollary 8. *Theorem 4 applies to the following families:*

(i) *LXoP$[\sigma, w]$ over $\mathbb{F}_2^d$, for every w-admissible σ and every width, including unbalanced ones. This is an explicit form of [BGGJ26, Theorem 7], whose method is followed here.*

(ii) *$\pi_1 + \pi_2$ over any finite abelian group: $\mathbf{Adv} \le 117\, q/N^{3/2}$. When $n = 1$, masks are absorbed into the permutations (Lemma 6), so masked variants are equivalent. [BGGJ26, Theorem 2] gives the same order with an unspecified constant, together with the sharper $O(q^2/N^2)$ below the birthday bound.*

(iii) *XORP with $w + 1$ independent permutations, $x \mapsto (\pi_j(x) - \pi_0(x))_{1\le j\le w}$, over any finite abelian group: $\mathbf{Adv} \le C_w q/N^{3/2}$, with $\Omega_1 = 0$. As far as we can see, this multi-block instance is not among the statements of [BGGJ26].*

(iv) *Mixed families, with several permutations, each called several times per input, where every same-class difference $\lambda_j^{-1} - \lambda_{j'}^{-1}$ is invertible. Examples are pairwise distinct nonzero field scalars within each class. Masks of different classes are unrestricted.*

Families with a rigid pair, such as XORP$[w]$ *on one permutation and the domain-separated sum, are not covered. They admit an $\Omega(q/N)$ attack, and no matching upper bound is claimed here.*