

Indifferentiability of the Sum of Two Permutations: Tight $3n/4$ Security for the Uniform Simulator and a New Simulator for $4n/5$ Security

Sunyeop Kim^{1,2}

¹ Korea University

`sunyeopkim@korea.ac.kr`

² Nanyang Technological University

Abstract. We study the regular indifferentiability of the sum of two n -bit permutations. Previous work proved $2n/3$ -bit security and gave a $5n/6$ -bit attack against the uniform simulator. We prove $3n/4$ -bit security for the simulator and give a matching attack. Our proof bounds the KL divergence between response distributions by tracking the bias in unrevealed construction values through their conditional distribution, which was not analyzed in detail in previous work. To exceed this threshold, we introduce the Gyroscope simulator, which adjusts inverse-query acceptance probabilities to compensate for the bias left by earlier responses. The Gyroscope simulator achieves $4n/5$ -bit security.

Keywords: Indifferentiability · sum of permutations · simulators · distinguishing attacks

1 Introduction

Constructing pseudorandom functions from block ciphers is a basic problem in symmetric cryptography. The sum of permutations (SoP), proposed by Bellare, Krovetz, and Rogaway in 1998 [1], is one of the most widely studied constructions. For two independent random permutations Π_0, Π_1 on $\mathbb{F}_2^n$, it computes

$$\text{SoP}(x) = \Pi_0(x) \oplus \Pi_1(x).$$

When the permutations are secret, the construction achieves n -bit security [15, 6, 8]. An optimal bound of $O(q/2^{3n/2})$ on the distinguishing advantage is also known [9, 7].

When the permutations are public, however, the exact security of SoP under regular indifferentiability [13, 5] remains unresolved.

Previous results. Mandal et al. [12] claimed $2n/3$ -bit indifferentiability of SoP from a random function. Mennink and Preneel [14] identified a flaw in their proof and gave an alternative proof of the same bound. Bhattacharya and Nandi [2]

later claimed n -bit indistinguishability. Gunesing et al. [10] clarified that the proof of Mennink and Preneel establishes only sequential indistinguishability and identified further flaws in the proof of Bhattacharya and Nandi [2]. They proved $2n/3$ -bit regular indistinguishability of SoP and gave a $2^{5n/6}$ -query attack against the uniform simulator. Bhaumik et al. [3] subsequently gave a simpler proof of the $2n/3$ -bit regular indistinguishability bound.

This leaves two questions: what is the exact security level of the uniform simulator, and can a different simulator achieve higher security? We resolve both questions.

1.1 Our results

We prove tight $3n/4$ -bit security for the uniform simulator (USim) and design the Gyroscope simulator (GSim), which achieves $4n/5$ -bit security. With $N = 2^n$, our results are summarized as follows:

Simulator	Result	Reference
USim	$q = o(N^{3/4}) : \text{Adv} = o(1)$	Theorem 1
USim	$\exists D : q = O(N^{3/4}), \text{Adv}(D) > 1/2$	Theorem 2
GSim	$q = o(N^{4/5}/n^{1/5}) : \text{Adv} = o(1)$	Theorem 3.

The Gyroscope simulator makes a constant number of construction queries per primitive query and admits an efficient implementation (Proposition 3).

1.2 Proof strategy and simulator design

The security and limit of uniform simulation. To answer an inverse query, the uniform simulator tries a random input. The construction oracle fixes what the other permutation must return at that input. If this would reuse an occupied output, the simulator rejects the candidate and tries again. This keeps each answer consistent with the permutation tables and the construction oracle. Even though candidate inputs are sampled uniformly, earlier inverse responses bias the construction values that remain unrevealed [10].

Earlier proofs lose precision in two places. Their lower bounds on transcript probability ratios charge the number of coincidences between construction values and XORs of permutation outputs. A sum-capture bound controls this number, but the calculation loses the cancellation with its average that is present in the exact response probabilities. The resulting loss limits these bounds to $2n/3$ -bit security [14,3]. The analyses of inverse queries in [14,3] introduce another loss. To lower-bound transcript probabilities, they retain only executions in which rejected candidates satisfy additional freshness conditions. Even with more trials, the lower bound on transcript probability does not account for these excluded executions if the same freshness restrictions are retained.

We retain this cancellation by bounding the divergence between response distributions. To include the dependencies caused by rejected candidates, we compute the conditional distribution of unrevealed construction values, summing over all rejected trials. Bounding the average squared bias gives a smaller

divergence bound than bounding the maximum bias alone. Section 3.5 illustrates both improvements for a fresh construction query.

We compare the real and ideal worlds through a hybrid, classifying queries according to their type and the values already revealed. In both comparisons, the conditional KL divergence between response distributions is $O(q^3/N^3)$ for every query type, including the contribution from simulator failure. Summing over q queries gives an $O(q^4/N^3)$ bound, yielding $3n/4$ -bit security. We construct a matching attack by using inverse queries to bias unrevealed construction values.

Design rationale for the Gyroscope simulator. We quantify the bias that the uniform simulator introduces in unrevealed construction values when answering inverse queries. We then design the Gyroscope simulator to counteract this bias through an additional random acceptance test. The test uses each earlier inverse query and the simulator’s state immediately before it to reduce the acceptance probability of candidates whose construction values earlier responses made more likely.

For $3n/4$ -bit security, it suffices to bound the KL divergence between response distributions uniformly over all possible sequences of preceding queries and responses. Our proof of $4n/5$ -bit security requires a sharper bound that depends on those queries and responses. For the comparison of the real and hybrid worlds, we bound the second moment of the discrepancy between their response distributions over interactions in the real world.

Section 2 defines the security model and introduces the probabilistic tools. Section 3 proves security for the uniform simulator, and Section 4 presents a matching attack. Section 5 presents the Gyroscope simulator, and Section 6 explains its security proof. The appendices give the detailed security proofs for the uniform and Gyroscope simulators and analyze the implementation of the Gyroscope simulator.

2 Preliminaries

2.1 Notation and the SoP construction

Let $\mathcal{G} = \mathbb{F}_2^n$ and $N = 2^n$. For $T \subseteq \mathcal{G}$ and $z \in \mathcal{G}$, we write $T \oplus z = \{t \oplus z : t \in T\}$. For permutations Π_0, Π_1 of $\mathcal{G}$, the sum-of-permutations construction is defined by $\text{SoP}^{\Pi_0, \Pi_1}(x) := \Pi_0(x) \oplus \Pi_1(x)$.

Let $\text{Unif}(E)$ denote the uniform distribution on a nonempty finite set E . Similarly, $\text{Perm}(\mathcal{G})$ and $\text{Func}(\mathcal{G}, \mathcal{G})$ denote the uniform distributions over permutations of $\mathcal{G}$ and functions from $\mathcal{G}$ to $\mathcal{G}$, respectively. For a random variable X , $\mathcal{L}(X)$ denotes its distribution, and $\mathcal{L}(X | C)$ denotes its conditional distribution given information C of positive probability. We write $\mathbf{1}[A]$ for the indicator of a condition A and $\mathbf{1}_C(x) = \mathbf{1}[x \in C]$ for a set C .

For a real function f on a nonempty finite set E , we define

$$\|f\|_{1,E} = \sum_{x \in E} |f(x)|, \quad \|f\|_{2,E} = \left(\sum_{x \in E} |f(x)|^2 \right)^{1/2}, \quad \|f\|_{\infty,E} = \max_{x \in E} |f(x)|.$$

For $E = \emptyset$, we define all three norms to be zero. We omit E when the domain is clear and use the same notation for vectors indexed by E . For $f : \mathcal{G} \rightarrow \mathbb{R}$, we write $f^\circ := f - N^{-1} \sum_{z \in \mathcal{G}} f(z)$. Unless stated otherwise, we assume $N \geq 64$ and $q \leq N/64$ in the security proof.

2.2 The (regular) indistinguishability game

Game. Fix a simulator $S = (S_0, S_1)$ with access to a uniformly random construction oracle. The two worlds are

$$\begin{aligned} \text{Real} : \Pi_0, \Pi_1 &\leftarrow \text{Perm}(\mathcal{G}) \text{ independently, } F = \text{SoP}^{\Pi_0, \Pi_1}, \\ \text{Ideal} : \Gamma &\leftarrow \text{Func}(\mathcal{G}, \mathcal{G}), \quad (S_0, S_1) = S^\Gamma. \end{aligned}$$

In **Real**, the distinguisher has access to $(F, \Pi_0^\pm, \Pi_1^\pm)$. In **Ideal**, it has access to $(\Gamma, S_0^\pm, S_1^\pm)$. For $b \in \{0, 1\}$ and $O \in \{\Pi_b, S_b\}$, $O^\pm$ denotes access to the forward and inverse interfaces. The simulator may query Γ , but it is not notified of construction queries that the distinguisher makes directly to Γ . The *construction value* at x is $F(x)$ in **Real** or $\Gamma(x)$ in **Ideal**. This value need not be known to the distinguisher. A *construction response* is the value returned by a construction query. The simulators may fail and return $\perp$ in **Ideal**; this never occurs in **Real**.

Distinguisher. A distinguisher D makes adaptive construction and primitive queries, in any order, and outputs one bit. We assume that D makes no query whose response is already determined by its previous queries and responses. We work in the information-theoretic setting and assume that D is deterministic. Let q_c and q_p denote the numbers of construction and primitive queries made by D , respectively, with $q_c + q_p \leq q$. The distinguishing advantage is defined as

$$\begin{aligned} \text{Adv}_{F,S}^{\text{indiff}}(D) &:= \left| \Pr[D^{F, \Pi_0^\pm, \Pi_1^\pm} = 1] - \Pr[D^{\Gamma, S_0^\pm, S_1^\pm} = 1] \right|, \\ \text{Adv}_{F,S}^{\text{indiff}}(q) &:= \sup_{D \text{ makes at most } q \text{ queries}} \text{Adv}_{F,S}^{\text{indiff}}(D). \end{aligned}$$

2.3 Transcripts and query notation

In a world W , the *transcript* is the sequence of all query-response pairs (Q_i, v_i) , where Q_i is the i th query and v_i its response. We pad shorter transcripts to length q with a fixed dummy pair. Its first k pairs form the *transcript prefix of length k* , denoted by

$$\text{Tr}_W^k := ((Q_1, v_1), \dots, (Q_k, v_k)).$$

A transcript prefix τ of length k is *W-attainable at step $k+1$* if $\Pr_W[\text{Tr}_W^k = \tau] > 0$. For a W -attainable prefix and a possible response v to query $k+1$, put

$$p_{W,k+1}^\tau(v) := \Pr_W[v_{k+1} = v \mid \text{Tr}_W^k = \tau].$$

We introduce notation for queries to the SoP construction and the permutation interfaces. For $b \in \{0, 1\}$, let O_b denote the permutation interface on side b in each world. Thus $O_b = \Pi_b$ in *Real* and $O_b = S_b$ in *Ideal*. We write

$$\text{Con}(x), \quad \text{Fwd}_b(x), \quad \text{Inv}_b(y)$$

for construction queries, forward primitive queries, and inverse primitive queries. In the case of a forward primitive query $\text{Fwd}_b(x)$ or a successful inverse primitive query $\text{Inv}_b(y)$ with *inverse output* $x = O_b^{-1}(y)$, we call $O_{1-b}(x)$ the *companion output*. It is the output at the same domain point of the other permutation interface, namely the one not addressed by that primitive query. In the analysis, the transcript includes

$$(x, O_0(x), O_1(x))$$

for every successful primitive query at domain point x . Construction responses and $\perp$ responses are unchanged. Projecting this triple to the requested output recovers the ordinary interface, so this augmentation cannot decrease the distinguisher's advantage. Write TV for total variation distance and $\text{Tr}_{W,\text{aug}}$ for the complete augmented transcript in world W . Then

$$\text{Adv}_{F,S}^{\text{indiff}}(D) \leq \text{TV}(\mathcal{L}(\text{Tr}_{\text{Real},\text{aug}}), \mathcal{L}(\text{Tr}_{\text{Ideal},\text{aug}})).$$

Henceforth, we assume that transcripts are augmented, except in the attack analysis of Section 4. We call the construction value at x *revealed* if the transcript contains that value or both permutation outputs at x , and *unrevealed* otherwise. An unrevealed value may already have been read privately by the simulator.

In the response analysis, we omit queries whose responses are already determined by the transcript.

A construction query is *fresh* if no earlier query involved its input. A forward primitive query is *fresh* if the transcript contains no permutation outputs at its input. An inverse primitive query $\text{Inv}_b(y)$ is *fresh* if y does not appear among the outputs of O_b in the transcript. Freshness is determined after augmenting the transcript.

2.4 Statistical tools for the proof

We introduce the statistical notation and tools used in the security proof. All distributions here are on finite sets and are identified with their probability mass functions. For a real function g on a nonempty finite set E , write $\text{osc}_E(g) := \max_{x \in E} g(x) - \min_{x \in E} g(x)$. We omit E when the domain is clear. A variable subscript, as in $\text{osc}_x g(x)$, means that both the maximum and minimum are taken over x .

Let P and Q be probability distributions on the same finite set. Their total variation distance is $\text{TV}(P, Q) = \frac{1}{2} \|P - Q\|_1$. In the divergences below, we call Q the *reference distribution*. When $\text{supp}(P) \subseteq \text{supp}(Q)$, their chi-squared and Kullback–Leibler (KL) divergences are

$$\chi^2(P \| Q) := \sum_{\omega: Q(\omega) > 0} Q(\omega) \left(\frac{P(\omega)}{Q(\omega)} - 1 \right)^2,$$

$$D_{\text{KL}}(P\|Q) := \sum_{\omega: P(\omega) > 0} P(\omega) \log \frac{P(\omega)}{Q(\omega)}.$$

For a fixed construction value z , either permutation output determines the other. Likewise, for a fixed inverse query $\text{Inv}_b(y)$, the construction value z and the companion output $y \oplus z$ determine each other. Thus, including the determined output leaves both divergences unchanged.

We use the bound $D_{\text{KL}}(P\|Q) \leq \chi^2(P\|Q)$ together with Pinsker's inequality and the triangle inequality:

$$\text{TV}(P, Q) \leq \sqrt{D_{\text{KL}}(P\|Q)/2}, \quad \text{TV}(P, R) \leq \text{TV}(P, Q) + \text{TV}(Q, R).$$

Both chi-squared divergence and KL divergence are convex in each argument when the other is fixed.

Write $\text{Var}_Q(Z) = \mathbb{E}_Q[(Z - \mathbb{E}_Q[Z])^2]$. For a distribution P on a nonempty set E , the chi-squared divergence from $U_E = \text{Unif}(E)$ is

$$\chi^2(P\|U_E) = \text{Var}_{X \leftarrow U_E}[|E|P(X)] = |E| \sum_{x \in E} (P(x) - 1/|E|)^2. \quad (1)$$

The KL chain rule. Let W_0, W_1 be two worlds with the same deterministic distinguisher. For $W \in \{W_0, W_1\}$, let $T_W := \mathcal{L}(\text{Tr}_W^q)$ denote the complete transcript distribution. If every transcript possible in W_0 is also possible in W_1 , the KL chain rule gives

$$D_{\text{KL}}(T_{W_0}\|T_{W_1}) = \sum_{i=1}^q \mathbb{E}_{\tau \leftarrow \text{Tr}_{W_0}^{i-1}}[D_{\text{KL}}(p_{W_0, i}^\tau \| p_{W_1, i}^\tau)]. \quad (2)$$

Both response distributions are conditioned on the same prefix, sampled from W_0 . This identity, $D_{\text{KL}} \leq \chi^2$, and Pinsker's inequality give the chi-squared method of Dai, Hoang, and Tessaro [6, Lemma 3].

Reweighting. The simulator's sampling and acceptance rules can bias the conditional response distribution. In our proofs, we express response distributions as weighted versions of a reference distribution, which need not be uniform. The following lemma bounds their divergence and total variation distance from that reference using the variance or range of the logarithms of the weights.

Lemma 1 (Reweight a finite distribution). *Let Q be a distribution on a finite set Ω and let $w : \Omega \rightarrow (0, \infty)$ assign a positive weight to each outcome. Multiplying each probability $Q(x)$ by its weight $w(x)$ and normalizing gives the distribution*

$$P(x) = \frac{Q(x)w(x)}{\sum_{y \in \Omega} Q(y)w(y)}.$$

Set $g(x) = \log w(x)$ and

$$d = \text{osc}_{\text{supp}(Q)}(g).$$

Then

$$\chi^2(P\|Q) \leq e^{2d} \text{Var}_Q(g), \quad (3)$$

$$\begin{aligned} D_{\text{KL}}(P\|Q) &\leq \min\{\tfrac{1}{2}e^d \text{Var}_Q(g), d^2/8\}, \\ D_{\text{KL}}(Q\|P) &\leq d^2/8, \quad \text{TV}(P, Q) \leq d/4. \end{aligned} \quad (4)$$

The proof is given in Appendix A.

2.5 Handling simulator failure

For each query, we condition the response distribution on success and account for failure separately. With the same assumptions on W_0, W_1 as in Section 2.4, fix a W_0 -attainable transcript prefix τ of length k and let $p_{\text{fail},\tau} := p_{W_1,k+1}^\tau(\perp) < 1$. The distribution of successful responses in W_1 is

$$p_{W_1,k+1}^{\tau,+}(v) := \frac{p_{W_1,k+1}^\tau(v)}{1 - p_{\text{fail},\tau}}, \quad v \neq \perp.$$

If $p_{W_0,k+1}^\tau(\perp) = 0$, then

$$\begin{aligned} D_{\text{KL}}(p_{W_0,k+1}^\tau \| p_{W_1,k+1}^\tau) &= D_{\text{KL}}(p_{W_0,k+1}^\tau \| p_{W_1,k+1}^{\tau,+}) \\ &\quad + \log \frac{1}{1 - p_{\text{fail},\tau}}. \end{aligned} \quad (5)$$

We bound the two terms separately, then average over prefixes and sum over queries using the KL chain rule.

Comparison with previous analyses. The analyses in [14,10,3] use probabilities of complete transcripts without conditioning on simulator success. Bhattacharya and Nandi [2] analyze response distributions for individual queries, but likewise do not condition on success. We instead condition the response distribution on success for each query and account for failure separately.

3 Security of the Uniform Simulator

In this section, we define the uniform simulator, compute the conditional response distributions, and introduce a hybrid world to compare the real and ideal worlds. Subsection 3.7 combines the divergence bounds to prove $3n/4$ security. The bias from earlier inverse responses will motivate the Gyroscope simulator.

3.1 The uniform simulator

For an integer $m \geq 1$, Algorithm 1 defines the uniform simulator USim_m , with access to a uniformly random construction oracle $\Gamma : \mathcal{G} \rightarrow \mathcal{G}$. It maintains partial permutations P_0, P_1 , with

$$\begin{aligned} \text{Dom} &:= \text{dom}(P_0) = \text{dom}(P_1), \quad \text{Ran}_b := \text{ran}(P_b), \quad U := \mathcal{G} \setminus \text{Dom}, \\ P_0(x) \oplus P_1(x) &= \Gamma(x) \quad (x \in \text{Dom}), \end{aligned}$$

where $b \in \{0, 1\}$. It also maintains a terminal flag **Failure**, initialized to zero. Once **Failure** = 1, every primitive query returns $\perp$.

Algorithm 1 The uniform simulator USim_m .

1: function $S_b(x)$ 2: if Failure = 1 then return $\perp$ 3: if $x \in \text{Dom}$ then return $P_b(x)$ 4: $z \leftarrow \Gamma(x)$ 5: $y_b \leftarrow \text{Unif}(\mathcal{G} \setminus (\text{Ran}_b \cup (\text{Ran}_{1-b} \oplus z)))$ 6: $P_b(x) \leftarrow y_b$; $P_{1-b}(x) \leftarrow y_b \oplus z$ 7: return y_b	1: function $S_b^{-1}(y)$ 2: if Failure = 1 then return $\perp$ 3: if $y \in \text{Ran}_b$ then return $P_b^{-1}(y)$ 4: for $j = 1, \dots, m$ do 5: $x \leftarrow \text{Unif}(U)$; $z \leftarrow \Gamma(x)$ 6: if $y \oplus z \notin \text{Ran}_{1-b}$ then 7: $P_b(x) \leftarrow y$; $P_{1-b}(x) \leftarrow y \oplus z$ 8: return x 9: Failure $\leftarrow 1$; return $\perp$
--	--

Let Ideal_m denote the ideal world with simulator USim_m . We fix $m = 3$ for the rest of this section.

Inverse responses with a fixed construction oracle. We first compute the distribution of inverse responses with Γ fixed. Their likelihoods determine the conditional distribution of construction oracle values not revealed in the transcript.

Proposition 1 (Inverse distribution for a fixed Γ). *Fix Γ and the current tables P_0, P_1 with **Failure** = 0, and let $y \notin \text{Ran}_b$. Put*

$$L := |U|, \quad E_{\text{leg}} := \{x \in U : y \oplus \Gamma(x) \notin \text{Ran}_{1-b}\},$$

$$I := L - |E_{\text{leg}}|.$$

Then

$$\Pr[S_b^{-1}(y) = x \mid \Gamma] = \frac{\mathbf{1}[x \in E_{\text{leg}}]}{L} \left(1 + \frac{I}{L} + \left(\frac{I}{L} \right)^2 \right),$$

$$\Pr[S_b^{-1}(y) = \perp \mid \Gamma] = \left(\frac{I}{L} \right)^3.$$

If $E_{\text{leg}} \neq \emptyset$, then conditional on success and on Γ , the inverse output is uniform on E_{leg} .

Proof. The simulator's three sampling trials are independent. Each rejects with probability I/L and selects a specified legal input with probability $1/L$, giving the first formula. The second follows because the simulator returns $\perp$ if and only if all three trials are rejected. $\square$

3.2 Query classification

To classify queries in each world, we partition $\mathcal{G}$ according to a transcript τ containing no failure response: inputs involved in primitive queries, inputs queried

only to the construction oracle, and inputs not involved in either type of query. Let

$$\text{Dom}_\tau^{\text{prim}} := \{x : x \text{ is a forward primitive input or inverse output in } \tau\}.$$

For a fixed transcript τ , let Ran_b be the range of the recorded permutation table on side b , and put

$$\begin{aligned} \text{Ran}_0 &:= O_0(\text{Dom}_\tau^{\text{prim}}), & \text{Ran}_1 &:= O_1(\text{Dom}_\tau^{\text{prim}}), \\ s_{\text{prim}} &:= |\text{Dom}_\tau^{\text{prim}}| = |\text{Ran}_0| = |\text{Ran}_1|. \end{aligned}$$

Let Cons denote the list of query–response pairs for inputs queried only to the construction oracle in τ , in the order in which they are queried. Define

$$s_{\text{con}} := |\text{Cons}|, \quad s := s_{\text{prim}} + s_{\text{con}}$$

and write

$$\text{Cons} = ((x_1, z_1), \dots, (x_{s_{\text{con}}}, z_{s_{\text{con}}})) .$$

Its input set is $\text{dom}(\text{Cons}) := \{x_1, \dots, x_{s_{\text{con}}}\}$. Define

$$E_{\text{new}} := \mathcal{G} \setminus (\text{Dom}_\tau^{\text{prim}} \cup \text{dom}(\text{Cons})).$$

The construction values are revealed at $\text{Dom}_\tau^{\text{prim}} \cup \text{dom}(\text{Cons})$ and unrevealed at E_{new} . In *Ideal*, the simulator state satisfies

$$\text{Dom} = \text{Dom}_\tau^{\text{prim}}, \quad \text{Ran}_b = \text{ran}(\text{P}_b), \quad U = \mathcal{G} \setminus \text{Dom}_\tau^{\text{prim}}.$$

These sets partition the input space as

$$\mathcal{G} = \text{Dom}_\tau^{\text{prim}} \cup \text{dom}(\text{Cons}) \cup E_{\text{new}}.$$

Figures 1–3 use this partition to illustrate the queries in each world.

Among queries whose responses are not already determined, we distinguish four cases:

1. A construction query $\text{Con}(x)$ with $x \in E_{\text{new}}$, in the green input region.
2. A forward primitive query $\text{Fwd}_b(x)$ with $x \in \text{dom}(\text{Cons})$, in the blue input region.
3. A forward primitive query $\text{Fwd}_b(x)$ with $x \in E_{\text{new}}$, in the green input region.
4. An inverse primitive query $\text{Inv}_b(y)$ with $y \notin \text{Ran}_b$. If the query succeeds, its inverse output lies in the blue or green input region.

3.3 Response distributions in Real

Fix a transcript τ in *Real*. Figure 1 illustrates construction and primitive queries.




Fig. 1. Fresh queries in *Real*, conditional on the sampled vector u . The blue output regions represent the coordinate sets of u and $u \oplus z$, where $z = (z_i)$. The two remaining bijections between the green regions are independent and uniform.

To obtain the response distributions in *Real*, we average over the compatible Π_0 -output vectors at $\text{dom}(\text{Cons})$. These form the set

$$\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons}) := \left\{ X \in \mathcal{G}^{s_{\text{con}}} : \begin{array}{l} X_i \notin \text{Ran}_0 \quad (1 \leq i \leq s_{\text{con}}), \\ X_i \oplus z_i \notin \text{Ran}_1 \quad (1 \leq i \leq s_{\text{con}}), \\ X_1, \dots, X_{s_{\text{con}}} \text{ distinct}, \\ X_1 \oplus z_1, \dots, X_{s_{\text{con}}} \oplus z_{s_{\text{con}}} \text{ distinct} \end{array} \right\}. \quad (6)$$

Each compatible vector u specifies both permutation outputs $(u_i, u_i \oplus z_i)$ at $\text{dom}(\text{Cons})$. Since Π_0 and Π_1 are independent uniform permutations, all permutation pairs consistent with the transcript are equally likely. Thus, conditioned on $\text{Ran}_0, \text{Ran}_1, \text{Cons}$, the output vector is uniform on $\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})$. Equivalently, the unexposed permutation values have conditional distribution

$$\mathcal{L}((\Pi_0(x_i))_{i=1}^{s_{\text{con}}} \mid \text{Ran}_0, \text{Ran}_1, \text{Cons}) = \text{Unif}(\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})). \quad (7)$$

For a fixed $u \in \text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})$, put

$$\text{Ran}_{0,u} = \text{Ran}_0 \cup \{u_i : 1 \leq i \leq s_{\text{con}}\}, \quad \text{Ran}_{1,u} = \text{Ran}_1 \cup \{u_i \oplus z_i : 1 \leq i \leq s_{\text{con}}\}.$$

Conditioned on u , the two permutations on E_{new} are independent uniform bijections onto $\mathcal{G} \setminus \text{Ran}_{0,u}$ and $\mathcal{G} \setminus \text{Ran}_{1,u}$, respectively. At $x_i \in \text{dom}(\text{Cons})$, the pair is already fixed as $(u_i, u_i \oplus z_i)$. A fresh inverse query $\Pi_0^{-1}(y)$ returns x_i if $y = u_i$; otherwise its inverse output is uniform on E_{new} , independently of a uniform companion output on $\mathcal{G} \setminus \text{Ran}_{1,u}$. The case of $\Pi_1^{-1}(y)$ is symmetric.

3.4 Response distributions in Ideal_3

We call the collection $(\Gamma(x))_{x \in \mathcal{G}}$ the table of Γ . Figure 2 illustrates the responses in Ideal_3 .

For fixed Γ and a fresh inverse query, define

$$\vartheta_L(I) := \frac{1}{L} \left(1 + \frac{I}{L} + \left(\frac{I}{L} \right)^2 \right),$$

where L is the number of available inputs and I the number of illegal ones. By Proposition 1, each input in E_{leg} is returned with probability $\vartheta_L(I)$. The Unif labels in Figure 2 describe the distribution conditional on the inverse output lying in the indicated region. The failure probability is $(I/L)^3$ and the probability of a successful response is $(L - I)\vartheta_L(I) = 1 - (I/L)^3$.

Fix an augmented transcript prefix τ containing no failure response. Let $\text{Dom}_\tau^{\text{pub}} = \text{Dom}_\tau^{\text{prim}} \cup \text{dom}(\text{Cons})$ be the inputs whose construction values are revealed. Let $\Phi_\tau : \mathcal{G} \rightarrow \mathcal{G}$ denote a table agreeing with these revealed values. Let μ_τ be the distribution on these tables under which all unrevealed values are independent and uniform on $\mathcal{G}$. Sums over Φ_τ range over these tables. For each earlier fresh inverse call j , let τ_j be its preceding prefix, b_j its side, and y_j its target. With $\text{Ran}_{b,j}$ denoting the range on side b at that prefix, put

$$E_j = \mathcal{G} \setminus \text{Dom}_{\tau_j}^{\text{prim}}, \quad L_j = |E_j|, \quad \text{Forb}_j = y_j \oplus \text{Ran}_{1-b_j,j}.$$

Define

$$I_j(\Phi_\tau) = \sum_{x \in E_j} \mathbf{1}[\Phi_\tau(x) \in \text{Forb}_j].$$

Lemma 2 (Conditional distribution of Γ). *In Ideal_3 , the construction oracle Γ conditioned on a transcript prefix τ containing no failure response has distribution π_τ , where*

$$W_\tau(\Phi_\tau) := \prod_j \vartheta_{L_j}(I_j(\Phi_\tau)), \quad \pi_\tau(\Phi_\tau) := \frac{W_\tau(\Phi_\tau)}{\mathbb{E}_{\mu_\tau}[W_\tau]} \mu_\tau(\Phi_\tau). \quad (8)$$

Proof. For a fixed table Φ_τ , the likelihood of the observed inverse response is its legality indicator times $\vartheta_{L_j}(I_j)$. Compatibility with the augmented response fixes that indicator to one. Conditional on the recorded construction values, the probabilities of the observed forward responses and the constraints on revealed construction values do not depend on the unrevealed values. Multiplying the remaining likelihoods and applying Bayes' rule gives the formula. This calculation sums over all rejected candidates, including candidates tested in earlier calls. $\square$




Fig. 2. Fresh queries in Ideal_3 , conditioned on the transcript τ . For the inverse query, the responses are conditional on the sampled construction-oracle table Φ_τ .

Construction and forward responses. A construction response at an input x with an unrevealed construction value has distribution

$$\Pr_{\text{Ideal}_3}[\Gamma(x) = z \mid \tau] = \sum_{\Phi_\tau} \pi_\tau(\Phi_\tau) \mathbf{1}[\Phi_\tau(x) = z].$$

A fresh forward primitive query uses the same distribution for its construction value, unless that value is already fixed by the transcript. Given this value, Algorithm 1 samples the two permutation outputs uniformly among the legal pairs. The construction value can be nonuniform because π_τ is conditioned on earlier inverse responses.

An inverse response conditioned on success. For a fresh $\text{Inv}_b(y)$, put $L = |U|$ and $\text{Forb} = y \oplus \text{Ran}_{1-b}$, and let $I(\Phi_\tau)$ be the number of inputs $x \in U$ with $\Phi_\tau(x) \in \text{Forb}$. For the table Φ_τ , the current call succeeds with probability $1 - (I(\Phi_\tau)/L)^3$. Conditioning on this success therefore changes the distribution of the table of Γ to

$$\pi_\tau^+(\Phi_\tau) = \frac{\pi_\tau(\Phi_\tau)(1 - (I(\Phi_\tau)/L)^3)}{\mathbb{E}_{\pi_\tau}[1 - (I/L)^3]}. \quad (9)$$

Given Φ_τ and success, the inverse output is uniform among the legal inputs by Proposition 1. Combining this conditional distribution with π_τ^+ gives the joint

distribution of the table of I and the inverse output:

$$Q_\tau^+(\Phi_\tau, x) = \frac{\pi_\tau(\Phi_\tau) \mathbf{1}[\Phi_\tau(x) \notin \text{Forb}] \vartheta_L(I(\Phi_\tau))}{\mathbb{E}_{\pi_\tau}[(L - I) \vartheta_L(I)]}. \quad (10)$$

Its marginal $\sum_{\Phi_\tau} Q_\tau^+(\Phi_\tau, x)$ is the distribution of the inverse output. The pair $(x, \Phi_\tau(x))$ determines the full augmented response. Thus this joint distribution determines the conditional response distribution $p_{\text{Ideal}_3, k+1}^{\tau, +}$. The normalization accounts for the different success probabilities of the tables compatible with τ .

3.5 Where the sharper bounds come from

For a fresh construction query at $x \in E_{\text{new}}$, let p_τ^{Real} and $p_\tau^{\text{Ideal}_3}$ denote the response distributions given a prefix τ possible in both worlds and containing no failure response. For $u \leftarrow \text{Unif}(\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons}))$, define

$$c_u(z) := |\text{Ran}_{0,u} \cap (\text{Ran}_{1,u} \oplus z)|, \quad p_u(z) := \frac{N - 2s + c_u(z)}{(N - s)^2}.$$

Here p_u is the response distribution in **Real** conditioned on u . The distributions above give the direct comparison

$$p_\tau^{\text{Real}}(z) - p_\tau^{\text{Ideal}_3}(z) = \frac{\mathbb{E}_u[c_u(z)] - s^2/N}{(N - s)^2} - \frac{1}{N} \left(\frac{\mathbb{E}_{\mu_\tau}[W_\tau \mid \Phi_\tau(x) = z]}{\mathbb{E}_{\mu_\tau}[W_\tau]} - 1 \right).$$

The first term comes from the occupied permutation ranges; the second comes from earlier inverse responses in **Ideal**₃.

Retaining cancellation in Real. The following bound illustrates the loss in earlier proofs [14, 3]:

$$\frac{1}{N p_u(z)} = 1 - \frac{c_u(z) - s^2/N}{N - 2s + c_u(z)} \geq 1 - \frac{c_u(z)}{N}.$$

This drops the subtraction of s^2/N , the average of $c_u(z)$ over uniform z . The average alone contributes s^2/N^2 to the loss for each query, giving $O(q^3/N^2)$ over q queries. We retain the cancellation in the divergence bound. Since $\sum_z c_u(z) = s^2$ and $c_u(z) \leq s$,

$$\chi^2(p_u \parallel \text{Unif}(\mathcal{G})) = \frac{N}{(N - s)^4} \sum_z (c_u(z) - s^2/N)^2 \leq \frac{N s^3}{(N - s)^4} = O(s^3/N^3).$$

Averaging over u preserves this bound.

*Including rejected trials in Ideal*₃. The analyses in [14, Section 4.4] and [3, Appendix D] lower-bound transcript probabilities by excluding certain executions. Mennink and Preneel require rejected candidates to avoid the inputs in the final transcript and all earlier rejected candidates. Bhaumik et al. require the

rejected candidate to avoid inputs queried to the construction oracle in the final transcript.

In these calculations, the freshness condition excludes $O(q)$ candidate inputs. Conditional on this condition, the rejection probability is s_{prim}/N , where s_{prim} is the size of each permutation range before the query. This restriction introduces a term of order qs_{prim}/N^2 in the lower bound on the ratio of transcript probabilities for that query.

More trials can reduce the s_{prim}^2/N^2 term caused by stopping after two trials. However, the separate $O(qs_{\text{prim}}/N^2)$ loss remains if the same freshness restriction is retained. Our weight W_τ accounts for rejected candidates without these restrictions. Conditional on the other unrevealed values, the probability of $\Gamma(x) = z$ is proportional to $e^{b(z)}$, where

$$b(z) = \sum_j \delta_j \mathbf{1}[z \in \text{Forb}_j], \quad \delta_j := \log \frac{\vartheta_{L_j}(I_j^- + 1)}{\vartheta_{L_j}(I_j^-)}, \quad 0 \leq \delta_j \leq \frac{4}{N},$$

and I_j^- is the number of illegal candidates other than x at query j . Using only $0 \leq b(z) \leq 4q/N$ gives a KL bound of order q^2/N^2 . Using $|\text{Forb}_j| \leq q$ also gives

$$\mathbb{E}_{z \leftarrow \text{Unif}(\mathcal{G})}[b(z)^2] \leq \left(\sum_j \delta_j \right) \sum_j \delta_j \frac{|\text{Forb}_j|}{N} \leq \frac{16q^3}{N^3}.$$

The reweighting lemma gives $\chi^2(\text{Unif}(\mathcal{G}) \| p_\tau^{\text{Ideal}_3}) = O(q^3/N^3)$, including averaging over the other unrevealed values; see Appendix D. The hybrid below separates these two contributions in the security proof.

3.6 The hybrid Hyb

The hybrid **Hyb** answers fresh construction queries uniformly and, given a construction value, samples forward outputs uniformly among legal pairs. We compare **Real** with **Hyb** to bound the effect of the permutation constraints, and **Hyb** with **Ideal**₃ to bound the bias from earlier inverse responses and account for simulator failure.


The hybrid uses the notation of Section 3.2. Recording a primitive response (x, y_0, y_1) adds x to $\text{Dom}_\tau^{\text{prim}}$ and y_b to Ran_b for both sides, and removes $(x, y_0 \oplus y_1)$ from Cons , if present. Figure 3 illustrates the hybrid's query procedures.

For a fresh inverse query $\text{Inv}_b(y)$, let

$$\begin{aligned} \text{Forb} &= y \oplus \text{Ran}_{1-b}, \\ E_{\text{con}}^{\text{leg}} &= \{x : (x, z) \in \text{Cons}, z \notin \text{Forb}\}, \\ k &= |E_{\text{con}}^{\text{leg}}|, \quad u = |E_{\text{new}}|, \quad \rho = (N - s_{\text{prim}})/N, \quad \mathcal{N}_{\text{Hyb}} = k + \rho u. \end{aligned}$$

A candidate with a revealed legal construction value is accepted with probability one; a uniform unrevealed construction value avoids Forb with probability ρ .

Construction and forward queries



Inverse query $\text{Inv}_0(y)$, $y \notin \text{Ran}_0$




Fig. 3. Fresh queries in Hyb, conditioned on the transcript τ . For the inverse query, red marks construction-only inputs whose revealed values are forbidden. All inputs in the green region remain eligible.

Algorithm 2 The hybrid Hyb: fresh queries.

1: function $\text{Con}(x)$	1: function $\text{Inv}_b(y)$
2: $z \leftarrow \text{Unif}(\mathcal{G})$	2: Sample x assigning probability
3: Append (x, z) to Cons ; return z	$1/\mathcal{N}_{\text{Hyb}}$ to each $x \in E_{\text{con}}^{\text{leg}}$
	and $\rho/\mathcal{N}_{\text{Hyb}}$ to each $x \in E_{\text{new}}$
4: function $\text{Fwd}_b(x)$	3: if $x \in E_{\text{con}}^{\text{leg}}$ then
5: if $x \in \text{dom}(\text{Cons})$ then	4: Read (x, z) from Cons
6: Read (x, z) from Cons	5: $y_{1-b} \leftarrow y \oplus z$
7: else	6: else
8: $z \leftarrow \text{Unif}(\mathcal{G})$	7: $y_{1-b} \leftarrow \text{Unif}(\mathcal{G} \setminus \text{Ran}_{1-b})$
9: Sample $y_0 \notin \text{Ran}_0$, $y_1 \notin \text{Ran}_1$	8: $y_b \leftarrow y$
uniformly with $y_0 \oplus y_1 = z$	9: Record and return (x, y_0, y_1)
10: Record and return (x, y_0, y_1)	

Accordingly, each input in $E_{\text{con}}^{\text{leg}}$ is returned with probability $1/\mathcal{N}_{\text{Hyb}}$, and each input in E_{new} with probability $\rho/\mathcal{N}_{\text{Hyb}}$. The Unif labels in Figure 3 describe the distributions conditional on the inverse output lying in these sets. Algorithm 2 specifies the query procedures.

We write X for the inverse output, $K_{\text{Hyb}}^{\text{out}}$ for its distribution, and $K_{\text{Hyb},x}^{\text{comp}}$ for that of y_{1-b} given $X = x$.

Every transcript prefix in Hyb is attainable in Ideal_3 , since its inverse responses can occur on their first trials. This gives the support condition for the KL chain rule.

The KL chain rule (2) averages the conditional divergences over prefixes of the first world in each comparison. Equation (5) accounts for simulator failure.

3.7 Security bound

The following proposition bounds the divergence between response distributions and the failure probability. Here $s = s_{\text{prim}} + s_{\text{con}}$ is the number of distinct inputs whose construction values are revealed in the prefix, so $s \leq q$.

Proposition 2 (Conditional divergence and failure probability). *The conditional KL divergences at an attainable prefix containing no failure response are bounded as follows. The last column conditions the next response in Ideal_3 on success.*

Response	Real to Hyb	Hyb to Ideal_3^+
(1) Fresh construction response	$4s^3/N^3$	$20q^3/N^3$
(2) Forward at a construction-only input	$320s^3/N^3$	0
(3) Forward with an unrevealed construction value	$324s^3/N^3$	$20q^3/N^3$
(4) Inverse response	$101s^3/N^3$	$21q^3/N^3$.

At each attainable prefix in Ideal_3 containing no failure response, a fresh inverse call fails with probability at most $10q^3/N^3$. Consequently,

$$\Pr_{\text{Ideal}_3} [\text{Failure} = 1] \leq 10q^4/N^3. \quad (11)$$

For (1), Subsection 3.5 gives the calculation for a fresh construction response. For (2), we compare the permutation outputs conditioned on the revealed construction value; their distributions in Hyb and Ideal_3 coincide. For (3), we first compare the construction values, then the permutation outputs conditioned on those values. For (4), we first compare the inverse outputs, then the companion outputs conditioned on them. The KL chain rule combines the two bounds in (3) and (4). Appendices B.1–C.3 prove the bounds from Real to Hyb. Appendix D uses the conditional distribution (8) to prove the bounds for the comparison with Ideal_3 , including the failure probability.

The following theorem bounds the indistinguishability advantage for USim_3 .

Theorem 1. *Let $N = 2^n \geq 64$ and $q \leq N/64$. Then*

$$\text{Adv}_{\text{SoP}, \text{USim}_3}^{\text{indiff}}(q) \leq 17 \frac{q^2}{N^{3/2}}. \quad (12)$$

Proof. Proposition 2 bounds the conditional KL divergence between response distributions by $324s^3/N^3$ from Real to Hyb, and by $21q^3/N^3$ from Hyb to the distribution of successful responses in Ideal_3 . It also bounds the failure probability

for each query by $10q^3/N^3$. Since $q/N \leq 1/64$, the failure contribution to KL is at most $\log(1/(1 - 10q^3/N^3)) \leq 11q^3/N^3$. By (5) and the KL chain rule,

$$D_{\text{KL}}(T_{\text{Real}} \| T_{\text{Hyb}}) \leq 324q^4/N^3, \quad D_{\text{KL}}(T_{\text{Hyb}} \| T_{\text{Ideal}_3}) \leq 32q^4/N^3.$$

Using Pinsker's inequality, we obtain

$$\text{TV}(T_{\text{Real}}, T_{\text{Ideal}_3}) \leq (\sqrt{162} + 4) \frac{q^2}{N^{3/2}} \leq 17 \frac{q^2}{N^{3/2}}. \quad (13)$$

□

4 A Matching Attack on the Uniform Simulator

The following attack distinguishes **Real** from the ideal world with USim_m using $O(N^{3/4})$ primitive queries, matching the security bound for USim_3 .

4.1 Distinguisher

Throughout this section, $N = 2^n \geq 64$, $m \geq 1$, and r is a power of two satisfying

$$64n \leq r \leq \min \left\{ \frac{N}{2^{11}}, \frac{N}{24n} \right\}. \quad (14)$$

Fix a subspace $V = \{v_1, \dots, v_r\} \leq \mathcal{G}$ of size r and a vector $c \notin V$. Let $y_i := c \oplus v_i$ and $C := c \oplus V = \{y_1, \dots, y_r\}$.

Phase 1 queries $O_0^{-1}(v)$ for every $v \in V$. For a fresh query $O_1^{-1}(v)$ in Phase 2, the forbidden set contains $v \oplus V = V$. These inverse responses make unrevealed construction values more likely to lie in V . In Phase 3, the distinguisher queries O_0^{-1} at values in C and tests whether the construction values lie in V more often than expected in **Real**. Algorithm 3 specifies the distinguisher.

4.2 Analysis of the attack

Each phase has r iterations, each making at most two primitive queries. Thus every execution makes at most $6r$ queries.

We first analyze the auxiliary world Ideal^∞ . A fresh inverse query returns a uniformly chosen input from E_{leg} if this set is nonempty. Otherwise it sets the terminal flag **Failure** = 1 and returns $\perp$. Assuming that $E_{\text{leg}} \neq \emptyset$, this is the response obtained by uniform sampling without a trial limit.

The attack statistic. In iteration k of Phase 3, suppose $O_0^{-1}(y_k)$ returns a fresh input x_k . Let τ_k be the transcript observed by D up to and including this response, before the query $O_1(x_k)$. Thus x_k is known, but $Z_k = y_k \oplus O_1(x_k)$ has not yet been revealed to D . Put $T_k = \mathbf{1}[Z_k \in V] - \pi_k$.

In **Real**, Z_k is uniform on $\mathcal{Z}_k$ conditional on τ_k , so T_k has conditional mean zero. In Ideal^∞ , earlier inverse responses change the distribution of the unrevealed

Algorithm 3 The distinguisher $D^{O_0^\pm, O_1^\pm}$.

<pre> 1: $\text{Ran}_0 \leftarrow \emptyset; \text{Ran}_1 \leftarrow \emptyset;$ $T_{\text{atk}} \leftarrow 0$ <i>Phase 1.</i> 2: for $i = 1, \dots, r$ do 3: $x \leftarrow O_0^{-1}(v_i)$ 4: if $x = \perp$ then return 1 5: $y \leftarrow O_1(x)$ 6: $\text{Ran}_0 \leftarrow \text{Ran}_0 \cup \{v_i\};$ $\text{Ran}_1 \leftarrow \text{Ran}_1 \cup \{y\}$ <i>Phase 2.</i> 7: for $i = 1, \dots, r$ do 8: if $v_i \in \text{Ran}_1$ then continue 9: $x \leftarrow O_1^{-1}(v_i)$ 10: if $x = \perp$ then return 1 11: $y \leftarrow O_0(x)$ 12: $\text{Ran}_0 \leftarrow \text{Ran}_0 \cup \{y\};$ $\text{Ran}_1 \leftarrow \text{Ran}_1 \cup \{v_i\}$ </pre>	<pre> <i>Phase 3.</i> 13: for $k = 1, \dots, r$ do 14: if $y_k \in \text{Ran}_0$ then continue 15: $x_k \leftarrow O_0^{-1}(y_k)$ 16: if $x_k = \perp$ then return 1 17: $\text{Ran}_{1,k} \leftarrow \text{Ran}_1;$ $\mathcal{Z}_k \leftarrow \mathcal{G} \setminus (y_k \oplus \text{Ran}_{1,k})$ 18: $\pi_k \leftarrow V \cap \mathcal{Z}_k / \mathcal{Z}_k$ 19: $w_k \leftarrow O_1(x_k);$ $Z_k \leftarrow y_k \oplus w_k$ 20: $T_{\text{atk}} \leftarrow T_{\text{atk}} + \mathbf{1}[Z_k \in V] - \pi_k$ 21: $\text{Ran}_0 \leftarrow \text{Ran}_0 \cup \{y_k\};$ $\text{Ran}_1 \leftarrow \text{Ran}_1 \cup \{w_k\}$ 22: return $\mathbf{1}\left[T_{\text{atk}} \geq \frac{r^3}{64N^2}\right].$ </pre>
--	--

construction values. With all other construction values fixed, assigning to x_k a value forbidden in an earlier call removes x_k from that call's legal candidates. This makes the observed inverse response more likely. Conditioning on that response therefore favors such values at x_k . Phase 1 ensures that V is contained in the set of construction values forbidden at each fresh inverse query in Phase 2, so these queries contribute a positive bias toward V . Appendix E.4 gives a lower bound on the conditional mean of T_k .

In Ideal^∞ , define

$$d_k = \mathbb{E}_{\text{Ideal}^\infty}[T_k \mid \tau_k], \quad M_\infty = T_{\text{atk}} - \sum_{k=1}^r d_k.$$

We set $T_k = d_k = 0$ when the inverse query is skipped because its response is already known, or when the algorithm has terminated on failure. Note that T_k in Real and $T_k - d_k$ in Ideal^∞ are martingale differences. We call $\sum_k d_k$ the *drift*.

Write $\theta = r^3/(64N^2)$ for the threshold in (15). Let Bad be the event that all three phases complete but $\sum_k d_k \leq 2\theta$. The bounds below, proved in Appendix E, control the martingale parts on the scale $r/\sqrt{N}$. On completed executions outside Bad , the drift in Ideal^∞ exceeds $2\theta = r^3/(32N^2)$. Its ratio to this scale is at least of order $r^2/N^{3/2}$, explaining why $r \asymp N^{3/4}$ suffices.

Lemma 3 (Bounds for the attack statistic). *For every power of two r satisfying (14),*

$$\mathbb{E}_{\text{Real}}[T_{\text{atk}}^2] \leq \frac{2r^2}{N}, \quad \mathbb{E}_{\text{Ideal}^\infty}[M_\infty^2] \leq \frac{5r^2}{N}, \quad (16)$$

$$\Pr_{\text{Ideal}^\infty}[\text{Bad}] \leq 128 \frac{r}{N} + \frac{r^2}{N^2}. \quad (17)$$

These bounds give the following advantage guarantee for every trial budget.

Theorem 2 (Indifferentiability attack on the uniform simulator). *Let $N = 2^n \geq 64$ and $m \geq 1$, and let r be a power of two satisfying*

$$64n \leq r \leq \min \left\{ \frac{N}{2^{11}}, \frac{N}{24n} \right\}.$$

The distinguisher D in Algorithm 3 makes at most $6r$ primitive queries and satisfies

$$\text{Adv}_{\text{SoP,USim}_m}^{\text{indiff}}(D) \geq 1 - 28672 \frac{N^3}{r^4} - 129 \frac{r}{N}.$$

Proof. We couple Ideal_m and Ideal^∞ using the same construction oracle Γ and sampling the same input at each of the first m trials of an inverse query. They agree until Ideal_m fails. If $E_{\text{leg}} = \emptyset$, both fail; otherwise Ideal^∞ continues sampling. Since D outputs 1 on failure, extending the trials cannot decrease the probability that $D = 0$. Hence $\Pr_{\text{Ideal}_m}[D = 1] \geq \Pr_{\text{Ideal}^\infty}[D = 1]$ for every $m \geq 1$.

In Ideal^∞ , $D = 0$ and $\neg \text{Bad}$ imply $T_{\text{atk}} < \theta$ and hence $M_\infty < -\theta$. Chebyshev's inequality and (16)–(17) give

$$\begin{aligned} \text{Adv}_{\text{SoP,USim}_m}^{\text{indiff}}(D) &\geq 1 - \Pr_{\text{Real}}[D = 1] - \Pr_{\text{Ideal}^\infty}[D = 0] \\ &\geq 1 - \frac{\mathbb{E}_{\text{Real}}[T_{\text{atk}}^2] + \mathbb{E}_{\text{Ideal}^\infty}[M_\infty^2]}{\theta^2} - \Pr_{\text{Ideal}^\infty}[\text{Bad}] \\ &\geq 1 - 28672 \frac{N^3}{r^4} - 128 \frac{r}{N} - \frac{r^2}{N^2} \\ &\geq 1 - 28672 \frac{N^3}{r^4} - 129 \frac{r}{N}. \quad \square \end{aligned}$$

For $n \geq 64$, taking $r = 2^{\lceil 3n/4 \rceil + 4}$ gives $6r < 162N^{3/4}$ and advantage greater than $1/2$ (Appendix E.6).

5 The Gyroscope Simulator

We derive an acceptance rule from the bias accumulated under uniform simulation, then use it to define the Gyroscope simulator.

5.1 Deriving the correction

Bias under uniform simulation. Fix an augmented transcript prefix τ in Ideal_3 containing no failure response. Earlier inverse responses bias the construction values that remain unrevealed. Let I_τ index the earlier fresh inverse queries in query order. For query j , let b_j be its side and y_j its target. Immediately before

that query, let d_j be the number of assigned inputs and let $\text{Ran}_{0,j}, \text{Ran}_{1,j}$ be the occupied ranges. Thus

$$L_j := N - d_j, \quad \text{Forb}_j := y_j \oplus \text{Ran}_{1-b_j,j}$$

are the number of available candidates and the set of forbidden construction values, as in Section 3.4.

For m trials with rejection probability t , the possibilities of zero through $m - 1$ rejections before the observed inverse response give the factor $H_m(t) := \sum_{\ell=0}^{m-1} t^\ell$. Lemma 2 weights the table by $\prod_j L_j^{-1} H_3(p_{\text{rej},j})$. A higher rejection probability increases H_3 , making the observed inverse response more likely. Thus the observed responses favor unrevealed values that would have been rejected in earlier calls. An input x whose construction value is still unrevealed was available at every earlier inverse call. Fix the other construction values as w , and let $u_j(w)$ be their contribution to the rejection probability in call j , with denominator L_j . If $\Gamma(x) = z$, then

$$p_{\text{rej},j} = u_j(w) + \frac{\mathbf{1}[z \in \text{Forb}_j]}{L_j}.$$

The coefficient $1/L_j$ is the probability of selecting x in one trial. Up to a term independent of z , the log likelihood contributed by call j is therefore

$$\log H_3\left(u_j(w) + \frac{\mathbf{1}[z \in \text{Forb}_j]}{L_j}\right) - \log H_3(u_j(w)).$$

Assuming that the rejection probability is small, its leading dependence on z is $L_j^{-1} \mathbf{1}[z \in \text{Forb}_j]$. The leading contributions add over the earlier inverse calls. We denote their sum by

$$\text{Corr}_\tau(z) := \sum_{j \in I_\tau} \frac{1}{L_j} \mathbf{1}[z \in \text{Forb}_j].$$

We write Corr for Corr_τ when the prefix τ is clear. This suggests approximating the conditional distribution of $\Gamma(x)$, given w , by a distribution proportional to $e^{\text{Corr}_\tau(z)}$.

Choosing the acceptance probability. For a fresh inverse query $\text{Inv}_b(y)$ in Real , conditional on an inverse output $x \in E_{\text{new}}$, the companion output is close to uniform over $\mathcal{G} \setminus \text{Ran}_{1-b}$ (Lemma 8). Since it equals $y \oplus z$, the corresponding construction value is close to uniform on $\mathcal{G} \setminus \text{Forb}$.

The uniform simulator accepts every legal candidate, so its acceptance rule does not remove the accumulated factor $e^{\text{Corr}_\tau(z)}$. To offset this factor in inverse responses, we instead accept a legal candidate with construction value z with probability

$$\alpha_\tau(z) := e^{-\text{Corr}_\tau(z)}.$$

Multiplication by this probability cancels the leading factor in the distribution of accepted values. Each earlier query j then multiplies the acceptance probability

by e^{-1/L_j} for construction values in Forb_j . This is the acceptance rule of the Gyroscope simulator, defined below. Appendix L bounds the error under the modified rule, including the effect of its additional rejections. The leading bias in unrevealed construction values also appears in *Real*; Appendix I quantifies the discrepancy.

5.2 Simulator definition and implementation

We define GSim_m with $m \geq 1$ trials per inverse query and analyze GSim_5 .

The simulator uses the tables P_0, P_1 , terminal flag, and forward procedure of Algorithm 1 without change. Compared with USim_m , the simulator changes the procedure for fresh inverse queries. It also stores enough information to reconstruct the permutation tables as they were immediately before each successful fresh inverse query.

Using the notation of Section 5.1, the simulator records its *inverse history* $\text{Hist}_\tau := ((b_j, y_j, L_j, d_j))_{j \in I_\tau}$ and uses the acceptance probability $\alpha_\tau(z) = e^{-\text{Corr}_\tau(z)}$ derived there.

Algorithm 4 defines the inverse primitive oracle of GSim_m .

Algorithm 4 The inverse primitive oracle of GSim_m .

```

1: function  $S_b^{-1}(y)$ 
2:   if Failure = 1 then return  $\perp$ 
3:   if  $y \in \text{Ran}_b$  then return  $P_b^{-1}(y)$ 
4:    $U \leftarrow \mathcal{G} \setminus \text{Dom}; L \leftarrow |U|; d \leftarrow |\text{Dom}|; \text{Forb} \leftarrow y \oplus \text{Ran}_{1-b}$ 
5:   for  $\ell = 1, \dots, m$  do
6:      $x \leftarrow \text{Unif}(U); z \leftarrow \Gamma(x)$ 
7:     if  $z \notin \text{Forb}$  then
8:       accept  $\leftarrow 0$ 
9:       With probability  $\alpha_\tau(z)$ , set accept  $\leftarrow 1$ 
10:    if accept = 1 then
11:       $P_b(x) \leftarrow y; P_{1-b}(x) \leftarrow y \oplus z$ 
12:      Append  $(b, y, L, d)$  to  $\text{Hist}_\tau$ 
13:    return  $x$ 
14:   Failure  $\leftarrow 1$ ; return  $\perp$ 
```

The m trials sample from U with replacement; $U, \text{Forb}, \text{Corr}_\tau$ stay fixed during the call. From here on we take $m = 5$, and let $\text{Ideal}^{\text{GSim}}$ denote Ideal with GSim_5 . Five trials make the failure contribution small enough for the $4n/5$ security bound; Appendix L.4 gives the calculation.

Implementation. GSim_5 answers q_p primitive queries using at most $5q_p$ queries to the construction oracle. The implementation with exact sampling procedures has expected total running time $O(q \text{poly}(n))$. For every fixed $A > 0$, an implementation with total variation error at most N^{-A} uses $q \text{poly}_A(n)$ worst-case time and $O(qn)$ bits of memory (Proposition 3).

We use the exponential Bernoulli sampler in [4, Algorithm 2] to accept each legal candidate with probability $e^{-\text{Corr}_\tau(z)}$ without summing $\text{Corr}_\tau(z)$ explicitly.

Limiting the sampling iterations gives the stated error and worst-case time. Appendix F gives the proof and details of implementing the sampling procedures.

5.3 Conditional response distributions

We use the notation of Section 3.2. We first fix Γ , then average over its conditional distribution.

Lemma 4 (Inverse-response distribution). *Fix Γ and the simulator state before a fresh inverse call with $\text{Failure} = 0$. For $x \in U$, put*

$$p_{\text{acc}}(x) := \mathbf{1}[\Gamma(x) \notin \text{Forb}]e^{-\text{Corr}_\tau(\Gamma(x))}, \quad p_{\text{rej}} := 1 - \frac{1}{L} \sum_{x \in U} p_{\text{acc}}(x),$$

with H_5 as defined in Section 5.1. Then

$$\Pr[\text{the call returns } x \mid \Gamma] = \frac{p_{\text{acc}}(x)}{L} H_5(p_{\text{rej}}), \quad \Pr[\text{Failure} = 1 \mid \Gamma] = p_{\text{rej}}^5.$$

Proof. The simulator's five sampling trials are independent, and each accepts a specified x with probability $p_{\text{acc}}(x)/L$. The identities follow by summing the geometric series and by requiring all five trials to reject. $\square$

For GSim_5 , we reuse the notation W_τ, π_τ from Lemma 2. Fix an augmented transcript prefix τ containing no failure response. As in Section 3.4, $\Phi_\tau : \mathcal{G} \rightarrow \mathcal{G}$ agrees with the construction values fixed by τ . Under μ_τ , the remaining values are independent and uniform on $\mathcal{G}$. For each earlier fresh inverse query j that succeeded, put

$$p_{\text{rej},j}(\Phi_\tau) = 1 - \frac{1}{L_j} \sum_{x \in U_j} \mathbf{1}[\Phi_\tau(x) \notin \text{Forb}_j] e^{-\text{Corr}_j(\Phi_\tau(x))},$$

where $U_j, \text{Forb}_j, L_j, \text{Corr}_j$ are taken immediately before query j . In $\text{Ideal}^{\text{GSim}}$, the table of Γ conditioned on τ has distribution

$$W_\tau(\Phi_\tau) := \prod_j H_5(p_{\text{rej},j}(\Phi_\tau)), \quad \pi_\tau(\Phi_\tau) := \frac{W_\tau(\Phi_\tau)}{\mathbb{E}_{\mu_\tau}[W_\tau]} \mu_\tau(\Phi_\tau).$$

The factors $H_5(p_{\text{rej},j})$ are the parts of the observed inverse-response likelihoods that depend on the unrevealed values. Lemma 25 proves the formula by Bayes' rule, integrating out rejected candidates.

For a construction or forward query at an input x with an unrevealed construction value, $\Gamma(x) = z$ with probability $\pi_\tau[\Phi_\tau(x) = z]$. Given z , the forward procedure samples unused permutation outputs y_0, y_1 uniformly subject to $y_0 \oplus y_1 = z$. For a fresh inverse query, conditional on this response being successful, the inverse output and its construction value have joint distribution

$$P_\tau^+(x, z) = \frac{\mathbb{E}_{\Gamma \leftarrow \pi_\tau}[\mathbf{1}[\Gamma(x) = z] p_{\text{acc}}(x) H_5(p_{\text{rej}})]}{L \mathbb{E}_{\Gamma \leftarrow \pi_\tau}[1 - p_{\text{rej}}^5]}.$$

Here $p_{\text{acc}}(x), p_{\text{rej}}$ are from Lemma 4. For the query $\text{Inv}_b(y)$, the permutation outputs at x are $y_b = y$ and $y_{1-b} = y \oplus z$. We write $R_{\text{inv}}^{\text{Ideal}^{\text{GSim}},+}$ for the resulting distribution of successful responses.

5.4 The hybrid Hyb^{GSim}

We define a hybrid world Hyb^{GSim} that retains the leading bias in unrevealed construction values and applies the same acceptance correction as the simulator.

For construction queries and forward queries at inputs with unrevealed construction values, the hybrid uses the leading factor derived in Section 5.1 and samples the value from

$$\nu_{\text{Corr}}(z) := \frac{e^{\text{Corr}(z)}}{Z_{\text{Corr}}}, \quad Z_{\text{Corr}} := \sum_{w \in \mathcal{G}} e^{\text{Corr}(w)}.$$

For an inverse output whose construction value is unrevealed, weighting this distribution by the acceptance probability gives

$$\nu_{\text{Corr}}(z) \mathbf{1}[z \notin \text{Forb}] e^{-\text{Corr}(z)} = \frac{\mathbf{1}[z \notin \text{Forb}]}{Z_{\text{Corr}}}.$$

Thus, conditioned on acceptance, the construction value is uniform on $\mathcal{G} \setminus \text{Forb}$. The mean acceptance probability is $\alpha_{\text{Corr}} := (N - |\text{Forb}|)/Z_{\text{Corr}}$.

Algorithm 5 gives the full query procedures for Hyb^{GSim} . The hybrid maintains the state of Hyb in Section 3.6 and Hist_τ . It records each primitive response (x, y_0, y_1) as in that section and computes Corr_τ from Hist_τ as for GSim_5 . For an inverse query, let $U = \mathcal{G} \setminus \text{Dom}_\tau^{\text{prim}}$ and $\text{Forb} = y \oplus \text{Ran}_{1-b}$. The weights for selecting its inverse output are

$$w_x = \begin{cases} \mathbf{1}[z_i \notin \text{Forb}] e^{-\text{Corr}(z_i)}, & x = x_i \in \text{dom}(\text{Cons}), \\ \alpha_{\text{Corr}}, & x \in U \setminus \text{dom}(\text{Cons}). \end{cases}$$

Write $R_{\text{inv}}^{\text{Hyb}^{\text{GSim}}}$ for its augmented inverse-response distribution. The hybrid uses Cons , which is unavailable to the simulator. For $q \leq N/64$, an input with unrevealed construction value remains and has positive weight, so the rules are well defined.

Every prefix in Hyb^{GSim} is attainable in $\text{Ideal}^{\text{GSim}}$ by the argument of Section 3.6: each inverse query can succeed on its first trial, since the extra acceptance probability $e^{-\text{Corr}_\tau(z)}$ is positive.

6 Security of the Gyroscope Simulator

We compare Real with Hyb^{GSim} , and then Hyb^{GSim} with $\text{Ideal}^{\text{GSim}}$, to obtain the following security bound.

Algorithm 5 The hybrid Hyb^{GSim} : fresh queries.

1: function $\text{Con}(x)$ 2: $z \leftarrow \nu_{\text{Corr}}$ 3: Append (x, z) to Cons 4: return z 5: function $\text{Fwd}_b(x)$ 6: if $x \in \text{dom}(\text{Cons})$ then 7: Read (x, z) from Cons 8: else 9: $z \leftarrow \nu_{\text{Corr}}$ 10: Sample $y_0 \notin \text{Ran}_0, y_1 \notin \text{Ran}_1$ uniformly with $y_0 \oplus y_1 = z$ 11: Record (x, y_0, y_1) 12: return (x, y_0, y_1)	1: function $\text{Inv}_b(y)$ 2: Sample $x \in U$ with probability $w_x / \sum_{u \in U} w_u$ 3: if $x \in \text{dom}(\text{Cons})$ then 4: Read (x, z) from Cons 5: else 6: $z \leftarrow \text{Unif}(\mathcal{G} \setminus \text{Forb})$ 7: $y_b \leftarrow y; y_{1-b} \leftarrow y \oplus z$ 8: Append $(b, y, U , \text{Dom}_\tau^{\text{prim}})$ to Hist_τ 9: Record (x, y_0, y_1) 10: return (x, y_0, y_1)
---	--

Theorem 3. Let $N = 2^n \geq 2^{20}$ and $q \leq N/64$. For the regular game with at most q direct construction and primitive queries,

$$\text{Adv}_{\text{SoP,GSim}_5}^{\text{indiff}}(q) \leq 11n^{3/2} \frac{q^{3/2}}{N^{3/2}} + 11\sqrt{n} \frac{q^{5/2}}{N^2} + \varepsilon_{\text{impl}}. \quad (18)$$

Here $\varepsilon_{\text{impl}} = 0$ for the exact sampler in Appendix F. For every fixed $\kappa \geq 20$, Proposition 3 gives an implementation with $\varepsilon_{\text{impl}} \leq N^{-\kappa}$ and worst-case time $q \text{poly}_\kappa(n)$.

The full proof is in Appendix G. We explain the two comparisons below. The comparison with Real uses a second-moment bound over its transcript distribution; the comparison with $\text{Ideal}^{\text{GSim}}$ bounds the divergence at each prefix.

6.1 Comparison with Real

Construction responses in Real. Fix a prefix τ after $r < q$ queries in Real , and let p_τ^{Real} denote the response distribution for a fresh construction query. For the calculation, we reveal permutation outputs as primitive queries occur. After query r , we also reveal the outputs at inputs queried only to the construction, conditioned on the transcript. Let X denote this assignment. Together with the recorded outputs, it fixes both permutations at $S \leq r$ inputs, with ranges $\widetilde{\text{Ran}}_{0,S}, \widetilde{\text{Ran}}_{1,S}$. The response distribution p_X for a fresh construction query, conditioned on this assignment, satisfies

$$c_S(z) = |\widetilde{\text{Ran}}_{0,S} \cap (\widetilde{\text{Ran}}_{1,S} \oplus z)|, \quad p_X(z) = \frac{N - 2S + c_S(z)}{(N - S)^2}.$$

Here $c_S(z)$ is the number of pairs of occupied outputs whose XOR is z . Each occupied range excludes S output pairs with XOR z , and $c_S(z)$ of them are excluded by both ranges.

An inverse query j fixes one output to y_j . Its XORs with the previously occupied outputs on the other side contribute exactly $\mathbf{1}[z \in \text{Forb}_j]$ to the count $c_S(z)$. Its coefficient in Np_X is $N/(N-S)^2$, whereas the correction derived in Section 5.1 uses $1/L_j$. Both have leading term $1/N$ when r/N is small. We bound their difference together with the contributions from the sampled outputs.

Averaging over compatible assignments gives $p_\tau^{\text{Real}} = \mathbb{E}[p_X \mid \tau]$. To compare this distribution with ν_{Corr_τ} , we use the centered log-likelihood ratio

$$\ell_\tau^{\text{res}} := (\log(Np_\tau^{\text{Real}}) - \text{Corr}_\tau)^\circ, \quad \varepsilon_\tau^{\text{res}} := \frac{\|\ell_\tau^{\text{res}}\|_2}{\sqrt{N}}.$$

Centering removes the normalization constant of ν_{Corr_τ} . Appendix I subtracts the leading correction from $Np_X - 1$ and applies Jensen's inequality to average over compatible assignments. It then bounds the error from taking logarithms. This gives

$$\begin{aligned} \sqrt{D_{\text{KL}}(p_\tau^{\text{Real}} \parallel \nu_{\text{Corr}_\tau})} &\leq 2\varepsilon_\tau^{\text{res}}, \\ \sqrt{\mathbb{E}_{\tau \leftarrow \text{Tr}_{\text{Real}}}[(\varepsilon_\tau^{\text{res}})^2]} &\leq 3\frac{r+1}{N^{3/2}} + 2\frac{(r+1)^2}{N^2}. \end{aligned}$$

The first bound holds at each prefix; the second averages over prefixes that arise in **Real**. The size bound $r+1$ also covers a forward query that reveals a construction value before its permutation outputs are compared.

Primitive responses and transcript distributions. For primitive queries, we also bound the probabilities that the required permutation outputs remain unused, conditional on the revealed construction values. Appendices J and K relate these probabilities to the same construction-response distribution.

The primitive-response bounds also depend on repeated construction values and pairs with a given XOR. List these values as $Z_1, \dots, Z_s$, in order of first revelation. Even if several construction or primitive queries involve the same input x , record its construction value only once: for example, $\text{Con}(x)$ followed by $\text{Fwd}_0(x)$ contributes one entry. Distinct inputs contribute separate entries even when their values coincide. We call τ *good* if, for every $1 \leq k \leq s$,

$$\begin{aligned} \max_{z \in \mathcal{G}} |\{i \leq k : Z_i = z\}| &\leq 2n, \\ \max_{a \in \mathcal{G}} |\{(i, j) : i \neq j, i, j \leq k, Z_i \oplus Z_j = a\}| &\leq 3k^2/N + 42n^2. \end{aligned} \tag{19}$$

The ordered pairs in the second line have distinct indices. The first bound controls how many inputs share a construction value; the second controls how many pairs of constraints involve a given XOR. At a good prefix, each construction value appears at most $2n$ times in **Cons**, so

$$\frac{1}{N} \sum_{(x, z) \in \text{Cons}} \ell_\tau^{\text{res}}(z)^2 \leq \frac{2n}{N} \sum_{z \in \mathcal{G}} \ell_\tau^{\text{res}}(z)^2 = 2n(\varepsilon_\tau^{\text{res}})^2.$$

This prevents one value with a large discrepancy from being counted too many times in the comparison of inverse outputs. The XOR bound plays the corresponding role when comparing the availability of permutation outputs. The

normalization and availability calculations are given in Appendices J and K. Lemma 13 shows that all prefixes in Real are good with probability at least $1 - N^{-8}$.

Lemma 22 bounds the conditional KL divergence in the excluded cases by one, so their total contribution to the transcript KL divergence is at most qN^{-8} .

Combining the query types and summing their conditional divergences by the KL chain rule gives

$$\sqrt{D_{\text{KL}}(T_{\text{Real}} \| T_{\text{Hyb}}^{\text{GSim}})} \leq 14n^{3/2} \frac{q^{3/2}}{N^{3/2}} + 13\sqrt{n} \frac{q^{5/2}}{N^2} + \sqrt{q} N^{-4}.$$

Lemma 24 and its proof in Appendix K.6 give the full calculation.

6.2 Comparison with Ideal^{GSim}

Here we bound the divergence at each prefix. The conditional table distribution π_τ from Section 5.3 includes the additional rejections introduced by the Gyroscope simulator.

Construction and forward responses. Fix an input x with an unrevealed construction value, and let w denote the other construction values. In an earlier inverse query j , the contribution of x to the rejection probability is $1/L_j$ when $z \in \text{Forb}_j$, and $(1 - e^{-\text{Corr}_j(z)})/L_j$ otherwise. The expansion of $\log H_5$ therefore retains the leading term $L_j^{-1} \mathbf{1}[z \in \text{Forb}_j]$. Summing these terms gives $\text{Corr}_\tau(z)$.

Appendix L.2 bounds the remaining variation by

$$\mathbb{E}_w \left[\text{osc}_z (\log \pi_\tau[\Gamma(x) = z \mid w] - \text{Corr}_\tau(z))^2 \right] \leq \frac{81q^4}{N^4},$$

where w follows its marginal distribution under π_τ at the fixed prefix. The reweighting lemma and averaging over w give

$$D_{\text{KL}}(\nu_{\text{Corr}_\tau} \| \mathcal{L}_{\text{Ideal}^{\text{GSim}}}(\Gamma(x) \mid \tau)) \leq \frac{11q^4}{N^4}.$$

This also bounds the divergence for a forward query at such an input: conditioned on the construction value, the permutation outputs have the same distribution in both worlds. At an input with a revealed construction value, the forward-response distributions coincide.

Inverse responses conditioned on success. We first compare the inverse output, then its construction value. The possible rejected trials before a return contribute the factor $H_5(p_{\text{rej}})$. For an input x with an unrevealed construction value, its mean acceptance probability after this reweighting is

$$\theta := \frac{\mathbb{E}_{\Gamma \leftarrow \pi_\tau} [H_5(p_{\text{rej}}) p_{\text{acc}}(x)]}{\mathbb{E}_{\Gamma \leftarrow \pi_\tau} [H_5(p_{\text{rej}})]}.$$

Symmetry makes θ independent of the choice of x . For selecting the inverse output, $\text{Ideal}^{\text{GSim}}$, conditioned on success, assigns weight θ to these inputs; Hyb^{GSim} assigns weight α_{Corr} . Both use the same weights at available inputs with revealed construction values. Appendix L.3 proves

$$|\theta - \alpha_{\text{Corr}}| \leq \frac{4q^2}{N^2}$$

and bounds the divergence after normalizing these weights.

Conditional on selecting an input with an unrevealed construction value, acceptance cancels the leading factor $e^{\text{Corr}_\tau(z)}$ on the legal set. The leading contribution from rejected trials in the current call is supported on Forb , so it vanishes on accepted values. The appendix bounds the remaining departure from uniformity, including the effect of conditioning on the inverse output. If the selected input already had a revealed construction value, the query and that value determine the companion output and add no divergence.

Failure and the transcript comparison. For every query type, the divergence from the response distribution in Hyb^{GSim} to the distribution of successful responses in $\text{Ideal}^{\text{GSim}}$ is at most $11q^4/N^4$. An inverse call fails only when all five trials reject. Candidates with construction values revealed in τ or read earlier in the current call occupy only a small part of the available set. For the other candidates, Lemma 26 bounds the probability of a value in Forb , and the bound on Corr_τ controls the extra rejection.

Conditional on the earlier rejected trials and the construction values read in those trials, the rejection probability of the next trial is at most $4q/N$. This bound includes repeated candidates. Multiplying the five conditional bounds gives $p_{\text{fail},\tau} \leq (4q/N)^5$.

At a fixed prefix τ , let P_τ be the response distribution in Hyb^{GSim} and Q_τ that in $\text{Ideal}^{\text{GSim}}$, including failure. Let Q_τ^+ denote Q_τ conditioned on success. Equation (5) gives

$$\begin{aligned} D_{\text{KL}}(P_\tau \| Q_\tau) &= D_{\text{KL}}(P_\tau \| Q_\tau^+) + \log \frac{1}{1 - p_{\text{fail},\tau}} \\ &\leq \frac{11q^4}{N^4} + \log \frac{1}{1 - (4q/N)^5}. \end{aligned}$$

The KL chain rule sums these bounds over at most q queries, giving

$$D_{\text{KL}}(T_{\text{Hyb}}^{\text{GSim}} \| T_{\text{Ideal}}^{\text{GSim}}) \leq \frac{11q^5}{N^4} + q \log \frac{1}{1 - (4q/N)^5}.$$

The full calculation is in Appendix L.4, with Lemma 30.

AI use disclosure. Most results of this paper are discovered with the assistance of ChatGPT 5.6 and ChatGPT 6.0. The manuscript is written by the author based on the initial draft generated by AI. The author verified the technical content and takes full responsibility for the manuscript.

References

1. Bellare, M., Krovetz, T., Rogaway, P.: Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible. In: Nyberg, K. (ed.) *Advances in Cryptology – EUROCRYPT ’98*. Lecture Notes in Computer Science, vol. 1403, pp. 266–280. Springer, Berlin, Heidelberg (1998). <https://doi.org/10.1007/BFb0054132>
2. Bhattacharya, S., Nandi, M.: Full indifferentiable security of the xor of two or more random permutations using the χ^2 method. In: Nielsen, J.B., Rijmen, V. (eds.) *Advances in Cryptology – EUROCRYPT 2018, Part I*. Lecture Notes in Computer Science, vol. 10820, pp. 387–412. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_15
3. Bhaumik, R., Datta, N., Dutta, A., Jha, A., Mandal, S., Mennink, B., Nandi, H., Shen, Y.: How to build a short-input random oracle from public random permutations. In: Daemen, J., Thomé, E. (eds.) *Advances in Cryptology – EUROCRYPT 2026, Part VI*. Lecture Notes in Computer Science, vol. 16546, pp. 394–423. Springer, Cham (2026). https://doi.org/10.1007/978-3-032-25333-0_14, <https://eprint.iacr.org/2026/336>, full version: Cryptology ePrint Archive, Report 2026/336
4. Canonne, C.L., Kamath, G., Steinke, T.: The discrete Gaussian for differential privacy. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 15676–15688. Curran Associates, Inc. (2020), https://proceedings.neurips.cc/paper_files/paper/2020/file/b53b3a3d6ab90ce0268229151c9bde11-Paper.pdf
5. Coron, J.S., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård revisited: How to construct a hash function. In: Shoup, V. (ed.) *Advances in Cryptology – CRYPTO 2005*. Lecture Notes in Computer Science, vol. 3621, pp. 430–448. Springer, Berlin, Heidelberg (2005). https://doi.org/10.1007/11535218_26
6. Dai, W., Hoang, V.T., Tessaro, S.: Information-theoretic indistinguishability via the chi-squared method. In: Katz, J., Shacham, H. (eds.) *Advances in Cryptology – CRYPTO 2017, Part III*. Lecture Notes in Computer Science, vol. 10403, pp. 497–523. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_17
7. Dinur, I.: Tight indistinguishability bounds for the XOR of independent random permutations by Fourier analysis. In: Joye, M., Leander, G. (eds.) *Advances in Cryptology – EUROCRYPT 2024, Part I*. Lecture Notes in Computer Science, vol. 14651, pp. 33–62. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-58716-0_2, <https://eprint.iacr.org/2024/338>, full version: Cryptology ePrint Archive, Report 2024/338
8. Dutta, A., Nandi, M., Saha, A.: Proof of Mirror Theory for $\xi_{\max} = 2$. *IEEE Transactions on Information Theory* **68**(9), 6218–6232 (2022). <https://doi.org/10.1109/TIT.2022.3171178>
9. Eberhard, S.: More on additive triples of bijections. *arXiv:1704.02407* (2017), <https://arxiv.org/abs/1704.02407>
10. Günsing, A., Bhaumik, R., Jha, A., Mennink, B., Shen, Y.: Revisiting the indifferentiability of the sum of permutations. In: Handschuh, H., Lysyanskaya, A. (eds.) *Advances in Cryptology – CRYPTO 2023, Part III*. Lecture Notes in Computer Science, vol. 14083, pp. 628–660. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38548-3_21
11. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* **58**(301), 13–30 (1963). <https://doi.org/10.1080/01621459.1963.10500830>

12. Mandal, A., Patarin, J., Nachev, V.: Indifferentiability beyond the birthday bound for the xor of two public random permutations. In: Gong, G., Gupta, K.C. (eds.) Progress in Cryptology – INDOCRYPT 2010. Lecture Notes in Computer Science, vol. 6498, pp. 69–81. Springer, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17401-8_6
13. Maurer, U.M., Renner, R., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In: Naor, M. (ed.) Theory of Cryptography – TCC 2004. Lecture Notes in Computer Science, vol. 2951, pp. 21–39. Springer, Berlin, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24638-1_2
14. Mennink, B., Preneel, B.: On the XOR of multiple random permutations. In: Malkin, T., Kolesnikov, V., Lewko, A.B., Polychronakis, M. (eds.) Applied Cryptography and Network Security – ACNS 2015. Lecture Notes in Computer Science, vol. 9092, pp. 619–634. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-28166-7_30
15. Patarin, J.: A proof of security in $O(2^n)$ for the xor of two random permutations. In: Safavi-Naini, R. (ed.) Information Theoretic Security – ICITS 2008. Lecture Notes in Computer Science, vol. 5155, pp. 232–248. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85093-9_22

Supplementary Material

Appendix A proves the reweighting lemma. Appendices B–D analyze uniform simulation, and Appendix E proves the attack. Appendix F implements the Gyroscope simulator. Appendix G proves Theorem 3 from the two transcript comparisons. Appendix H bounds multiplicities and XOR counts; Appendices I–J analyze construction and primitive responses in Real. Appendix K combines these bounds in the comparison of Real and Hyb^{GSim} , and Appendix L compares Hyb^{GSim} with $\text{Ideal}^{\text{GSim}}$, including simulator failure.

A Proof of the Reweighting Lemma

We prove Lemma 1 with $g = \log w$ and

$$P(x) = \frac{Q(x)e^{g(x)}}{\mathbb{E}_Q[e^g]}, \quad d = \text{osc}_{\text{supp}(Q)}(g).$$

Proof (Lemma 1). For any constant c , replacing g by $g - c$ multiplies all weights by e^{-c} and leaves P , d , and $\text{Var}_Q(g)$ unchanged.

Chi-squared divergence. Subtract $\min_{x \in \text{supp}(Q)} g(x)$ from g , so $0 \leq g \leq d$ on $\text{supp}(Q)$. Since $|e^a - e^b| \leq e^d|a - b|$ for $a, b \in [0, d]$, the identity $\text{Var}_Q(h) = \frac{1}{2} \sum_{x, y} Q(x)Q(y)(h(x) - h(y))^2$ gives

$$\chi^2(P\|Q) = \frac{\text{Var}_Q(e^g)}{(\mathbb{E}_Q[e^g])^2} \leq e^{2d} \text{Var}_Q(g).$$

KL divergence in terms of $\text{Var}_Q(g)$. Subtract $\mathbb{E}_Q[g]$ from g , so $\mathbb{E}_Q[g] = 0$ and $|g| \leq d$ on $\text{supp}(Q)$, and put $Z = \mathbb{E}_Q[e^g] \geq 1$. The exponential power series gives, for $|u| \leq d$,

$$ue^u - e^u + 1 = \sum_{k \geq 2} \frac{(k-1)u^k}{k!} \leq \frac{1}{2}e^d u^2.$$

Using $\log Z \geq 1 - 1/Z$, we obtain

$$D_{\text{KL}}(P\|Q) \leq \frac{1}{Z} \sum_x Q(x)(g(x)e^{g(x)} - e^{g(x)} + 1) \leq \frac{1}{2}e^d \text{Var}_Q(g).$$

KL divergence in terms of d . Hoeffding's bounded-variable inequality [11, (4.16)], applied to $-g(X)$ for $X \leftarrow P$, gives

$$D_{\text{KL}}(P\|Q) = \log \sum_x P(x) \exp(-g(x) + \mathbb{E}_P[g]) \leq d^2/8.$$

Weights $1/w$ give $D_{\text{KL}}(Q\|P) \leq d^2/8$ by the same argument.

Total variation distance. Pinsker's inequality gives $\text{TV}(P, Q) \leq d/4$. $\square$

B Uniform Simulation: Construction and Forward Responses in Real

We use $\text{Ran}_0, \text{Ran}_1, \text{Cons}$ from Section 3.2, with $s = s_{\text{prim}} + s_{\text{con}} \leq N/4$. We bound the divergences between response distributions for construction and forward queries in Proposition 2.

B.1 Construction query at a fresh input

Fixing the unexposed pairs expresses the response bias through the overlap of occupied ranges. Averaging gives the following bound.

Lemma 5 (Construction query at a fresh input). *At an input x with no previous construction or primitive query, the response $Z = \Pi_0(x) \oplus \Pi_1(x)$ in Real satisfies*

$$\chi^2(\mathcal{L}(Z \mid \text{Ran}_0, \text{Ran}_1, \text{Cons}) \parallel \text{Unif}(\mathcal{G})) \leq \frac{Ns^3}{(N-s)^4} \leq 4\frac{s^3}{N^3}.$$

Proof (Lemma 5). Let $X = (\Pi_0(x_1), \dots, \Pi_0(x_{s_{\text{con}}}))$ and condition on $X = u \in \text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})$. Put

$$\text{Ran}_{0,u} := \text{Ran}_0 \cup \{u_1, \dots, u_{s_{\text{con}}}\}, \quad \text{Ran}_{1,u} := \text{Ran}_1 \cup \{u_1 \oplus z_1, \dots, u_{s_{\text{con}}} \oplus z_{s_{\text{con}}}\}.$$

Then $|\text{Ran}_{0,u}| = |\text{Ran}_{1,u}| = s$, and at a fresh input x ,

$$\mathcal{L}(\Pi_0(x), \Pi_1(x) \mid \text{Ran}_0, \text{Ran}_1, \text{Cons}, X = u) = \text{Unif}(\mathcal{G} \setminus \text{Ran}_{0,u}) \otimes \text{Unif}(\mathcal{G} \setminus \text{Ran}_{1,u}).$$

For $z \in \mathcal{G}$, define

$$\begin{aligned} c_u(z) &:= |\text{Ran}_{0,u} \cap (\text{Ran}_{1,u} \oplus z)|, \\ p_u(z) &:= \Pr[\Pi_0(x) \oplus \Pi_1(x) = z \mid \text{Ran}_0, \text{Ran}_1, \text{Cons}, X = u]. \end{aligned}$$

Thus

$$\begin{aligned} p_u(z) &= \frac{|\mathcal{G} \setminus (\text{Ran}_{0,u} \cup (\text{Ran}_{1,u} \oplus z))|}{(N-s)^2} = \frac{N-2s+c_u(z)}{(N-s)^2}, \\ p_u(z) - \frac{1}{N} &= \frac{c_u(z) - s^2/N}{(N-s)^2}. \end{aligned}$$

For $W \leftarrow \text{Unif}(\mathcal{G})$,

$$\sum_z c_u(z) = s^2, \quad 0 \leq c_u(z) \leq s, \quad \sum_z c_u(z)^2 \leq s^3,$$

and hence

$$\mathbb{E}[c_u(W)] = \frac{s^2}{N}, \quad \text{Var}[c_u(W)] \leq \frac{s^3}{N}.$$

Identity (1) yields

$$\chi^2(p_u \parallel \text{Unif}(\mathcal{G})) = \frac{N^2}{(N-s)^4} \text{Var}[c_u(W)] \leq \frac{Ns^3}{(N-s)^4}.$$

By convexity, averaging over X given $(\text{Ran}_0, \text{Ran}_1, \text{Cons})$ preserves the bound. Since $N-s \geq 3N/4$, we have $Ns^3/(N-s)^4 \leq 4s^3/N^3$. $\square$

B.2 Forward query at a construction-only input

The forward-response distribution is obtained by reweighting the uniform distribution on legal outputs, according to the probability that the corresponding pair avoids all unexposed outputs. We bound the variance of these weights.

Suppose $(x, z_0) \in \text{Cons}$ and put

$$\text{Cons}_{-x} := \text{Cons} \setminus \{(x, z_0)\}, \quad s_{\text{con}}^{(-x)} := |\text{Cons}_{-x}|, \quad s_{-x} := s_{\text{prim}} + s_{\text{con}}^{(-x)},$$

$$E_0 := \{a \in \mathcal{G} : a \notin \text{Ran}_0, a \oplus z_0 \notin \text{Ran}_1\}.$$

We use the same definition of **Comp** with the indicated occupied ranges and list, replacing s_{con} by the length of that list. For $a \in E_0$, let

$$w_a := |\text{Comp}(\text{Ran}_0 \cup \{a\}, \text{Ran}_1 \cup \{a \oplus z_0\}, \text{Cons}_{-x})|.$$

The conditional distribution of $\Pi_0(x)$ assigns mass $w_a / \sum_{c \in E_0} w_c$ to a , while the corresponding distribution in **Hyb** is uniform on E_0 .

Lemma 6 (Forward query after $F(x) = z_0$). *For every compatible state with $\sum_{c \in E_0} w_c > 0$ and $s_{-x} \leq N/4$, including states with repeated construction values in Cons_{-x} ,*

$$\chi^2(\mathcal{L}(\Pi_0(x) \mid \text{Ran}_0, \text{Ran}_1, \text{Cons}) \parallel \text{Unif}(E_0)) \leq \frac{40(s_{\text{con}}^{(-x)})^2 s_{-x}}{(N - 2s_{-x})^3} \leq 320 \frac{s_{-x}^3}{N^3}.$$

Proof (Lemma 6).

Reduction to a variance bound. Write

$$\text{Cons}_{-x} = ((u_i, z_i))_{i=1}^{s_{\text{con}}^{(-x)}}, \quad X \leftarrow \text{Unif}(\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons}_{-x})),$$

let $d_i = z_i \oplus z_0$, and put

$$\text{Occ}(X) := \{X_1, \dots, X_{s_{\text{con}}^{(-x)}}\} \cup \{X_1 \oplus d_1, \dots, X_{s_{\text{con}}^{(-x)}} \oplus d_{s_{\text{con}}^{(-x)}}\}.$$

For $a \in E_0$, compatibility is equivalent to $a \notin \text{Occ}(X)$. For all $a \in \mathcal{G}$, put

$$p_a := \Pr[a \in \text{Occ}(X)], \quad \pi_a := 1 - p_a, \quad \bar{p} := \frac{1}{|E_0|} \sum_{a \in E_0} p_a, \quad \bar{\pi} := 1 - \bar{p}.$$

Let $R_{\text{Cons}_{-x}}^{\text{fwd}}$ be the conditional distribution of $\Pi_0(x)$ and let $U_{\text{Cons}_{-x}}^{\text{fwd}} = \text{Unif}(E_0)$. On E_0 ,

$$\frac{R_{\text{Cons}_{-x}}^{\text{fwd}}(a)}{U_{\text{Cons}_{-x}}^{\text{fwd}}(a)} = \frac{\pi_a}{\bar{\pi}}, \quad \chi^2(R_{\text{Cons}_{-x}}^{\text{fwd}} \parallel U_{\text{Cons}_{-x}}^{\text{fwd}}) = \frac{\text{Var}_{W_0 \leftarrow \text{Unif}(E_0)}[p_{W_0}]}{\bar{\pi}^2}.$$

Bounding the variance of each term. Fix i and put $X_{-i} = (X_1, \dots, X_{i-1}, X_{i+1}, \dots)$. Conditional on X_{-i} , the legal set for X_i is

$$D_i = \mathcal{G} \setminus (\text{Ran}_0 \cup (\text{Ran}_1 \oplus z_i) \cup \{X_j : j \neq i\} \cup \{X_j \oplus z_j \oplus z_i : j \neq i\}).$$

Every element of D_i extends the other coordinates in exactly one way, so

$$\mathcal{L}(X_i \mid X_{-i}) = \text{Unif}(D_i), \quad |D_i| \geq N - 2s_{-x}.$$

Let $W \leftarrow \text{Unif}(\mathcal{G})$, independently of X , and define $p_i^{(0)}(a) = \Pr[X_i = a]$. For every nonempty $D \subseteq \mathcal{G}$,

$$\mathbb{E}_W[\text{Unif}(D)(W)] = \frac{1}{N}, \quad \text{Var}_W[\text{Unif}(D)(W)] = \frac{N - |D|}{N^2 |D|}.$$

Since the first expectation is $1/N$, Jensen's inequality gives

$$\text{Var}_W[p_i^{(0)}(W)] \leq \frac{2s_{-x}}{N^2(N - 2s_{-x})} \leq \frac{2s_{-x}}{N(N - 2s_{-x})^2}. \quad (20)$$

The values X_i are distinct, as are the values $X_i \oplus d_i$. Counting their union without duplication gives

$$\mathbf{1}[a \in \text{Occ}(X)] = \sum_i \mathbf{1}[X_i = a] + \sum_i \mathbf{1}[X_i \oplus d_i = a, X_i \oplus d_i \notin \{X_1, \dots, X_{s_{\text{con}}^{(-x)}}\}].$$

Let $p_i^{(1)}(a)$ be the probability of the second event. Conditional on X_{-i} , let $\ell_i = |D_i|$. If $d_i = 0$, put $E_i = \emptyset$; otherwise put $E_i = (D_i \oplus d_i) \setminus \{X_j : j \neq i\}$. Define

$$e_i := |E_i|, \quad \varphi_i(a) := \frac{\mathbf{1}[a \in E_i]}{\ell_i}.$$

Then

$$p_i^{(1)}(a) = \mathbb{E}_{X_{-i}}[\varphi_i(a)], \quad \text{Var}_W[\varphi_i(W)] = \frac{e_i(N - e_i)}{\ell_i^2 N^2}.$$

When $d_i = 0$, this variance is zero. Otherwise, $N - e_i \leq 2s_{-x} + (s_{\text{con}}^{(-x)} - 1) \leq 3s_{-x}$. Applying Jensen's inequality after subtracting each function's uniform mean gives

$$\text{Var}_W[p_i^{(1)}(W)] \leq \mathbb{E}_{X_{-i}}[\text{Var}_W[\varphi_i(W)]] \leq \frac{3s_{-x}}{N(N - 2s_{-x})^2}. \quad (21)$$

Combining the variance bounds. Thus

$$p_a = \sum_{i=1}^{s_{\text{con}}^{(-x)}} p_i^{(0)}(a) + \sum_{i=1}^{s_{\text{con}}^{(-x)}} p_i^{(1)}(a).$$

Put $\bar{p}_G = \mathbb{E}_{W \leftarrow \text{Unif}(G)}[p_W]$. Applying Cauchy–Schwarz to the centered versions of these $2s_{\text{con}}^{(-x)}$ summands and using (20)–(21) gives

$$\begin{aligned} \sum_{a \in G} (p_a - \bar{p}_G)^2 &= N \text{Var}_{W \leftarrow \text{Unif}(G)}[p_W] \\ &\leq 2N s_{\text{con}}^{(-x)} \sum_i (\text{Var}_W[p_i^{(0)}(W)] + \text{Var}_W[p_i^{(1)}(W)]) \\ &\leq \frac{10(s_{\text{con}}^{(-x)})^2 s_{-x}}{(N - 2s_{-x})^2}. \end{aligned}$$

Since $|E_0| \geq N - 2s_{-x}$ and the average over E_0 minimizes the squared deviation on E_0 ,

$$\text{Var}_{W_0 \leftarrow \text{Unif}(E_0)}[p_{W_0}] \leq \frac{10(s_{\text{con}}^{(-x)})^2 s_{-x}}{(N - 2s_{-x})^3}.$$

Since $|\text{Occ}(X)| \leq 2s_{\text{con}}^{(-x)}$, we have $1 - \bar{\pi} \leq 1/2$. These bounds give

$$\begin{aligned} \chi^2(\mathcal{L}(\Pi_0(x) \mid \text{Ran}_0, \text{Ran}_1, \text{Cons}) \parallel \text{Unif}(E_0)) \\ \leq \frac{40(s_{\text{con}}^{(-x)})^2 s_{-x}}{(N - 2s_{-x})^3} \leq 320 \frac{s_{-x}^3}{N^3}. \end{aligned}$$

□

C Uniform Simulation: Inverse Responses in Real

C.1 Distribution of the inverse output

We separate the inverse output from its companion output. For the inverse output, we bound the contribution of construction-only inputs and then that of inputs with unrevealed construction values.

By symmetry, fix a fresh query $\Pi_0^{-1}(y)$ with $y \notin \text{Ran}_0$, and let $X = \Pi_0^{-1}(y)$. Use $u, k, \rho, \mathcal{N}_{\text{Hyb}}, K_{\text{Hyb}}^{\text{out}}$ from Section 3.6 for this query, with $u = N - s$.

For $U \leftarrow \text{Unif}(\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons}))$, put

$$\iota_i := \mathbf{1}[y \oplus z_i \notin \text{Ran}_1], \quad r_i := \Pr[U_i = y].$$

The distribution of the inverse output in Real is

$$R_{\text{Cons}}(x_i) := r_i, \quad R_{\text{Cons}}(x) := \left(1 - \sum_{i=1}^{s_{\text{con}}} r_i\right) \text{Unif}(E_{\text{new}})(x) \quad (x \in E_{\text{new}}).$$

Lemma 7 (Inverse output for a fresh query). *For $s = s_{\text{prim}} + s_{\text{con}} \leq N/4$,*

$$\chi^2(R_{\text{Cons}} \parallel K_{\text{Hyb}}^{\text{out}}) \leq 93 \frac{s^3}{N^3}.$$

Proof (Lemma 7). Conditional on the other coordinates, U_i is uniform on at least $d = N - 2s$ values. If $\iota_i = 1$, then $y \notin \text{Ran}_0$ and $y \oplus z_i \notin \text{Ran}_1$, so y is excluded only by one of $2(s_{\text{con}} - 1)$ possible collisions with the other coordinates. Each collision has probability at most $1/d$. Averaging over those coordinates gives

$$\frac{1 - 2(s_{\text{con}} - 1)/d}{N} \leq r_i \leq \frac{1}{d}, \quad |r_i - 1/N| \leq \frac{2s}{Nd} \leq \frac{4s}{N^2}.$$

If $\iota_i = 0$, then $r_i = 0$.

The normalizing constant in **Hyb** satisfies

$$N - \mathcal{N}_{\text{Hyb}} = 2s_{\text{prim}} + s_{\text{con}} - k - \frac{s_{\text{prim}}(s_{\text{prim}} + s_{\text{con}})}{N}, \quad N - 2s \leq \mathcal{N}_{\text{Hyb}} \leq N.$$

For each legal construction-only input x_i , the probability of returning x_i in **Hyb** satisfies $|K_{\text{Hyb}}^{\text{out}}(x_i) - 1/N| \leq 4s/N^2$. The inputs in $\text{dom}(\text{Cons})$ therefore contribute at most $64s_{\text{con}}s^2/N^3$ to $\chi^2(R_{\text{Cons}} \| K_{\text{Hyb}}^{\text{out}})$.

Let $\Delta = \sum_i (R_{\text{Cons}}(x_i) - K_{\text{Hyb}}^{\text{out}}(x_i))$. Both distributions are uniform within E_{new} , where their total masses differ by $-\Delta$, and $|\Delta| \leq 8ss_{\text{con}}/N^2$. Moreover, $u \geq 3N/4$ and every $x \in E_{\text{new}}$ has $K_{\text{Hyb}}^{\text{out}}(x) \geq 3/(4N)$, so the coordinates in E_{new} contribute at most

$$\frac{\Delta^2}{uK_{\text{Hyb}}^{\text{out}}(x)} \leq 114 \frac{s_{\text{con}}^2 s^2}{N^4} \leq 29 \frac{s^3}{N^3}.$$

Summing the contributions from $\text{dom}(\text{Cons})$ and E_{new} gives $\chi^2(R_{\text{Cons}} \| K_{\text{Hyb}}^{\text{out}}) \leq 93s^3/N^3$. $\square$

C.2 Companion output

For an inverse output whose construction value is unrevealed, the distribution of the companion output depends on which outputs are occupied. Its divergence is controlled by their occupancy probabilities.

For the companion output, condition on $X \in E_{\text{new}}$ and put

$$V_{\text{comp}} := \Pi_1(X), \quad R_{\text{Cons}}^{\text{comp}}(v) := \Pr[V_{\text{comp}} = v \mid X \in E_{\text{new}}] \quad (v \notin \text{Ran}_1).$$

Lemma 8 (Companion output when $X \in E_{\text{new}}$). *Let $U_{\text{Ran}_1} := \text{Unif}(\mathcal{G} \setminus \text{Ran}_1)$. For $s \leq N/4 - 1$,*

$$\chi^2(R_{\text{Cons}}^{\text{comp}} \| U_{\text{Ran}_1}) \leq 8 \frac{s^3}{N^3}.$$

Proof (Lemma 8). Write $t = s_{\text{prim}}$, $h = s_{\text{con}}$, and $M = N - t$. If $h = 0$, the companion output is uniform on $\mathcal{G} \setminus \text{Ran}_1$ and the claim is immediate. Otherwise put $\ell = N - 2s + 1$. Condition on $X \in E_{\text{new}}$, equivalently $U_i \neq y$ for every i . For $v \notin \text{Ran}_1$, let

$$p_v := \Pr[\text{some } U_i \oplus z_i = v \mid X \in E_{\text{new}}].$$

Given U and $X \in E_{\text{new}}$, the companion output is uniform over the $M - h$ unoccupied outputs. Thus

$$R_{\text{Cons}}^{\text{comp}}(v) = \frac{1 - p_v}{M - h}.$$

Conditional on the other coordinates and $X \in E_{\text{new}}$, U_i is uniform on its allowed values. Its forbidden set is contained in

$$\text{Ran}_0 \cup (\text{Ran}_1 \oplus z_i) \cup \{y\} \cup \{U_j : j \neq i\} \cup \{U_j \oplus z_j \oplus z_i : j \neq i\}.$$

At most $2t + 1 + 2(h - 1) = 2s - 1$ values are excluded, leaving at least ℓ choices. Hence $\Pr[U_i \oplus z_i = v \mid X \in E_{\text{new}}] \leq 1/\ell$. Since the h outputs $U_i \oplus z_i$ are distinct, for $W \leftarrow U_{\text{Ran}_1}$ we have

$$0 \leq p_v \leq \frac{h}{\ell}, \quad \mathbb{E}[p_W] = \frac{h}{M}.$$

Using $p_W^2 \leq (h/\ell)p_W$ gives directly

$$\text{Var}[p_W] \leq \frac{h}{\ell} \mathbb{E}[p_W] - (\mathbb{E}[p_W])^2 = \frac{h^2(M - \ell)}{\ell M^2}.$$

Consequently,

$$\begin{aligned} \chi^2(R_{\text{Cons}}^{\text{comp}} \| U_{\text{Ran}_1}) &= \frac{\text{Var}[p_W]}{(1 - h/M)^2} \\ &\leq \frac{h^2(t + 2h - 1)}{\ell(N - s)^2} \leq 8 \frac{s^3}{N^3}, \end{aligned}$$

where we used $h \leq s$, $t + 2h - 1 \leq 2s$, $\ell \geq N/2$, and $N - s \geq 3N/4$. $\square$

C.3 The augmented inverse response

We combine the divergence bounds for the inverse output and companion output using the KL chain rule. For the augmented inverse response, let $V_{\text{comp}} = \Pi_1(X)$ and put $R_{\text{Cons}}^{\text{inv}} = \mathcal{L}((X, V_{\text{comp}}) \mid \text{Ran}_0, \text{Ran}_1, \text{Cons})$. Write $K_{\text{Hyb}}^{\text{inv}}$ for the joint response distribution in **Hyb** from Section 3.6, whose marginal distribution of the inverse output is $K_{\text{Hyb}}^{\text{out}}$.

Lemma 9 (Augmented inverse response). *For a fresh query $\Pi_0^{-1}(y)$ and $s \leq N/4 - 1$,*

$$D_{\text{KL}}(R_{\text{Cons}}^{\text{inv}} \| K_{\text{Hyb}}^{\text{inv}}) \leq 101 \frac{s^3}{N^3}.$$

Proof. When $X \in \text{dom}(\text{Cons})$, the companion output is the same fixed value in both worlds. For each $x \in E_{\text{new}}$, its conditional distributions given $X = x$ are $R_{\text{Cons}}^{\text{comp}}$ and U_{Ran_1} by exchangeability. The KL chain rule thus gives exactly

$$D_{\text{KL}}(R_{\text{Cons}}^{\text{inv}} \| K_{\text{Hyb}}^{\text{inv}}) = D_{\text{KL}}(R_{\text{Cons}} \| K_{\text{Hyb}}^{\text{out}}) + R_{\text{Cons}}(E_{\text{new}}) D_{\text{KL}}(R_{\text{Cons}}^{\text{comp}} \| U_{\text{Ran}_1}).$$

Bound the branch probability by one and use $D_{\text{KL}} \leq \chi^2$ with Lemmas 7 and 8. The two bounds sum to $101s^3/N^3$. $\square$

D Responses and Failure in $\mathbf{Ideal}_3$

D.1 Conditional distribution of an unrevealed construction value

We control how each inverse-response likelihood changes with one unrevealed construction value, then apply the reweighting lemma.

Lemma 10 (Conditional distribution of one construction value). *Consider the conditional distribution π_τ of the table of Γ in $\mathbf{Ideal}_3$, given by (8), optionally reweighted by the current factor $\vartheta_L(I)$. Fix all but one of the unrevealed construction values. The conditional distribution P of the remaining value satisfies*

$$\frac{1}{N}e^{-4q/N} \leq P(z) \leq \frac{1}{N}e^{4q/N}.$$

For $V = \mathcal{G}$ or any fixed set of legal construction values with $|V| \geq N - q$,

$$\chi^2(\text{Unif}(V) \| P(\cdot | V)) \leq 20q^3/N^3.$$

The same chi-squared bound holds after averaging over the other coordinates under their conditional distribution, including after the legal restriction.

Proof (Lemma 10). Write $\tilde{\vartheta}_L(I) := L\vartheta_L(I)$. For integers $0 \leq I \leq L-1$, $\tilde{\vartheta}_L(I) \geq 1$ and

$$\tilde{\vartheta}_L(I+1) - \tilde{\vartheta}_L(I) = \frac{1}{L} + \frac{2I+1}{L^2} \leq \frac{3}{L} \leq \frac{4}{N}.$$

The last inequality follows from $L \geq N - q \geq 63N/64$ and $N \geq 64$. Thus $0 \leq \log(\vartheta_L(I+1)/\vartheta_L(I)) \leq 4/N$.

Fix an input x with an unrevealed construction value and fix all the other coordinates. Let I_j^- denote the number of illegal candidates other than x at query j . The illegal count is $I_j^- + \mathbf{1}[z \in \text{Forb}_j]$ when $x \in E_j$, and otherwise is constant. The current factor, when included, has the same form. Taking logarithms of these factors shows that the conditional distribution of $\Gamma(x)$ is obtained by reweighting $\text{Unif}(\mathcal{G})$ by $e^{b(z)}$, where

$$b(z) = \sum_j \delta_j \mathbf{1}[z \in \text{Forb}_j], \quad 0 \leq \delta_j \leq 4/N, \quad 0 \leq b(z) \leq A := 4q/N \leq 1/16.$$

There are at most q factors, including the current one when used, and each $|\text{Forb}_j| \leq q$. Weighted Cauchy–Schwarz gives

$$\mathbb{E}_{z \leftarrow \text{Unif}(\mathcal{G})}[b(z)^2] \leq \left(\sum_j \delta_j \right) \sum_j \delta_j \frac{|\text{Forb}_j|}{N} \leq 16q^3/N^3.$$

Reweight $U = \text{Unif}(V)$ by e^b to obtain U_b . Since $|V| \geq N - q$, the second moment of b under U is at most $N/(N - q)$ times its second moment under $\text{Unif}(\mathcal{G})$. Since

$\text{osc}_V(-b) = \text{osc}_V(b) \leq A$, apply (3) to U_b with weights e^{-b} , which gives U . Using $U_b(z)/U(z) \leq e^A$, we obtain

$$\chi^2(U \| U_b) \leq e^{2A} \text{Var}_{U_b}(b) \leq e^{3A} \mathbb{E}_U[b^2] \leq \frac{16e^{3/16}}{1 - 1/64} \frac{q^3}{N^3} < 20q^3/N^3.$$

Before restricting to V , the probability ratio to $\text{Unif}(\mathcal{G})$ lies between e^{-A} and e^A . For fixed V , convexity in the second argument preserves the chi-squared bound when the other coordinates are averaged under their conditional distribution. If the construction value is restricted to V , this average uses the distribution of the other coordinates conditioned on $\Gamma(x) \in V$. $\square$

D.2 Distributions of successful responses

We now use the bound for unrevealed construction values to compare successful responses. For inverse queries, we separate the inverse output from its companion output.

Proof (Response bounds in Proposition 2). Lemma 10 bounds the divergence between the distributions of a fresh construction response. For a forward primitive query at an input with an unrevealed construction value, the same bound applies to the distributions of that value. Conditional on it, the permutation pair is uniform among the legal pairs in both worlds. If the construction value is revealed, the conditional pair distributions agree.

For a fresh inverse call, write $t = |\text{Dom}|$, $h = |\text{Cons}|$, $L = N - t$, and $u = N - t - h$. Let Forb be the set of forbidden construction values, $|\text{Forb}| = t$, and let k be the number of legal inputs in U with revealed construction values. Reweight π_τ by the current factor $\vartheta_L(I)$ and call the normalized distribution π_τ^ϑ . Its unrevealed entries are exchangeable: their inputs belonged to the available set of every earlier inverse call, and every likelihood factor is invariant under permuting these entries. Put

$$\theta = \Pr_{\pi_\tau^\vartheta}[\Gamma(x) \notin \text{Forb}], \quad \rho = 1 - t/N$$

for any input x with an unrevealed construction value. Summing $Q_\tau^+(\Phi_\tau, x)$ in (10) over Φ_τ gives the distribution of the inverse output

$$P_{\text{ret}}(x) = \begin{cases} 1/(k + u\theta), & x \in \text{dom}(\text{Cons}) \text{ and legal,} \\ \theta/(k + u\theta), & x \text{ with unrevealed construction value.} \end{cases}$$

The hybrid distribution $K_{\text{Hyb}}^{\text{out}}$ uses this formula with θ replaced by ρ . Direct summation over the two groups gives

$$\chi^2(K_{\text{Hyb}}^{\text{out}} \| P_{\text{ret}}) = \frac{ku(\theta - \rho)^2}{\theta(k + u\rho)^2}. \quad (22)$$

When $k = 0$, both distributions are uniform over the inputs with unrevealed construction values.

With $A = 4q/N$, the density bounds in Lemma 10 imply

$$|\theta - \rho| \leq (e^A - 1)t/N \leq 8qt/N^2, \quad \theta \geq e^{-A}\rho.$$

Since $k \leq q$, $u \geq N - q$, and $\rho \geq 63/64$, the coefficient of $(\theta - \rho)^2$ in (22) is at most $q/(u\theta\rho^2) \leq e^{1/16}(64/63)^4q/N < 2q/N$. Hence

$$\chi^2(K_{\text{Hyb}}^{\text{out}} \| P_{\text{ret}}) \leq 128q^5/N^5 \leq q^3/N^3.$$

Conditional on returning a specified input x whose construction value was unrevealed, the distribution of the table of Γ is precisely π_τ^ϑ conditioned on $\Gamma(x) \notin \text{Forb}$. Lemma 10, averaged over the other coordinates under this conditional distribution, bounds the divergence from $\text{Unif}(\text{Forb}^c)$ to this distribution of the construction value by $20q^3/N^3$. XOR with the target gives the companion output, whose distribution in **Hyb** is uniform over the unused outputs of the other permutation. If the inverse output has a revealed construction value, the companion output is deterministic. We apply the KL chain rule to the inverse output and its companion output. Bounding the probability in **Hyb** of returning an input with an unrevealed construction value by one gives the KL bound $21q^3/N^3$. $\square$

D.3 Failure probability

We control rejection at each of the three trials, then sum the failure probabilities over all queries.

Proof (Failure bound in Proposition 2). Write $t = |\text{Dom}|$, $h = |\text{Cons}|$, and $L = N - t$, so $h + t \leq q$. For $q = 1$, a fresh inverse call has $\text{Forb} = \emptyset$ and cannot fail. Assume $q \geq 2$ and condition on the candidate identities and construction values tested so far in the current call. The rejection indicators and sampling probabilities add no factor depending on the remaining unrevealed construction values, so each remaining construction value has conditional probability at most $e^{1/16}/N$ by Lemma 10. Earlier calls are represented by π_τ , which averages over their rejected candidates.

There are at most h inputs with revealed illegal values and at most two additional inputs whose construction values were read during earlier rejected trials in this call. Sampling with replacement may revisit these inputs. Every trial is uniform on the same set of size L , so its conditional rejection probability is at most

$$\frac{h + 2 + e^{1/16}t}{L} \leq \frac{16q/15 + 2}{63N/64}.$$

Multiplying the three conditional rejection bounds and using $q \geq 2$ gives

$$\Pr[\text{the call fails} \mid \tau] \leq \left(\frac{16q/15 + 2}{63N/64} \right)^3 < 10q^3/N^3.$$

A union bound over at most q calls gives $\Pr_{\text{ideal}_3}[\text{Failure} = 1] \leq 10q^4/N^3$. $\square$

E Analysis of the Uniform Simulator Attack

We prove the bounds on the attack statistic in Lemma 3. Throughout, r is a power of two satisfying

$$64n \leq r \leq \min \left\{ \frac{N}{2^{11}}, \frac{N}{24n} \right\}.$$

There are at most $3r$ fresh successful inverse responses, and each forbidden-value set has size at most $3r$.

E.1 The attack statistic and the good event

Let Rep_2 and Rep_3 denote the numbers of iterations skipped in Phases 2 and 3, respectively, because the target of the inverse query is already in the recorded range. Let $\text{Ran}_{1,\text{fin}}$ be the value of Ran_1 when the query procedure terminates, and let Bot_∞ be the event that at least one inverse response in Ideal^∞ is $\perp$. Put

$$N_C^{\text{hit}} := |\text{Ran}_{1,\text{fin}} \cap C|,$$

$$\text{Good}_{\text{atk}} := (\neg \text{Bot}_\infty) \cap \{\text{Rep}_2 \leq r/8, \text{Rep}_3 \leq r/8, N_C^{\text{hit}} \leq r/8\}.$$

Let K be the indices of fresh successful inverse responses in Phase 3, and set

$$T_k := \mathbf{1}[Z_k \in V] - \pi_k \quad (k \in K), \quad T_k := 0 \quad (k \notin K).$$

Thus $T_{\text{atk}} = \sum_{k=1}^r T_k$. Write τ_k for the ordinary transcript after the range check and any inverse query in iteration k of Phase 3, before the companion query. Membership in K is determined by this prefix.

For $W \in \{\text{Real}, \text{Ideal}^\infty\}$ and $k \in K$, put

$$\rho_k^W := \Pr_W[Z_k \in V \mid \tau_k], \quad \Delta_k^W := \mathbf{1}[Z_k \in V] - \rho_k^W,$$

and set both variables to zero for $k \notin K$. The Δ_k^W are the martingale differences from Section 4.2, with

$$M_\infty := \sum_{k=1}^r \Delta_k^{\text{Ideal}^\infty}.$$

The cumulative drift is $\sum_k d_k$, where

$$d_k = \rho_k^{\text{Ideal}^\infty} - \pi_k \quad (k \in K), \quad d_k = 0 \quad (k \notin K),$$

and

$$T_{\text{atk}} = \sum_{k=1}^r d_k + M_\infty.$$

E.2 Legal candidate counts

Fix an ordinary transcript τ of Ideal^∞ with no $\perp$ response. Let R_τ be the inputs whose construction values remain unrevealed, and let J_τ index the fresh successful inverse responses in τ . If τ ends just after a fresh inverse response, these sets include the inverse output and the current response, respectively. For $j \in J_\tau$, let R_j and Forb_j be the available inputs and forbidden construction-value set at call j .

Let $\Phi_\tau : \mathcal{G} \rightarrow \mathcal{G}$ agree with the construction values fixed by τ . Let μ_τ^{atk} be the distribution on these tables with independent uniform values in $\mathcal{G}$ at inputs in R_τ , subject to the inverse output's legality when its value is still unrevealed. For every event E determined by Γ , let $\Pr_\tau^\infty[E]$ denote its conditional probability in Ideal^∞ given the ordinary transcript τ . Write

$$V_j(\Phi_\tau) = V_j^{\text{pub}} + \sum_{y \in R_\tau} \mathbf{1}[y \in R_j] \mathbf{1}[\Phi_\tau(y) \notin \text{Forb}_j].$$

Here V_j^{pub} is the fixed contribution of coordinates that have been revealed since call j . For a table Φ_τ with $\mu_\tau^{\text{atk}}(\Phi_\tau) > 0$, the observed inverse responses contribute the likelihood factor $\prod_{j \in J_\tau} 1/V_j(\Phi_\tau)$. The conditional distribution of Γ given τ is therefore

$$\Pr_\tau^\infty[\Gamma = \Phi_\tau] = \frac{\mu_\tau^{\text{atk}}(\Phi_\tau) \prod_{j \in J_\tau} V_j(\Phi_\tau)^{-1}}{\mathbb{E}_{\Phi'_\tau \leftarrow \mu_\tau^{\text{atk}}} [\prod_{j \in J_\tau} V_j(\Phi'_\tau)^{-1}]}. \quad (23)$$

For $T \subseteq R_\tau$, $|T| \leq 1$, let $V_j^{(-T)}$ omit the coordinates in T . If the inverse output has an unrevealed value restricted by legality, require that it belong to T . Define

$$\text{Dense}_{\tau, T}^{\text{atk}} := \{V_j^{(-T)} \geq N/2 \text{ for every } j \in J_\tau\}.$$

Lemma 11 (Lower bounds on the numbers of legal candidates). *Uniformly over τ and every admissible fixed T of size at most one,*

$$\Pr_\tau^\infty[(\text{Dense}_{\tau, T}^{\text{atk}})^c] \leq \frac{r^2}{32N^2}. \quad (24)$$

Proof. Every coordinate whose construction value is still unrevealed was available at each earlier call. After omitting T , at least $N - 3r - 1$ independent uniform unrevealed values of Γ remain, and each is legal for call j with probability $1 - |\text{Forb}_j|/N \geq 1 - 3r/N$. The range assumptions imply that the mean is at least $7N/8$. Hoeffding's inequality [11, Theorem 2] bounds the probability that this sum is below $N/2$ by $e^{-9N/32} \leq e^{-N/4}$. A union bound over at most $3r$ factors gives

$$\mu_\tau^{\text{atk}}((\text{Dense}_{\tau, T}^{\text{atk}})^c) \leq 3re^{-N/4} < \frac{1}{2}.$$

Put $W = \prod_{j \in J_\tau} 1/V_j$. On the dense event, $W \geq N^{-|J_\tau|}$. For every assignment compatible with τ , each $V_j \geq 1$, so $W \leq 1$. Comparing numerator and

denominator in (23) gives

$$\Pr_{\tau}[(\text{Dense}_{\tau,T}^{\text{atk}})^c] \leq 2N^{|J_{\tau}|} \mu_{\tau}^{\text{atk}}((\text{Dense}_{\tau,T}^{\text{atk}})^c) \leq 6rN^{3r}e^{-N/4}.$$

Since $rn \leq N/24$, we have $N^{3r} \leq e^{N/8}$, so the last expression is at most $6r/e^{N/8}$. Since $1 \leq r \leq N/2^{11}$, the exponential series gives

$$\frac{192N^2}{re^{N/8}} \leq \frac{192 \cdot 5! \cdot 8^5}{rN^3} < 1.$$

Thus $6r/e^{N/8} \leq r^2/(32N^2)$. $\square$

E.3 Construction value at the inverse output

Write $\text{Dense}_{\tau,x}^{\text{atk}}$ for the event with $T = \{x\}$. The following calculation applies to every fresh inverse response, using $x = x_k$ for an iteration of Phase 3. Fix all construction oracle values except $\Gamma(x)$, and write $z = \Gamma(x)$. Let J_k^{hist} index the fresh successful inverse responses strictly before the current inverse query, so $|J_k^{\text{hist}}| \leq 3r$. For $j \in J_k^{\text{hist}}$, let Forb_j be its forbidden construction-value set and $V_j^{(-x)}$ its number of legal candidates other than x . Suppose these values satisfy $\text{Dense}_{\tau_k,x}^{\text{atk}}$. An earlier inverse response contributes the likelihood factor

$$\frac{1}{V_j^{(-x)} + \mathbf{1}[z \notin \text{Forb}_j]} = \frac{1}{V_j^{(-x)} + 1} \left(1 + \frac{1}{V_j^{(-x)}} \right)^{\mathbf{1}[z \in \text{Forb}_j]}.$$

For $z \in \mathcal{Z}_k$, the likelihood factor for the current inverse response is constant in z and cancels on normalization. Let $\tilde{\nu}_k$ be the conditional distribution of $\Gamma(x)$ with the other values fixed, and put $\nu_k = \text{Unif}(\mathcal{Z}_k)$. Then

$$\frac{\tilde{\nu}_k(z)}{\nu_k(z)} = \frac{e^{\Psi_k(z)}}{\mathbb{E}_{w \leftarrow \nu_k}[e^{\Psi_k(w)}]}, \quad \Psi_k(z) = \sum_{j \in J_k^{\text{hist}}} \lambda_j \mathbf{1}[z \in \text{Forb}_j], \quad (25)$$

where $\lambda_j = \log(1 + 1/V_j^{(-x)})$. We average $\tilde{\nu}_k$ over the other construction values under their distribution given τ_k , which records the current inverse response. On $\text{Dense}_{\tau_k,x}^{\text{atk}}$, $N/2 \leq V_j^{(-x)} \leq N$, and hence

$$\frac{1}{2N} \leq \lambda_j \leq \frac{2}{N}, \quad \text{osc}_z(\Psi_k) \leq \frac{6r}{N}. \quad (26)$$

Hit probabilities and the good event. On the dense event, (25)–(26) bound the conditional probability of each construction value by $e^{6r/N} < 2$ times its uniform probability. After averaging over the other values and applying (24), we obtain the following bound on the conditional probability that a new companion output lies in any fixed set of size r :

$$\frac{4r}{N} + \frac{r^2}{32N^2}. \quad (27)$$

At most r companion outputs contribute to each of Rep_2 and Rep_3 , and at most $2r$ contribute to N_C^{hit} . Markov's inequality at threshold $r/8$ and a union bound give

$$\Pr_{\text{Ideal}^\infty} [(\neg \text{Bot}_\infty) \cap \text{Good}_{\text{atk}}^c] \leq \frac{128r}{N} + \frac{r^2}{N^2}. \quad (28)$$

E.4 Drift of the attack statistic

For $k \in K$, suppose that Phase 2 contains at least $7r/8$ fresh successful inverse responses and $|V \setminus \mathcal{Z}_k| \leq r/8$. These conditions depend only on τ_k and hold on Good_{atk} . Fix the other construction oracle values so that $\text{Dense}_{\tau_k, x_k}^{\text{atk}}$ holds. For each fresh inverse query j in Phase 2, $V \subseteq \text{Forb}_j$. We also have $|\mathcal{G} \setminus \mathcal{Z}_k| \leq 3r$. Therefore, for every such index j from Phase 2,

$$\frac{7r}{8N} \leq \pi_k \leq \frac{2r}{N}, \quad \nu_k(V \cap \text{Forb}_j) - \pi_k \nu_k(\text{Forb}_j) \geq \frac{21r}{32N}. \quad (29)$$

There are at most $2r$ forbidden sets from inverse calls outside Phase 2. For each such index j , positivity of the intersection term and $|\text{Forb}_j| \leq 3r$ give

$$\nu_k(V \cap \text{Forb}_j) - \pi_k \nu_k(\text{Forb}_j) \geq -\frac{8r^2}{N^2}. \quad (30)$$

Since

$$\mathbb{E}_{z \leftarrow \nu_k} [(\mathbf{1}[z \in V] - \pi_k) \Psi_k(z)] = \sum_{j \in J_k^{\text{hist}}} \lambda_j (\nu_k(V \cap \text{Forb}_j) - \pi_k \nu_k(\text{Forb}_j)),$$

combining (26)–(30) and $r \leq N/2^{11}$ gives

$$\mathbb{E}_{z \leftarrow \nu_k} [(\mathbf{1}[z \in V] - \pi_k) \Psi_k(z)] \geq \frac{3r^2}{16N^2}. \quad (31)$$

Let

$$\text{inc}_k(z) := \sum_{j \in J_k^{\text{hist}}} \mathbf{1}[z \in \text{Forb}_j].$$

On the dense event, $0 \leq \Psi_k(z) \leq 2 \text{inc}_k(z)/N$. The sum of these counts is at most $9r^2$, each count is at most $3r$, and $|\mathcal{Z}_k| \geq N - 3r$. Thus

$$\mathbb{E}_{W \leftarrow \nu_k} [\text{inc}_k(W)^2] \leq \frac{54r^3}{N}.$$

Hence

$$\text{Var}_{W \leftarrow \nu_k} [\Psi_k(W)] \leq \frac{216r^3}{N^3}.$$

To bound the nonlinear contribution in (25), let g and a have mean zero under ν , and put $\omega = \text{osc}(g)$, $\sigma^2 = \text{Var}_\nu(g)$, $Z = \mathbb{E}_\nu[e^g]$, $h = e^g - 1 - g$, and $R_g = e^g/Z - 1 - g$. Since $\mathbb{E}_\nu[a] = 0$,

$$\mathbb{E}_\nu[aR_g] = \frac{\mathbb{E}_\nu[ah] - (Z-1)\mathbb{E}_\nu[ag]}{Z}.$$

Here $Z \geq 1$, $|g| \leq \omega$, and the exponential power series gives

$$(\mathbb{E}_\nu[h^2])^{1/2} \leq \frac{1}{2}e^\omega \omega \sigma, \quad 0 \leq Z - 1 \leq \frac{1}{2}e^\omega \sigma^2.$$

Cauchy–Schwarz and $\sigma \leq \omega$ therefore imply

$$|\mathbb{E}_\nu[aR_g]| \leq \frac{1}{2}e^\omega (1 + \omega) \omega \sqrt{\text{Var}_\nu(a) \text{Var}_\nu(g)}.$$

Apply this with $g = \Psi_k - \mathbb{E}_{\nu_k}[\Psi_k]$ and $a = \mathbf{1}_V - \pi_k$. By (26), $\omega \leq 6r/N \leq 3/1024$. Using $\text{Var}_{\nu_k}(a) \leq 2r/N$ and the variance bound for Ψ_k above gives

$$|\mathbb{E}_{z \leftarrow \nu_k}[(\mathbf{1}[z \in V] - \pi_k)R_g(z)]| \leq 64 \frac{r^3}{N^3} \leq \frac{r^2}{32N^2}. \quad (32)$$

Equations (24), (31), and (32), after adding the contribution of the complement of $\text{Dense}_{\tau_k, x_k}^{\text{atk}}$, show that, under these prefix conditions, the conditional mean for each fresh successful inverse response in Phase 3 satisfies

$$d_k \geq \frac{r^2}{16N^2}.$$

On Good_{atk} , at least $7r/8$ such increments contribute to the drift, so

$$\sum_{k \in K} d_k > \frac{r^3}{32N^2}. \quad (33)$$

E.5 Second moments of the martingale parts

For $W \in \{\text{Real}, \text{Ideal}^\infty\}$, martingale orthogonality gives

$$\mathbb{E}_W \left[\left(\sum_{k=1}^r \Delta_k^W \right)^2 \right] = \sum_{k=1}^r \mathbb{E}_W[(\Delta_k^W)^2] \leq \sum_{k=1}^r \mathbb{E}_W[\rho_k^W].$$

In Real , Z_k is exactly uniform on $\mathcal{Z}_k$, so $\rho_k^{\text{Real}} = \pi_k$ and $T_{\text{atk}} = \sum_k \Delta_k^{\text{Real}}$. Since $|\text{Ran}_{1,k}| \leq 3r$, one has $\pi_k \leq r/(N - 3r) < 2r/N$, and hence

$$\mathbb{E}_{\text{Real}}[T_{\text{atk}}^2] \leq 2r^2/N.$$

Equation (27) gives $\rho_k^{\text{Ideal}^\infty} \leq 5r/N$, so $\mathbb{E}_{\text{Ideal}^\infty}[M_\infty^2] \leq 5r^2/N$. By (33), a completed execution with drift $\sum_k d_k \leq r^3/(32N^2)$ lies in $(\neg \text{Bot}_\infty) \cap \text{Good}_{\text{atk}}^c$. Equation (28) bounds its probability by $128r/N + r^2/N^2$.

E.6 Choice of parameters

Set $r = 2^{\lceil 3n/4 \rceil + 4}$ and suppose $n \geq 64$. Then

$$\frac{r}{N^{3/4}} = 2^{\lceil 3n/4 \rceil - 3n/4 + 4} \leq 2^{19/4}, \quad 6r < 162N^{3/4}.$$

Moreover, for $n \geq 64$,

$$\begin{aligned} r &\geq 2^{3n/4+4} \geq 64n, \\ \frac{N}{r} &= 2^{\lfloor n/4 \rfloor - 4} \geq 24n, \\ \frac{N}{r} &\geq 2^{12}, \end{aligned}$$

so both upper bounds on r in (14) hold. Finally,

$$\frac{N^3}{r^4} \leq 2^{-16}, \quad 28672 \frac{N^3}{r^4} \leq \frac{7}{16}.$$

Since $r/N \leq 2^{-12}$,

$$28672 \frac{N^3}{r^4} + 129 \frac{r}{N} < \frac{1}{2}.$$

F Implementing the Gyroscope Simulator

We describe the data structures and exact acceptance sampler for GSim_5 , then prove the memory, time, and implementation-error bounds stated in Section 5.2.

The sets Forb_j need not be stored explicitly. Index the points of Dom in the order in which they are added. When the k th point is added with outputs y_0, y_1 , store

$$\text{birth}_0(y_0) = k, \quad \text{birth}_1(y_1) = k.$$

Then membership in Forb_j is recovered exactly from entry j of Hist_τ :

$$z \in \text{Forb}_j \iff \text{birth}_{1-b_j}(y_j \oplus z) \text{ exists and is at most } d_j.$$

Exact acceptance sampling. Algorithm 6 implements the exponential Bernoulli sampler of [4, Algorithm 2] using Hist_τ . For a candidate with construction value z , it returns one with probability $e^{-\text{Corr}_\tau(z)}$.

Algorithm 6 Exact acceptance sampling for the Gyroscope simulator.

```

1: function ACCEPT( $z, \text{Hist}_\tau$ )
2:    $m_{\text{inv}} \leftarrow |\text{Hist}_\tau|$ 
3:   if  $m_{\text{inv}} = 0$  then return 1
4:   for  $k = 1, 2, \dots$  do
5:      $J \leftarrow \text{Unif}(I_\tau)$ 
6:      $V \leftarrow \text{Unif}(\{1, \dots, kL_J\})$ 
7:     if  $z \notin \text{Forb}_J$  or  $V > m_{\text{inv}}$  then
8:       return  $\mathbf{1}[k \text{ is odd}]$ 
```

For $m_{\text{inv}} > 0$, write $\text{Corr} = \text{Corr}_\tau(z)$. Since $m_{\text{inv}} \leq q \leq N/64$ and $L_j \geq N - q$, we have $m_{\text{inv}} \leq kL_j$. The probability of continuing at iteration k is therefore

$$\frac{1}{m_{\text{inv}}} \sum_{j \in I_\tau} \mathbf{1}[z \in \text{Forb}_j] \frac{m_{\text{inv}}}{kL_j} = \frac{\text{Corr}}{k}.$$

The integer draws use unbiased random bits and rejection sampling.

The data structures and sampler give the following implementation guarantee.

Proposition 3 (Efficiency of GSim_5). *For interactions of at most $q \leq N/64$ direct queries in total, including at most q_p primitive queries, GSim_5 stores $O(q)$ entries in its permutation tables and Hist_τ and makes at most $5q_p$ internal Γ -queries. The implementation with exact sampling procedures has expected total running time $O(q \text{poly}(n))$.*

For every constant $A > 0$, there is an implementation of GSim_5 with total variation error $\varepsilon_{\text{impl}} \leq N^{-A}$, worst-case running time $q \text{poly}_A(n)$, and $O(qn)$ bits of memory.

Proof (Proposition 3). The tables P_0, P_1 , their inverse maps, the dictionaries recording when outputs were assigned, and Hist_τ contain $O(q)$ entries. A forward call reads at most one new Γ -entry and a fresh inverse call reads at most five. Since there are at most q_p primitive queries, the total number of internal construction queries is at most $5q_p$.

Fix a candidate with construction value z and abbreviate $\text{Corr} = \text{Corr}_\tau(z) \leq 1/63$. The acceptance sampler reaches iteration k with probability $\text{Corr}^{k-1}/(k-1)!$. It therefore terminates almost surely, and its probability of returning one is

$$\sum_{j \geq 0} \left(\frac{\text{Corr}^{2j}}{(2j)!} - \frac{\text{Corr}^{2j+1}}{(2j+1)!} \right) = e^{-\text{Corr}}.$$

It uses $e^{\text{Corr}} < 2$ iterations in expectation. Iteration k requires a constant number of dictionary operations and draws from sets of size at most kN . Uniform integers are sampled by rejecting out-of-range bit strings; each attempt uses $O(n + \log k)$ bits and succeeds with probability greater than $1/2$. Since the probability of reaching iteration k decreases exponentially, the acceptance sampler runs in expected $\text{poly}(n)$ time. Dictionary lookups also cost $\text{poly}(n)$, proving the expected running-time bound. These data structures use $O(qn)$ bits, and iteration k uses $O(n + \log k)$ additional working bits.

For the worst-case implementation, limit every acceptance sampler and every uniform rejection sampler to R iterations. An acceptance sampler exceeds this limit with probability

$$\frac{\text{Corr}^R}{R!} \leq 63^{-R}.$$

There are at most $5q$ acceptance calls and at most $10qR$ integer draws within them. Each integer sampler exceeds its limit with probability at most 2^{-R} . Sampling a new forward primitive output rejects with probability at most $2q/N \leq$

1/32 per trial, and sampling an unused inverse-query input rejects with probability at most $q/N \leq 1/64$. These samplers are called at most q and $5q$ times, respectively.

Couple the exact and bounded implementations using the same random bits. They agree until a sampling limit is reached. On that event the bounded implementation returns $\perp$ on the current and all subsequent primitive queries, while leaving the construction oracle available. A union bound gives

$$\varepsilon_{\text{impl}} \leq q32^{-R} + 5q64^{-R} + 5q63^{-R} + 10qR2^{-R}.$$

Taking $R = \lceil C_A n \rceil$ for a sufficiently large constant C_A makes this at most N^{-4} . Every integer then has $O_A(n)$ bits, and each primitive query requires at most $O(R^2)$ sampling trials and a constant number of updates to the stored data. $\square$

G Proof of Theorem 3

We prove Theorem 3 by comparing Real with Hyb^{GSim} , and then Hyb^{GSim} with $\text{Ideal}^{\text{GSim}}$. The hybrid is defined in Section 5.4. This appendix combines the two transcript bounds. Appendices H–K prove the comparison with Real , and Appendix L proves the comparison with $\text{Ideal}^{\text{GSim}}$, including simulator failure. Throughout these comparisons, $N = 2^n \geq 2^{20}$ and $q \leq N/64$.

Lemma 24 bounds the divergence from T_{Real} to $T_{\text{Hyb}^{\text{GSim}}}$, and Lemma 30 bounds that from $T_{\text{Hyb}^{\text{GSim}}}$ to $T_{\text{Ideal}^{\text{GSim}}}$. Their statements and proofs are in Appendices K.6 and L.4, respectively. These bounds give the security theorem as follows.

Proof (Theorem 3). Using $q/N \leq 1/64$ in Lemma 30 gives

$$D_{\text{KL}}(T_{\text{Hyb}^{\text{GSim}}} \| T_{\text{Ideal}^{\text{GSim}}}) \leq \frac{28q^5}{N^4}.$$

Combining this with Lemma 24, Pinsker's inequality and the triangle inequality gives, for $n \geq 20$,

$$\text{TV}(T_{\text{Real}}, T_{\text{Ideal}^{\text{GSim}}}) \leq 11n^{3/2} \frac{q^{3/2}}{N^{3/2}} + 11\sqrt{n} \frac{q^{5/2}}{N^2},$$

where the term $\sqrt{q/2} N^{-4}$ is absorbed into the first term. This bounds the advantage with exact sampling; Proposition 3 adds at most $\varepsilon_{\text{impl}}$ when sampling loops are limited. $\square$

H Bounds on Repeated Values and Pairs with a Given XOR in Real

Assume $N = 2^n \geq 2^{20}$ and $q \leq N/64$, and put

$$B_{\text{occ}} := 2n, \quad B_{\text{xor}}(s) := 3s^2/N + 42n^2.$$

H.1 The sequence of newly revealed construction values

Run `Real` with augmented transcripts. List, in chronological order, every construction value when it is first revealed:

$$Z_1, Z_2, \dots, Z_\rho, \quad \rho \leq q.$$

The list contains construction values first revealed by construction queries, fresh forward primitive queries, and fresh inverse primitive queries returning an input whose construction value was unrevealed. Distinct inputs may have equal construction values. If only ρ construction values are revealed, keep all counts unchanged for $\rho < j \leq q$.

The convention for recording each input once and the definition of a good prefix are given in Section 6.1.

We first bound the probability of each newly revealed value.

Lemma 12 (Conditional distribution of a revealed construction value). *Suppose $q < N/2$. Let $\mathcal{F}_{j-1}^{\text{rev}}$ contain all publicly revealed information immediately before the j th revelation of a previously unrevealed construction value, but no unused permutation entry. Then*

$$\max_{z \in \mathcal{G}} \Pr[Z_j = z \mid \mathcal{F}_{j-1}^{\text{rev}}] \leq \frac{1}{N-q} < \frac{2}{N}. \quad (34)$$

Proof. Condition on the unexposed permutation outputs at inputs whose construction values have already been revealed. Condition also on the input at which Z_j is revealed and on one of the two permutation outputs there, taking the target for an inverse query. The other output is uniform on a remaining range of size at least $N - q$. Each XOR value determines at most one of these outputs, so its probability is at most $1/(N - q)$. Averaging over the additional conditioning preserves this bound, and $1/(N - q) \leq 2/N$. $\square$

The one-value bound gives simultaneous control of repeated construction values and pairs with a given XOR. The next lemma proves that the good-prefix conditions (19) hold throughout the interaction with high probability.

Lemma 13 (Bounds on repeated values and pairs with a given XOR). *Except with probability at most N^{-8} , the first k newly revealed construction values satisfy, simultaneously for every $k \leq \rho$,*

$$\begin{aligned} \max_z |\{i \leq k : Z_i = z\}| &\leq 2n, \\ \max_d |\{(i, j) : i \neq j, i, j \leq k, Z_i \oplus Z_j = d\}| &\leq \frac{3k^2}{N} + 42n^2. \end{aligned}$$

The same bounds hold for any subset of the first k indices, with the original value of k on the right-hand side.

Proof. For a fixed value z , let $\text{Mult}_k(z)$ denote the number of its occurrences among the first $\min\{k, \rho\}$ revealed construction values. By (34), the sum of the conditional expectations of its increments is at most $2q/N \leq 1/32$. Iterating the conditional exponential-moment bound for Bernoulli increments gives, for every $\theta > 0$,

$$\Pr[\text{Mult}_q(z) > B] \leq \exp((e^\theta - 1)/32 - \theta B).$$

The counts are nondecreasing, so this also controls every prefix. For $B = 2n$, choose $\theta = 6$. Since $(e^6 - 1)/32 < 13 < 12n(1 - \log 2)$ for $n \geq 20$, a union bound over z gives failure probability at most N^{-11} .

For a fixed difference d , put $S_k = \sum_{j \leq k} Y_j$, where

$$Y_j = \mathbf{1}[\text{all earlier multiplicities are at most } B] \sum_{i < j} \mathbf{1}[Z_i \oplus Z_j = d],$$

with $Y_j = 0$ for $j > \rho$. Here $\mathbb{E}_{j-1}[\cdot]$ conditions on $\mathcal{F}_{j-1}^{\text{rev}}$. Then $0 \leq Y_j \leq B$, including the first revelation that violates the multiplicity bound, and $\mathbb{E}_{j-1}[Y_j] \leq 2(j-1)/N$. Convexity on $[0, B]$ yields

$$\mathbb{E}_{j-1}[e^{\theta Y_j/B}] \leq \exp((e^\theta - 1)\mathbb{E}_{j-1}[Y_j]/B).$$

For each fixed k , iterating this inequality gives

$$\mathbb{E}[e^{\theta S_k/B}] \leq \exp\left(\frac{(e^\theta - 1)k(k-1)}{BN}\right).$$

Markov's inequality therefore gives

$$\Pr\left[S_k > \frac{e^\theta - 1}{\theta} \frac{k(k-1)}{N} + \frac{B}{\theta} x\right] \leq e^{-x}.$$

Choose $\theta = 3/4$ and $x = 11 \log N$. Since $2(e^{3/4} - 1)/(3/4) < 3$, this gives, for each fixed (d, k) , except with probability N^{-11} ,

$$2S_k \leq \frac{3k(k-1)}{N} + 30B \log N \leq \frac{3k^2}{N} + 42n^2.$$

A union bound over all $Nq \leq N^2$ pairs (d, k) gives failure probability at most N^{-9} . When no construction value occurs more than B times, the indicator in Y_j is always one, and $2S_k$ is the number of ordered pairs of distinct indices whose construction values have XOR d . The total failure probability is at most $N^{-11} + N^{-9} \leq N^{-8}$. Deleting indices can only decrease all counts. $\square$

H.2 Matching XOR values and squared sums

For $C \subseteq \mathcal{G}$, write $h_C := \mathbf{1}_C - |C|/N$ for its centered indicator.

For a set $C \subseteq \mathcal{G}$ and construction values $(z_j)_{j \in J}$, allowing repetitions, put

$$r_C(d) = |\{(c, c') \in C^2 : c \oplus c' = d\}|, \quad r_J(d) = |\{(j, k) \in J^2 : z_j \oplus z_k = d\}|.$$

Writing $F_{C,J}(v) = \sum_{j \in J} h_{C \oplus z_j}(v)$, we have

$$\|F_{C,J}\|_2^2 = \sum_{d \in \mathcal{G}} r_C(d)r_J(d) - \frac{|J|^2|C|^2}{N}. \quad (35)$$

Indeed, if $M(v) = \sum_{j \in J} \mathbf{1}_{C \oplus z_j}(v)$, then $\sum_v M(v)^2$ is the number of tuples (c, c', j, k) satisfying $c \oplus c' = z_j \oplus z_k$, and hence equals $\sum_d r_C(d)r_J(d)$. Centering M subtracts $|J|^2|C|^2/N$ from this squared sum.

Corollary 1 (Bound on matching XOR values for a subset of indices).

On the event of Lemma 13, let $C \subseteq \mathcal{G}$ have size $t \leq s$, and let J contain $h \leq s$ indices among the first s newly revealed construction values. Then

$$\sum_{d \in \mathcal{G}} r_C(d)r_J(d) \leq th + B_{\text{xor}}(s)t^2. \quad (36)$$

Proof. Separate the h diagonal index pairs at difference zero. Every off-diagonal difference count is at most $3s^2/N + 42n^2$, including at zero. Now use $r_C(0) = t$ and $\sum_d r_C(d) = t^2$. $\square$

I Construction Responses in Real

We compare the conditional distribution of a fresh construction response with ν_{Corr_r} in Hyb^{GSim} .

Let $N \geq 2^{20}$ and fix a query stage $0 \leq r \leq q \leq N/64$ in **Real**. Let $0 \leq s \leq q$ be a deterministic upper bound on the number S_r of inputs with revealed construction values, so $S_r \leq s$ on every path. This count is nondecreasing, so the same bound holds at every earlier state. We use the norms from Section 2.1. Put $h_C = \mathbf{1}_C - |C|/N$ for $C \subseteq \mathcal{G}$, $a_s := s/N^{3/2}$, and $b_s := s^2/N^2$. The empty case is immediate; below assume $s \geq 1$ and put $\eta = s/N$.

I.1 Marginal distributions of permutation outputs

We first control the marginal bias of unexposed permutation outputs; this will bound the drift terms when pairs are revealed.

Lemma 14 (Unexposed permutation outputs). *Suppose primitive queries have fixed t permutation pairs, and h other inputs have revealed construction values but unexposed permutation outputs, with $t + h \leq s$. At a construction-only input, conditional on the permutation pairs at all other such inputs, either output is uniform on a legal set of size at least $N - 2s$. Thus each value has conditional probability at most $(N - 2s)^{-1}$, also when only a subset of those pairs is fixed. The marginal distribution ρ of either permutation output satisfies*

$$\|\rho - \text{Unif}(\mathcal{G})\|_1 \leq 4s/N. \quad (37)$$

For a permutation output at a new input, the bound is $2s/N$. For the companion output of a fresh inverse query, without conditioning on which input it returns, the bound is $5s/N$. These assertions remain valid when further permutation pairs are exposed and removed from the list of unexposed pairs.

Proof. At a construction-only input x_i , conditional on the permutation pairs at the other such inputs, $\Pi_0(x_i)$ is uniform on its legal set D_i , where $|D_i| \geq N - 2s$. For any nonempty $D \subseteq \mathcal{G}$,

$$\|\text{Unif}(D) - \text{Unif}(\mathcal{G})\|_1 = 2 \left(1 - \frac{|D|}{N}\right).$$

Applying this identity to D_i and then averaging the conditional distributions proves (37). Translation gives the same bound for the other permutation output.

For an output at a new input, condition on the unexposed permutation outputs X at the construction-only inputs. Let Ran_b be the exposed range on side b and $X^{(b)}$ the set of these unexposed outputs on that side. The new output has distribution

$$\Pr[Y = y \mid X] = \frac{\mathbf{1}[y \notin \text{Ran}_b \cup X^{(b)}]}{N - t - h}.$$

The same identity now bounds its distance from uniform by $2(t+h)/N \leq 2s/N$, before and after averaging over X . For an inverse query, the target is fixed by the current prefix. Its probability of matching an unexposed output at a construction-only input on the queried side is at most $h/(N - 2s)$, by the conditional probability bound. Conditional on avoiding these outputs, every compatible assignment still leaves $N - t - h$ unused outputs on the opposite side, so the companion output has distance at most $2s/N$ from uniform. On the remaining event, use the general bound 2 for the distance between probability distributions. The unconditional companion distribution therefore satisfies

$$\|\rho - \text{Unif}(\mathcal{G})\|_1 \leq \frac{2s}{N} + \frac{2h}{N - 2s} < \frac{5s}{N}.$$

Exposing further coordinates only fixes more permutation entries; all conditional support bounds continue to hold with the same upper bound s on the total number of inputs. $\square$

Let $\mathcal{E}$ be the information immediately before a permutation output is exposed, and let Y be an output covered by the lemma. If C is $\mathcal{E}$ -measurable, $|C| = v$, and $V = h_{C \oplus Y}$, then

$$\mathbb{E}[\|V - \mathbb{E}[V \mid \mathcal{E}]\|_{2,\mathcal{G}}^2 \mid \mathcal{E}] \leq v, \quad \|\mathbb{E}[V \mid \mathcal{E}]\|_{2,\mathcal{G}} \leq 5(s/N)\sqrt{v}. \quad (38)$$

Indeed, every translate has squared norm $v - v^2/N$, and its expectation under uniform Y is zero. For two translates from one permutation pair, $\|v_1 + v_2\|_2^2 \leq 2\|v_1\|_2^2 + 2\|v_2\|_2^2$.

I.2 Decomposing the overlap of the occupied ranges

Let $\tau = \text{Tr}_{\text{Real}}^r$. Conditional on τ , let X be uniform on $\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})$. Its coordinates are the unexposed Π_0 -outputs at inputs in $\text{dom}(\text{Cons})$. For $u \in \text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})$, its joint distribution is

$$\Pr[\text{Tr}_{\text{Real}}^r = \tau, X = u] = \frac{\Pr_{\text{Real}}[\text{Tr}_{\text{Real}}^r = \tau]}{|\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})|}. \quad (39)$$

Let $\mathcal{F}_k$ contain the information revealed through step k . The first r steps follow the distinguisher's queries. During the next r steps, the remaining permutation pairs at inputs in $\text{dom}(\text{Cons})$ are revealed in construction-query order; unused steps reveal nothing. Each step uses the conditional distribution given $\mathcal{F}_{k-1}$. Let $(u_1, v_1), \dots, (u_S, v_S)$, with $S = S_r \leq s$, denote the pairs in this order. Put $\widetilde{\text{Ran}}_{0,\ell} = \{u_1, \dots, u_\ell\}$, $\widetilde{\text{Ran}}_{1,\ell} = \{v_1, \dots, v_\ell\}$, and $c_\ell(z) = |\widetilde{\text{Ran}}_{0,\ell} \cap (\widetilde{\text{Ran}}_{1,\ell} \oplus z)|$. The additional pairs counted by c_ℓ combine u_ℓ with each previous right output, each previous left output with v_ℓ , and u_ℓ with v_ℓ . Each pair contributes at its XOR value. Subtracting the mean gives

$$c_\ell^\circ - c_{\ell-1}^\circ = h_{u_\ell \oplus \widetilde{\text{Ran}}_{1,\ell-1}} + h_{v_\ell \oplus \widetilde{\text{Ran}}_{0,\ell-1}} + h_{\{u_\ell \oplus v_\ell\}}. \quad (40)$$

Conditional on $\widetilde{\text{Ran}}_{0,S}, \widetilde{\text{Ran}}_{1,S}$, the two outputs at a fresh input are independent and uniform on their respective unused ranges. There are N output pairs with XOR z . The two occupied ranges each exclude S of them, and $c_S(z)$ pairs are excluded by both ranges. Thus the distribution of the fresh construction value satisfies

$$p_X(z) = \frac{N - 2S + c_S(z)}{(N - S)^2}, \quad (41)$$

$$v_X(z) := Np_X(z) - 1 = \alpha_S c_S^\circ(z), \quad \alpha_S = \frac{N}{(N - S)^2}.$$

For every fresh inverse call j , let Forb_j be its set of forbidden construction values and $L_j = N - |\text{Forb}_j|$. One of the first two terms in (40) is h_{Forb_j} : it combines the query's target with the previous outputs on the other side. The sum of these terms is c_S^{tgt} . The other contributions from the first two terms sum to c_S^{rem} , and the terms $h_{\{u_\ell \oplus v_\ell\}}$ sum to c_S^{diag} . In particular, c_S^{rem} includes the contribution from the companion output of every fresh inverse query, regardless of which input it returns. Using the index set I_τ of the fresh inverse queries in τ , we have

$$c_S^\circ = c_S^{\text{tgt}} + c_S^{\text{rem}} + c_S^{\text{diag}}, \quad \text{Corr}_\tau^\circ = \sum_{j \in I_\tau} L_j^{-1} h_{\text{Forb}_j}.$$

The coefficient of each h_{Forb_j} in v_X is α_S , whereas the correction uses $1/L_j$. Both have leading term $1/N$ when s/N is small. We bound their difference along with the contributions of the sampled outputs.

The final number S of inputs with revealed construction values may not be known when an earlier query is answered. We therefore separate the coefficient α_S into the fixed part $1/N$ and the remaining part $\alpha_S - 1/N$, which gives

$$v_X - \text{Corr}_\tau^\circ = \sum_{j \in I_\tau} (N^{-1} - L_j^{-1}) h_{\text{Forb}_j} + N^{-1} (c_S^{\text{rem}} + c_S^{\text{diag}}) + (\alpha_S - N^{-1}) c_S^\circ. \quad (42)$$

I.3 Bounding the centered log-likelihood ratio

We bound the martingale and drift parts of c_S^{rem} and the other terms in this decomposition to control the centered log-likelihood ratio.

Lemma 15 (Centered log-likelihood ratio). *Given τ , let p_τ^{Real} be the response distribution for a fresh construction query in Real. Put*

$$\ell_\tau^{\text{res}} := (\log(Np_\tau^{\text{Real}}) - \text{Corr}_\tau)^\circ, \quad \varepsilon_\tau^{\text{res}} := \|\ell_\tau^{\text{res}}\|_{2,\mathcal{G}}/\sqrt{N}.$$

Centering removes the normalization constant, so ℓ_τ^{res} is also the centered log-likelihood ratio of p_τ^{Real} to ν_{Corr_τ} . Then

$$\sqrt{\frac{1}{N} \mathbb{E}[\|\ell_\tau^{\text{res}}\|_{2,\mathcal{G}}^2]} \leq 3a_s + 2b_s. \quad (43)$$

Proof. Write $\|V\|_{L_2}^2 = \mathbb{E}[\|V\|_{2,\mathcal{G}}^2]$, where the expectation uses the joint distribution of the transcript and X in (39), and put $\psi = s^{5/2}/N^{5/2} = \sqrt{\eta} b_s$. Let V_k be the sum of the one or two translates contributing to c_S^{rem} at step k , or zero if no pair is exposed. In $c_S^{\text{rem}} = c_S^{\text{rem},\text{mg}} + c_S^{\text{rem},\text{drift}}$, the martingale part and drift are

$$c_S^{\text{rem},\text{mg}} = \sum_{k=1}^{2r} (V_k - \mathbb{E}[V_k \mid \mathcal{F}_{k-1}]), \quad c_S^{\text{rem},\text{drift}} = \sum_{k=1}^{2r} \mathbb{E}[V_k \mid \mathcal{F}_{k-1}].$$

Since N^{-1} is deterministic, martingale orthogonality and (38), together with $\sum_{j=0}^{s-1} \sqrt{j} \leq (2/3)s^{3/2}$, give

$$\frac{1}{\sqrt{N}} \|N^{-1} c_S^{\text{rem},\text{mg}}\|_{L_2} \leq \sqrt{2} a_s, \quad \frac{1}{\sqrt{N}} \|N^{-1} c_S^{\text{rem},\text{drift}}\|_{L_2} \leq 7\psi.$$

The martingale bound sums conditional variances of at most $4j$ when j pairs have already been exposed. The drift bound sums at most two conditional means per pair, each of norm at most $5(s/N)\sqrt{j}$. For inverse queries, the moments are taken over the companion-output distribution without conditioning on the inverse output. The remaining components satisfy, for every transcript and compatible assignment X ,

$$\begin{aligned} \frac{1}{\sqrt{N}} \|N^{-1} c_S^{\text{diag}}\|_2 &\leq a_s, \\ \frac{1}{\sqrt{N}} \left\| \sum_{j \in I_\tau} (N^{-1} - L_j^{-1}) h_{\text{Forb}_j} \right\|_2 &\leq \frac{\psi}{1-\eta}, \\ \frac{1}{\sqrt{N}} \|(\alpha_S - N^{-1}) c_S^\circ\|_2 &\leq \frac{2-\eta}{(1-\eta)^2} \psi. \end{aligned}$$

For the first, n_z is the number of indices ℓ with $u_\ell \oplus v_\ell = z$. Then

$$\|c_S^{\text{diag}}\|_2^2 \leq \sum_z n_z^2 \leq s^2.$$

For the second, there are at most s calls, $L_j \geq N - s$, and $\|h_{\text{Forb}_j}\|_2 \leq \sqrt{s}$. For the third, use $\|c_S^\circ\|_2^2 \leq S^3$ and $\alpha_S - N^{-1} \leq (2 - \eta)s/[N^2(1 - \eta)^2]$.

Averaging over X conditional on a fixed prefix gives

$$p_\tau^{\text{Real}} = \mathbb{E}[p_X \mid \tau], \quad v_\tau = N p_\tau^{\text{Real}} - 1 = \mathbb{E}[v_X \mid \tau].$$

Since Corr_τ is determined by τ , Jensen's inequality gives

$$\|v_\tau - \text{Corr}_\tau^\circ\|_2^2 \leq \mathbb{E}[\|v_X - \text{Corr}_\tau^\circ\|_2^2 \mid \tau].$$

Averaging over transcripts in Real transfers the second-moment bounds for the terms in (42) to $v_\tau - \text{Corr}_\tau^\circ$. Moreover, (41) gives

$$\|v_\tau\|_\infty \leq \frac{\eta}{(1 - \eta)^2}, \quad \frac{1}{\sqrt{N}} \|v_\tau\|_2 \leq \frac{s^{3/2}}{N^{3/2}(1 - \eta)^2}.$$

The inequality $|\log(1 + x) - x| \leq x^2/[2(1 - |x|)]$ and the contraction of centering imply

$$\frac{1}{\sqrt{N}} \|(\log(1 + v_\tau) - v_\tau)^\circ\|_2 \leq \frac{\psi}{2[1 - \eta/(1 - \eta)^2](1 - \eta)^4}.$$

Combining these bounds, the coefficient of a_s is at most $\sqrt{2} + 1 < 3$. The coefficient of b_s is at most

$$\sqrt{\eta} \left[7 + \frac{1}{1 - \eta} + \frac{2 - \eta}{(1 - \eta)^2} + \frac{1}{2[1 - \eta/(1 - \eta)^2](1 - \eta)^4} \right] \Big|_{\eta=1/64} < 2.$$

All summands increase with $\eta \in [0, 1/64]$. $\square$

Reweightings bounds the KL divergence between p_τ^{Real} and ν_{Corr_τ} ; the second-moment bound then controls its expectation.

Corollary 2 (KL bound for a fresh construction query). *For every compatible prefix,*

$$\sqrt{D_{\text{KL}}(p_\tau^{\text{Real}} \parallel \nu_{\text{Corr}_\tau})} \leq \frac{2}{\sqrt{N}} \|\ell_\tau^{\text{res}}\|_{2, \mathcal{G}}. \quad (44)$$

Consequently, averaging over $\tau = \text{Tr}_{\text{Real}}^r$,

$$\sqrt{\mathbb{E}[D_{\text{KL}}(p_\tau^{\text{Real}} \parallel \nu_{\text{Corr}_\tau})]} \leq 2(3a_s + 2b_s).$$

Proof. By the definition of ℓ_τ^{res} and normalization,

$$p_\tau^{\text{Real}}(z) = \frac{\nu_{\text{Corr}_\tau}(z) e^{\ell_\tau^{\text{res}}(z)}}{\mathbb{E}_{\nu_{\text{Corr}_\tau}}[e^{\ell_\tau^{\text{res}}}]}.$$

Since $0 \leq c_S(z) \leq S$ and $0 \leq \text{Corr}_\tau(z) \leq s/(N-s)$,

$$\text{osc}(\ell_\tau^{\text{res}}) \leq \log \frac{N-s}{N-2s} + \frac{s}{N-s} < 0.032, \quad \frac{\nu_{\text{Corr}_\tau}}{\text{Unif}(\mathcal{G})} \leq e^{1/63}.$$

Apply (3) with $\text{Var}_{\nu_{\text{Corr}}}(\ell_\tau^{\text{res}}) \leq (e^{1/63}/N)\|\ell_\tau^{\text{res}}\|_{2,\mathcal{G}}^2$, then use $D_{\text{KL}} \leq \chi^2$. The resulting multiplier $e^{0.032+1/126}$ is less than 2. Average over transcripts and apply Lemma 15. $\square$

J Permutation Outputs in Real

Fix a compatible augmented prefix with state $\text{Ran}_0, \text{Ran}_1, \text{Cons}$. Put $t := |\text{Ran}_0| = |\text{Ran}_1|$, $h := |\text{Cons}|$, and $s := t + h$. Write the construction-only list as $\text{Cons} = ((x_j, z_j))_{j \in I}$, where $I = \{1, \dots, h\}$. Assume $N \geq 2^{20}$ and $1 \leq s \leq q \leq N/64$. Put $\eta = s/N$, $a_s = s/N^{3/2}$, and $b_s = s^2/N^2$. For $C \subseteq \mathcal{G}$, write $h_C = \mathbf{1}_C - |C|/N$. The identities below hold at every compatible prefix. Bounds on squared sums additionally assume the event in Lemma 13, with

$$B_{\text{occ}} = 2n, \quad B_{\text{xor}}(s) = 3s^2/N + 42n^2.$$

These bounds remain valid after deleting indices, with $B_{\text{xor}}(s)$ computed using the original value of s .

J.1 Conditional output distributions

For $J \subseteq I$ put

$$E_j = \mathcal{G} \setminus (\text{Ran}_0 \cup (\text{Ran}_1 \oplus z_j)), \quad L_j = |E_j| \geq N - 2t,$$

and let

$$\text{Comp}_J = \left\{ u \in \prod_{j \in J} E_j : \begin{array}{l} u_j \neq u_k, \\ u_j \oplus z_j \neq u_k \oplus z_k \end{array} (j \neq k) \right\}.$$

Write $\mu_J = \text{Unif}(\text{Comp}_J)$, let $X = (X_j)_{j \in J} \leftarrow \mu_J$, and put $L_*(J) = N - 2t - 2(|J| - 1) \geq N - 2s$.

Lemma 16 (Conditional and marginal output distributions). *Under μ_J , conditional on all other coordinates, X_j is uniform on*

$$D_j(u_{-j}) = E_j \setminus \bigcup_{k \neq j} \{u_k, u_k \oplus z_k \oplus z_j\}, \quad |D_j(u_{-j})| \geq L_*(J). \quad (45)$$

After fixing any subset of the other coordinates, each value of X_j therefore has conditional probability at most $L_*(J)^{-1}$. Moreover,

$$\left| \Pr_{\mu_J}[X_j = c] - L_j^{-1} \right| \leq \frac{2(|J| - 1)}{L_*(J)L_j} \quad (c \in E_j), \quad (46)$$

$$\Pr_{\mu_J}[X_j = c, X_k = d] \leq L_*(J)^{-2} \quad (j \neq k). \quad (47)$$

Proof. Conditional on all other coordinates, every value in (45) determines one compatible assignment. Averaging gives the conditional probability bound, and applying it twice gives (47). Put $m = |J| - 1$ and $\ell = L_*(J)$. For $c \in E_j$, at most $2m$ collision events can remove c from D_j , each with probability at most $1/\ell$. Also $L_j - 2m \leq |D_j| \leq L_j$. Hence

$$\frac{1 - 2m/\ell}{L_j} \leq \Pr[X_j = c] = \mathbb{E} \left[\frac{\mathbf{1}[c \in D_j]}{|D_j|} \right] \leq \frac{1}{L_j - 2m}.$$

Since $L_j - 2m \geq \ell$, the distance from $1/L_j$ is at most $2m/(\ell L_j)$. Repeated construction values can merge collision events and do not increase their number. $\square$

J.2 Forward response at a construction-only input

Fix $i \in I$, let $J = I \setminus \{i\}$, and take $a \in E_i$. Put $d_{ij} = z_i \oplus z_j$, $\mathcal{B}_j(a) = \{X_j = a\} \cup \{X_j = a \oplus d_{ij}\}$, and

$$W_i(a) = \Pr_{\mu_J} \left[\bigcap_{j \in J} \mathcal{B}_j(a)^c \right].$$

Let R_i^{fwd} be the conditional distribution of the Π_0 -output at input x_i . Inserting $(a, a \oplus z_i)$ into an assignment in Comp_J gives

$$\frac{R_i^{\text{fwd}}(a)}{U_i(a)} = \frac{W_i(a)}{\mathbb{E}_{b \leftarrow U_i}[W_i(b)]}, \quad U_i = \text{Unif}(E_i). \quad (48)$$

For $Y \leftarrow \text{Unif}(E_j)$ define $q_{ij}(a) = \Pr[Y = a \text{ or } Y = a \oplus d_{ij}]$.

Lemma 17 (Logarithm of a collision avoidance probability). *Uniformly over legal a ,*

$$\left| \log W_i(a) + \sum_{j \in J} q_{ij}(a) \right| \leq \frac{10h^2}{(N - 2s)^2}.$$

If each collision event specifies a single value of X_j , the bound is $4h^2/\ell^2$, provided both $|E_j|$ and every full conditional support size $|D_j|$ are at least $\ell \geq N - 2s$.

Proof. Let $\beta_j = \Pr_{\mu_J}[\mathcal{B}_j(a)]$ and $p = \Pr_{\mu_J}[\cup_j \mathcal{B}_j(a)]$. Equation (46) gives total marginal error $\sum_j |\beta_j - q_{ij}| \leq 4h^2/(N - 2s)^2$. For $j \neq k$, the intersection of $\mathcal{B}_j(a)$ and $\mathcal{B}_k(a)$ specifies at most four possible pairs (X_j, X_k) , so $0 \leq \sum_j \beta_j - p \leq 2h^2/(N - 2s)^2$. Finally $p \leq 2h/(N - 2s) < 1/16$ and $|\log(1 - p) + p| \leq p^2 \leq 4h^2/(N - 2s)^2$. The three coefficients sum to 10. With one possible value per event they become 2, $1/2$, and 1, whose sum is less than 4. $\square$

Membership in E_j gives, when $d_{ij} \neq 0$,

$$q_{ij}(a) = \frac{2 - \mathbf{1}_{\text{Ran}_1 \oplus z_j}(a) - \mathbf{1}_{\text{Ran}_0 \oplus z_i \oplus z_j}(a)}{L_j}.$$

When the construction values repeat, the two specified values coincide:

$$d_{ij} = 0 \implies q_{ij}(a) = L_i^{-1}.$$

Define

$$T_i(a) = \frac{1}{N} \sum_{j \in J} (h_{\text{Ran}_1 \oplus z_j}(a) + h_{\text{Ran}_0 \oplus z_i \oplus z_j}(a)), \quad \text{Ctr}_i f = f - \mathbb{E}_{U_i}[f].$$

A summand for repeated construction values is constant on E_i and vanishes after Ctr_i . Since $|L_j^{-1} - N^{-1}| \leq 2t/[N(N - 2t)]$,

$$\text{Ctr}_i \log W_i = \text{Ctr}_i T_i + \mathcal{R}_i, \quad \|\mathcal{R}_i\|_\infty \leq 26b_s.$$

Centering bounds the supremum norm of the logarithmic remainder by $20b_s/(1 - 2\eta)^2$. The nonconstant error from replacing L_j^{-1} by N^{-1} is a sum of at most two indicators per index, with nonnegative weights. Using $\max_{a \in E_i} |\text{Ctr}_i f(a)| \leq \text{osc}_{E_i}(f)$, its centered supremum is at most $4b_s/(1 - 2\eta)$. The two coefficients sum to less than 26 for $\eta \leq 1/64$.

In the notation of (35), $T_i = (F_{\text{Ran}_1, J} + F_{\text{Ran}_0, J}(\cdot \oplus z_i))/N$. Translation preserves the ℓ_2 norm. Thus (35) and (36) give

$$\frac{1}{\sqrt{N}} \|T_i\|_{2, \mathcal{G}} \leq 2\sqrt{1 + 42n^2} a_s + 2\sqrt{3} b_s.$$

Since $N/|E_i| \leq 32/31$, $\text{Var}_{U_i}(T_i) \leq \frac{32}{31N} \|T_i\|_{2, \mathcal{G}}^2$. The union bound gives $W_i \geq 1 - 2h/(N - 2s) \geq 30/31$. Since $W_i \leq 1$, we have $\text{osc}_{E_i}(\log W_i) \leq \log(31/30)$. Applying (3) with $\|\mathcal{R}_i\|_\infty \leq 26b_s$ and substituting the preceding bounds yields

$$\begin{aligned} \sqrt{\chi^2(R_i^{\text{fwd}} \|U_i)} &\leq \frac{31}{30} \sqrt{\frac{32}{31}} 2\sqrt{1 + 42n^2} a_s \\ &\quad + \frac{31}{30} \left(2\sqrt{\frac{96}{31}} + 26 \right) b_s \\ &\leq 4n^{3/2} a_s + 7\sqrt{n} b_s. \end{aligned} \tag{49}$$

J.3 Companion output at a new input

Fix an inverse query $\text{Inv}_0(y)$ with $y \notin \text{Ran}_0$, and let $X \leftarrow \mu_I$. Conditional on y avoiding the unexposed Π_0 -outputs, X is uniform over the compatible assignments with supports

$$E'_j = \mathcal{G} \setminus (\text{Ran}_0 \cup \{y\} \cup (\text{Ran}_1 \oplus z_j)).$$

Their full conditional supports have size at least $N - 2s + 1$. Conditioning further on a particular inverse output with unrevealed construction value leaves this distribution unchanged: each compatible assignment assigns probability $1/(N - s)$ to that return. For $v \notin \text{Ran}_1$ define the events $\mathcal{B}_j^{\text{opp}}(v) = \{X_j \oplus z_j = v\}$, and let $W^{\text{opp}}(v)$ be the probability that none occurs. Let R^{opp} be the companion-output distribution under this conditioning. Counting the compatible assignments gives

$$\frac{R^{\text{opp}}(v)}{U_{\text{Ran}_1}(v)} = \frac{W^{\text{opp}}(v)}{\mathbb{E}_{w \leftarrow U_{\text{Ran}_1}}[W^{\text{opp}}(w)]}, \quad U_{\text{Ran}_1} = \text{Unif}(\mathcal{G} \setminus \text{Ran}_1). \quad (50)$$

The corresponding collision probability under $\text{Unif}(E'_j)$ is

$$q_j^{\text{opp}}(v) = \frac{1 - \mathbf{1}_{\text{Ran}_0 \oplus z_j}(v) - \mathbf{1}_{\{y \oplus z_j\}}(v)}{|E'_j|}.$$

Define the centered sum

$$T^{\text{opp}}(v) = \frac{1}{M} \sum_{j \in I} (h_{\text{Ran}_0 \oplus z_j}(v) + h_{\{y \oplus z_j\}}(v)),$$

$$M = N - t, \quad \text{Ctr}_{\text{Ran}_1} f = f - \mathbb{E}_{U_{\text{Ran}_1}}[f].$$

The two indicator sets are disjoint because $y \notin \text{Ran}_0$. The single-value collision bound in Lemma 17 gives

$$\left| \log W^{\text{opp}}(v) + \sum_j q_j^{\text{opp}}(v) \right| \leq \frac{4h^2}{(N - 2s + 1)^2}.$$

Moreover

$$||E'_j|^{-1} - M^{-1}| \leq \frac{t + 1}{(N - 2t - 1)(N - t)}.$$

When $h \geq 1$, $h(t + 1) \leq s^2$. Centering increases the supremum bound on the logarithmic remainder by at most a factor of two. For the weighted indicator sum, use $\max_{v \notin \text{Ran}_1} |\text{Ctr}_{\text{Ran}_1} f(v)| \leq \text{osc}_{\mathcal{G} \setminus \text{Ran}_1}(f)$. Consequently

$$\text{Ctr}_{\text{Ran}_1} \log W^{\text{opp}} = \text{Ctr}_{\text{Ran}_1} T^{\text{opp}} + \mathcal{R}^{\text{opp}}, \quad \|\mathcal{R}^{\text{opp}}\|_{\infty} \leq 10b_s,$$

because $8/(1 - 2\eta)^2 + [(1 - 2\eta - 1/N)(1 - \eta)]^{-1} < 10$.

Here $T^{\text{opp}} = F_{\text{Ran}_0 \cup \{y\}, I}/M$. Since $t + 1 \leq s$ when $h \geq 1$, applying (35) and (36) with $C = \text{Ran}_0 \cup \{y\}$ yields

$$\frac{1}{\sqrt{N}} \|T^{\text{opp}}\|_{2, \mathcal{G}} \leq \frac{64}{63} (\sqrt{1 + 42n^2} a_s + \sqrt{3} b_s).$$

Since $N/M \leq 64/63$, $\text{Var}_{U_{\text{Ran}_1}}(T^{\text{opp}}) \leq \frac{64}{63N} \|T^{\text{opp}}\|_{2, \mathcal{G}}^2$. The union bounds give, without multiplicity or collision assumptions,

$$30/31 \leq W_i \leq 1, \quad 61/62 \leq W^{\text{opp}} \leq 1, \quad (51)$$

so $\text{osc}_{G \setminus \text{Ran}_1}(\log W^{\text{opp}}) \leq \log(62/61)$. Applying (3) with $\|\mathcal{R}^{\text{opp}}\|_\infty \leq 10b_s$ gives

$$\begin{aligned} \sqrt{\chi^2(R^{\text{opp}}\|U_{\text{Ran}_1})} &\leq \frac{62}{61} \left(\sqrt{\frac{64}{63N}} \|T^{\text{opp}}\|_{2,\mathcal{G}} + 10b_s \right) \\ &\leq \frac{62}{61} \left(\frac{64}{63} \right)^{3/2} \sqrt{1 + 42n^2} a_s \\ &\quad + \frac{62}{61} \left(\sqrt{3} \left(\frac{64}{63} \right)^{3/2} + 10 \right) b_s \\ &\leq 2n^{3/2} a_s + 3\sqrt{n} b_s. \end{aligned} \tag{52}$$

If $h = 0$, the sum, remainder and divergence all vanish.

J.4 Collision avoidance for an inverse query

Fix $y \notin \text{Ran}_0$. A construction-only input x_i is legal when $\iota_i = \mathbf{1}[y \oplus z_i \notin \text{Ran}_1] = 1$. Returning x_i requires the pair $(y, y \oplus z_i)$ to be compatible with the other unexposed pairs, whereas returning an input with an unrevealed construction value requires y to avoid the unexposed Π_0 -outputs. We compare these two probabilities; Appendix K uses their ratio to compare the distributions of the inverse output. Define

$$\pi_i(y) = \Pr_{\mu_{I \setminus \{i\}}} [X_j \neq y, X_j \oplus z_j \neq y \oplus z_i \text{ for all } j \neq i], \tag{53}$$

$$\pi_0(y) = \Pr_{\mu_I} [X_j \neq y \text{ for all } j]. \tag{54}$$

Put $v = N - s$, $M = N - t$, $d = v + 1$, and, for legal i ,

$$A_i(y) = \log \pi_i(y) - \log \pi_0(y) + \log \frac{vM}{d^2}. \tag{55}$$

Let $J_i^\neq = \{j \neq i : z_j \neq z_i\}$, $m_i = |\{j \neq i : z_j = z_i\}|$, and

$$F_i(y) = \frac{1}{N} \sum_{j \in J_i^\neq} h_{\text{Ran}_0 \oplus z_i \oplus z_j}(y).$$

Lemma 18 (Comparison of collision avoidance probabilities). *For every legal i ,*

$$A_i(y) = F_i(y) + m_i/N + \epsilon_i(y), \quad |\epsilon_i(y)| \leq 23b_s. \tag{56}$$

Proof. Apply the two bounds of Lemma 17 to π_i and π_0 , respectively, and subtract. Their total error is at most $14h^2/(N-2s)^2$. For $j \in J_i^\neq$, the terms involving Ran_1 cancel:

$$-q_{ij}(y) + q_{0j}(y) = [-1 + \mathbf{1}_{\text{Ran}_0 \oplus z_i \oplus z_j}(y)]/L_j, \quad q_{0j}(y) = \mathbf{1}[y \in E_j]/L_j.$$

For repeated construction values, $q_{ij}(y) = q_{0j}(y) = 1/L_i$, and the omitted index contributes $q_{0i}(y) = 1/L_i$. Therefore

$$A_i(y) = \log \frac{vM}{d^2} + \frac{1}{L_i} - \sum_{j \in J_i^\neq} \frac{1}{L_j} + \sum_{j \in J_i^\neq} \frac{\mathbf{1}_{\text{Ran}_0 \oplus z_i \oplus z_j}(y)}{L_j} + e_i,$$

$$|e_i| \leq \frac{14b_s}{(1-2\eta)^2}.$$

The power series for the three logarithms in $\log(vM/d^2)$ gives

$$\left| \log \frac{vM}{d^2} - \frac{h-2}{N} \right| \leq \frac{2b_s}{1-\eta}.$$

Replacing each remaining reciprocal by $1/N$ changes the expression by at most $4b_s/(1-2\eta)$. Its constant part is $(2-h+m_i)/N$, so the leading constants cancel and leave m_i/N . Finally, changing the indicator sum into F_i changes it by $|J_i^\neq|t/N^2 \leq b_s$. The sum of the four error coefficients is

$$\frac{14}{(1-2\eta)^2} + \frac{2}{1-\eta} + \frac{4}{1-2\eta} + 1 < 23.$$

The term m_i/N accounts for repeated construction values, for which the two collision events coincide. $\square$

We next bound the sum of $A_i(y)^2$ over legal construction-only inputs.

Lemma 19 (Bound on the sum of $A_i(y)^2$). *On the event of Lemma 13, the following bound holds simultaneously for every $y \notin \text{Ran}_0$:*

$$\left(\frac{1}{N} \sum_{i:\iota_i=1} A_i(y)^2 \right)^{1/2} \leq 10n^{3/2}a_s + 4\sqrt{n}b_s. \quad (57)$$

Proof. Let $I_y = \{i \in I : y \oplus z_i \notin \text{Ran}_1\}$. We bound the sum over I_y using sums of F_i^2 and m_i^2 over all of I . Put $C_y = y \oplus \text{Ran}_0$ and $\deg_i = |\{j \in J_i^\neq : z_i \oplus z_j \in C_y\}|$. Then $F_i = (\deg_i - |J_i^\neq|t/N)/N$. The multiplicity and collision bounds give

$$\max_i \deg_i \leq tB_{\text{occ}}, \quad \sum_i \deg_i \leq tB_{\text{xor}}(s), \quad \sum_i \deg_i^2 \leq B_{\text{occ}}B_{\text{xor}}(s)t^2.$$

Thus

$$\|F\|_{2,I} \leq \frac{\sqrt{B_{\text{occ}}B_{\text{xor}}(s)}t}{N} + \frac{h^{3/2}t}{N^2}, \quad \|(m_i/N)_i\|_{2,I} \leq \frac{\sqrt{B_{\text{occ}}B_{\text{xor}}(s)}}{N}.$$

For the second inequality, use $\max_i m_i \leq B_{\text{occ}}$ and $\sum_i m_i \leq B_{\text{xor}}(s)$, with the diagonal h removed from the zero-difference count. Since $t+1 \leq s$ when $h \geq 1$,

(56) gives

$$\begin{aligned} \left(\frac{1}{N} \sum_{i \in I_y} A_i(y)^2 \right)^{1/2} &\leq 10n^{3/2}a_s + (3\sqrt{n} + 3)b_s \\ &\leq 10n^{3/2}a_s + 4\sqrt{n}b_s. \end{aligned}$$

Here $h^{3/2}t/N^{5/2} \leq b_s/8$ and $23\sqrt{h/N}b_s \leq 23b_s/8$, since $\sqrt{s/N} \leq 1/8$. The final rounding uses $n \geq 20$. For $h = 0$ the sum is empty. $\square$

K Real to Hyb^{GSim}

This appendix proves the transcript comparison in Lemma 24. We combine the construction-response bounds in Appendix I with the bounds on permutation outputs in Appendix J to bound the divergence between the distributions of the next response. We first work at good prefixes and then include the contribution of bad cases.

Suppose $N \geq 2^{20}$ and $q \leq N/64$. Fix $r \in \{0, \dots, q-1\}$, and let $\tau = \text{Tr}_{\text{Real}}^r$ be the transcript after r queries. Let R_τ and K_τ^{GSim} be the conditional distributions of the response to query $r+1$ in Real and Hyb^{GSim}, respectively, and put $d_r(\tau) = D_{\text{KL}}(R_\tau \| K_\tau^{\text{GSim}})$.

We use the notation of Section 3.2, with

$$|\text{Ran}_0| = |\text{Ran}_1| = t, \quad \text{Cons} = ((x_i, z_i))_{i=1}^h, \quad s = t + h, \quad \eta = s/N \leq 1/64.$$

Put $a_s = s/N^{3/2}$ and $b_s = s^2/N^2$. Use the good prefixes defined in (19): among the first k newly revealed construction values, each value occurs at most $2n$ times and each XOR value occurs for at most $3k^2/N + 42n^2$ ordered pairs of distinct indices, for every $k \leq s$. We use $B_{\text{occ}} = 2n$. At bad prefixes, $d_r^{\text{good}} = 0$. At a good prefix, $d_r^{\text{good}} = d_r$ except for forward queries at $x \in E_{\text{new}}$. In that case, we retain the full KL term for the construction value. The conditional KL term for the permutation outputs includes only values that preserve the good-prefix conditions. Their probability weights are unchanged. The omitted contribution is $d_r^{\text{bad}} = d_r - d_r^{\text{good}}$. We first bound $d_r^{\text{good}}(\tau)$ at each prefix, then average over transcripts in Real.

At a good prefix, each construction value appears at most $2n$ times in Cons, so

$$\frac{1}{N} \sum_{(x,z) \in \text{Cons}} \ell_\tau^{\text{res}}(z)^2 \leq \frac{2n}{N} \sum_{z \in \mathcal{G}} \ell_\tau^{\text{res}}(z)^2 = 2n(\varepsilon_\tau^{\text{res}})^2.$$

We use the bounds in Appendix J on the probabilities that the required permutation outputs remain unused. Lemma 13 shows that all prefixes in Real are good with probability at least $1 - N^{-8}$.

The current state has size at most r . A forward query may reveal one additional construction value, so every state used below has size at most $r+1 \leq q$. Empty states are handled directly.

K.1 Removing one construction constraint

For a primitive query at an input with a revealed construction value, the response must satisfy that value as well as the permutation constraints. For an inverse query, the inverse output must also be chosen. The comparison therefore uses both the construction-response distribution and the probabilities that the required outputs are still available.

Write $\mu_{\text{Cons}} = \text{Unif}(\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons}))$ for the distribution of compatible assignments, and let $\text{Cons} - i$ be the list Cons with its i th constraint removed. For $X = (X_1, \dots, X_h) \leftarrow \mu_{\text{Cons}}$, overlap counting gives

$$\begin{aligned} d_X(z) &= N - 2s + c_X(z), \\ c_X(z) &= |(\text{Ran}_0 \cup \{X_i : 1 \leq i \leq h\}) \\ &\quad \cap ((\text{Ran}_1 \cup \{X_i \oplus z_i : 1 \leq i \leq h\}) \oplus z)|. \end{aligned}$$

There are $(N - s)^2$ ordered pairs of unused permutation outputs, so

$$p_X(z) = \frac{d_X(z)}{(N - s)^2}, \quad p_{\text{Cons}}^{\text{Real}}(z) = \mathbb{E}_{X \leftarrow \mu_{\text{Cons}}} [p_X(z)].$$

For the current list, $p_{\text{Cons}}^{\text{Real}} = p_{\tau}^{\text{Real}}$. Since $0 \leq c_X(z) \leq s$,

$$\frac{N - 2s}{(N - s)^2} \leq p_{\text{Cons}}^{\text{Real}}(z) \leq \frac{1}{N - s}. \quad (58)$$

Define the log-likelihood ratios

$$\begin{aligned} \ell_{\text{Cons}} &= \log(p_{\text{Cons}}^{\text{Real}} / \nu_{\text{Corr}}), & \ell_{-i} &= \log(p_{\text{Cons} - i}^{\text{Real}} / \nu_{\text{Corr}}), \\ \ell_{\text{Cons}}^{\text{res}} &= \ell_{\tau}^{\text{res}} = \ell_{\text{Cons}}^{\circ} = (\log(N p_{\text{Cons}}^{\text{Real}}) - \text{Corr})^{\circ}. \end{aligned}$$

The function Corr and the distribution $\nu_{\text{Corr}} = e^{\text{Corr}} / Z_{\text{Corr}}$ remain fixed when the constraint is removed.

Fix a fresh inverse query $\text{Inv}_0(y)$ with $y \notin \text{Ran}_0$; the other side is symmetric. For a construction-only input x_i with $y \oplus z_i \notin \text{Ran}_1$, let r_i be the probability of returning x_i in Real , and set $d = N - s + 1$. Under $\mu_{\text{Cons} - i}$, let $\pi_i(y)$ be the probability that no remaining pair uses left output y or right output $y \oplus z_i$. Then

$$r_i = \frac{\pi_i(y)}{d^2 p_{\text{Cons} - i}^{\text{Real}}(z_i)}.$$

The denominator is the expected number of legal pairs with XOR z_i under $\mu_{\text{Cons} - i}$; the numerator is the probability that the particular pair $(y, y \oplus z_i)$ is legal. Indeed, let C_{-i} denote the number of compatible assignments for $\text{Cons} - i$. The count with left output y is $C_{-i} \pi_i(y)$, while the number of compatible assignments for Cons is

$$C_{-i} d^2 p_{\text{Cons} - i}^{\text{Real}}(z_i).$$

A forward query to Π_0 at x_i returns y exactly when the assigned pair is $(y, y \oplus z_i)$. Counting the compatible assignments containing this pair gives the response distribution in (48).

The following lemma bounds the change in the distribution of a fresh construction response when one constraint is removed. Together with the availability bounds in Appendix J, it connects the distribution of the inverse output to $\varepsilon_\tau^{\text{res}}$.

Lemma 20 (Removing one construction constraint). *For every compatible Cons and $1 \leq i \leq h$,*

$$\begin{aligned} \|p_{\text{Cons}}^{\text{Real}} - p_{\text{Cons}-i}^{\text{Real}}\|_\infty &\leq \frac{2}{(N-s)^2} + \frac{s^2}{4(N-s)^2(N-2s)}, \\ \|\ell_{-i} - \ell_{\text{Cons}}\|_\infty &\leq \frac{3}{N} + \frac{s^2}{N^2}. \end{aligned} \quad (59)$$

Proof. Write $m = s-1$ and $d = N-m$, and sample $X \leftarrow \mu_{\text{Cons}-i}$. Let $p_X^{(m)}$ be the response distribution for a fresh construction query, given X and before inserting the removed pair. Let $\mathcal{L}_i(X)$ be the set of first outputs b for which $(b, b \oplus z_i)$ is a legal insertion. Then $N - 2m \leq |\mathcal{L}_i(X)| \leq N - m$ and $\text{osc}_X |\mathcal{L}_i(X)| \leq m$. For $b \in \mathcal{L}_i(X)$, let $p_{X,b}^{(m+1)}$ be the corresponding response distribution after fixing $(b, b \oplus z_i)$. Since μ_{Cons} is uniform on $\text{Comp}(\text{Ran}_0, \text{Ran}_1, \text{Cons})$, we obtain

$$p_{\text{Cons}}^{\text{Real}}(w) = \frac{\mathbb{E}_X[|\mathcal{L}_i(X)| \bar{p}_{X,i}(w)]}{\mathbb{E}_X[|\mathcal{L}_i(X)|]}, \quad \bar{p}_{X,i}(w) = \frac{1}{|\mathcal{L}_i(X)|} \sum_{b \in \mathcal{L}_i(X)} p_{X,b}^{(m+1)}(w). \quad (60)$$

Fixing the additional pair removes at most two legal output pairs with XOR w . If their original number is $c \leq d$, the change in the probability of w is bounded by

$$\left| \frac{c - \delta}{(d-1)^2} - \frac{c}{d^2} \right| = \frac{|(2d-1)c/d^2 - \delta|}{(d-1)^2} \leq \frac{2}{(N-s)^2}, \quad 0 \leq \delta \leq 2.$$

Also $\text{osc}_X p_X^{(m)}(w) \leq m/d^2$. Decomposing (60) into the change caused by fixing the pair and the covariance of $|\mathcal{L}_i(X)|$ with $p_X^{(m)}$, and using $|\text{Cov}(U, V)| \leq \text{osc}(U) \text{osc}(V)/4$, proves (59). Both distributions assign probability at least $(N-2s)/(N-s)^2$ to every value. The common factor ν_{Corr} cancels in the log-ratio difference. The inequality $|\log a - \log b| \leq |a - b|/\min\{a, b\}$ for $a, b > 0$ gives

$$\|\ell_{-i} - \ell_{\text{Cons}}\|_\infty \leq \frac{2}{N-2s} + \frac{s^2}{4(N-2s)^2} \leq \frac{3}{N} + \frac{s^2}{N^2}.$$

The last rounding uses $\eta \leq 1/64$. □

K.2 Log-likelihood ratios at revealed construction values

Normalization gives $p_{\text{Cons}}^{\text{Real}} = \nu_{\text{Corr}} e^{\ell_{\text{Cons}}^{\text{res}}}/Z$ and $\ell_{\text{Cons}} = \ell_{\text{Cons}}^{\text{res}} - \log Z$, where $Z = \sum_z \nu_{\text{Corr}}(z) e^{\ell_{\text{Cons}}^{\text{res}}(z)}$. Jensen's inequality, applied to $e^{\ell_{\text{Cons}}^{\text{res}}}$ under ν_{Corr} and to $e^{-\ell_{\text{Cons}}^{\text{res}}}$ under $p_{\text{Cons}}^{\text{Real}}$, gives

$$\mathbb{E}_{\nu_{\text{Corr}}}[\ell_{\text{Cons}}^{\text{res}}] \leq \log Z \leq \mathbb{E}_{p_{\text{Cons}}^{\text{Real}}}[\ell_{\text{Cons}}^{\text{res}}].$$

Both densities relative to $\text{Unif}(\mathcal{G})$ are less than 1.1: $N\nu_{\text{Corr}}(z) \leq e^{1/63} < 1.1$ and $Np_{\text{Cons}}^{\text{Real}}(z) \leq 64/63 < 1.1$. Thus Cauchy–Schwarz gives

$$|\log Z| \leq 1.1\mathbb{E}_{\text{Unif}(\mathcal{G})}[|\ell_{\text{Cons}}^{\text{res}}|] \leq 1.1\varepsilon_{\tau}^{\text{res}}.$$

Since $\mathbb{E}_{\text{Unif}(\mathcal{G})}[\ell_{\text{Cons}}^{\text{res}}] = 0$, it follows that

$$\frac{1}{N}\|\ell_{\text{Cons}}\|_{2,\mathcal{G}}^2 = (\varepsilon_{\tau}^{\text{res}})^2 + (\log Z)^2 \leq (1 + 1.1^2)(\varepsilon_{\tau}^{\text{res}})^2 \leq 1.5^2(\varepsilon_{\tau}^{\text{res}})^2.$$

Let $I_y = \{i \in \{1, \dots, h\} : y \oplus z_i \notin \text{Ran}_1\}$. At a good prefix, the multiplicity bound gives

$$\frac{1}{N} \sum_{i=1}^h \ell_{\text{Cons}}(z_i)^2 \leq \frac{B_{\text{occ}}}{N} \|\ell_{\text{Cons}}\|_{2,\mathcal{G}}^2.$$

The stability lemma, $\sqrt{h} \leq s$, and $\sqrt{h/N} \leq 1/8$ imply

$$\frac{1}{\sqrt{N}} \|(\ell_{-i}(z_i))_i\|_{2,I_y} \leq 1.5\sqrt{2n}\varepsilon_{\tau}^{\text{res}} + 3a_s + b_s. \quad (61)$$

K.3 Distribution of the inverse output

Put $\iota_i = \mathbf{1}[y \oplus z_i \notin \text{Ran}_1]$, $M = N - t$, and $u = N - s$. The probabilities of the inverse output in Hyb^{GSim} are given by

$$\begin{aligned} Z_{\text{Corr}} &= \sum_z e^{\text{Corr}(z)}, & \alpha_{\text{Corr}} &= M/Z_{\text{Corr}}, \\ Z_{\text{ret}} &= \sum_i \iota_i e^{-\text{Corr}(z_i)} + u\alpha_{\text{Corr}}, & k_i &= \frac{\iota_i e^{-\text{Corr}(z_i)}}{Z_{\text{ret}}}, & u_0 &= \frac{\alpha_{\text{Corr}}}{Z_{\text{ret}}}. \end{aligned}$$

Here k_i is the mass of input x_i , whose construction value is revealed, and u_0 is the mass of each input with an unrevealed construction value. Let $\pi_0(y)$ be the probability under μ_{Cons} that no unexposed left output equals y . By exchangeability, the probability that a fixed input with an unrevealed construction value is the inverse output in **Real** is

$$r_0 = \pi_0(y)/u.$$

Indeed, conditional on y avoiding the unexposed Π_0 -outputs, the inverse output is uniform among the u inputs with unrevealed construction values. Set $k_f = u u_0$ and $r_f = u r_0$. For $i \in I_y$, define

$$b_i = \log \frac{r_i/r_f}{k_i/k_f}.$$

Substituting the identity for r_i and the weights in Hyb^{GSim} gives

$$b_i = A_i - \ell_{-i}(z_i), \quad A_i = \log \pi_i - \log \pi_0 + \log \frac{uM}{(u+1)^2}, \quad (62)$$

Here $e^{b_i} = r_i \alpha_{\text{Corr}} / (r_0 e^{-\text{Corr}(z_i)})$ and $\alpha_{\text{Corr}} = M/Z_{\text{Corr}}$; the normalizing factor Z_{Corr} is included in ℓ_{-i} .

The following bounds also hold at bad prefixes. Pair availability, the identity for r_i , and the conditional probability bound in Lemma 16 imply

$$\frac{30}{31N} \leq r_i \leq \frac{32}{31N}, \quad \frac{1-s/(N-2s)}{N-s} \leq r_0 \leq \frac{1}{N-s}. \quad (63)$$

Writing $\beta = 1/63$, one has $\alpha_{\text{Corr}} \geq (1-\eta)e^{-\beta}$ and $Z_{\text{ret}} \leq N$. In particular

$$k_i \leq \frac{e^\beta}{N(1-\eta)^2} < \frac{1.1}{N}, \quad |b_i| < 0.1.$$

For the second assertion, multiply r_i/r_0 by $\alpha_{\text{Corr}} e^{\text{Corr}(z_i)}$. This factor lies between $(1-\eta)e^{-\beta}$ and e^β . The resulting lower and upper bounds on e^{b_i} are respectively

$$\frac{30}{31}(1-\eta)^2 e^{-\beta} > e^{-0.1}, \quad \frac{32}{31} \frac{1-\eta}{1-\eta/(1-2\eta)} e^\beta < e^{0.1}.$$

Lemma 21 (KL divergence between distributions of the inverse output). *Combine all inputs with unrevealed construction values into one outcome $\mathbf{f}$. Let Q assign masses $k_i, k_{\mathbf{f}}$ to the legal indices $i \in I_y$ and $\mathbf{f}$, and set $b_{\mathbf{f}} = 0$. Multiplying these masses by e^{b_i} and normalizing gives the corresponding distribution in Real:*

$$P(i) = \frac{k_i e^{b_i}}{D}, \quad P(\mathbf{f}) = \frac{k_{\mathbf{f}}}{D}, \quad D = k_{\mathbf{f}} + \sum_i k_i e^{b_i}. \quad (64)$$

Moreover,

$$\sqrt{D_{\text{KL}}(P\|Q)} \leq \sqrt{\frac{1.1e^{0.2}}{2N}} \|(b_i)_i\|_{2, I_y} \leq \frac{1}{\sqrt{N}} \|(b_i)_i\|_{2, I_y}. \quad (65)$$

Proof. The definition of b_i and normalization give (64). Apply (4) with $\text{osc } b \leq 0.2$ and $\text{Var}_Q(b) \leq \sum_i k_i b_i^2 \leq (1.1/N) \sum_i b_i^2$. $\square$

K.4 The augmented inverse response

Let $R_{\text{inv}}^{\text{Real}}$ denote the augmented inverse-response distribution in Real, and use $R_{\text{inv}}^{\text{Hyb}^{\text{GSim}}}$ from Section 5.4 for the hybrid. Conditional on $\mathbf{f}$, the inverse output is uniform over the u inputs with unrevealed construction values in both worlds, so revealing it adds zero KL divergence. Conditional on any fixed inverse output with an unrevealed construction value, the companion-output distribution is R^{opp} from (50); when the inverse output has a revealed construction value, the companion output is deterministic in both worlds. The KL chain rule therefore gives exactly

$$D_{\text{KL}}(R_{\text{inv}}^{\text{Real}} \| R_{\text{inv}}^{\text{Hyb}^{\text{GSim}}}) = D_{\text{KL}}(P\|Q) + r_{\mathbf{f}} D_{\text{KL}}(R^{\text{opp}} \| \text{Unif}(\mathcal{G} \setminus \text{Ran}_1)).$$

Use $r_{\mathbf{f}} \leq 1$, the companion-output bound (52), the bound on $\sum_{i: \iota_i=1} A_i(y)^2$ in (57), and (62). Together with (61) and (65), they give at a good prefix

$$\sqrt{D_{\text{KL}}(R_{\text{inv}}^{\text{Real}} \| R_{\text{inv}}^{\text{Hyb}^{\text{GSim}}})} \leq 1.5\sqrt{2n} \varepsilon_{\tau}^{\text{res}} + 13n^{3/2} a_s + 8\sqrt{n} b_s.$$

K.5 Combining response bounds and handling bad prefixes

For a forward query at a construction-only input, the divergence between the response distributions satisfies (49). For a forward query at an input whose construction value is still unrevealed, we first compare the distributions of that value: $p_{\text{Cons}}^{\text{Real}}$ in Real and ν_{Corr} in Hyb^{GSim} . Conditional on the construction value, the pair has the forward-response distribution for the enlarged list. Write $R_{z,\text{pair}}$ and $K_{z,\text{pair}}$ for the joint distributions of the construction value and the two permutation outputs in Real and Hyb^{GSim} , respectively. Let $R_{\text{pair}|z}$ and $U_{\text{pair}|z}$ be the corresponding distributions of the two outputs given z . The KL chain rule gives

$$D_{\text{KL}}(R_{z,\text{pair}} \| K_{z,\text{pair}}) = D_{\text{KL}}(p_{\text{Cons}}^{\text{Real}} \| \nu_{\text{Corr}}) + \mathbb{E}_{z \leftarrow p_{\text{Cons}}^{\text{Real}}} [D_{\text{KL}}(R_{\text{pair}|z} \| U_{\text{pair}|z})]. \quad (66)$$

The reweighting bound (44) gives $\sqrt{D_{\text{KL}}(p_{\text{Cons}}^{\text{Real}} \| \nu_{\text{Corr}})} \leq 2\varepsilon_{\tau}^{\text{res}}$. In the second term, conditioning on the newly revealed construction value gives an enlarged state of size at most $r+1$. The bound for the forward response applies whenever adding z preserves the good-prefix conditions.

In d_r^{good} , the second term retains only values z that preserve the good-prefix conditions, with unchanged probability weights. The bound for a forward query is no larger than the inverse-response bound above, and $2 \leq 1.5\sqrt{2n}$. Thus every fixed prefix satisfies

$$\sqrt{d_r^{\text{good}}(\tau)} \leq 1.5\sqrt{2n}\varepsilon_{\tau}^{\text{res}} + 13n^{3/2}a_{r+1} + 8\sqrt{n}b_{r+1}.$$

Averaging over $\tau \leftarrow \text{Tr}_{\text{Real}}^r$ gives

$$\sqrt{\mathbb{E}_{\tau}[d_r^{\text{good}}(\tau)]} \leq 1.5\sqrt{2n}\sqrt{\mathbb{E}_{\tau}[(\varepsilon_{\tau}^{\text{res}})^2]} + 13n^{3/2}a_{r+1} + 8\sqrt{n}b_{r+1}.$$

Using (43) with size bound $r+1$ gives coefficients

$$13 + \frac{4.5\sqrt{2}}{n} < 14, \quad 8 + 3\sqrt{2} < 13. \quad (67)$$

To include the omitted bad cases, we need a divergence bound that holds also at bad prefixes. The following bound on likelihood ratios provides it.

Lemma 22 (Bounded response likelihood ratios). *At every compatible prefix of size at most $N/64$, the likelihood ratios of the conditional distributions in Real relative to those in Hyb^{GSim} are less than 2 on their common legal support. Thus every conditional KL is less than $\log 2 < 1$.*

Proof. For a fresh construction query, (58) and $0 \leq \text{Corr} \leq 1/63$ bound the response likelihood ratio by $e^{1/63}/(1-\eta) < 1.1$. Given the construction value, the likelihood ratio for the permutation pair is at most $31/30$ by (51). For a legal return at a construction-only input x_i , $k_i \geq e^{-1/63}/N$, so (63) gives $r_i/k_i < 1.1$. For an inverse output with an unrevealed construction value, $u_0 \geq (1 -$

$\eta)e^{-1/63}/N$ gives $r_0/u_0 < 1.1$. Finally, the probability of avoiding all collision events satisfies $61/62 \leq W^{\text{opp}} \leq 1$, so (50) bounds the likelihood ratio of the companion output by $62/61$. The entire forward and inverse response ratios are thus bounded by $1.1(31/30)$ and $1.1(62/61)$, respectively, both less than 2. Availability lower bounds ensure common legal support. $\square$

We now average over prefixes in Real and include the bad cases.

Lemma 23 (Conditional divergence from Real to Hyb^{GSim}). *Suppose $N \geq 2^{20}$ and $q \leq N/64$. Fix $r \in \{0, \dots, q-1\}$, and let $\tau = \text{Tr}_{\text{Real}}^r$ be the transcript after r queries. Let R_τ and K_τ^{GSim} be the conditional distributions of the response to query $r+1$ in Real and Hyb^{GSim}, respectively, and put $d_r(\tau) = D_{\text{KL}}(R_\tau \| K_\tau^{\text{GSim}})$. Averaging over τ , we have*

$$\sqrt{\mathbb{E}_{\tau \leftarrow \text{Tr}_{\text{Real}}^r}[d_r(\tau)]} \leq 14n^{3/2} \frac{r+1}{N^{3/2}} + 13\sqrt{n} \frac{(r+1)^2}{N^2} + N^{-4}.$$

Proof (Lemma 23). Lemma 13 bounds multiplicities and collisions simultaneously at every prefix. The omitted contribution is $d_r(\tau)$ at a bad prefix. When a forward query makes a good prefix bad, only the KL divergence between the distributions of the two permutation outputs, conditional on the newly revealed construction value, is omitted. These cases are disjoint and their total probability is at most N^{-8} . Lemma 22 therefore gives $\mathbb{E}_\tau[d_r^{\text{bad}}(\tau)] \leq N^{-8}$. Combining this with (67) and substituting the definitions of a_{r+1} and b_{r+1} yields

$$\begin{aligned} \sqrt{\mathbb{E}_\tau[d_r(\tau)]} &\leq \sqrt{\mathbb{E}_\tau[d_r^{\text{good}}(\tau)]} + N^{-4} \\ &\leq 14n^{3/2} \frac{r+1}{N^{3/2}} + 13\sqrt{n} \frac{(r+1)^2}{N^2} + N^{-4}. \end{aligned}$$

For the empty state, the distributions of fresh construction, forward, and inverse responses coincide in Real and Hyb^{GSim}. $\square$

K.6 From response distributions to transcript distributions

The conditional divergence bound now gives the following transcript bound.

Lemma 24 (KL divergence between transcript distributions). *For $N \geq 2^{20}$ and $q \leq N/64$,*

$$\sqrt{D_{\text{KL}}(T_{\text{Real}} \| T_{\text{Hyb}^{\text{GSim}}})} \leq 14n^{3/2} \frac{q^{3/2}}{N^{3/2}} + 13\sqrt{n} \frac{q^{5/2}}{N^2} + \sqrt{q} N^{-4}. \quad (68)$$

Proof (Lemma 24). For $r = 0, \dots, q-1$, Lemma 23 gives

$$\sqrt{\mathbb{E}_{\tau \leftarrow \text{Tr}_{\text{Real}}^r}[d_r(\tau)]} \leq 14n^{3/2} \frac{r+1}{N^{3/2}} + 13\sqrt{n} \frac{(r+1)^2}{N^2} + N^{-4}.$$

By the KL chain rule,

$$D_{\text{KL}}(T_{\text{Real}} \| T_{\text{Hyb}^{\text{GSim}}}) = \sum_{r=0}^{q-1} \mathbb{E}_{\tau \leftarrow \text{Tr}_{\text{Real}}^r} [d_r(\tau)].$$

Therefore,

$$\begin{aligned} \sqrt{D_{\text{KL}}(T_{\text{Real}} \| T_{\text{Hyb}^{\text{GSim}}})} &\leq \frac{14n^{3/2}}{N^{3/2}} \left(\sum_{r=0}^{q-1} (r+1)^2 \right)^{1/2} \\ &\quad + \frac{13\sqrt{n}}{N^2} \left(\sum_{r=0}^{q-1} (r+1)^4 \right)^{1/2} + \sqrt{q} N^{-4} \\ &\leq 14n^{3/2} \frac{q^{3/2}}{N^{3/2}} + 13\sqrt{n} \frac{q^{5/2}}{N^2} + \sqrt{q} N^{-4}, \end{aligned}$$

where $\sum_{r < q} (r+1)^2 \leq q^3$ and $\sum_{r < q} (r+1)^4 \leq q^5$. $\square$

L Hyb^{GSim} to Ideal^{GSim}

This appendix proves the transcript comparison in Lemma 30. We first derive the conditional response distributions, then bound their divergence and the failure probability. Throughout, let $N = 2^n \geq 2^{20}$ and $q \leq N/64$, and put

$$\beta := \frac{q}{N-q} \leq \frac{1}{63}, \quad c_* := \frac{32}{31}.$$

L.1 Conditional distribution of the construction oracle

We use the exact inverse-response distribution from Lemma 4 and the notation of Section 5.3.

For GSim_5 , we reuse the notation W_τ, π_τ from Lemma 2. Fix an attainable augmented transcript prefix τ in $\text{Ideal}^{\text{GSim}}$ containing no failure response, immediately before the next query. Let $\text{Dom}_\tau^{\text{pub}} = \text{Dom}_\tau^{\text{prim}} \cup \text{dom}(\text{Cons})$ be the inputs with revealed construction values, so $|\text{Dom}_\tau^{\text{pub}}| \leq q$. As in Section 3.4, $\Phi_\tau : \mathcal{G} \rightarrow \mathcal{G}$ agrees with these values. Under μ_τ , they are fixed and the remaining values are independent and uniform on $\mathcal{G}$.

For each past fresh successful inverse call j , fix its available set U_j , size L_j , set of forbidden construction values Forb_j , and correction term Corr_j . Define

$$q_j(z) = 1 - \mathbf{1}[z \notin \text{Forb}_j] e^{-\text{Corr}_j(z)}, \quad p_{\text{rej},j}(\Phi_\tau) = \frac{1}{L_j} \sum_{x \in U_j} q_j(\Phi_\tau(x)).$$

Use index 0 for the analogous quantities in a current inverse call. Write $H_5(v) = 1 + v + v^2 + v^3 + v^4$ and $f(v) = \log H_5(v)$. The factors $H_5(p_{\text{rej},j})$ account for the rejected candidates preceding each observed inverse response. The following lemma gives their effect on the conditional distribution of the table of Γ .

Lemma 25 (Conditional distribution of the table of Γ). *In $\text{Ideal}^{\text{GSim}}$, the table of Γ conditioned on a prefix τ containing no failure response has distribution π_τ given by*

$$W_\tau(\Phi_\tau) := \prod_j H_5(p_{\text{rej},j}(\Phi_\tau)), \quad \pi_\tau(\Phi_\tau) := \frac{W_\tau(\Phi_\tau)}{\mathbb{E}_{\mu_\tau}[W_\tau]} \mu_\tau(\Phi_\tau).$$

Proof. By Lemma 4, conditional on Γ , an observed successful response at x_j has likelihood $p_{\text{acc},j}(x_j)H_5(p_{\text{rej},j})/L_j$. The augmented response determines $p_{\text{acc},j}(x_j)/L_j$, so only $H_5(p_{\text{rej},j})$ depends on the remaining values. The rejected candidate inputs, their construction values, and the random coins used in those trials have been summed out. Construction and forward responses fix revealed coordinates. Conditional on the recorded construction values, the probabilities of the observed forward responses do not depend on the remaining values. Query choices are determined by the fixed visible prefix and add no factor depending on the remaining construction oracle values. $\square$

For $u, v \geq 0$ with $u + v \leq 1$, the function $f = \log H_5$ satisfies

$$v \leq f(u + v) - f(u) \leq \min\{2v, v + 2uv + v^2\}. \quad (69)$$

Indeed, direct differentiation gives

$$1 \leq f'(t) = \frac{1 + 2t + 3t^2 + 4t^3}{H_5(t)} \leq \min\{2, 1 + 2t\} \quad (0 \leq t \leq 1).$$

Integrate over $[u, u + v]$.

The increment bound controls how W_τ changes when one unrevealed construction value changes. Normalization then gives the following conditional probability bounds.

Lemma 26 (Conditional distribution of an unrevealed construction value). *For each input x with an unrevealed construction value, write Γ_{-x} for the other unrevealed construction values. For every setting of Γ_{-x} ,*

$$\frac{1}{c_* N} \leq \pi_\tau[\Gamma(x) = z \mid \Gamma_{-x}] \leq \frac{c_*}{N}.$$

The same bounds hold after fixing any subset of coordinates other than x , and after multiplying π_τ by $H_5(p_{\text{rej},0})$ and normalizing.

Proof. Changing one construction value changes $p_{\text{rej},j}$ by at most $1/L_j$. Thus, with the other coordinates fixed,

$$\text{osc}_z(\log \pi_\tau[\Gamma(x) = z \mid \Gamma_{-x}]) \leq 2 \sum_j L_j^{-1} \leq 2\beta.$$

Since $e^{2\beta} < c_*$, normalization relative to the uniform distribution gives the claimed bounds.

Averaging over unfixed coordinates preserves the resulting bounds. Including the current call, there are at most q factors, so the same bound holds after including $H_5(p_{\text{rej},0})$. $\square$

L.2 Comparing construction responses

We compare ν_{Corr} with the distribution of a fresh construction response in $\text{Ideal}^{\text{GSim}}$ given τ . We first bound the second moments of the rejection probabilities that determine the conditional distribution of the table of Γ .

Lemma 27 (Second moment of the rejection probability). *For the rejection probabilities in earlier calls and the current probability $p_{\text{rej},0}$,*

$$\mathbb{E}_{\pi_\tau}[p_{\text{rej},j}^2] \leq 11\beta^2.$$

The same bound holds for $p_{\text{rej},j,-x}$, obtained by omitting one coordinate while retaining L_j , and after reweighting by $H_5(p_{\text{rej},0})$.

Proof. Let B_j denote the number of inputs in U_j with an unrevealed construction value in Forb_j . For their indicators I_y , Lemma 26 gives $\mathbb{E}[I_y] \leq c_*|\text{Forb}_j|/N$ and, by conditioning on all coordinates except y , $\mathbb{E}[I_y I_v] \leq (c_*|\text{Forb}_j|/N)^2$ for $y \neq v$. Consequently

$$\mathbb{E}[B_j] \leq c_*q, \quad \mathbb{E}[B_j^2] \leq c_*q + c_*^2q^2.$$

There are at most q revealed coordinates, and $q_j(z) \leq \mathbf{1}[z \in \text{Forb}_j] + \text{Corr}_j(z) \leq \mathbf{1}[z \in \text{Forb}_j] + \beta$. Hence $p_{\text{rej},j} \leq (2q + B_j)/(N - q)$ and

$$\mathbb{E}[p_{\text{rej},j}^2] \leq (4 + 5c_* + c_*^2)\beta^2 \leq 11\beta^2.$$

Omitting a nonnegative summand decreases the average. The same conditional probability bound holds after reweighting by $H_5(p_{\text{rej},0})$, so the calculation also applies under that distribution. $\square$

We reuse the likelihood calculation of Section 5.1, now with five trials and the additional acceptance probability $e^{-\text{Corr}_j(z)}$. For an input x with an unrevealed construction value, fixing $w = \Gamma_{-x}$ gives the exact increments

$$\Delta_{j,x}(z) = f\left(p_{\text{rej},j,-x} + \frac{q_j(z)}{L_j}\right) - f(p_{\text{rej},j,-x}).$$

For $z \in \text{Forb}_j$, the contribution of x to the rejection probability is $1/L_j$. Outside this set it is $(1 - e^{-\text{Corr}_j(z)})/L_j$, arising from the additional acceptance test. The leading term is still $L_j^{-1}\mathbf{1}[z \in \text{Forb}_j]$; the earlier correction and $p_{\text{rej},j,-x}$ determine the error. Every input whose construction value is currently unrevealed belonged to the available set of every earlier inverse call. Define

$$r_x(z) = \sum_j \Delta_{j,x}(z) - \text{Corr}_\tau(z), \quad \text{Corr}_\tau(z) = \sum_j L_j^{-1}\mathbf{1}[z \in \text{Forb}_j].$$

Applying (69) gives

$$0 \leq \Delta_{j,x}(z) - \frac{\mathbf{1}[z \in \text{Forb}_j]}{L_j} \leq \frac{2p_{\text{rej},j,-x}}{L_j} + \frac{\beta}{L_j} + \frac{1}{L_j^2}. \quad (70)$$

Indeed, take $u = p_{\text{rej},j,-x}$, $v = q_j(z)/L_j$ in that inequality and use $0 \leq q_j(z) - \mathbf{1}[z \in \text{Forb}_j] \leq \beta$ and $0 \leq q_j(z) \leq 1$. The lower bound follows from $v \geq L_j^{-1}\mathbf{1}[z \in \text{Forb}_j]$.

Lemma 28 (Oscillation of the log-likelihood remainder). *With oscillation over z at fixed w ,*

$$\text{osc}(r_x) \leq 2\beta, \quad (\mathbb{E}_{\pi_\tau}[\text{osc}(r_x)^2])^{1/2} \leq 9q^2/N^2.$$

The same bounds hold after reweighting by $H_5(p_{\text{rej},0})$.

Proof. Every summand of r_x is nonnegative by (70), and $r_x \leq \sum_j \Delta_{j,x} \leq 2\beta$. Thus its oscillation is bounded by the sum of the upper bounds in that equation. The second-moment bound gives

$$\begin{aligned} (\mathbb{E}_{\pi_\tau}[\text{osc}(r_x)^2])^{1/2} &\leq 2 \sum_j \frac{(\mathbb{E}_{\pi_\tau}[p_{\text{rej},j,-x}^2])^{1/2}}{L_j} + \beta \sum_j L_j^{-1} + \sum_j L_j^{-2} \\ &\leq (2\sqrt{11} + 2)\beta^2. \end{aligned}$$

Here $\sum_j L_j^{-1} \leq \beta$ and $\sum_j L_j^{-2} \leq q/(N-q)^2 \leq \beta^2$. Finally $(2\sqrt{11}+2)(64/63)^2 < 9$ and $\beta \leq (64/63)q/N$. The same calculation applies after reweighting by $H_5(p_{\text{rej},0})$. $\square$

For $\nu_{\text{Corr}}(z) = e^{\text{Corr}_\tau(z)}/Z_{\text{Corr}}$, Bayes' rule gives

$$\frac{\pi_\tau[\Gamma(x) = z \mid w]}{\nu_{\text{Corr}}(z)} = \frac{e^{r_x(z)}}{\mathbb{E}_{\nu_{\text{Corr}}}[e^{r_x}]}$$

Write P_w for this conditional distribution. Using weights e^{-r_x} in Lemma 1 takes P_w to ν_{Corr} and gives $D_{\text{KL}}(\nu_{\text{Corr}} \| P_w) \leq \text{osc}(r_x)^2/8$. Convexity in the second argument and averaging over the w -marginal of π_τ now give

$$D_{\text{KL}}(\nu_{\text{Corr}} \| \mathcal{L}(\Gamma(x) \mid \tau)) \leq \frac{9^2 q^4}{8N^4} < \frac{11q^4}{N^4}.$$

This bounds the divergence between the distributions of the construction value revealed by a fresh construction or forward primitive query. Conditional on that value, the pair of permutation outputs has the same distribution in $\text{Ideal}^{\text{GSim}}$ and Hyb^{GSim} . At an input with a revealed construction value, the two forward-response distributions coincide.

L.3 Inverse responses conditioned on success

For a successful inverse query, we first compare the inverse output and then its construction value. Given the query, these determine the augmented response. We use the joint distribution P_τ^+ and the resulting distribution $R_{\text{inv}}^{\text{Ideal}^{\text{GSim}},+}$ from Section 5.3.

For the current inverse query, write the acceptance probability as a function of the candidate's construction value z :

$$\begin{aligned} p_{\text{acc}}^{\text{val}}(z) &= \mathbf{1}[z \notin \text{Forb}]e^{-\text{Corr}_\tau(z)}, \quad q_0(z) = 1 - p_{\text{acc}}^{\text{val}}(z), \\ p_{\text{rej},0} &= \frac{1}{L} \sum_{x \in U} q_0(\Gamma(x)), \quad w_{\text{cur}} = H_5(p_{\text{rej},0}). \end{aligned}$$

The possible rejected trials before the return contribute w_{cur} . We include this factor in the auxiliary distribution $\pi^{\text{cur}} = w_{\text{cur}}\pi_\tau / \mathbb{E}_{\pi_\tau}[w_{\text{cur}}]$. For a specified input x with an unrevealed construction value, define its mean acceptance probability and the distribution of the table of Γ conditional on returning x :

$$\theta = \mathbb{E}_{\pi^{\text{cur}}}[p_{\text{acc}}^{\text{val}}(\Gamma(x))], \quad \pi^{\text{sel}}(\Phi_\tau) = \frac{p_{\text{acc}}^{\text{val}}(\Phi_\tau(x))}{\theta} \pi^{\text{cur}}(\Phi_\tau).$$

Lemma 26 applies to π^{cur} , so

$$\theta \geq e^{-\beta} \left(1 - \frac{c_* q}{N}\right) \geq \frac{61}{63}, \quad \pi^{\text{sel}} \leq \frac{63}{61} \pi^{\text{cur}}. \quad (71)$$

Indeed $e^{-\beta} \geq 1 - \beta \geq 62/63$ and $c_* q/N \leq 1/62$. The same pointwise comparison holds for the marginal distributions of the other coordinates. In particular $\mathbb{E}_{\pi_\tau}[p_{\text{acc}}^{\text{val}}(\Gamma(x))w_{\text{cur}}] = \theta \mathbb{E}_{\pi_\tau}[w_{\text{cur}}] \geq 61/63$.

For the current inverse call, the contribution of $\Gamma(x) = z$ to $\log H_5(p_{\text{rej},0})$ is

$$\Delta_{0,x}(z) = f\left(p_{\text{rej},0,-x} + \frac{q_0(z)}{L}\right) - f(p_{\text{rej},0,-x}).$$

Put $\ell(z) = L^{-1}\mathbf{1}[z \in \text{Forb}]$ and $\zeta_x = r_x + \Delta_{0,x} - \ell$. By (70), this is a sum of nonnegative remainders from at most q past and current calls. The proof of Lemma 28 gives

$$\begin{aligned} \text{osc}(\zeta_x) &\leq 2\beta, \\ \mathbb{E}_\rho[\text{osc}(\zeta_x)^2] &\leq \frac{9^2 q^4}{N^4} \quad (\rho = \pi_\tau, \pi^{\text{cur}}), \\ \mathbb{E}_{\pi^{\text{sel}}}[\text{osc}(\zeta_x)^2] &\leq \frac{63}{61} \frac{9^2 q^4}{N^4}. \end{aligned}$$

The final line applies (71) to the w -marginals of π^{sel} and π^{cur} .

Distribution of the inverse output. Partition U according to whether the construction value is revealed or unrevealed. Let u denote the number of inputs with unrevealed values, and put

$$w_{\text{pub}} = \sum_{x' \in U \cap \text{Dom}_\tau^{\text{pub}}} p_{\text{acc}}^{\text{val}}(\Gamma(x')), \quad \alpha_{\text{Corr}} = \mathbb{E}_{\nu_{\text{Corr}}}[p_{\text{acc}}^{\text{val}}] = \frac{N - |\text{Forb}|}{Z_{\text{Corr}}}.$$

Exchangeability makes θ the same for every input with an unrevealed construction value. By Lemma 4, the probability of returning x , before conditioning on success, is

$$\frac{1}{L} \mathbb{E}_{\pi_\tau}[w_{\text{cur}} p_{\text{acc}}^{\text{val}}(\Gamma(x))].$$

Dividing these masses by the total success probability gives

$$P_{\text{ret}}(x) = \begin{cases} p_{\text{acc}}^{\text{val}}(\Gamma(x)) / (w_{\text{pub}} + u\theta), & x \text{ with revealed construction value,} \\ \theta / (w_{\text{pub}} + u\theta), & x \text{ with unrevealed construction value.} \end{cases}$$

Hyb^{GSim} replaces θ by α_{Corr} and keeps the same weights for inputs with revealed construction values. This changes the total weight of all available inputs. Direct summation gives

$$\chi^2(P_{\text{ret}}^{\text{Hyb}^{\text{GSim}}} \| P_{\text{ret}}) = \frac{w_{\text{pub}} u (\theta - \alpha_{\text{Corr}})^2}{\theta (w_{\text{pub}} + u \alpha_{\text{Corr}})^2}. \quad (72)$$

When $w_{\text{pub}} = 0$ both distributions are uniform over the inputs with unrevealed construction values.

Lemma 29 (Acceptance weights for inputs with unrevealed construction values). *The mean acceptance probabilities θ and α_{Corr} satisfy*

$$|\theta - \alpha_{\text{Corr}}| \leq \frac{4q^2}{N^2}. \quad (73)$$

Proof. For fixed w , the distribution of the construction value under π^{cur} is proportional to $\nu_{\text{Corr}} e^{\ell + \zeta_x}$. Since $p_{\text{acc}}^{\text{val}} = 0$ on **Forb**, reweighting ν_{Corr} by e^{ℓ} and normalizing gives expectation

$$\frac{\mathbb{E}_{\nu_{\text{Corr}}} [p_{\text{acc}}^{\text{val}} e^{\ell}]}{\mathbb{E}_{\nu_{\text{Corr}}} [e^{\ell}]} = \frac{\alpha_{\text{Corr}}}{1 + (e^{1/L} - 1) \nu_{\text{Corr}}(\text{Forb})}.$$

Its difference from α_{Corr} is at most $(e^{1/L} - 1) \nu_{\text{Corr}}(\text{Forb})$. Now

$$\nu_{\text{Corr}}(\text{Forb}) \leq \frac{63q}{62N}, \quad e^{1/L} - 1 \leq \frac{1}{L-1} \leq \frac{32}{31N},$$

since $e^{\beta} \leq 63/62$ and $L \geq N - q \geq 63N/64 \geq 63$. Thus the change is less than $1.1q/N^2$. Further multiplication by e^{ζ_x} and normalization change the expectation of $p_{\text{acc}}^{\text{val}} \in [0, 1]$ by at most $\text{osc}(\zeta_x)/4$, by (4). Average under the π^{cur} marginal of w and use Cauchy–Schwarz to obtain

$$|\theta - \alpha_{\text{Corr}}| \leq \frac{1.1q}{N^2} + \frac{9q^2}{4N^2} < \frac{4q^2}{N^2}.$$

□

Here $w_{\text{pub}} \leq q$, $u \geq N - q$, and $\alpha_{\text{Corr}} \geq e^{-\beta}(1 - q/N) \geq 31/32$. Hence

$$\frac{q}{(N - q)\theta\alpha_{\text{Corr}}^2} < \frac{9q}{8N}.$$

Equations (72) and (73), with $q/N \leq 1/64$, imply

$$D_{\text{KL}}(P_{\text{ret}}^{\text{Hyb}^{\text{GSim}}} \| P_{\text{ret}}) \leq \frac{9}{32} \frac{q^4}{N^4}.$$

Construction value at the inverse output. Conditional on returning an input x with an unrevealed construction value, the table of Γ has distribution π^{sel} . The leading factor $e^{\text{Corr}_\tau(z)}$ is canceled by the acceptance probability: on Forb^c , $\nu_{\text{Corr}}(z)p_{\text{acc}}^{\text{val}}(z) = 1/Z_{\text{Corr}}$. The leading contribution $\ell(z)$ from rejected trials in the current call also vanishes on this set. Thus, with the other coordinates fixed, the conditional distribution of $\Gamma(x)$ under π^{sel} is proportional to $e^{\zeta_x(z)} \text{Unif}(\text{Forb}^c)(z)$. Lemma 1 with weights $e^{-\zeta_x}$, the second-moment bound under π^{sel} , and convexity in the second argument give

$$D_{\text{KL}}(\text{Unif}(\text{Forb}^c) \parallel \mathcal{L}(\Gamma(x) \mid \tau, x \text{ is the inverse output})) \leq \frac{63}{61} \frac{9^2 q^4}{8N^4}.$$

Here the average over w is conditional on returning x . When the inverse output has a revealed construction value, that value and the companion output are deterministic. We apply the KL chain rule to the inverse output and its construction value. Bounding the probability in Hyb^{GSim} of returning an input with an unrevealed construction value by one gives

$$D_{\text{KL}}(R_{\text{inv}}^{\text{Hyb}^{\text{GSim}}} \parallel R_{\text{inv}}^{\text{Ideal}^{\text{GSim}},+}) \leq \left(\frac{63}{61} \frac{9^2}{8} + \frac{9}{32} \right) \frac{q^4}{N^4} < \frac{11q^4}{N^4}. \quad (74)$$

L.4 Failure probability and divergence of transcript distributions

The bounds on divergence between distributions of successful responses hold uniformly over attainable prefixes containing no failure response. We now include failure and prove the bound for transcript distributions.

Lemma 30 (KL divergence from Hyb^{GSim} to $\text{Ideal}^{\text{GSim}}$). *For $N \geq 2^{20}$ and $q \leq N/64$,*

$$D_{\text{KL}}(T_{\text{Hyb}^{\text{GSim}}} \parallel T_{\text{Ideal}^{\text{GSim}}}) \leq \frac{11q^5}{N^4} + q \log \frac{1}{1 - (4q/N)^5}. \quad (75)$$

Proof. Fix an attainable prefix τ containing no failure response and its next query. Sections L.2 and L.3 bound the KL divergence from the response distribution in Hyb^{GSim} to the distribution of successful responses in $\text{Ideal}^{\text{GSim}}$ by $11q^4/N^4$ for every query type. Only inverse queries can fail.

Conditional on Γ , a current inverse call fails with probability $p_{\text{rej},0}^5$. Fix the candidates and construction values tested earlier in this call. Conditioning on their rejection contributes likelihood factors independent of the remaining unrevealed values, so Lemma 26 still bounds the probability of each possible construction value at a newly tested input by $32/(31N)$. Rejected candidates from earlier calls are averaged out.

Before the first query, $\text{Forb} = \emptyset$ and $\text{Corr} \equiv 0$, so failure is impossible. Otherwise let $r \geq 1$ denote the number of previous queries, so $r \leq q - 1$. At each of the five trials, at most $r + 4$ available inputs have a construction value revealed in τ or read in an earlier trial of the current call. At any other input,

the construction value lies in **Forb** with probability at most $32r/(31N)$, and its acceptance coin rejects with probability at most $r/(N-r)$. Thus, conditional on the earlier rejections and the construction values read in those trials, the next rejection probability is at most

$$\frac{r+4}{N-r} + \frac{32r}{31N} + \frac{r}{N-r} < \frac{4q}{N}.$$

Multiplying the five conditional bounds gives

$$p_{\text{fail},\tau} = \mathbb{E}_{\pi_\tau}[p_{\text{rej},0}^5] \leq (4q/N)^5. \quad (76)$$

Let P_τ be the response distribution in Hyb^{GSim} and Q_τ that in $\text{Ideal}^{\text{GSim}}$, including failure. Let Q_τ^+ denote Q_τ conditioned on success. Equation (5) gives

$$D_{\text{KL}}(P_\tau \| Q_\tau) = D_{\text{KL}}(P_\tau \| Q_\tau^+) + \log \frac{1}{1 - p_{\text{fail},\tau}} \leq \frac{11q^4}{N^4} + \log \frac{1}{1 - (4q/N)^5}.$$

Every prefix in Hyb^{GSim} is attainable in $\text{Ideal}^{\text{GSim}}$ by Section 5.4. Applying the KL chain rule with prefixes sampled from Hyb^{GSim} yields

$$\begin{aligned} D_{\text{KL}}(T_{\text{Hyb}^{\text{GSim}}} \| T_{\text{Ideal}^{\text{GSim}}}) &= \sum_{r=0}^{q-1} \mathbb{E}_{\tau \leftarrow \text{Tr}_{\text{Hyb}^{\text{GSim}}}^r} [D_{\text{KL}}(P_\tau \| Q_\tau)] \\ &\leq \frac{11q^5}{N^4} + q \log \frac{1}{1 - (4q/N)^5}. \end{aligned} \quad \square$$