

Towards Defeating Backdoored Random Oracles: Indifferentiability with Bounded Adaptivity

Yevgeniy Dodis¹, Pooya Farshim², Sogol Mazaheri³, and Stefano Tessaro⁴

¹ New York University

`dodis@cs.nyu.edu`

² University of York

`pooya.farshim@gmail.com`

³ Technische Universität Darmstadt

`sogol.mazaheri@cryptoplexity.de`

⁴ University of Washington

`tessaro@cs.washington.edu`

Abstract. In the backdoored random-oracle (BRO) model, besides access to a random function H , adversaries are provided with a backdoor oracle that can compute arbitrary leakage functions f of the function table of H . Thus, an adversary would be able to invert points, find collisions, test for membership in certain sets, and more. This model was introduced in the work of Bauer, Farshim, and Mazaheri (Crypto 2018) and extends the auxiliary-input idealized models of Unruh (Crypto 2007), Dodis, Guo, and Katz (Eurocrypt 2017), Coretti et al. (Eurocrypt 2018), and Coretti, Dodis, and Guo (Crypto 2018). It was shown that certain security properties, such as one-wayness, pseudorandomness, and collision resistance can be re-established by combining two independent BROs, even if the adversary has access to both backdoor oracles.

In this work we further develop the technique of combining two or more independent BROs to render their backdoors useless in a more general sense. More precisely, we study the question of building an *indifferentiable* and backdoor-free random function by combining multiple BROs. Achieving full indifferentiability in this model seems very challenging at the moment. We however make progress by showing that the xor combiner goes well beyond security against preprocessing attacks and offers indifferentiability as long as the adaptivity of queries to different backdoor oracles remains logarithmic in the input size of the BROs. We even show that an extractor-based combiner of three BROs can achieve indifferentiability with respect to a linear adaptivity of backdoor queries. Furthermore, a natural restriction of our definition gives rise to a notion of *indifferentiability with auxiliary input*, for which we give two positive feasibility results.

To prove these results we build on and refine techniques by Gös et al. (STOC 2015) and Kothari et al. (STOC 2017) for decomposing distributions with high entropy into distributions with more structure and show how they can be applied in the more involved adaptive settings.

Keywords: Hash functions · Indifferentiability · Backdoors · Auxiliary input · Communication complexity

1 Introduction

Hash functions are one of the most fundamental building blocks in protocol design. For this reason, both the cryptanalysis and provable security of hash functions have been active areas of research in recent years. The first known instances of collisions and chosen-prefix collisions in SHA-1 were recently demonstrated by Stevens et al. [26] and Leurent and Peyrin [20], respectively. Furthermore, feasibility of built-in adversarial weaknesses (aka. backdoors) in efficient hash functions have been demonstrated by Fischlin, Janson, and Mazaheri [13]. A practical way to provide safeguards against similar failures of hash functions is to *combine* a number of independent hash functions so that the resulting function is at least as secure as their strongest. Most works in this area have focused attention on a setting where at least one of the hash functions is secure, although positive results when *all* underlying hash functions have weaknesses have also been demonstrated [22,15].

In this work we are interested in protecting hash functions against a variety of attacks that may arise due to backdoors, cryptanalytic advances, or preprocessing attacks. We carry out our study in the recent backdoored random-oracle (BRO) model, which uniformly treats these settings and also permits strong adversarial settings where all hash functions may be weak.

1.1 The BRO model

Bauer, Farshim, and Mazaheri (BFM) [3] at Crypto 2018 formulated a new model for the analysis of hash functions that substantially weakens the traditional random-oracle (RO) model. Here an adversary, on top of direct access to the random oracle, is able to obtain arbitrary functions of the function table of the random oracle.¹ The implications of this weakening are manifold. To start with, positive results in this model imply positive results in the traditional setting where all but one of the hash functions is weak. Second, this model captures arbitrary preprocessing attacks on hash functions, another highly active area of research [27,10,7,6]. Finally, it allows to model unrestricted adversarial capabilities, which can adaptively depend on input instances, and thus captures built-in as well as inadvertent weaknesses that may or may not be discovered in course of time.

BFM studied three natural combiners in this setting: those of concatenation, cascade, and xor combiners:

$$\begin{aligned} C_{|}^{H_1, H_2}(x) &:= H_1(x) \| H_2(x) & C_{\circ}^{H_1, H_2}(x) &:= H_2(H_1(x)) \\ C_{\oplus}^{H_1, H_2}(x) &:= H_1(x) \oplus H_2(x) . \end{aligned}$$

They showed, using new types of reductions to problems with high communication complexity, that central cryptographic security properties, such as one-way

¹ The model allows for a parameterization of the class of functions that can be computed. Both BFM and we here work with respect to the full set of functions.

security, pseudorandomness, and collision resistance are indeed achievable by these combiners.

The reductions to communication complexity problems are at times tedious and very specific to the combiner. Moreover, the hardness of the communication complexity problem underlying collision resistance is conjectural and still remains to be proven. Furthermore, a number of deployed protocols have only been shown to be secure in the random-oracle model, and thus may rely on properties beyond one-wayness, pseudorandomness, or collision resistance.

This raises the question whether or not other cryptographic properties expected from a good hash function are also met by these combiners. In other words:

Can combining two or more backdoored random oracles render access to independent but adaptive auxiliary information useless?

We formalize and study this question in the indistinguishability framework, which has been immensely successful in justifying the soundness of hash-function designs.

1.2 Indistinguishability

A common paradigm in the design of hash functions is to start with some underlying primitive, and through some construction build a more complex one. The provable security of such constructions have been analyzed through two main approaches. One formulates specific goals (such as collision resistance) and goes on to show that the construction satisfies them if its underlying primitives satisfy their own specific security properties. Another is a general approach, whose goal is to show that a (wide) class of security goals are simultaneously met.

The latter has been formalized in a number of frameworks, notably in the UC framework of Canetti [5], the reactive systems framework of Pfitzmann and Waidner [24], and the indistinguishability framework of Maurer, Renner, and Holenstein [23]. The latter is by now a standard methodology to study the soundness of cryptographic constructions, particularly symmetric ones such as hash functions [8,4] and block-ciphers [9,16,1,12] in idealized models of computation.

In the MRH framework, a public primitive H is available and the goal is to build another primitive, say a random oracle RO , from H through a construction C^H . Indistinguishability formalizes a set of necessary and sufficient conditions for the construction C^H to securely replace its ideal counterpart RO in a wide range of environments: for a *simulator* Sim , the systems (C^H, H) and (RO, Sim^{RO}) should be indistinguishable. The composition theorem proved by MRH states that, if C^H is indistinguishable from RO , then C^H can securely replace RO in arbitrary single-stage contexts. A central corollary of this composition theorem is that indistinguishability implies any single-stage security goal, which includes among others, one-wayness, collision resistance, PRG/PRF security, and more.

1.3 Contributions

With the above terminology in hand, the central question tackled in this work is whether or not combiners that are *indifferentiable* from a conventional (backdoor-free) random oracle exist, when the underlying primitives are two (or more) backdoored random oracles.

Let us consider the concatenation combiner $H_1(x)|H_2(x)$, where H_1 and H_2 are both backdoored. This construction was shown to be one-way, collision resistant, and PRG secure if both underlying functions are highly compressing. Despite this, the concatenation combiner cannot be indifferentiable from a random oracle: using the backdoor oracle for H_1 an attacker can compute two inputs x and x' such that $H_1(x) = H_1(x')$, query them to the construction and return 1 iff the left sides of the outputs match. However, any simulator attempting to find such a pair with respect to a backdoor-free random oracle must place an exponentially large number of queries. Attacks on the cascade combiner $H_2(H_1(x))$ were also given in [3, Section D.2] for a wider range of parameter regimes, leaving only the expand-then-compress case as potentially indifferentiable. Finally, the xor combiner $H_1(x) \oplus H_2(x)$, which is simpler, more efficient, and one of the most common ways to combine hash functions, resists these.²

DECOMPOSITION OF DISTRIBUTIONS. When proving results in the presence of auxiliary input, Uhrh [27] observed that pre-computation (or leakage) on a random oracle can reveal a significant amount of information only on restricted parts of its support. The problem of dealing with auxiliary input was later revised in a number of works [10,7,6]. In particular Coretti et al. [7], building on work in communication complexity, employed a *pre-sampling technique* to prove a number of positive results in the RO model with auxiliary input with tighter bounds. At a high level, this method permits writing a high min-entropy distribution (here, over a set of functions) as the convex combination of a (large) number of distributions which are fixed on a certain number (p) of points and highly unpredictable on the rest, the so-called $(p, 1 - \delta)$ -dense distributions. This technique was originally introduced in the work of G  os et al. [14].

THE SIMULATOR. Our simulator for the xor combiner builds on this technique to decompose distributions into a convex combination of $(p, 1 - \delta)$ -dense distributions. Simulation of backdoor oracles is arguably quite natural and proceeds as follows. Starting with uniform random oracles H_1 and H_2 , on each backdoor query f for H_1 the simulator computes $z = f(H_1)$ and updates the distribution of the random oracle H_1 to be uniform conditioned on the output of f being z . This distribution is then decomposed into a convex combination of $(p, 1 - \delta)$ -dense distributions, from which one function is sampled. For all of the p fixed points, the simulator sets the value of H_2 consistently with the random oracle and the distribution of H_2 is updated accordingly. An analogous procedure is implemented as the simulator for the second backdoored random oracle.

² Further, an indifferentiability proof of the expand-then-compress cascade combiner would closely follow that of the xor combiner and thus we focus on the latter here.

TECHNICAL ANALYSIS. The first technical contribution of our work is a refinement of the decomposition technique which can be used to *adaptively* decompose distributions after backdoor queries. We show that this refinement is sufficiently powerful to allow proving indifferenciability up to a logarithmic (in the input size of the BROs) number of switches between the backdoor queries. We prove this via a sequence of games which are carefully designed so as to be compatible with the decomposition technique. A key observation is that in contrast to previous works in the AI-RO model, we do not replace the dense (intuitively, unpredictable) part of the distribution of random oracles with uniform: backdoor functions “see” the entire table of the random oracle and this replacement would result in a noticeable change. Second, we modify the number of fixed points in the (partially) dense distributions so that progressively smaller sets of points are fixed. Even though each leakage corresponds to fixing a large number of points, it is proportionally smaller than the previous number of fixed points. Thus the overall bound remains small.

SIMULATOR EFFICIENCY. Our simulator runs in doubly exponential time in the bit-length of the random oracle and thus is of use in information-theoretic settings. These include the vast majority of symmetric constructions. Protocols based on computational assumptions (such as public-key encryption) escape this treatment: the overall adversary obtained via the composition would run the decomposition algorithm and hence will not be poly-time. This observation, however, also applies to the BRO model as the backdoor oracles also allow for non-polynomial time computation, trivially breaking any computational assumption if unrestricted. Despite this, in a setting where the computational assumption holds relative to the backdoor oracles, positive results may hold. We can for example restrict the backdoor capability to achieve this. Another promising avenue is to rely on an independent idealized model such as the generic-group model (GGM) and for instance, prove IND-CCA security of Hashed ElGamal in the BRO and (backdoor-free) GGM models. We leave exploring these solutions to future work.

AN EXTRACTOR-BASED COMBINER WITH IMPROVED SECURITY. We apply the above proof technique to the analysis of an alternative combiner for three independent backdoored random oracles, which relies on 2-out-of-3-source extractors that output good randomness as long as two out of the three of the inputs have sufficient min-entropy. Given such an extractor Ext , our combiner is

$$C_{\text{3ext}}^{H_1, H_2, H_3}(x) := \text{Ext}(H_1(x), H_2(x), H_3(x)) .$$

As mentioned above, our simulator for the xor combiner programs H_2 on the fixed points for H_1 (and vice versa) using the random oracle. This results in a loss since dense values are replaced with uniform values. In contrast, here the extractor ensures that image values are closer to uniform and thus the overall loss is lower. We show that a 2-out-of-3-source extractor can tolerate even a number of switches between the backdoor oracles which is slightly sub-linear in the size of the BRO inputs. This gives us more hope for unbounded adaptivity, in case improved decomposition techniques are found.

COMPOSITION. Let c denote the number of times the adversary switches between one backdoor oracle to the other. Regarding the query complexities of our simulators, each query to the backdoor oracle translates to roughly $N^{1-2^{-c}}$ queries to the random oracle for the xor combiner and roughly $N^{1-3/(c+3)}$ queries to the random oracle for the extractor combiner. This in particular means that, for a wide range of parameters, composition is only meaningful with respect to security notions whereby the random oracle can tolerate a large number of queries. This, for example, would be the case for one-way, PRG, and PRF security notions where the security bounds are of the form $\mathcal{O}(q/N)$. However, with respect to a smaller number of switches (as well as in the auxiliary-input setting with no adaptivity), collision resistance can still be achieved.

INDIFFERENTIABILITY WITH AUXILIARY INPUT. When our definition of indifferntiability is restricted so that only a single backdoor query to each hash function at the onset is allowed, we obtain a notion that formalizes *indifferentiability with auxiliary input*. This definition is interesting as it is sufficiently strong to allow for the generic replacement of random oracles with iterative constructions even in the presence of preprocessing attacks. Accordingly, our positive results in the BRO model when considered with no adaptivity translate to indifferntiability with independent preprocessing attacks. To complement this picture, we also discuss the case of auxiliary-input indifferntiability with a single BRO and show, as expected, that a salted indifferntiable construction is also indifferntiable with auxiliary input.

OPEN PROBLEMS. In order to overcome the bounded adaptivity restriction and prove full indifferntiability, one would require an improved decomposition technique which fixes considerably less points after each leakage. This, at the moment, seems (very) challenging and is left as an open question. In particular, such a result would simultaneously give new proofs of known communication complexity lower bounds for a host of problems, such as set-disjointness and intersection, potentially a proof of the conjecturally hard problem stated in [3], and many others. (We note that improved decomposition techniques can potentially also translate to improved bounds.) Indeed the xor combiner may achieve security well beyond what we establish here (and indeed the original work of BFM does so for specific games). Finally, as the extractor combiner suggests, the form of the combiner and the number of underlying BROs can also affect the overall bounds.

2 Preliminaries

Throughout the paper, when we write $[N]$ for any uppercase letter N , we use the convention that N is an integer and a power of two, i.e., $N = 2^n$ for some $n \in \mathbb{N}$. Let $[N] := \{0, \dots, N-1\}$ denote the set of all n -bit strings. We use $[M]^N$ to denote the set of all bit-strings of length $N \cdot \log M$, which corresponds to the set of all functions $F : [N] \rightarrow [M]$. We denote the uniform distribution over an arbitrary finite set S by \mathcal{U}_S .

For $F \in [M]^N$ and $I \subseteq [N]$ we denote by F_I the projection of F onto the points in I . Let μ be a probability density function over $[M]^N$. We define $\mu(D) := \Pr_{F \sim \mu}[F \in D]$ as the probability that a sample randomly drawn from μ falls into the domain $D \subseteq [M]^N$. By $\mu|_D$ we denote the density μ conditioned on the domain D . For a function $f : [M]^N \rightarrow \{0, 1\}^\ell$ and $z \in \{0, 1\}^\ell$, by $\mu|_{f(\cdot)=z}$ we denote μ conditioned on $f(F) = z$ for all $F \sim \mu|_{f(\cdot)=z}$.

For a set of assignments $A \subseteq \{(a, b) : (a, b) \in [N] \times [M]\}$, by $\mu|_A$ we denote μ conditioned on $F_{\{a\}} = b$ for all $(a, b) \in A$ and all $F \sim \mu|_A$. We further let $A_{.1} \subseteq [N]$ (resp. $A_{.2} \subseteq [M]$) denote the set containing the first (resp. second) coordinates of all elements in A .

For an algorithm Alg we denote by $\text{Alg}[param](input)$ a call of the algorithm with (constant) parameters $param$ and variable inputs $input$. This is to increase clarity among multiple calls to the algorithm about the main input, while the parameters remain unchanged.

2.1 Backdoored random oracles

We recall the definition of the backdoored random-oracle model from [3]. The $\text{BRO}(N_1, M_1, \dots, N_k, M_k)$ model (for some $k \in \mathbb{N}$) defines a setting where all parties have access to k functions H_1, \dots, H_k , where H_i 's are chosen uniformly and independently at random from $[M_i]^{N_i}$, while the adversarial parties also have access to the corresponding backdoor oracles BD_i 's. A backdoor oracle BD_i can be queried on functions f and return $f(H_i)$. If for all $i \in [k]$ we have $N_i = N$ and $M_i = M$, we simply refer to this model as $k\text{-BRO}(N, M)$ and when N and M are clear from the context, we simply use $k\text{-BRO}$.

These models may be weakened by restricting the adversary to query BD_i only on functions f in some capability class \mathcal{F}_i . However our results as well as those in [3] hold for arbitrary backdoor capabilities. In other words an adversary can (adaptively) query arbitrary functions f to any of the backdoor oracles.

2.2 Indifferentiability in the BRO model

We follow the indifferentiability framework of Maurer, Renner, and Holenstein (MRH) [23]. Here the underlying honest interfaces are k random oracles H_i and respective adversarial interfaces BD_i . We define the advantage of a differentiator \mathcal{D} with respect to a construction C^{H_i} and a simulator $\text{Sim}^{\text{RO}} := (\text{SimH}_i^{\text{RO}}, \text{SimBD}_i^{\text{RO}})$ as

$$\text{Adv}_{\text{C}^{H_i}, \text{Sim}}^{\text{indiff}}(\mathcal{D}) := \left| \Pr \left[\mathcal{D}^{\text{C}^{H_i}, H_i, \text{BD}_i} \right] - \Pr \left[\mathcal{D}^{\text{RO}, \text{SimH}_i^{\text{RO}}, \text{SimBD}_i^{\text{RO}}} \right] \right|,$$

where RO is a random oracle whose domain and co-domain match those of C .

We emphasize that the simulators do not get access to any backdoor oracles. This ensures that any attack against a construction with backdoors translates to one against the underlying random oracles *without* any backdoors.

2.3 Randomness extractors

Let X be a random variable. The min-entropy of X is defined as $\mathbf{H}_\infty(X) := -\log \max_x \Pr[X = x]$. The random variable X is called a (weak) k -source if $\mathbf{H}_\infty(X) \geq k$, i.e., $\Pr[X = x] \leq 2^{-k}$. The min-entropy of a distribution typically determines how many bits can be extracted from it which are close to uniform. The notion of closeness is formalized by the statistical distance. For two random variables X and Y over a common support D , their statistical distance is defined as $\text{SD}(X, Y) := \frac{1}{2} \sum_{z \in D} |\Pr[X = z] - \Pr[Y = z]|$.

In this paper we are interested in extractors that do not require seeds but rather rely on multiple weak sources.

Definition 1 (Multi-source extractors). *An efficient function $\text{Ext} : [N_1] \times \dots \times [N_t] \rightarrow [M]$ is a $(k_1, \dots, k_t, \varepsilon)$ -extractor if for all weak k_i -sources X_i over domains $[N_i]$, we have:*

$$\text{SD}(\text{Ext}(X_1, \dots, X_t), \mathcal{U}_{[M]}) \leq \varepsilon,$$

where ε is usually defined as a function of k_1, \dots, k_t . We call Ext an s -out-of- t $(k_1, \dots, k_t, \varepsilon)$ -extractor if $\text{Ext}(X_1, \dots, X_t)$ is ε -close to uniform even if only s sources fulfill the min-entropy condition.

Below we define useful classes of distributions, the so-called (partially) dense distributions, resp. dense probability density functions. Intuitively, bit strings from a dense distribution are unpredictable not only as a whole but also in any of their substrings and any combination of those substrings.

Definition 2 (Dense distributions). *Let μ be a probability density function over $[M]^N$. Then*

- μ is called $(1 - \delta)$ -dense if for $\mathbf{F} \sim \mu$, it holds that for every subset $I \subseteq [N]$ we have $\mathbf{H}_\infty(\mathbf{F}_I) \geq (1 - \delta) \cdot |I| \cdot \log M$.
- μ is called $(p, 1 - \delta)$ -dense if for $\mathbf{F} \sim \mu$ there exists a set $I \subseteq [N]$ of size $|I| \leq p$ such that $\mathbf{H}_\infty(\mathbf{F}_I) = 0$, while for every subset $J \subseteq [N] \setminus I$ we have $\mathbf{H}_\infty(\mathbf{F}_J) \geq (1 - \delta) \cdot |J| \cdot \log M$. That is, μ is fixed on at most p coordinates and $(1 - \delta)$ -dense on the rest.

We call a distribution dense, if the corresponding density function is dense.

3 Decomposition of High Min-Entropy Distributions

Any high min-entropy distribution can be written as a convex combination of distributions that are fixed on a number of points and dense on the rest (i.e., $(p, 1 - \delta)$ -dense distributions for some p and $\delta > 0$).³ The decomposition technique introduced by G600s et al. [14] has its origins in communication complexity

³ A convex combination of distributions μ_1, \dots, μ_n is a distribution that can be written as $\alpha_1 \cdot \mu_1 + \dots + \alpha_n \cdot \mu_n$, where $\alpha_1, \dots, \alpha_n$ are non-negative real numbers that sum up to 1.

theory. We generalize this technique, with a terminology closer to that of Kothari et al. [18], in order to allow for adaptive leakage. The original lemma, also used by Coretti et al. [7], can be easily derived as a special case of our lemma. For this, one assumes that the starting distribution before the leakage was uniform, in other words $(0, 1)$ -dense.

When proving results in the auxiliary-input random-oracle (AI-RO) model, Uhrh [27] observed that pre-computation (or leakage) on a random oracle can cause a significant decrease of its min-entropy only on restricted parts of its support (i.e., on p points), causing that part to become practically fixed, while the rest remains indistinguishable from random to a bounded-query distinguisher. This means that after fixing p coordinates of the random oracle, the rest can be lazily sampled from a uniform distribution. Coretti et al. [7] recently gave a different and tighter proof consisting of two main steps. First, the decomposition technique is used to show that the distribution of a random oracle given some leakage is *statistically* close to a $(p, 1 - \delta)$ -dense distribution. Second, they prove that no *bounded-query* algorithm can distinguish a $(p, 1 - \delta)$ -dense distribution from one that is fixed on the same p points and is otherwise uniform (a so-called p -bit-fixing distribution), as suggested by Uhrh [27].

Since in the BRO model adaptive queries are allowed, a function queried to the backdoor oracle is able to “see” the entire random oracle, rather than a restricted part of it. Hence, when analyzing the distribution of a random oracle after adaptive leakage, it is crucial that we keep the distributions *statistically close*. In other words we use $(p, 1 - \delta)$ -dense distributions instead of p -bit-fixing.

In the k -BRO model, we are concerned with multiple queries to the backdoor oracles, i.e., continuous and adaptive leakage that can depend on previously leaked information about both hash functions. Intuitively, since the leakage function can be arbitrary, it can in particular depend on the previously leaked values. We still need to argue that the distribution obtained after leakage about a $(p_{\text{prv}}, 1 - \delta_{\text{prv}})$ distribution, which is not necessarily uniform, is also close to a convex combination of $(p, 1 - \delta)$ distributions. Naturally, we have $\delta \geq \delta_{\text{prv}}$, since min-entropy decreases after new leakage, and $p \geq p_{\text{prv}}$, since additional points are fixed. Looking ahead, in the indistinguishability proofs, this refined decomposition lemma allows us to simply fix a new portion p_{frsh} of the simulated hash function after each leakage (i.e., backdoor query) and not to worry about the rest, which still has high entropy and can be lazily sampled (from a dense distribution) upon receiving the next query.

Lemma 1 (Refined decomposition after leakage). *Let μ be a $(p_{\text{prv}}, 1 - \delta_{\text{prv}})$ -dense density function over $[M]^N$ for some $p_{\text{prv}}, \delta_{\text{prv}} \geq 0$. Let $f : [M]^N \rightarrow \{0, 1\}^\ell$ be an arbitrary function and $z \in \{0, 1\}^\ell$ be a bit string. Then for any $p_{\text{frsh}}, \gamma > 0$, the density function conditioned on the leakage $\mu|_{f(\cdot)=z}$ is γ -close to a convex combination of finitely many $(p, 1 - \delta)$ -dense density functions for some p and δ such that*

$$p_{\text{prv}} \leq p \leq p_{\text{prv}} + p_{\text{frsh}} \quad \text{and} \quad \delta_{\text{prv}} \leq \delta \leq \frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{p_{\text{frsh}} \cdot \log M},$$

where $\ell_z := \mathbf{H}_\infty(\mathbf{G}) - \mathbf{H}_\infty(\mathbf{F})$ is the min-entropy deficiency of $\mathbf{F} \sim \mu|_{f(\cdot)=z}$ compared to $\mathbf{G} \sim \mu$.

Proof. This refined decomposition lemma differs from the original lemma in that the starting density function μ is $(p_{\text{prv}}, 1 - \delta_{\text{prv}})$ -dense. As a first step, we modify the original decomposition algorithm from [14, 18] so that it additionally gets the set of p_{prv} indices $I_{\text{prv}} \subseteq [N]$ that are already fixed in μ from the start.

Our refined decomposition algorithm **RefinedDecomp**, given below, recursively decomposes the domain $[M]^N$, according to the density function after leakage $\mu_z := \mu|_{f(\cdot)=z}$, into $d + 1$ partitions $D_1, \dots, D_d, D_{\text{err}} \subseteq [M]^N$ such that $(\bigcup_{i=1}^d D_i) \cup D_{\text{err}} = [M]^N$, where **err** stands for erroneous. For all i with $1 \leq i \leq d$ the partition D_i defines a $(p, 1 - \delta)$ -dense density function $\mu_z|_{D_i}$.

Each recursive call on a domain D to **RefinedDecomp** (other than the call leading to D_{err} , which we will discuss shortly) returns a pair (D_i, I_i) , where D_i represents a subset of $[M]^N$, where the images of all points in the set $I_i \subset [N]$ are fixed to the same values under all functions $\mathbf{H} \in D_i$. In other words, we have $\mathbf{H}_{I_i} = \alpha_i$ for some $\alpha_i \in [M]^{|I_i|}$. The algorithm finds such a pair (D_i, I_i) by considering the biggest set I_i (excluding those points fixed from the start, i.e., I_{prv}) such that the min-entropy of \mathbf{F}_{I_i} (for $\mathbf{F} \sim \mu_z|_D$) is too small (as determined by the rate δ) and then finding some α_i which is a very likely value of \mathbf{F}_{I_i} . Then I_i is returned with some D_i as the partition that contains all \mathbf{H} with $\mathbf{H}_{I_i} = \alpha_i$. The next recursive call will exclude D_i from the considered domain.

Decomposition halts either if the probability of a sample falling into the current domain is smaller than γ (i.e., $\mu_z(D) \leq \gamma$) or the current distribution is already $(p_{\text{prv}}, 1 - \delta)$ -dense. In both cases the algorithm returns the current domain D together with an empty set. In the former case the returned domain is marked as an erroneous domain $D_{\text{err}} := D$, since it may not define a $(p, 1 - \delta)$ -dense distribution. Let us without loss of generality assume that μ_z is not $(p_{\text{prv}} + p_{\text{frsh}}, 1 - \delta)$ -dense, as otherwise the claim holds trivially.

The formal definition of the algorithm **RefinedDecomp** is given below. We initialize the desired density rate as $\delta := \frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{p_{\text{frsh}} \cdot \log M}$ before calling **RefinedDecomp**.

RefinedDecomp $[\mu_z, \delta, \gamma, I_{\text{prv}}](D)$

if $\mu_z(D) \leq \gamma$ **then return** $(D_{\text{err}} \leftarrow D, \emptyset)$
if $\mu_z|_D$ is $(|I_{\text{prv}}|, 1 - \delta)$ -dense **then return** (D, \emptyset)
for $\mathbf{F} \sim \mu_z|_D$ let $I \subseteq [N]$ be a maximal set such that
 $\mathbf{H}_\infty(\mathbf{F}_I) < (1 - \delta) \cdot |I| \cdot \log M$ and $I \cap I_{\text{prv}} = \emptyset$.
let $\alpha \in [M]^{|I|}$ be such that $\Pr[\mathbf{F}_I = \alpha] > 2^{-(1-\delta) \cdot |I| \cdot \log M}$.
 $D_\alpha \leftarrow D \cap \{\mathbf{F} \in [M]^N \mid \mathbf{F}_I = \alpha\}$
 $D_{\neq \alpha} \leftarrow D \cap \{\mathbf{F} \in [M]^N \mid \mathbf{F}_I \neq \alpha\}$
return $((D_\alpha, I), \text{RefinedDecomp}[\mu_z, \delta, \gamma, I_{\text{prv}}](D_{\neq \alpha}))$

Now we turn our attention to proving that every partition D_i (other than D_{err}) returned by the above decomposition algorithm defines a density function $\mu_z|_{D_i}$ which is $(p, 1 - \delta)$ -dense.

Claim 1. *For all values of i with $1 \leq i \leq d$ it holds that the density function $\mu_z|_{D_i}$ is $(p, 1 - \frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{p_{\text{frsh}} \cdot \log M})$ -dense, where $p_{\text{prv}} \leq p \leq p_{\text{prv}} + p_{\text{frsh}}$.*

Proof. Let $\delta := \frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{p_{\text{frsh}} \cdot \log M}$. Let I be the set of freshly fixed points in $\mu_z|_{D_i}$ and $\overline{I \cup I_{\text{prv}}} := [N] \setminus (I \cup I_{\text{prv}})$. Let $\alpha_{\cup} \in [N]^{|I \cup I_{\text{prv}}|}$ be such that $F_{I \cup I_{\text{prv}}} = \alpha_{\cup}$ for $F \sim \mu_z|_{D_i}$. We first argue for the $(1 - \delta)$ -density of $\mu_z|_{D_i}$ on values projected to $\overline{I \cup I_{\text{prv}}}$ and afterwards bound the size of I .

1. Suppose $\mu_z|_{D_i}$ is not $(1 - \delta)$ -dense on $\overline{I \cup I_{\text{prv}}}$. Then there exists a non-empty set which violates the density property. That is, there exists a non-empty set $J \subseteq \overline{I \cup I_{\text{prv}}}$ and some $\beta \in [N]^{|J|}$ such that, with the probability taken over $F \sim \mu_z|_{D_i}$, we have:

$$\Pr[F_J = \beta] > 2^{-(1-\delta) \cdot |J| \cdot \log M}.$$

Now the union of the three sets $I^* := I \cup I_{\text{prv}} \cup J$ forms a new set such that for some $\beta^* \in [N]^{|I \cup I_{\text{prv}} \cup J|}$ we have

$$\begin{aligned} \Pr[F_{I^*} = \beta^*] &= \Pr[F_{I \cup I_{\text{prv}}} = \alpha_{\cup} \wedge F_J = \beta] \\ &= \Pr[F_{I \cup I_{\text{prv}}} = \alpha_{\cup}] \cdot \Pr[F_J = \beta | F_{I \cup I_{\text{prv}}} = \alpha_{\cup}] \\ &> 2^{-(1-\delta) \cdot |I \cup I_{\text{prv}}| \cdot \log M} \cdot 2^{-(1-\delta) \cdot |J| \cdot \log M} \\ &= 2^{-(1-\delta) \cdot |I \cup I_{\text{prv}} \cup J| \cdot \log M}. \end{aligned}$$

Since J was assumed to be non-empty and disjoint from $I \cup I_{\text{prv}}$ (and in particular with I), its existence violates the maximality of I . Therefore, $F_{\overline{I \cup I_{\text{prv}}}}$ is $(1 - \delta)$ dense.

2. We now bound the size of I , given that $\delta = \frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{p_{\text{frsh}} \cdot \log M}$. Let $F \sim \mu_z$ and $G \sim \mu$. We have $\mathbf{H}_{\infty}(F) = \mathbf{H}_{\infty}(G) - \ell_z \geq (1 - \delta_{\text{prv}}) \cdot (N - p_{\text{prv}}) \cdot \log M - \ell_z$, where the inequality holds, since μ is $(1 - \delta_{\text{prv}})$ -dense in $N - p_{\text{prv}}$ rows. Let $\beta \in [M]^{|I|}$. Then we have:

$$\begin{aligned} \Pr_{\mu_z|_{D_i}}[F_I = \beta] &\leq \Pr_{\mu_z}[F_I = \beta] / \mu_z(D_i) \\ &\leq \Pr_{\mu_z}[F_I = \beta] / \gamma \\ &= \sum_{\beta' \in [M]^{N-|I|-|I_{\text{prv}}|}} \Pr_{\mu_z}[F_I = \beta \wedge F_{[N] \setminus (I \cup I_{\text{prv}})} = \beta'] / \gamma \\ &\leq 2^{(N-|I|-p_{\text{prv}}) \cdot \log M} \cdot 2^{-\mathbf{H}_{\infty}(F)} / \gamma \\ &\leq 2^{(N-|I|-p_{\text{prv}}) \cdot \log M} \cdot 2^{-((1-\delta_{\text{prv}}) \cdot (N-p_{\text{prv}}) \cdot \log M - \ell_z)} / \gamma \\ &= 2^{\delta_{\text{prv}} \cdot N \cdot \log M - \delta_{\text{prv}} \cdot p_{\text{prv}} \cdot \log M - |I| \cdot \log M + \ell_z} / \gamma \end{aligned}$$

$$= 2^{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) - |I| \cdot \log M + \ell_z + \log \gamma^{-1}}.$$

Since by definition of the decomposition algorithm, there exists an $\alpha \in [M]^I$ such that $\Pr_{\mu_z|D_i}[\mathbf{F}_I = \alpha] > 2^{-(1-\delta) \cdot |I| \cdot \log M}$, we obtain

$$|I| \leq \frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{\delta \cdot \log M}.$$

Substituting δ by $\frac{\delta_{\text{prv}} \cdot \log M \cdot (N - p_{\text{prv}}) + \ell_z + \log \gamma^{-1}}{p_{\text{frsh}} \cdot \log M}$, we obtain $|I| \leq p_{\text{frsh}}$ and therefore, for the total number of fixed points $p := |I \cup I_{\text{prv}}|$ we get $p_{\text{prv}} \leq p \leq p_{\text{prv}} + p_{\text{frsh}}$, as stated in the claim. \square

Therefore, μ_z can be written as a convex combination of $\mu_z|_{D_1}, \dots, \mu_z|_{D_d}$ and $\mu_z|_{D_{\text{err}}}$, i.e., $\mu_z = \sum_{i=1}^d \mu_z(D_i) \cdot \mu_z|_{D_i} + \mu_z(D_{\text{err}}) \cdot \mu_z|_{D_{\text{err}}}$. Since $\mu_z(D_{\text{err}}) \leq \gamma$ when the algorithm **RefinedDecomp** terminates, the distribution μ_z is γ -close to a convex combination of $(p, 1 - \delta)$ distributions. \square

A special case of the above lemma for a uniform (i.e., $(0, 1)$ -dense) starting distribution μ , where $p_{\text{prv}} = 0$ and $\delta_{\text{prv}} = 0$, implies the bound $\delta \leq (\ell_z + \log \gamma^{-1}) / (p_{\text{frsh}} \cdot \log M)$ used by Coretti et al. [7].

REMARK. Note that the coefficient of δ_{prv} in the right hand side of the inequality established in the lemma is of the order $\mathcal{O}(N/p_{\text{frsh}})$. Looking ahead (see discussions on parameter estimation) this results in an increase in the number of points that the simulator needs to set. Thus any improvement in the bound established in this lemma would translate to tolerating a higher level of adaptivity and/or obtaining an improved bound.

Below we show that the expected min-entropy deficiency after leaking ℓ bits of information can be upper-bounded by ℓ bits.

Lemma 2. *Let \mathbf{F} be a random variable over $[M]^N$ and $f : [M]^N \rightarrow \{0, 1\}^\ell$ be an arbitrary function. Let $\ell_z := \mathbf{H}_\infty(\mathbf{F}) - \mathbf{H}_\infty(\mathbf{F}|f(\mathbf{F}) = z)$ be the min-entropy deficiency of $\mathbf{F}|f(\mathbf{F}) = z$. Then, we have $\mathbb{E}_{z \in f(\text{supp}(\mathbf{F}))}[\ell_z] \leq \ell$.*

Proof. Recall that $\tilde{\mathbf{H}}_\infty(A|B) := -\log(\mathbb{E}_b[\max_a \Pr[A = a|B = b]])$ defines the average min-entropy of A , given B .

$$\begin{aligned} \mathbb{E}_{z \in f(\text{supp}(\mathbf{F}))}[\ell_z] &= \mathbf{H}_\infty(\mathbf{F}) - \mathbb{E}_{z \in f(\text{supp}(\mathbf{F}))}[\mathbf{H}_\infty(\mathbf{F}|f(\mathbf{F}) = z)] \\ &\leq \mathbf{H}_\infty(\mathbf{F}) - \tilde{\mathbf{H}}_\infty(\mathbf{F}|f(\mathbf{F}) = z) \\ &\leq \mathbf{H}_\infty(\mathbf{F}) - \mathbf{H}_\infty(\mathbf{F}) + \log |f(\text{supp}(\mathbf{F}))| \leq \ell, \end{aligned}$$

where for deriving the second line we have used Jensen's inequality and for the third line we have used [11, Lemma 2.2.b].⁴ \square

⁴ The lemma is as follows. Let A, B be random variables. Then we have $\tilde{\mathbf{H}}_\infty(A|B) \geq \mathbf{H}_\infty(A, B) - n \geq \mathbf{H}_\infty(A) - n$, where B has at most 2^n possible values.

4 The xor Combiner

In this section, we study the indistinguishability of the xor combiner $C_{\oplus}^{H_1, H_2}(x) := H_1(x) \oplus H_2(x)$ in the 2-BRO model from a random oracle RO. We show indistinguishability against adversaries that switch between the two backdoor oracles BD_1 and BD_2 only a logarithmic number of times, while arbitrarily interleaving queries to the underlying BROs H_1 and H_2 , as well as to the random oracle RO.

To prove indistinguishability we need to show that there exists a simulator $\text{Sim} := (\text{Sim}H_1^{\text{RO}}, \text{Sim}H_2^{\text{RO}}, \text{Sim}BD_1^{\text{RO}}, \text{Sim}BD_2^{\text{RO}})$ such that no distinguisher placing a “reasonable” number of queries can distinguish

$$(C_{\oplus}^{H_1, H_2}, H_1, H_2, BD_1, BD_2) \quad \text{and} \quad (RO, \text{Sim}H_1^{\text{RO}}, \text{Sim}H_2^{\text{RO}}, \text{Sim}BD_1^{\text{RO}}, \text{Sim}BD_2^{\text{RO}}).$$

Such a simulator is described in Figure 1. Simulating the evaluation queries to H_1 and H_2 is straightforward. In simulating the backdoor queries, we take advantage of the decomposition technique (discussed in Section 3) for transforming high min-entropy distributions into distributions that have a number of fixed points and are dense otherwise. The backdoor simulator $\text{Sim}BD_1$ (resp. $\text{Sim}BD_2$) computes the queried function f on the truth table of H_1 (resp. H_2), where H_1 and H_2 are initialized by picking two functions uniformly at random. For the sake of simplicity, we consider an adversary that makes Q consecutive queries, ignoring evaluation and RO-queries in between, to one backdoor oracle before moving to the other. After the i -th sequence of Q queries to one of the backdoor oracles, the leaked backdoor information is translated into fixing p_i rows of the hash function such that the rest is dense and the resulting distribution is statistically close to the true one. In other words, the distribution conditioned on the leakage is γ -close (for some $\gamma > 0$) to a convex combination of $(p, 1 - \delta)$ -dense distributions obtained after decomposition.

Regarding the density rates δ_i ’s, we use odd values of i for the distributions obtained after backdoor queries on H_1 and even values of i for distributions of H_2 . Note that is crucial for the statistical distance of these two distributions on the entire table to remain small, since the distinguisher can adaptively query a backdoor oracle which sees and can depend on the *entire* hash function table (as opposed to a limited number of rows).

Finding a distribution, which is partly fixed and partly dense, is performed by the `FixRows` algorithm from 1. On input of a distribution μ_z , integer $p \in \mathbb{N}$, and a set $I_{\text{prv}} \in [N]$, the algorithm `FixRows` returns a new distribution which is fixed on points in a set I of size at most $p + |I_{\text{prv}}|$ and is for some δ , $(1 - \delta)$ -dense on the rest, together with a set of assignments A for elements in I according to the output distribution. The `FixRows` algorithm internally calls the refined decomposition algorithm, whose existence is guaranteed by Lemma 1 and its output distribution is one of the distributions in the convex combination returned by `RefinedDecomp`.

Upon fixing rows of one simulated BRO, the same rows in the other simulated BRO have to be fixed in a way that consistency with RO is assured. More precisely, for any x if $H_1(x)$ is fixed, the simulator $\text{Sim}BD_1$ will immediately

$\text{RO}(x)$ <hr/> if $\exists y \in [M]$ s.t. $(x, y) \in \text{hst}_{\text{RO}}$ then return y $y \leftarrow [M]$ $\text{hst}_{\text{RO}} \leftarrow \text{hst}_{\text{RO}} \cup \{(x, y)\}$ return y	
$\text{SimH}_1^{\text{RO}}(x)$ <hr/> $y_1 \leftarrow H_1(x)$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, \text{RO}(x) \oplus y_1)\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}; \mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $H_2 \leftarrow \mu_2$ return y_1	$\text{SimH}_2^{\text{RO}}(x)$ <hr/> $y_2 \leftarrow H_2(x)$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, \text{RO}(x) \oplus y_2)\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}; \mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ $H_1 \leftarrow \mu_1$ return y_2
$\text{SimBD}_1^{\text{RO}}[\bar{p}, \gamma](f)$ <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_1)$ $\mu_1 \leftarrow \mu_1 _{f(\cdot)=z}$ if $q = Q$ do $(\mu_1, A_1) \leftarrow \text{FixRows}[\gamma](\mu_1, p_{2s+1}, \text{hst}_{1.1})$ $H_1 \leftarrow \mu_1$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup A_1$ for $(x, y_1) \in A_1$ do $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, \text{RO}(x) \oplus y_1)\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $H_2 \leftarrow \mu_2$ $q \leftarrow 0$ return z	$\text{SimBD}_2^{\text{RO}}[\bar{p}, \gamma](f)$ <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_2)$ $\mu_2 \leftarrow \mu_2 _{f(\cdot)=z}$ if $q = Q$ then $(\mu_2, A_2) \leftarrow \text{FixRows}[\gamma](\mu_2, p_{2s+2}, \text{hst}_{2.1})$ $H_2 \leftarrow \mu_2$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup A_2$ for $(x, y_2) \in A_2$ do $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, \text{RO}(x) \oplus y_2)\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ $H_1 \leftarrow \mu_1$ $q \leftarrow 0$ $s \leftarrow s + 1$ return z
$\text{FixRows}[\gamma](\mu_z, p, I_{\text{prv}})$ <hr/> $((D_1, I_1), \dots, (D_d, I_d), (D_{\text{err}}, I_{\text{err}})) \leftarrow \text{RefinedDecomp}[\mu_z, p, \gamma, I_{\text{prv}}]([M]^N)$ $D_{\text{err}} \leftarrow [M]^N$ $(D_i, I_i) \leftarrow \{(D_1, I_1), \dots, (D_d, I_d), (D_{\text{err}}, I_{\text{err}})\}$ with probability $\mu_z(D_i)$, where $i \in \{1, \dots, d, \text{err}\}$ $A \leftarrow \emptyset; F \leftarrow D_i$ for $x \in I_i$ do $A \leftarrow A \cup \{(x, F(x))\}$ return $(\mu _{D_i}, A)$	

Fig. 1: Indifferentiability simulator for the xor combiner. We assume initial values $\text{hst}_1 = \text{hst}_2 = \text{hst}_{\text{RO}} := \emptyset$, $\mu_1 = \mu_2 := \mathcal{U}_{[M]^N}$, $H_1, H_2 \leftarrow \mathcal{U}_{[M]^N}$, $q := 0$, and $s := 0$.

set $H_2(x) := \text{RO}(x) \oplus H_1(x)$ (and, analogously, so does SimBD_2). The simulator specifies the number of points that it can afford to fix (since every such query requires a call to RO) and the statistical distance that it wants. Such a strategy to assure consistency with RO is also followed by evaluation simulators SimH_1 and SimH_2 , where only one coordinate of each BRO is fixed.

Note that the simulator SimBD_1 programs values of H_2 , which were supposed to be dense (after a first SimBD_2 query), to values that are uniform instead. Hence, we need to argue later that the statistical distance between a uniform and a dense distribution is small for the number of points that are being treated this way. This is formalized in Claim 2, below. Looking ahead, the need to keep the advantage of the differentiator small is the reason why the simulator adapts the number of fixed points with a differentiator's switch to the other backdoor oracle. Finally, via a hybrid argument we can upper bound the total number of random oracle queries by the simulator and the advantage of the differentiator.

Claim 2. *Let \mathcal{U} be the uniform distribution and \mathcal{V} be a $(1-\delta)$ -dense distribution, both over the domain $[M]^t$. Then we have $\text{SD}(\mathcal{U}, \mathcal{V}) \leq t \cdot \delta \cdot \log M$.*

Proof. This proof follows that of [7, Claim 3]. Let V_+ be the set of all values $z \in [M]^t$ for which $\Pr[\mathcal{V} = z] > 0$ holds. We can write the statistical distance between \mathcal{U} and \mathcal{V} as:

$$\begin{aligned} \text{SD}(\mathcal{U}, \mathcal{V}) &= \sum_{z \in [M]^t} \max \{0, \Pr[\mathcal{V} = z] - \Pr[\mathcal{U} = z]\} \\ &= \sum_{z \in V_+} \max \{0, \Pr[\mathcal{V} = z] - \Pr[\mathcal{U} = z]\} \\ &= \sum_{z \in V_+} \Pr[\mathcal{V} = z] \cdot \max \left\{0, 1 - \frac{\Pr[\mathcal{U} = z]}{\Pr[\mathcal{V} = z]}\right\}. \end{aligned}$$

Now, observe that for any value $z \in [M]^t$, we have $\Pr[\mathcal{V} = z] \leq M^{-(1-\delta) \cdot t}$ and $\Pr[\mathcal{U} = z] = M^{-t}$. Hence we have:

$$\text{SD}(\mathcal{U}, \mathcal{V}) \leq 1 - M^{-\delta \cdot t} \leq t \cdot \delta \cdot \log M,$$

where the last inequality uses the fact that for all $x \geq 0$, it holds that $2^{-x} \geq 1 - x$ (and hence, $x \geq 1 - 2^{-x}$). \square

The following theorem states our indistinguishability result for xor.

Theorem 1 (Indistinguishability of xor in 2-BRO with bounded adaptivity). *Consider the xor combiner $C_{\oplus}^{H_1, H_2}(x) := H_1(x) \oplus H_2(x)$ in the 2-BRO model with backdoored hash functions $H_1, H_2 \in [M]^N$. It holds that for any $\bar{p} := (p_1, \dots, p_{c+1}) \in \mathbb{N}^{c+1}$, $0 < \gamma < 1$, and an integer $c \geq 0$, there exists a simulator $\text{Sim}[\bar{p}, \gamma] := (\text{SimH}_1^{\text{RO}}, \text{SimH}_2^{\text{RO}}, \text{SimBD}_1^{\text{RO}}[\bar{p}, \gamma], \text{SimBD}_2^{\text{RO}}[\bar{p}, \gamma])$ such that for any differentiator \mathcal{D} that always makes Q queries to a backdoor oracle (starting from BD_1 and always receiving an ℓ -bit response) before switching to the other, with a total number of c switches, while being allowed to arbitrarily interleave up to q_H primitive queries as well as q_C construction queries, we have*

$$\begin{aligned} \text{Adv}_{C_{\oplus}^{H_1, H_2}, \text{Sim}[\bar{p}, \gamma]}^{\text{indiff}}(\mathcal{D}) &\leq (c+1) \cdot \gamma \\ &\quad + \log M \cdot \left(\sum_{i=1}^c p_i \cdot \delta_{i-1} + q_H \cdot \delta_{c+1} + q_C \cdot (\delta_c + \delta_{c+1}) \right), \end{aligned}$$

where $\delta_{-1} := \delta_0 := 0$ and the density rate after the i -th sequence of Q -many backdoor queries is $\delta_i := (\delta_{i-2} \cdot (N - \sum_{j=1}^{i-2} p_j) \cdot \log M + Q \cdot \ell + \log \gamma^{-1}) / (p_i \cdot \log M)$. The simulator places at most $q_{\text{Sim}} \leq q_{\text{H}} + \sum_{i=1}^{c+1} p_i$ queries to the random oracle RO.

Proof. We prove indistinguishability by (1) defining a simulator, (2) upper bounding the advantage of any differentiator in distinguishing the real and the simulated worlds, and (3) upper bounding the number of queries that the simulator makes to the random oracle.

Simulator. All four sub-algorithms of the simulator are described in Figure 1. They share state, in particular, variables to keep track of the fixed history and the current distribution of the hash functions. Two sets $\text{hst}_1, \text{hst}_2$ are used to keep track of the fixed coordinates of the simulated hash functions H_1 and H_2 , respectively. The density functions, from which the simulated backdoored hash functions will be sampled, are denoted by μ_1 and μ_2 . Furthermore, the simulator uses a counter s to recognize switches from one backdoor oracle to the other in order to use the appropriate number of points to fix from the list \bar{p} . It also maintains a counter q for counting the number of consecutive queries to a backdoor oracle in order to decompose, i.e., substitute the current distribution with a partially fixed and partially dense distribution, only when necessary which is the case after each set of Q backdoor queries. We assume the initial values $\mu_1 = \mu_2 := \mathcal{U}_{[M]^N}$, $H_1, H_2 \leftarrow \mathcal{U}_{[M]^N}$, $\text{hst}_1 = \text{hst}_2 = \text{hst}_{\text{RO}} := \emptyset$, $q := 0$, and $s := 0$.

Security analysis. Here we analyze the indistinguishability of the xor combiner using a sequence of eight games $\text{Game}_0, \dots, \text{Game}_7$, where Game_0 and Game_7 are the real and ideal indistinguishability games, respectively. In the following we use the shorthand notation $\Pr[\mathcal{D}^{\text{Game}_i}] := \Pr[\mathcal{D}^{\text{Game}_i} = 1]$, where $\mathcal{D}^{\text{Game}_i}$ indicates the interaction of an adversary \mathcal{D} with a game Game_i . We define the intermediate games Game_1 through Game_6 by gradually modifying the oracles and highlighting the changes in each step. Unchanged oracles are omitted in games and correspond to those from their direct predecessor. We bound the advantage of differentiators in distinguishing every two consecutive games.

$\text{Game}_0 : \mathbb{C}_{\oplus}^{H_1, H_2}(x)$			
$y_1 \leftarrow H_1(x); y_2 \leftarrow H_2(x)$ $y \leftarrow y_1 \oplus y_2$ return y			
$\text{Game}_0 : H_1(x)$	$\text{Game}_0 : H_2(x)$	$\text{Game}_0 : \text{BD}_1(f)$	$\text{Game}_0 : \text{BD}_2(f)$
$y_1 \leftarrow H_1(x)$	$y_2 \leftarrow H_2(x)$	$z \leftarrow f(H_1)$	$z \leftarrow f(H_2)$
return y_1	return y_2	return z	return z

Game_1 . We next update the distributions of hash functions based on past evaluation queries, backdoor queries, and the history of coordinates that are fixed

through construction queries. The distributions μ_i are conditioned on these updates, but are never actually used (i.e., sampled from) in the game. Hence it is easy to see that these two games are identical, i.e., $\text{SD}(\text{Game}_0, \text{Game}_1) = 0$.

$\text{Game}_1 : \mathbb{C}_{\oplus}^{\text{H}_1, \text{H}_2}(x)$			
$y_1 \leftarrow \text{H}_1(x); y_2 \leftarrow \text{H}_2(x)$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}; \text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}; \mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $y \leftarrow y_1 \oplus y_2$ return y			
$\text{Game}_1 : \text{H}_1(x)$	$\text{Game}_1 : \text{H}_2(x)$	$\text{Game}_1 : \text{BD}_1(f)$	$\text{Game}_1 : \text{BD}_2(f)$
$y_1 \leftarrow \text{H}_1(x)$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ return y_1	$y_2 \leftarrow \text{H}_2(x)$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ return y_2	$z \leftarrow f(\text{H}_1)$ $\mu_1 \leftarrow \mu_1 _{f(\cdot)=z}$ return z	$z \leftarrow f(\text{H}_2)$ $\mu_2 \leftarrow \mu_2 _{f(\cdot)=z}$ return z

Game_2 . Here, after each sequence of Q queries to a backdoor oracle, i.e., right before a switch, a $(p, 1 - \delta)$ -dense distribution μ'_i is obtained using the algorithm `FixRows` by decomposing the distribution of the corresponding hash function after responding to the last query (i.e., $\mu_i|_{f(\cdot)=z}$). However, since the new distributions μ'_i are never actually used elsewhere, Game_2 remains identical to Game_1 , i.e., $\text{SD}(\text{Game}_1, \text{Game}_2) = 0$.

$\text{Game}_2 : \text{BD}_1(f)$	$\text{Game}_2 : \text{BD}_2(f)$
$q \leftarrow q + 1$ $z \leftarrow f(\text{H}_1); \mu_1 \leftarrow \mu_1 _{f(\cdot)=z}$ if $q = Q$ then $(\mu'_1, A_1) \leftarrow \text{FixRows}[\gamma](\mu_1, p_{2s+1}, \text{hst}_{1,1})$ $q \leftarrow 0$ return z	$q \leftarrow q + 1$ $z \leftarrow f(\text{H}_2); \mu_2 \leftarrow \mu_2 _{f(\cdot)=z}$ if $q = Q$ then $(\mu'_2, A_2) \leftarrow \text{FixRows}[\gamma](\mu_2, p_{2s+2}, \text{hst}_{2,1})$ $q \leftarrow 0$ $s \leftarrow s + 1$ return z

Game_3 . In this game, evaluation queries on a value x , fix the image of both functions, i.e., to $\text{H}_1(x)$ and $\text{H}_2(x)$. Similarly, in backdoor simulation the rows in the assignments A_1 (resp. A_2) are fixed for the other hash function H_2 (resp. H_1) according to its current distribution. In both games, the oracles' responses are at all times consistent with their past responses (and the construction) and we still do not sample from the updated distributions. Hence, it does not matter, if more or less of the hash function tables are fixed in each query and therefore the two games are identical, i.e., $\text{SD}(\text{Game}_2, \text{Game}_3) = 0$.

Game₃ : H₁(x) <hr/> $y_1 \leftarrow H_1(x)$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, H_2(x))\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}; \mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ return y_1	Game₃ : H₂(x) <hr/> $y_2 \leftarrow H_2(x)$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, H_1(x))\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}; \mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ return y_2
Game₃ : BD₁(f) <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_1); \mu_1 \leftarrow \mu_1 _{f(\cdot)=z}$ if $q = Q$ then $(\mu'_1, A_1) \leftarrow \text{FixRows}[\gamma](\mu_1, p_{2s+1}, \text{hst}_{1.1})$ for $x \in A_{1.1}$ do $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, H_2(x))\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $q \leftarrow 0$ return z	Game₃ : BD₂(f) <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_2); \mu_2 \leftarrow \mu_2 _{f(\cdot)=z}$ if $q = Q$ then $(\mu'_2, A_2) \leftarrow \text{FixRows}[\gamma](\mu_2, p_{2s+2}, \text{hst}_{2.1})$ for $x \in A_{2.1}$ do $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, H_1(x))\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ $q \leftarrow 0$ $s \leftarrow s + 1$ return z

Game₄. In this game the distributions obtained by decomposition actually replace the distributions conditioned on leakage. Hence, the histories are also updated and a new hash function H_i is later sampled for potential usage in the construction. According to Lemma 1, there is a convex combination of $(p, 1 - \delta)$ -dense distributions which is γ -close to the real distribution, one of such distributions being the one returned by **FixRows**. Hence, the distinguishing advantage can increase by γ for every Q sequence of backdoor queries. I.e., $|\Pr[\mathcal{D}^{\text{Game}_3}] - \Pr[\mathcal{D}^{\text{Game}_4}]| \leq (c + 1) \cdot \gamma$.

Game₄ : BD₁(f) <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_1); \mu_1 \leftarrow \mu_1 _{f(\cdot)=z}$ if $q = Q$ then $(\mu_1, A_1) \leftarrow \text{FixRows}[\gamma](\mu_1, p_{2s+1}, \text{hst}_{1.1})$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup A_1$ $H_1 \leftarrow \mu_1$ for $x \in A_{1.1}$ do $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, H_2(x))\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $q \leftarrow 0$ return z	Game₄ : BD₂(f) <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_2); \mu_2 \leftarrow \mu_2 _{f(\cdot)=z}$ if $q = Q$ then $(\mu_2, A_2) \leftarrow \text{FixRows}[\gamma](\mu_2, p_{2s+2}, \text{hst}_{2.1})$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup A_2$ $H_2 \leftarrow \mu_2$ for $x \in A_{2.1}$ do $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, H_1(x))\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ $q \leftarrow 0$ $s \leftarrow s + 1$ return z
---	---

Game₅. This game behaves exactly as **Game₄** except when fixing the same rows for the distribution of the other BRO. It fixes those points by calling C_\oplus (rather than directly) and then redundantly updates the history with e.g.,

some $(x, H_1(x) \oplus C_\oplus(x))$ and samples a new BRO from the updated distribution. However, since the construction C_\oplus itself calls the BROs, Game_5 is only taking a detour and the two games are perfectly indistinguishable. Hence $\text{SD}(\text{Game}_4, \text{Game}_5) = 0$.

<p>Game₅ : H₁(x)</p> <hr/> $y_1 \leftarrow H_1(x)$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, C_\oplus(x) \oplus y_1)\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}; \mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $H_2 \leftarrow \mu_2$ return y_1	<p>Game₅ : H₂(x)</p> <hr/> $y_2 \leftarrow H_2(x)$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, C_\oplus(x) \oplus y_2)\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}; \mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ $H_1 \leftarrow \mu_1$ return y_2
<p>Game₅ : BD₁(f)</p> <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_1); \mu_1 \leftarrow \mu_1 _{f(\cdot)=z}$ if $q = Q$ then $(\mu_1, A_1) \leftarrow \text{FixRows}[\gamma](\mu_1, p_{2s+1}, \text{hst}_{1.1})$ $\text{hst}_1 \leftarrow \text{hst}_1 \cup A_1$ $H_1 \leftarrow \mu_1$ for $(x, y_1) \in A_1$ do $\text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, C_\oplus(x) \oplus y_1)\}$ $\mu_2 \leftarrow \mu_2 _{\text{hst}_2}$ $H_2 \leftarrow \mu_2$ $q \leftarrow 0$ return z	<p>Game₅ : BD₂(f)</p> <hr/> $q \leftarrow q + 1$ $z \leftarrow f(H_2); \mu_2 \leftarrow \mu_2 _{f(\cdot)=z}$ if $q = Q$ then $(\mu_2, A_2) \leftarrow \text{FixRows}[\gamma](\mu_2, p_{2s+2}, \text{hst}_{2.1})$ $\text{hst}_2 \leftarrow \text{hst}_2 \cup A_2$ $H_2 \leftarrow \mu_2$ for $(x, y_2) \in A_2$ do $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, C_\oplus(x) \oplus y_2)\}$ $\mu_1 \leftarrow \mu_1 _{\text{hst}_1}$ $H_1 \leftarrow \mu_1$ $q \leftarrow 0$ $s \leftarrow s + 1$ return z

Game₆. We now modify C_\oplus to start to resemble a lazily sampled random oracle. In the new construction oracle, a query is stored together with its image in the history hst_{RO} . In case a query is repeated, its stored image is simply returned. Otherwise, there are three cases to consider: the corresponding row to the current query x is fixed in both hash functions, in one of them, or in neither one. In the first case the output of the construction is computed by xoring the individual images stored in hst_1 and hst_2 . In the second case, a uniformly random value is chosen (and later stored in hst_{RO}). In the final case, Game_6 behaves exactly as Game_5 . So, the distinguishing advantage is bounded by distinguishing uniform points (set to uniform when xoring with the returned uniform value of C_\oplus) from dense points. In fact, according to Claim 2, for each evaluation query it adds at most $\delta_{c+1} \cdot \log M$, since δ_i 's are increasing. Further, for all points that are fixed upon a backdoor query this adds $p_i \cdot \delta_{i-1} \cdot \log M$, except for the last one, since there will be no backdoor query after that which can see the entire p_{c+1} points.

$$|\Pr[\mathcal{D}^{\text{Game}_5}] - \Pr[\mathcal{D}^{\text{Game}_6}]| \leq \log M \cdot \left(\sum_{i=1}^c p_i \cdot \delta_{i-1} + q_H \cdot \delta_{c+1} \right)$$

```

Game6 :  $\mathcal{C}_{\oplus}^{H_1, H_2}(x)$ 

if  $\exists y \in [M]$  s.t.  $(x, y) \in \text{hst}_{\text{RO}}$  then return  $y$ 
if  $\exists y_1, y_2 \in [M]$  s.t.  $(x, y_1) \in \text{hst}_1 \wedge (x, y_2) \in \text{hst}_2$  then return  $y_1 \oplus y_2$ 
if  $\exists y' \in [M]$  s.t.  $(x, y') \in \text{hst}_1 \vee (x, y') \in \text{hst}_2$  then
   $y \leftarrow [M]$ 
else
   $y_1 \leftarrow H_1(x); y_2 \leftarrow H_2(x)$ 
   $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}; \text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ 
   $\mu_1 \leftarrow \mu_1|_{\text{hst}_1}; \mu_2 \leftarrow \mu_2|_{\text{hst}_2}$ 
   $y \leftarrow y_1 \oplus y_2$ 
   $\text{hst}_{\text{RO}} \leftarrow \text{hst}_{\text{RO}} \cup \{(x, y)\}$ 
return  $y$ 

```

Game₇. The construction oracle in this game differs from **Game**₆ in that it never evaluates the individual hash functions anymore. Here, we can safely remove the second case distinction, where x is in both hst_1 and hst_2 , since this case is covered by the first case where x has been queried to the construction itself. It remains to bound the distinguisher's advantage in distinguishing the two games while making queries x to the construction that are prior to the query fixed for neither hash function.

Claim. Let X and Y be two independent $(1 - \delta)$ and $(1 - \delta')$ -dense distributions over a domain $[M]^N$. Then the xor distribution $X \oplus Y$ is $(1 - (\delta + \delta'))$ -dense over the same domain $[M]^N$.

Proof. Let $I \subseteq [N]$ and $z \in [M]^{|I|}$ be arbitrary. Then we have:

$$\begin{aligned}
\Pr[X_I \oplus Y_I = z] &= \sum_x \Pr[X_I = x \wedge Y_I = x \oplus z] = \sum_x \Pr[X_I = x] \cdot \Pr[Y_I = x \oplus z] \\
&\leq 2^{|I| \cdot \log M} \cdot 2^{-(1-\delta) \cdot |I| \cdot \log M} \cdot 2^{-(1-\delta') \cdot |I| \cdot \log M} \\
&= 2^{-(1-(\delta+\delta')) \cdot |I| \cdot \log M} . \quad \square
\end{aligned}$$

We can now bound the distinguisher's advantage by computing the distance between the sum of two dense distributions from uniform, given that only $q_{\mathcal{C}}$ queries to \mathcal{C}_{\oplus} are allowed. Below, in the second line, we use the fact that according to Lemma 1, δ 's should increase.

$$\left| \Pr[\mathcal{D}^{\text{Game}_6}] - \Pr[\mathcal{D}^{\text{Game}_7}] \right| \leq q_{\mathcal{C}} \cdot \log M \cdot \max_{0 \leq i \leq c} \{\delta_i + \delta_{i+1}\} = q_{\mathcal{C}} \cdot \log M \cdot (\delta_c + \delta_{c+1}) .$$

```

Game7 :  $C_{\oplus}^{H_1, H_2}(x)$ 

if  $\exists y \in [M]$  s.t.  $(x, y) \in \text{hst}_{\text{RO}}$  then return  $y$ 
if  $\exists y_1, y_2 \in [M]$  s.t.  $(x, y_1) \in \text{hst}_1 \wedge (x, y_2) \in \text{hst}_2$  then return  $y_1 \oplus y_2$ 
if  $\exists y' \in [M]$  s.t.  $(x, y') \in \text{hst}_1 \vee (x, y') \in \text{hst}_2$  then
   $y \leftarrow [M]$ 
else
   $y_1 \leftarrow H_1(x); y_2 \leftarrow H_2(x)$ 
   $\text{hst}_1 \leftarrow \text{hst}_1 \cup \{(x, y_1)\}; \text{hst}_2 \leftarrow \text{hst}_2 \cup \{(x, y_2)\}$ 
   $\mu_1 \leftarrow \mu_1|_{\text{hst}_1}; \mu_2 \leftarrow \mu_2|_{\text{hst}_2}$ 
   $y \leftarrow y_1 \oplus y_2$ 
 $\text{hst}_{\text{RO}} \leftarrow \text{hst}_{\text{RO}} \cup \{(x, y)\}$ 
return  $y$ 

```

The last game Game_7 is identical to the simulated world. Therefore, the overall advantage of \mathcal{D} is as stated in the theorem.

Query complexity. The queries made by the simulator to RO consist of those made when simulating evaluation queries and those made when simulating backdoor queries. Responding to each evaluation query requires exactly one query to RO , which makes a total of q_H queries. Right after the Q -th consecutive backdoor query (i.e., right before a switch), the simulator fixes some rows of the other BRO , where for each fixed row one query to the random oracle RO is made. The maximum number of rows that should be fixed after each sequence of Q queries to BD_1 (resp. BD_2) is predetermined by the simulator's parameter \bar{p} . Hence we obtain the claimed query complexity $q_H + \sum_{i=1}^{c+1} p_i$. \square

We now provide estimates for the involved parameters.

Corollary 1. *Let the number of switches be $c \geq 1$. Then for any $\alpha_1 > 1 - 1/F_{c+1}$, where F_i are the Fibonacci numbers, there is an indistinguishability simulator Sim for the C_{\oplus} construction in the 2-BRO model which has query complexity $q_H + (c+1) \cdot N^{\alpha_1}$ for any distinguisher with q_H queries to the underlying BROs . Furthermore, any such distinguisher which places q_C construction queries and Q consecutive queries to the same backdoor oracle before switching has advantage at most*

$$(c+1) \cdot \gamma + \log M \cdot (c^2 B + 2q_H + 2q_C) \cdot N^{(1-\alpha_1) \cdot F_{c+1}/F_{c+2} - 1/F_{c+2}},$$

against the simulator, where $B := (Q\ell + \log \gamma^{-1})/\log M$. Asymptotically the query complexity is $q_H + \mathcal{O}(N^{1-1/F_{c+2}})$ and the advantage $\mathcal{O}((q_H + q_C) \cdot Q \cdot \ell / N^{0.38/F_{c+2}})$.

Proof. From Lemma 1 we have that

$$\delta_i \leq (\delta_{i-2} \cdot A + B)/p_i,$$

where $A := N$ and $B := (Q\ell + \log \gamma^{-1})/\log M$. Recursively applying the equation we get for odd i

$$\delta_i \leq \frac{B}{p_i} + \frac{AB}{p_i p_{i-2}} + \cdots + \frac{A^{(i-1)/2} B}{p_i p_{i-2} \cdots p_1}$$

Using $p_i < A$, the terms progressively get larger. Thus, in general

$$\delta_i \leq \frac{c \cdot N^{(i-2+i \bmod 2)/2} B}{p_i p_{i-2} \cdots p_{1+(i+1) \bmod 2}}.$$

For the indifferentiability advantage to be small, we would need to minimize

$$\sum_{i=1}^c p_i \cdot \delta_{i-1} + (q_H + q_C)(\delta_c + \delta_{c+1}).$$

Let's assume $p_i = N^{\alpha_i}$ for some $\alpha_i \in [0, 1)$. Then the i -th summand for $i > 1$ is

$$c \cdot B \cdot N^{\alpha_i - \alpha_{i-1} - \alpha_{i-3} - \cdots - \alpha_{1+i \bmod 2} + (i-3+(i-1) \bmod 2)/2}.$$

To minimize, we set all terms equal to a common value $c \cdot B \cdot N^\theta$. We obtain

$$\alpha_i - \alpha_{i-1} - \cdots - \alpha_{1+i \bmod 2} + (i-3+(i-1) \bmod 2)/2 = \theta,$$

Solving this system of linear equations gives

$$\alpha_i = F_i \cdot \theta + F_{i-1} \cdot (\alpha_1 - 1) + 1,$$

where F_i are the Fibonacci numbers with $F_0 = 0$ and $F_1 = 1$.

We may arrange the terms so that $(\delta_c + \delta_{c+1}) = 2 \cdot N^\theta$ (not including the $(q_H + q_C)$ factor). To this end, we set $\alpha_{c+2} = 0$ so that $\delta_{c+1} = N^\theta/p_{c+2} = N^\theta$ and $\delta_c = N^\theta/p_{c+1} \leq N^\theta/p_{c+2} = N^\theta$. Thus we set $\alpha_{c+2} = 0$. This gives $\theta = (1 - \alpha_1) \cdot F_{c+1}/F_{c+2} - 1/F_{c+2}$. Now for $\theta < 0$ we would need that $\alpha_1 > 1 - 1/F_{c+1}$. This means that the query complexity of the simulator is $q_H + (c+1) \cdot N^{\alpha_1}$ and its advantage is

$$(c+1) \cdot \gamma + \log M \cdot (c^2 B + 2q_H + 2q_C) \cdot N^{(1-\alpha_1) \cdot F_{c+1}/F_{c+2} - 1/F_{c+2}}.$$

We obtain the bound stated in the asymptotic part of the corollary by setting $\alpha_1 := 1 - 1/F_{c+2} > 1 - 1/F_{c+1}$. \square

We note that in the special case where $c = 1$, we must have that $\alpha_1 > 1 - 1/F_2 = 0$. In particular we can set $\alpha_1 := 1/4$ to obtain a simulator that places $N^{\alpha_1} = N^{1/4} \leq \sqrt{N}$ queries. Thus in this case we obtain collision resistance. Note, however, that as soon as $c \geq 2$ we would need to have that $\alpha_1 > 1 - 1/F_3 = 1/2$, which means the simulator places at least \sqrt{N} queries, and we do not get collision resistance.

The above corollary shows that the xor combiner can only tolerate a logarithmic number of switches in $\log N$, which we think of as the security parameter. This is due to the fact that the simulator complexity needs to be less than $N/2$ for it to be non-trivial. Although our bounds are arguably weak, they are still meaningful, and we conjecture that much better bounds in reality hold.

5 An Extractor-Based Combiner

In this section we study the indistinguishability of extractor-based combiners and show that they can give better security parameters compared to the xor combiner of Section 4. Recall that in the k -BRO model one considers adversaries that have access to all k backdoor oracles. A query to the backdoor oracle BD_i reveals some information about the underlying BRO H_i . The resulting distribution conditioned on the leakage can, using the decomposition technique, be translated into a distribution with a number of fixed coordinates, while the distribution of the rest remains dense. An indistinguishability simulator then fixes the same rows of the other BRO(s) in a way that consistency with the random oracle (which is to be indistinguishable from the construction) is ensured.

We demonstrated this idea for the xor combiner, where, before a switch to the other backdoor oracle, the simulator substituted p images of that BRO by uniformly random values, i.e., the result of the random oracle values xored with the ones just fixed. This causes a security loss of $p \cdot \delta \cdot \log M$ per switch, which corresponds to the advantage of an adversary distinguishing p uniform values from $(1 - \delta)$ -dense ones. Now consider a multi-source $(k_1, \dots, k_t, \varepsilon)$ -extractor as the combiner in t -BRO. The hope would be that as long as the images of the BROs have high min-entropy, the output of the extractor is ε -close to uniform. This makes it possible for us to express the loss described above in terms of a negligible ε and forgo the requirement on δ to be negligible.

In this section we focus on 2-out-of-3-source extractors as combiners, i.e., extractors that only require a minimal amount of min-entropy from two of the sources. More formally, let $\text{Ext} : [M]^3 \rightarrow [2]$ be a 2-out-of-3-source $(k_1, k_2, k_3, \varepsilon)$ -extractor. For three functions $H_1, H_2, H_3 : [N] \rightarrow [M]$, the combiner $C_{3\text{ext}}^{H_1, H_2, H_3} : [N] \rightarrow [2]$ is defined as $C_{3\text{ext}}^{H_1, H_2, H_3}(x) := \text{Ext}(H_1(x), H_2(x), H_3(x))$. Here we show that in the 3-BRO model the construction $C_{3\text{ext}}^{H_1, H_2, H_3}$ is indistinguishable from a random oracle.

WHY NOT A TWO-SOURCE EXTRACTOR? Note that we cannot guarantee that images which are being fixed by the simulator in some H_i as a result of a BD_i -query have *any* min-entropy whatsoever. To understand why, simply consider an adversary that makes a backdoor query to BD_1 requesting a preimage of the zero-string $y^* := 0^{\log M}$ under H_1 . Suppose BD_1 responds to this query with $x^* \in [N]$. In this case $H_1(x^*)$ has no min-entropy, since $y^* = H_1(x^*)$ was chosen by the adversary and is, therefore, completely predictable. Hence, $H_1(x^*)$ cannot be used in a (k_1, k_2, ε) -two-source extractor, i.e., $\text{Ext}(H_1(x^*), H_2(x^*))$, which relies on min-entropy from both sources for its output to be ε -close to uniform. Overall, using a two-source extractor does not seem to have any advantage over the xor combiner in the 2-BRO model. On the contrary, when using a 2-out-of-3-source extractor, assuming that the rows under consideration are not already fixed in the function tables of all three BROs due to some previous query, there will be two images with high min-entropy, from which we can extract a value ε -close to uniform.

Theorem 2 (Indifferentiability of 2-out-of-3-source extractors in the 3-BRO model with bounded adaptivity). *Let $\text{Ext} : [M]^3 \rightarrow [2]$ be a $(k_1, k_2, k_3, \varepsilon)$ -2-out-of-3-source randomness extractor, where ε is a function of k_1, k_2, k_3 . Consider the combiner $C_{\text{3ext}}^{\text{H}_1, \text{H}_2, \text{H}_3}(x) := \text{Ext}(\text{H}_1(x), \text{H}_2(x), \text{H}_3(x))$ in the 3-BRO model with backdoored hash functions $\text{H}_1, \text{H}_2, \text{H}_3 \in [M]^N$. It holds that for all values of $\bar{p} := (p_1, \dots, p_{c+1}) \in \mathbb{N}^{c+1}$, $0 < \gamma < 1$, and an integer $c \geq 0$, there exists a simulator $\text{Sim}[\bar{p}, \gamma] := (\text{SimH}_1^{\text{RO}}, \text{SimH}_2^{\text{RO}}, \text{SimH}_3^{\text{RO}}, \text{SimBD}_1^{\text{RO}}[\bar{p}, \gamma], \text{SimBD}_2^{\text{RO}}[\bar{p}, \gamma], \text{SimBD}_3^{\text{RO}}[\bar{p}, \gamma])$ such that for any differentiator \mathcal{D} that always makes Q queries to one backdoor oracle (always receiving an ℓ -bit response) before switching to the next, with a total number of c switches, while arbitrarily interleaving up to q_{H} primitive queries and q_{C} construction queries, we have*

$$\begin{aligned} \text{Adv}_{C_{\text{3ext}}^{\text{H}_1, \text{H}_2, \text{H}_3}, \text{Sim}[\bar{p}, \gamma]}^{\text{indiff}}(\mathcal{D}) &\leq (c+1) \cdot \gamma \\ &+ \sum_{i=1}^c \text{SD}(E_1 | \dots | E_{p_i}, \mathcal{U}_{[2]^{p_i}}) + q_{\text{H}} \cdot \text{SD}(E_1, \mathcal{U}_{[2]}) \\ &+ q_{\text{C}} \cdot \varepsilon((1-\delta_{c-1}) \cdot \log M, (1-\delta_c) \cdot \log M, (1-\delta_{c+1}) \cdot \log M), \end{aligned}$$

where for all $n \in \mathbb{N}$, we define $E_n := \text{Ext}(X, Y, Z)$ for some random variables X, Y, Z over $[M]$ such that at least 2 of them have min-entropy $(1-\delta_c) \cdot \log M$. Furthermore, we let $\delta_{-2} := \delta_{-1} := \delta_0 := 0$ and for other values of $i \leq c+1$ let $\delta_i := (\delta_{i-3} \cdot (N - \sum_{j=1}^{i-3} p_j) \cdot \log M + Q \cdot \ell + \log \gamma^{-1}) / (p_i \cdot \log M)$ be the density rate after the i -th sequence of Q -many backdoor queries. The simulator places at most $q_{\text{Sim}} \leq q_{\text{H}} + \sum_{i=1}^{c+1} p_i$ queries to the random oracle RO.

Proof. The proof structure closely follows the proof of Theorem 1. We show indifferentiability by (1) defining a simulator, (2) upper bounding the advantage of any differentiator in distinguishing the real world from the simulated world, and (3) upper-bounding the number RO-queries made by the given simulator.

Simulator. The simulator is described in Figure 2 by algorithms SimH_i and SimBD_i (for $i = 1, 2, 3$). The simulator sub-algorithms share state and keep track of the current distribution of the backdoored hash functions. The histories $\text{hst}_1, \text{hst}_2, \text{hst}_3$, initialized as empty sets, are used to keep track of the fixed coordinates of the simulated BROs. The distributions, according to which the simulated backdoored hash functions are sampled, are denoted by μ_1, μ_2 , and μ_3 and initialized as $\mathcal{U}_{[M]^N}$, since the hash functions without the backdoors are supposed to behave like random oracles. The corresponding hash functions are initialized as uniform random functions $\text{H}_i \leftarrow \mathcal{U}_{[M]^N}$. Furthermore, the simulator uses a counter q to keep track of the number of consecutive queries to a backdoor oracle and use this information to substitute the current distribution with a partially fixed and partially dense distribution, only when necessary (i.e., when $q = Q$), as opposed to doing so upon every backdoor query. Each time images of a simulated H_i are fixed by the simulator BD_i , images of the same rows must be fixed for H_j and H_k (i.e., the other two functions) to provide consistency with the random oracle RO. For this, images of H_j are fixed truthfully according to

the currently sampled function, while H_k is tweaked in a way that the extracted values match images of RO. Note that the simulators need to re-sample H_i and H_k if their distribution is modified in a non-trivial way, i.e., not just fixing more values, but either through FixRows or to force consistency with RO.

$\text{SimH}_i^{\text{RO}}(x)$	$\text{RO}(x)$
$y_i \leftarrow H_i(x); \text{hst}_i \leftarrow \text{hst}_i \cup \{(x, y_i)\}$ $\mu_i \leftarrow \mu_i _{\text{hst}_i}$ $j \leftarrow (i \bmod 3) + 1; k \leftarrow (j \bmod 3) + 1$ $y \leftarrow \text{RO}(x)$ if $i = 1$ then $H_k \leftarrow \mu_k _{\text{Ext}(y_i, H_j(x), H_k(x))=y}$ elseif $i = 2$ then $H_k \leftarrow \mu_k _{\text{Ext}(H_k(x), y_i, H_j(x))=y}$ else $H_k \leftarrow \mu_k _{\text{Ext}(H_j(x), H_k(x), y_i)=y}$ $\text{hst}_j \leftarrow \text{hst}_j \cup \{(x, H_j(x))\}; \text{hst}_k \leftarrow \text{hst}_k \cup \{(x, H_k(x))\}$ $\mu_j \leftarrow \mu_j _{\text{hst}_j}; \mu_k \leftarrow \mu_k _{\text{hst}_k}$ $H_k \leftarrow \mu_k$ return y_i	if $\exists y \in [2]$ s.t. $(x, y) \in \text{hst}_{\text{RO}}$ then return y $y \leftarrow [2]$ $\text{hst}_{\text{RO}} \leftarrow \text{hst}_{\text{RO}} \cup \{(x, y)\}$ return y
$\text{SimBD}_i^{\text{RO}}[\bar{p}, \gamma](f)$	
$q \leftarrow q + 1$ $z \leftarrow f(H_i)$ $\mu_i \leftarrow \mu_i _{f(\cdot)=z}$ $j \leftarrow (i \bmod 3) + 1; k \leftarrow (j \bmod 3) + 1$ if $q = Q$ then $(\mu_i, A_i) \leftarrow \text{FixRows}[\gamma](\mu_i, p_{3s+i}, \text{hst}_{i,1})$ $H_i \leftarrow \mu_i$ $\text{hst}_i \leftarrow \text{hst}_i \cup A_i$ for $x \in A_{i,1}$ do $r_x \leftarrow \text{RO}(x)$ if $i = 1$ then $H_k \leftarrow \mu_k _{\forall (x, y_i) \in A_i. \text{Ext}(y_i, H_j(x), H_k(x))=r_x}$ elseif $i = 2$ then $H_k \leftarrow \mu_k _{\forall (x, y_i) \in A_i. \text{Ext}(H_k(x), y_i, H_j(x))=r_x}$ else $H_k \leftarrow \mu_k _{\forall (x, y_i) \in A_i. \text{Ext}(H_j(x), H_k(x), y_i)=r_x}$ for $x \in A_{i,1}$ do $\text{hst}_j \leftarrow \text{hst}_j \cup \{(x, H_j(x))\}; \text{hst}_k \leftarrow \text{hst}_k \cup \{(x, H_k(x))\}$ $\mu_j \leftarrow \mu_j _{\text{hst}_j}; \mu_k \leftarrow \mu_k _{\text{hst}_k}$ $H_k \leftarrow \mu_k$ $q \leftarrow 0$ if $i = 3$ then $s \leftarrow s + 1$ return z	

Fig. 2: Indifferentiability simulator for the 2-out-of-3-source extractor. We assume for $i = 1..3$ initialization values $\text{hst}_i = \text{hst}_{\text{RO}} := \emptyset$, $\mu_i := \mathcal{U}_{[M]^N}$, $H_i \leftarrow \mathcal{U}_{[M]^N}$, $q := 0$, and $s := 0$. The FixRows algorithm is identical to that of Figure 1.

Security Analysis. We analyze indistinguishability of 3ext-combiner using a sequence of eight games $\text{Game}_0, \dots, \text{Game}_7$, where Game_0 and Game_7 are the real and the ideal indistinguishability games, respectively. The modified lines in each game are highlighted. Oracles are omitted in some games if they have not changed since the previous game.

$\text{Game}_0 : \mathcal{C}_{3\text{ext}}^{\text{H}_1, \text{H}_2, \text{H}_3}(x)$	
for $i = 1..3$ do $y_i \leftarrow \text{H}_i(x)$ $y \leftarrow \text{Ext}(y_1, y_2, y_3)$ return y	
$\text{Game}_0 : \text{H}_i(x)$	$\text{Game}_0 : \text{BD}_i(f)$
$y_i \leftarrow \text{H}_i(x)$ return y_i	$z \leftarrow f(\text{H}_i)$ return z

We use the shorthand notation $\Pr[\mathcal{D}^{\text{Game}}] := \Pr[\mathcal{D}^{\text{Game}} = 1]$, where $\mathcal{D}^{\text{Game}}$ indicates the interaction of an adversary \mathcal{D} with a game Game . In each game hop, we bound the adversary's advantage in distinguishing any two consecutive games from one another. The first game Game_0 is the real game, where the adversary interacts with the 3ext-combiner and the backdoor oracles of the underlying BROs.

$\text{Game}_1 : \mathcal{C}_{3\text{ext}}^{\text{H}_1, \text{H}_2, \text{H}_3}(x)$	
for $i = 1..3$ do $y_i \leftarrow \text{H}_i(x); \text{hst}_i \leftarrow \text{hst}_i \cup \{(x, y_i)\}; \mu_i \leftarrow \mu_i _{\text{hst}_i}$ $y \leftarrow \text{Ext}(y_1, y_2, y_3)$ return y	
$\text{Game}_1 : \text{H}_i(x)$	$\text{Game}_1 : \text{BD}_i(f)$
$y_i \leftarrow \text{H}_i(x)$ $\text{hst}_i \leftarrow \text{hst}_i \cup \{(x, y_i)\}$ $\mu_i \leftarrow \mu_i _{\text{hst}_i}$ return y_i	$z \leftarrow f(\text{H}_i)$ $\mu_i \leftarrow \mu_i _{f(\cdot)=z}$ return z

Game_1 . Game Game_1 updates the distribution of hash functions based on evaluation queries, backdoor queries, and the history of coordinates that are fixed through construction queries. The distributions μ_i are conditioned on these values but are never actually sampled from in the game. Hence the two games are identical, i.e., $\text{SD}(\text{Game}_0, \text{Game}_1) = 0$.

```

Game2 : BDi(f)


---


 $q \leftarrow q + 1$ 
 $z \leftarrow f(H_i)$ 
 $\mu_i \leftarrow \mu_i|_{f(\cdot)=z}$ 
if  $q = Q$  then
   $(\mu'_i, A_i) \leftarrow \text{FixRows}[\gamma](\mu_i, p_{3s+i}, \text{hst}_{i.1})$ 
   $q \leftarrow 0$ 
  if  $i \equiv 3$  then  $s \leftarrow s + 1$ 
return  $z$ 

```

Game₂. In game **Game₂**, after each sequence of Q queries to a backdoor oracle, i.e., right before switching to a different one, a $(p, 1 - \delta)$ -dense distribution μ'_i is obtained from the real distribution using the algorithm **FixRows** by decomposing the distribution of the corresponding hash function after responding to the last query (i.e., $\mu_i|_{f(\cdot)=z}$). The number of fixed points p is a parameter determined by the simulator and the density rate δ can be obtained by applying Lemma 1. However, since the new distributions μ'_i are never used elsewhere, **Game₂** remains identical to the previous **Game₁**, i.e., $\text{SD}(\text{Game}_1, \text{Game}_2) = 0$.

Game₃ : H_i(x)	Game₃ : BD_i(f)
$y_i \leftarrow H_i(x)$ $\text{hst}_i \leftarrow \text{hst}_i \cup \{(x, y_i)\}$ $\mu_i \leftarrow \mu_i _{\text{hst}_i}$ $j \leftarrow (i \bmod 3) + 1; k \leftarrow (j \bmod 3) + 1$ $\text{hst}_j \leftarrow \text{hst}_j \cup \{(x, H_j(x))\}$ $\text{hst}_k \leftarrow \text{hst}_k \cup \{(x, H_k(x))\}$ $\mu_j \leftarrow \mu_j _{\text{hst}_j}; \mu_k \leftarrow \mu_k _{\text{hst}_k}$ return y_i	$q \leftarrow q + 1$ $z \leftarrow f(H_i)$ $\mu_i \leftarrow \mu_i _{f(\cdot)=z}$ $j \leftarrow (i \bmod 3) + 1; k \leftarrow (j \bmod 3) + 1$ if $q = Q$ then $(\mu'_i, A_i) \leftarrow \text{FixRows}[\gamma](\mu_i, p_{3s+i}, \text{hst}_{i.1})$ for $x \in A_{i.1}$ do $\text{hst}_j \leftarrow \text{hst}_j \cup \{(x, H_j(x))\}$ $\text{hst}_k \leftarrow \text{hst}_k \cup \{(x, H_k(x))\}$ $\mu_j \leftarrow \mu_j _{\text{hst}_j}; \mu_k \leftarrow \mu_k _{\text{hst}_k}$ $q \leftarrow 0$ if $i \equiv 3$ then $s \leftarrow s + 1$ return z

Game₃. In this game, the fixed rows in one simulated BRO are also fixed for the other two BROs. E.g., in backdoor simulation, the rows in the assignment A_i are fixed for H_j and H_k . In both games, the oracles' behaviors are at all times consistent with their past responses as well as the construction. Hence, it does not matter, if more or less of the hash function tables are fixed in each query. The two games are again perfectly indistinguishable, i.e., we have $\text{SD}(\text{Game}_2, \text{Game}_3) = 0$.

Game₄ : BD_i(f)

```

q ← q + 1
z ← f(Hi)
μi ← μi |f(·)=z
j ← (i mod 3) + 1; k ← (j mod 3) + 1
if q = Q then
  (μi, Ai) ← FixRows[γ](μi, p3s+i, hsti.1)
  Hi ← μi
  hsti ← hsti ∪ Ai
  for x ∈ Ai.1 do
    hstj ← hstj ∪ {(x, Hj(x))}
    hstk ← hstk ∪ {(x, Hk(x))}
  μj ← μj |hstj; μk ← μk |hstk
  q ← 0
  if i = 3 then s ← s + 1
return z

```

Game₅ : H_i(x)

```

yi ← Hi(x)
hsti ← hsti ∪ {(x, yi)}
μi ← μi |hsti
j ← (i mod 3) + 1; k ← (j mod 3) + 1
y ← RO(x)
if i = 1 then Hk ← μk |Ext(yi, Hj(x), Hk(x))=y
elseif i = 2 then Hk ← μk |Ext(Hk(x), yi, Hj(x))=y
else Hk ← μk |Ext(Hj(x), Hk(x), yi)=y
hstj ← hstj ∪ {(x, Hj(x))}
hstk ← hstk ∪ {(x, Hk(x))}
μj ← μj |hstj; μk ← μk |hstk
Hk ← μk
return yi

```

Game₅ : BD_i(f)

```

q ← q + 1
z ← f(Hi)
μi ← μi |f(·)=z
j ← (i mod 3) + 1; k ← (j mod 3) + 1
if q = Q then
  (μi, Ai) ← FixRows[γ](μi, p3s+i, hsti.1)
  Hi ← μi
  hsti ← hsti ∪ Ai
  for x ∈ Ai.1 do rx ← C3ext(x)
  if i = 1 then
    Hk ← μk |∀(x, yi) ∈ Ai. Ext(yi, Hj(x), Hk(x))=rx
  elseif i = 2 then
    Hk ← μk |∀(x, yi) ∈ Ai. Ext(Hk(x), yi, Hj(x))=rx
  else Hk ← μk |∀(x, yi) ∈ Ai. Ext(Hj(x), Hk(x), yi)=rx
  for x ∈ Ai.1 do
    hstj ← hstj ∪ {(x, Hj(x))}
    hstk ← hstk ∪ {(x, Hk(x))}
  μj ← μj |hstj; μk ← μk |hstk
  Hk ← μk
  q ← 0
  if i = 3 then s ← s + 1
return z

```

Game₄. In Game₄ the distribution obtained by FixRows finally replaces the true distribution, i.e., the one conditioned on the recent backdoor responses. Hence, the history is updated. Notably, a new function H_i must be sampled for future references, since its distribution has changed in a non-trivial way. According to Lemma 1, there is a convex combination of (p, 1 - δ)-dense distributions

which is γ -close to the real distribution, one of such distributions being the one returned by **FixRows**. Thus, the distinguishing advantage increases by γ after each sequence of backdoor queries, i.e., $|\Pr[\mathcal{D}^{\text{Game}_3}] - \Pr[\mathcal{D}^{\text{Game}_4}]| \leq (c+1) \cdot \gamma$.

Game₅. Contrary to **Game₄**, the next game **Game₅** somewhat indirectly fixes images of rows x in A_i (and x queries to **SimH_i**) for the other functions H_j and H_k . More precisely, the simulator calls the construction $C_{3\text{ext}}$ on freshly fixed rows according to A_i and samples a H_k in such a way that it is consistent with those construction images, and aligned H_i and H_k images. Notice that a query to the construction already fixes the images for the underlying BROs and therefore, sampling H_k in a consistent way and fixing coordinates of H_j and H_k in the simulator is simply redundant. Hence $\text{SD}(\text{Game}_4, \text{Game}_5) = 0$.

```

Game6 :  $C_{3\text{ext}}^{H_1, H_2, H_3}(x)$ 


---


if  $\exists y \in [M]$  s.t.  $(x, y) \in \text{hst}_{\text{RO}}$  then return  $y$ 
if  $\exists y_1, y_2, y_3 \in [M]$  s.t.  $(x, y_1) \in \text{hst}_1 \wedge (x, y_2) \in \text{hst}_2 \wedge (x, y_3) \in \text{hst}_3$  then return  $\text{Ext}(y_1, y_2, y_3)$ 
if  $\exists y' \in [M]$  s.t.  $(x, y') \in \text{hst}_1 \vee (x, y') \in \text{hst}_2 \vee (x, y') \in \text{hst}_3$  then
   $y \leftarrow [M]$ 
else
  for  $i = 1..3$  do
     $y_i \leftarrow H_i(x)$ ;  $\text{hst}_i \leftarrow \text{hst}_i \cup \{(x, y_i)\}$ ;  $\mu_i \leftarrow \mu_i|_{\text{hst}_i}$ 
   $y \leftarrow \text{Ext}(y_1, y_2, y_3)$ 
   $\text{hst}_{\text{RO}} \leftarrow \text{hst}_{\text{RO}} \cup \{(x, y)\}$ 
return  $y$ 

```

Game₆. In this game we modify $C_{3\text{ext}}$ so that it starts to resemble a lazily sampled random oracle. Query-response pairs of the construction are kept in a set hst_{RO} and in case a query is repeated the stored image is simply returned. Otherwise, we distinguish three cases: (a) the corresponding row to the current query x is fixed in all hash functions, (b) in one of them, or (c) in none of them. In case (a), **Game₆** computes the output of the construction by extracting from the individual images stored in histories of the BROs. Note, however, that this case is never reached, since if the current x is in all individual histories, then the construction must have already been called on x in some previous evaluation or backdoor query. Hence, x must also be in hst_{RO} . In case (b), a uniformly random value is chosen (and stored in hst_{RO}). In the final case (c), **Game₆** behaves exactly as **Game₅**.

Overall, the distinguishing advantage is bounded by distinguishing p uniform (chosen by the construction) points each time a backdoor query fixes p points from values that were supposed to be extracted from three sources, from which one is not guaranteed to have any min-entropy, as well as q_H many times distinguishing a single extracted value from random. Let $E_n := \text{Ext}(H_1(x_n), H_2(x_n))$,

$H_3(x_n)$), where $x_n \in A_{i,1}$ is a row being fixed. Then we have:

$$|\Pr[\mathcal{D}^{\text{Game}_5}] - \Pr[\mathcal{D}^{\text{Game}_6}]| \leq \sum_{i=1}^c \text{SD}(E_1 | \dots | E_{p_i}, \mathcal{U}_{[2]^{p_i}}) + q_H \cdot \text{SD}(E_1, \mathcal{U}_{[2]}) ,$$

Game₇ : $C_{3\text{ext}}^{H_1, H_2, H_3}(x)$

```

if  $\exists y \in [M]$  s.t.  $(x, y) \in \text{hst}_{\text{RO}}$  then return  $y$ 
if  $\exists y_1, y_2, y_3 \in [M]$  s.t.  $(x, y_1) \in \text{hst}_1 \wedge (x, y_2) \in \text{hst}_2 \wedge (x, y_3) \in \text{hst}_3$  then return  $\text{Ext}(y_1, y_2, y_3)$ 
if  $\exists y' \in [M]$  s.t.  $(x, y') \in \text{hst}_1 \vee (x, y') \in \text{hst}_2 \vee (x, y') \in \text{hst}_3$  then
   $y \leftarrow [M]$ 
else
  for  $i = 1..3$  do
     $y_i \leftarrow H_i(x); \text{hst}_i \leftarrow \text{hst}_i \cup \{(x, y_i)\}; \mu_i \leftarrow \mu_i|_{\text{hst}_i}$ 
   $y \leftarrow \text{Ext}(y_1, y_2, y_3)$ 
 $\text{hst}_{\text{RO}} \leftarrow \text{hst}_{\text{RO}} \cup \{(x, y)\}$ 
return  $y$ 

```

Game₇. The $C_{3\text{ext}}$ oracle in **Game₇** differs from **Game₆** in that it never evaluates the underlying BROs any more and rather acts as a lazily sampled random oracle. We can safely remove the case distinction (a), where x is included in all histories hst_1 , hst_2 , and hst_3 , since this x would also be in hst_{RO} . It remains to bound the adversary's advantage in distinguishing the two games while making up to q_C fresh queries x to the construction $C_{3\text{ext}}$ that are not fixed for any of the BROs. While the outputs of the construction are uniformly random in **Game₇**, they are extracted from three dense images in **Game₆**. The distinguisher can only try to maximize the distance between q_C uniform values vs. values extracted from three dense images of BROs by querying the construction on values and at times which it can choose freely.

$$\begin{aligned}
|\Pr[\mathcal{D}^{\text{Game}_6}] - \Pr[\mathcal{D}^{\text{Game}_7}]| &\leq \sum_{t=1}^{q_C} \max_{x_t, H_1, H_2, H_3} \left(\text{SD}(\text{Ext}(H_1(x_t), H_2(x_t), H_3(x_t)), \mathcal{U}_{[2]}) \right) \\
&\leq q_C \cdot \max_{x, H_1, H_2, H_3} \left(\text{SD}(\text{Ext}(H_1(x), H_2(x), H_3(x)), \mathcal{U}_{[2]}) \right) \\
&\leq q_C \cdot \varepsilon((1 - \delta_{c-1}) \cdot \log M, (1 - \delta_c) \cdot \log M, \\
&\quad (1 - \delta_{c+1}) \cdot \log M) ,
\end{aligned}$$

where according to Lemma 1 we have δ_i as defined in the theorem statement with ℓ_i being the min-entropy deficiency after the i -th sequence of Q -many backdoor queries. Note that the maximum statistical distance corresponds to minimum entropy of the BRO-images, which is in turn given for the last three $(c-1, c, c+1)$ values of the density rate.

Query complexity. The simulator makes queries to the random oracle RO to set images of the other BROs each time one point of some BRO is fixed, either caused

by evaluation queries or by backdoor queries right after the Q -th consecutive backdoor query (i.e., before a switch). Hence we obtain the bound $q_{\text{Sim}} \leq q_{\text{H}} + \sum_{i=1}^{c+1} p_i$ on the number of queries that the simulator makes to the random oracle. \square

5.1 Instantiation with the pairwise inner-product extractor

Next we investigate a concrete instantiation of such a 2-out-of-3-source extractor. General multi-source extractors such as those from [2,25,21] which require a minimal amount of min-entropy from *every* source are inapplicable in our setting. We can, however, use the pairwise inner-product extractor as introduced by Lee et al. [19], which roughly speaking needs the sum of min-entropies to be sufficient. Formally a pairwise inner-product extractor $\text{Ext}_{\text{pip}} : [M]^t \rightarrow [2]$ is defined as:

$$\text{Ext}_{\text{pip}}(x_1, \dots, x_t) := \sum_{1 \leq i < j \leq t} x_i \cdot x_j .$$

This extractor is proven ([19], Corollary 1) to be a $(k_1, \dots, k_t, \varepsilon)$ -extractor with $\varepsilon = 2^{-(k+k'-\log M+1)/2}$, where k and k' are the two largest values among k_1, \dots, k_t . Hence, Ext_{pip} is also a 2-out-of- t extractor.

Corollary 2. *Let $\text{Ext}_{\text{pip}} : [M]^t \rightarrow [2]$ be a pairwise inner-product extractor. Then the construction $\mathbf{C}_{\text{pip}}^{\text{H}_1, \text{H}_2, \text{H}_3}(x) := \text{Ext}_{\text{pip}}(\text{H}_1(x), \text{H}_2(x), \text{H}_3(x))$ in the 3-BRO model is indifferentiable from a random oracle, where*

$$\begin{aligned} \text{Adv}_{\mathbf{C}_{\text{ext}}^{\text{H}_1, \text{H}_2, \text{H}_3}, \text{Sim}[p, \gamma]}^{\text{indiff}}(\mathcal{D}) &\leq (c+1) \cdot \gamma \\ &\quad + c \cdot \sqrt{(e^{p \cdot M^{-(1-2\delta_c)}} - 1)/2} \\ &\quad + (q_{\text{H}} + q_{\text{C}}) \cdot 2^{-((1-2\delta_{c+1}) \cdot \log M + 1)/2} , \end{aligned}$$

while the simulator makes up to $q_{\text{Sim}} \leq q_{\text{H}} + (c+1) \cdot p$ queries to RO.

Proof. The differentiator's advantage stated in the corollary is easily obtained by upper bounding the term $\text{SD}(E_1, \mathcal{U}_{[2]})$ by $2^{-((1-2\delta_{c+1}) \cdot \log M + 1)/2}$ and upper bounding the term $\text{SD}(E_1 | \dots | E_p, \mathcal{U}_{[2]^p})$ from the advantage in Theorem 2, using the following claim.

Claim. Let $x_i \in [N]$ and $E_i := \text{Ext}_{\text{pip}}(\text{H}_1(x_i), \text{H}_2(x_i), \text{H}_3(x_i))$ for $i = 1..n$. Suppose that for all x_i , at least two of the (distributions of the) functions H_1 , H_2 , H_3 are $(1 - \delta)$ -dense. Then for all $n \in \mathbb{N}$ we have:

$$\text{SD}(E_1 | \dots | E_n, \mathcal{U}_{[2]^n}) \leq \sqrt{(e^{n \cdot M^{-(1-2\delta)}} - 1)/2} .$$

Proof. In the proof below we use the parity lemma ⁵ (1) and the fact that the pairwise inner product (in $[L]$) is linear, i.e., $\sum_{n \in I} E_n = \sum_{n \in I} \text{Ext}_{\text{pip}}(\text{H}_1(x_n),$

⁵ Let X be a random variable over $[2^\ell]$. Then we have

$$\text{SD}(X, \mathcal{U}_{[2^\ell]}) \leq \sqrt{\sum_{0^\ell \neq a \in [2^\ell]} (\text{SD}(X \cdot a, \mathcal{U}_{[2]})^2} .$$

$$\mathbf{H}_2(x_n), \mathbf{H}_3(x_n)) = \text{Ext}_{\text{pip}}(\mathbf{H}_1(x_1) | \cdots | \mathbf{H}_1(x_{|I|}), \mathbf{H}_2(x_1) | \cdots | \mathbf{H}_2(x_{|I|}), \mathbf{H}_3(x_1) | \cdots | \mathbf{H}_3(x_{|I|})) = E_I \quad (2).$$

$$\begin{aligned} \text{SD}(E_1 | \cdots | E_n, \mathcal{U}_{[2]^n}) &\leq \sqrt{\sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} (\text{SD}(E_1 | \cdots | E_n \cdot a, \mathcal{U}_{[2]^n}))^2} \quad (1) \\ &= \sqrt{\sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} \left(\text{SD}\left(\sum_{i \in I} E_i, \mathcal{U}_{[2]^n}\right) \right)^2} \\ &= \sqrt{\sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} (\text{SD}(E_I, \mathcal{U}_{[2]^n}))^2} \quad (2) \\ &\leq \sqrt{\sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} 2^{-(|I| \cdot \log M \cdot (1-2\delta) + 2 - \log 2)}} \\ &= \sqrt{2^{-2 + \log 2} \cdot \sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} 2^{-|I| \cdot \log M \cdot (1-2\delta)}} \\ &= \sqrt{2^{-1} \cdot \sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} (M^{-(1-2\delta)})^{|I|}} \\ &= \sqrt{\left((1 + M^{-(1-2\delta)})^n - 1 \right) / 2} \\ &\leq \sqrt{(e^{n \cdot M^{-(1-2\delta)}} - 1) / 2} \quad \square \end{aligned}$$

Hence, the claim about the advantage holds. The query complexity of the simulator is bounded by the sum of $q_{\mathbf{H}}$ and $(c+1) \cdot p$. \square

We now provide estimates for the involved parameters.

Corollary 3. *Let the number of switches be $c \geq 1$ and assume the range size of the three random oracles are $M \geq N^9$. Then there is an indistinguishability simulator Sim for the \mathbf{C}_{pip} construction in the 3-BRO model that places at most*

$$q_{\mathbf{H}} + (c+1) \cdot \left(\frac{6Q\ell}{\log M} \right)^{1/\alpha(c)} \cdot N^{1-1/\alpha(c)}$$

queries to RO, where $\alpha(c) := \lfloor \frac{c}{3} \rfloor + 1$, against any distinguisher with $q_{\mathbf{H}}$ queries to the underlying BROs. Further, any such distinguisher with $q_{\mathbf{C}}$ construction queries and Q consecutive queries to the same backdoor oracle before switching, has advantage at most $(c+1) \cdot \gamma + (c + q_{\mathbf{H}} + q_{\mathbf{C}})/N$ against this simulator.

Proof. The recurrence relations for δ_i in the statement of Theorem 2 can be written as

$$\delta_i \leq A \cdot \delta_{i-3} + B,$$

where $A := N/p$ and $B := (Q\ell + \log \gamma^{-1})/p \log M$. Solving this recurrence relation we get

$$\delta_i \leq \frac{A^{\lfloor \frac{i-1}{3} \rfloor + 1} - 1}{A - 1} \cdot B.$$

We set $\delta_{c+1} \leq 1/3$ so that the term $1 - 2\delta_{c+1}$ is positive. To this end, it is sufficient to have that

$$\frac{A^{\lfloor \frac{c}{3} \rfloor + 1} - 1}{A - 1} \cdot B \leq \frac{1}{3}.$$

Substituting A and B and removing the -1 in the numerator we need to have that

$$\left(\frac{N}{p}\right)^{\lfloor \frac{c}{3} \rfloor + 1} \leq \frac{A - 1}{3B} = \frac{(N/p - 1)p \log M}{3Q\ell} = \frac{N \log M - p \log M}{3Q\ell} \leq \frac{N \log M}{6Q\ell},$$

where for the last inequality we have assumed that $p \leq N/2$. Thus,

$$p \geq \left(\frac{6Q\ell}{\log M}\right)^{1/\alpha(c)} \cdot N^{1-1/\alpha(c)},$$

where $\alpha(c) := \lfloor \frac{c}{3} \rfloor + 1$. For sufficiently large c , the factor above is at most 2.

The advantage stated in Corollary 2 is

$$(c + 1) \cdot \gamma + c \cdot \sqrt{p/M^{1-2\delta_c}} + (q_H + q_C) \cdot \sqrt{1/M^{1-2\delta_{c+1}}}.$$

Since $1 - 2\delta_{c+1} \leq 1 - 2/3 = 1/3$, $\delta_c \leq \delta_{c+1}$, $p \leq N$ and $M \geq N^9$, the advantage is upper-bounded by $(c + 1) \cdot \gamma + (c + q_H + q_C)/N$. \square

Note that for $c = 1, 2$ the query complexity of the simulator does not involve the $N^{1-1/\alpha(c)}$ factor, and hence we obtain collision resistance. For $c \geq 3$, however there is a factor of at least $N^{1/2}$.

The above corollary shows that the extractor combiner can tolerate a *linear* number of switches in $\log N$ (which can be thought of as the security parameter) for the simulator query complexity to be less than $N/2$. As for the xor combiner we conjecture that (much) better bounds for the extractor combiner are possible.

6 Indifferentiability with Auxiliary Input

In this section we discuss indifferentiability in a setting where there is no adaptivity and the backdoor oracles are called only once at the onset. Although this may seem overly restrictive, the resulting definition is sufficiently strong to model indifferentiability in the presence of auxiliary input, whereby we would like to securely replace random oracles in generic applications even in the presence of auxiliary input.

In this setting we can view an indifferentiability simulator as operating in two stages: An off-line stage which responds to the single backdoor queries for

each BRO, and an on-line stage which simulates direct evaluation calls to the underlying BROs. As defined, the off-line phase of the simulator can pass an arbitrary state to its on-line phase. Further, both stages have access to the reference object oracles (although the query complexities of both stages need to be small). More precisely, this definition in the 2-BRO requires that for any $(\mathcal{D}_{0,1}, \mathcal{D}_{0,2}, \mathcal{D}_1)$ in the real world with two BROs H_1 and H_2 with

$$z_1 \leftarrow \mathcal{D}_{0,1}(H_1); z_2 \leftarrow \mathcal{D}_{0,2}(H_2, z_1); b \leftarrow \mathcal{D}_1^{C^{H_1, H_2}, H_1, H_2}(z_1, z_2) ,$$

there exists some $(\text{Sim}_{0,1}^{\text{RO}}, \text{Sim}_{0,2}^{\text{RO}}, \text{Sim}_{1,1}^{\text{RO}}, \text{Sim}_{1,2}^{\text{RO}})$ in the ideal (simulated) world

$$(z_1, st) \leftarrow \text{Sim}_{0,1}^{\text{RO}}(); (z_2, st) \leftarrow \text{Sim}_{0,2}^{\text{RO}}(st); b \leftarrow \mathcal{D}_1^{\text{RO}, \text{Sim}_{1,1}^{\text{RO}}[st], \text{Sim}_{1,2}^{\text{RO}}[st]}(z_1, z_2) ,$$

with indistinguishable outputs b . The on-line simulators can also share state.

Let us now take a step back and define indifferentiability with auxiliary input driven by a composition theorem: for any game \mathcal{G} and any attacker \mathcal{A}_1 in this game against C^{H_1, H_2} which receives auxiliary input on H_1 and H_2 , there is an attacker \mathcal{B}_1 on RO in the same game \mathcal{G} but now *without* auxiliary input. More explicitly, the real world

$$z \leftarrow \mathcal{A}_0(H_1, H_2); b \leftarrow \mathcal{G}^{C^{H_1, H_2}, \mathcal{A}_1^{H_1, H_2}(z)}$$

and the ideal world

$$(z, st) \leftarrow \mathcal{B}_0^{\text{RO}}(); b \leftarrow \mathcal{G}^{\text{RO}, \mathcal{B}_1^{\text{RO}}(z, st)}$$

are indistinguishable. Once again the query complexity of \mathcal{B}_0 should be small (or even zero) to obtain a definition which meaningfully formalizes indifferentiability from random oracles without auxiliary input. This definition, however, turns out to be unachievable: \mathcal{A}_0 can simply encode a pair of collisions for the construction, which \mathcal{B}_0 will not be able to match (with respect to RO) without an exponentially large number of queries to RO.⁶

There are two natural ways to overcome this: (1) restrict the interface of the construction; or (2) restrict the form of preprocessing. The former is motivated by use of salting as a means to defeat preprocessing, and the latter by independence of preprocessing for BROs.

A final question arises here: is it possible to simplify this definition further by removing the quantification over \mathcal{A}_1 (as done for standard indifferentiability)? This could be done in the standard way by absorbing \mathcal{A}_1 into \mathcal{G} to form a differentiator \mathcal{D} . However, this means that \mathcal{D} must receive the auxiliary information z .

⁶ One can formulate an intermediate notion of indifferentiability from random oracle *with* auxiliary input. Without salting, this notion would not be of great help. Consider, for example, the case of domain extension via an iterative hashing mode. Due to Joux's multi-collision attack [17] one can encode exponentially many collisions for the construction in a small auxiliary input, whereas this would not be possible for the random oracle.

The resulting notion is stronger and models composition with respect to *games* that also depend on preprocessing. Thus, due to its simplicity, strength, and the fact that we can establish positive results for it, we focus on this definitional approach. We now make the two definitions arising from (1) and (2) explicit.

SALTED AI-INDIFFERENTIABILITY. We call a construction C^H *salted* if the construction takes a salt $hk \in \{0, 1\}^k$ as input and prepends all calls to H with hk . We define salted AI-indifferentiability from a random oracle by requiring that for any $(\mathcal{D}_0, \mathcal{D}_1)$ in the real world

$$z \leftarrow \mathcal{D}_0(H); hk \leftarrow \{0, 1\}^k; b \leftarrow \mathcal{D}_1^{C^H(hk, \cdot), H}(hk, z)$$

there is a simulator $(\text{Sim}_0^{\text{RO}}, \text{Sim}_1^{\text{RO}})$ in the ideal world

$$(z, st) \leftarrow \text{Sim}_0^{\text{RO}}(); hk \leftarrow \{0, 1\}^k; b \leftarrow \mathcal{D}_1^{\text{RO}(hk, \cdot), \text{Sim}_1^{\text{RO}}[st]}(hk, z)$$

resulting in indistinguishable outputs b . We denote the advantage of \mathcal{D} in the salted AI-indifferentiability game with simulator Sim for a construction C^H by $\text{Adv}_{C^H, \text{Sim}}^{\text{s-ai-indiff}}(\mathcal{D})$. Notice that in the above definition, the distinguisher gets access to a *salted* RO. A different definition arises when the distinguisher gets access to an unsalted RO instead. However, since the simulated auxiliary information is computed given access to an unsalted RO (which can be interpreted as having implicit access to the salt), such a definition calls for the existence of a more powerful simulator. In particular, such Sim_0 and \mathcal{D}_1 can easily call RO on common points. The practical implications of such a definition are unclear to us, and moreover, it is strictly weaker than our definition.

AI-INDIFFERENTIABILITY WITH INDEPENDENT PREPROCESSING. We define AI-indifferentiability with independent preprocessing by requiring that for any adversary $(\mathcal{D}_{0,1}, \mathcal{D}_{0,2}, \mathcal{D}_1)$ in the real world

$$z_1 \leftarrow \mathcal{D}_{0,1}(H_1); z_2 \leftarrow \mathcal{D}_{0,2}(H_2); b \leftarrow \mathcal{D}_1^{C^{H_1, H_2}, H_1, H_2}(z_1, z_2)$$

there is a simulator $(\text{Sim}_{0,1}^{\text{RO}}, \text{Sim}_{0,2}^{\text{RO}}, \text{Sim}_{1,1}^{\text{RO}}, \text{Sim}_{1,2}^{\text{RO}})$ in the ideal world

$$(z_1, st) \leftarrow \text{Sim}_{0,1}^{\text{RO}}(); (z_2, st) \leftarrow \text{Sim}_{0,2}^{\text{RO}}(st); b \leftarrow \mathcal{D}_1^{\text{RO}, \text{Sim}_{1,1}^{\text{RO}}[st], \text{Sim}_{1,2}^{\text{RO}}[st]}(z_1, z_2)$$

resulting in indistinguishable outputs b . Note that this is slightly weaker than the definition of indifferentiability in 2-BRO since z_2 is fully independent of z_1 , whereas BRO indifferentiability allows for a limited amount of dependence. We denote by $\text{Adv}_{C^H, \text{Sim}}^{\text{ai-indiff}}(\mathcal{D})$ the advantage of \mathcal{D} in the AI-indifferentiability game with independent preprocessing with respect to a simulator Sim and a construction C^{H_1, H_2} in the 2-BRO model.

We are now ready to prove our feasibility results for AI-indifferentiability.

Theorem 3 (AI-Indifferentiability). *Any construction C^{H_1, H_2} that is indifferentiable with backdoors from a random oracle with no adaptive backdoor queries*

is also AI-indifferentiable from a random oracle with respect to independent preprocessing attacks. More precisely, for any auxiliary-input differentiator $\mathcal{D} := (\mathcal{D}_{0,1}, \mathcal{D}_{0,2}, \mathcal{D}_1)$ with independent preprocessing for two random oracles there is a 2-BRO differentiator $\tilde{\mathcal{D}}$ with one-time non-adaptive access to each backdoor oracle such that for any 2-BRO indifferntiability simulator $\tilde{\text{Sim}}$ there is an auxiliary-input simulator $\text{Sim} := (\text{Sim}_{0,1}, \text{Sim}_{0,2}, \text{Sim}_{1,1}, \text{Sim}_{1,2})$ such that

$$\text{Adv}_{\mathcal{C}^{\text{H}_1, \text{H}_2}, \text{Sim}}^{\text{ai-indiff}}(\mathcal{D}) = \text{Adv}_{\mathcal{C}^{\text{H}_1, \text{H}_2}, \tilde{\text{Sim}}}^{\text{indiff}}(\tilde{\mathcal{D}}) .$$

Further, any salted construction \mathcal{C}^{H} that is indifferentiable (in the standard sense) from a random oracle is also salted AI-indifferentiable from a random oracle. More precisely, for any auxiliary-input differentiator $\mathcal{D} := (\mathcal{D}_0, \mathcal{D}_1)$, with an auxiliary input of size ℓ , there is a (standard) differentiator $\tilde{\mathcal{D}}$ such that for any indifferntiability simulator $\tilde{\text{Sim}}$ there is an auxiliary-input simulator $\text{Sim} := (\text{Sim}_0, \text{Sim}_1)$ such that for any $p \in \mathbb{N}$ and any $\gamma > 0$

$$\text{Adv}_{\mathcal{C}^{\text{H}}, \text{Sim}}^{\text{s-ai-indiff}}(\mathcal{D}) \leq \text{Adv}_{\mathcal{C}^{\text{H}}, \tilde{\text{Sim}}}^{\text{indiff}}(\tilde{\mathcal{D}}) + \frac{\ell + \log \gamma^{-1}}{p} + \frac{p}{2^k} + \gamma .$$

Proof. The first part of the theorem follows directly from the discussion above that indifferntiability with backdoors and no adaptivity is stronger than indifferntiability with auxiliary input for independent preprocessing.

We now prove the second part of the theorem.

Game₀: We start with the real game in the salted AI-indifferentiability game:

$$z \leftarrow \mathcal{D}_0(\text{H}); \text{hk} \leftarrow \{0, 1\}^k; b \leftarrow \mathcal{D}_1^{\mathcal{C}^{\text{H}(\text{hk}, \cdot)}(\text{hk}, \cdot), \text{H}}(\text{hk}, z) .$$

Game₁: We now move to the bit-fixing RO model

$$(z, A) \leftarrow \tilde{\mathcal{D}}_0(); \text{hk} \leftarrow \{0, 1\}^k; b \leftarrow \mathcal{D}_1^{\mathcal{C}^{\text{H}[A]}(\text{hk}, \cdot), \text{H}[A]}(\text{hk}, z) .$$

Here $\tilde{\mathcal{D}}_0$ runs \mathcal{D}_0 by simulating an H for it and then runs the decomposition algorithm to get a set of assignments A for p fixed points (for any $p \in \mathbb{N}$). We may now apply [7, Theorem 5] to deduce that for any $\gamma > 0$,

$$\Pr[\text{Game}_1] - \Pr[\text{Game}_0] \leq \frac{\ell + \log \gamma^{-1}}{p} + \gamma ,$$

where ℓ is the size of auxiliary information.

Game₂: We now move to a setting where \mathcal{C} uses H rather than $\text{H}[A]$

$$(z, A) \leftarrow \tilde{\mathcal{D}}_0(); \text{hk} \leftarrow \{0, 1\}^k; b \leftarrow \mathcal{D}_1^{\mathcal{C}^{\text{H}(\text{hk}, \cdot)}(\text{hk}, \cdot), \text{H}[A]}(\text{hk}, z) .$$

This modification is justified by the fact that the probability that a uniform hk is (the prefix of the first component of some point) in A is at most $p/2^k$. We have that $\Pr[\text{Game}_2] - \Pr[\text{Game}_1] \leq p/2^k$.

Game₃:. We now move to a world where \mathcal{D}_1 is replaced by a differentiator $\tilde{\mathcal{D}}_1$ that gets the list A and does not query H on points in A :

$$(z, A) \leftarrow \tilde{\mathcal{D}}_0(); \quad hk \leftarrow \{0, 1\}^k; \quad b \leftarrow \tilde{\mathcal{D}}_1^{\mathcal{C}^{H(hk, \cdot)}(hk, \cdot), H}(hk, z, A) .$$

Here $\tilde{\mathcal{D}}_1(hk, z, A)$ runs $\mathcal{D}_1(hk, z)$ relaying its queries to the first oracle to its own first oracle and the second oracle queries to its own second oracle except when a queried point appears as a prefix of the first component of an entry in A in which case $\tilde{\mathcal{D}}_1$ uses A to answer the query. We have that $\Pr[\text{Game}_3] - \Pr[\text{Game}_2] = 0$.

Game₄:. We now absorb $\tilde{\mathcal{D}}_0$ and $\tilde{\mathcal{D}}_1$ into a single differentiator $\tilde{\mathcal{D}}$:

$$b \leftarrow \tilde{\mathcal{D}}^{\mathcal{C}^{H(hk, \cdot)}(hk, \cdot), H} .$$

Here $\tilde{\mathcal{D}}$ simply runs $\tilde{\mathcal{D}}_0$, followed by picking $hk \leftarrow \{0, 1\}^k$, and then running $\tilde{\mathcal{D}}_1$. We have that $\Pr[\text{Game}_4] - \Pr[\text{Game}_3] = 0$.

Game₅:. We now use the standard indistinguishability of the construction to move to the world

$$b \leftarrow \tilde{\mathcal{D}}^{\text{RO}(hk, \cdot), \tilde{\text{Sim}}^{\text{RO}}} ,$$

where $\tilde{\text{Sim}}$ is an indistinguishability simulator. We have that $\Pr[\text{Game}_3] - \Pr[\text{Game}_2] \leq \text{Adv}_{\mathcal{C}^H, \tilde{\text{Sim}}}^{\text{indiff}}(\tilde{\mathcal{D}})$.

Game₆:. We now syntactically unroll $\tilde{\mathcal{D}}$ into $(\tilde{\mathcal{D}}_0, \tilde{\mathcal{D}}_1)$:

$$(z, A) \leftarrow \tilde{\mathcal{D}}_0(); \quad hk \leftarrow \{0, 1\}^k; \quad b \leftarrow \tilde{\mathcal{D}}_1^{\text{RO}(hk, \cdot), \tilde{\text{Sim}}^{\text{RO}}}(hk, z, A) .$$

We have that $\Pr[\text{Game}_6] - \Pr[\text{Game}_5] = 0$.

Game₇:. We further unroll $\tilde{\mathcal{D}}_1$ into \mathcal{D}_1 and define $\text{Sim}_1[A]$ to be $\tilde{\text{Sim}}$ except that it uses A to answers queries in A :

$$(z, A) \leftarrow \tilde{\mathcal{D}}_0(); \quad hk \leftarrow \{0, 1\}^k; \quad b \leftarrow \mathcal{D}_1^{\text{RO}(hk, \cdot), \text{Sim}_1^{\text{RO}}[A]}(hk, z) .$$

We have that $\Pr[\text{Game}_7] - \Pr[\text{Game}_6] = 0$.

Game₈:. Finally we define $\text{Sim}_0 := \tilde{\mathcal{D}}_0$ and arrive at the simulated world

$$(z, A) \leftarrow \text{Sim}_0(); \quad hk \leftarrow \{0, 1\}^k; \quad b \leftarrow \mathcal{D}_1^{\text{RO}(hk, \cdot), \text{Sim}_1^{\text{RO}}[A]}(hk, z) .$$

We have that $\Pr[\text{Game}_8] - \Pr[\text{Game}_7] = 0$.

The second part of theorem follows by summing the (in)equalities established above; that is for any $p \in \mathbb{N}$ and any $\gamma > 0$ we get that

$$\begin{aligned} \text{Adv}_{\mathcal{C}^H, (\text{Sim}_0, \text{Sim}_1)}^{\text{s-ai-indiff}}(\mathcal{D}_0, \mathcal{D}_1) &= \Pr[\text{Game}_0] - \Pr[\text{Game}_8] \\ &\leq \text{Adv}_{\mathcal{C}^H, \tilde{\text{Sim}}}^{\text{indiff}}(\tilde{\mathcal{D}}) + \frac{\ell + \log \gamma^{-1}}{p} + \frac{p}{2^k} + \gamma . \end{aligned}$$

□

We can instantiate the first part of the above theorem with the xor combiner, which gives us the following corollary.

Corollary 4. *The xor combiner $C_{\oplus}^{H_1, H_2}(x) := H_1(x) \oplus H_2(x)$ is AI-indifferentiable from a random oracle with respect to independent preprocessing attacks for hash functions $H_1, H_2 \in [M]^N$. More precisely, for any $p \in \mathbb{N}$ and $0 < \gamma < 1$, there exists a simulator $\text{Sim}[p] := (\text{Sim}_{0,1}[p], \text{Sim}_{0,2}[p], \text{Sim}_{1,1}, \text{Sim}_{1,2})$ with oracle access to RO, such that for any auxiliary-input differentiator $\mathcal{D} := (\mathcal{D}_{0,1}, \mathcal{D}_{0,2}, \mathcal{D}_1)$ with auxiliary input of size ℓ for each hash function, where \mathcal{D}_1 makes up to q_H evaluation queries to H_1 and H_2 as well as q_C construction queries, we have*

$$\text{Adv}_{C_{\oplus}^{H_1, H_2}, \text{Sim}[p]}^{\text{ai-indiff}}(\mathcal{D}) \leq 2\gamma + \frac{(q_H + 2q_C) \cdot (\ell + \log \gamma^{-1})}{p},$$

while the simulator places at most $q_H + 2p$ queries to the random oracle RO.

Proof. The claim follows from the first part of Theorem 3 together with our indistinguishability result for xor (given in Theorem 1). However, deriving the concrete bounds using Corollary 1 results in somewhat suboptimal bounds with simulator query complexity $\mathcal{O}(p)$ and advantage $\mathcal{O}(1/\sqrt{p})$ with $p = N^{\alpha_1}$.

Here we directly use Theorem 1 for a simulator which fixes p points while simulating an ℓ -bit response of BD_1 and the same number of points while simulating an ℓ -bit response of BD_2 . Note that in the auxiliary-input setting we only consider one query to each backdoor oracle and therefore we have $Q = 1$. Overall we will have a simulator $\text{Sim}[p]$ for the above corollary, such that its off-line phase (i.e., $\text{Sim}_{0,1}[p]$ and $\text{Sim}_{0,2}[p]$) makes no queries to the RO and it simulates the auxiliary inputs by randomly choosing the hash functions and computing the output of the desired auxiliary-input functions (similar to a queried backdoor function) on them. This off-line phase then can use the refined decomposition algorithm of Lemma 1 for some small γ to come up with and (in addition to the auxiliary input) output two sets of pre-set points, each of size p , as its state. The state will be shared with the on-line phase of simulation, i.e., $\text{Sim}_{1,1}$ and $\text{Sim}_{1,2}$. Now this on-line simulator is a simple xor indistinguishability simulator which ensures consistency with the pre-set points. Note that our on-line simulator fixes p points for H_1 and again p points for H_2 . This results in simulator query complexity of $q_H + 2p$.

In this case, since $\delta_{-1} = \delta_0 = 0$ we obtain that

$$\delta_1 = \delta_2 = \frac{\ell + \log \gamma^{-1}}{p \log M}.$$

Plugging these back into the advantage bound in Theorem 1 we obtain the bound claimed above. \square

Note that for $p = o(\sqrt{N})$ we get a bound that is meaningful for collision resistance. As a result, we get that the xor combiner is collision resistant in the presence of independent auxiliary input (with no-salting). We note that the xor construction comes with added advantage that its security goes beyond AI-indistinguishability, and is also more domain efficient. Strictly speaking, however, the two settings are incomparable as the form of auxiliary information changes.

Acknowledgments

Dodis was partially supported by gifts from VMware Labs, Facebook and Google, and NSF grants 1314568, 1619158, 1815546. Mazaheri was supported by the German Federal Ministry of Education and Research (BMBF) and by the Hessian State Ministry for Higher Education, Research and the Arts, within ATHENE. Tessaro was partially supported by NSF grants CNS-1930117 (CAREER), CNS-1926324, CNS-2026774, a Sloan Research Fellowship, and a JP Morgan Faculty Award.

References

1. E. Andreeva, A. Bogdanov, Y. Dodis, B. Mennink, and J. P. Steinberger. On the indistinguishability of key-alternating ciphers. In *CRYPTO 2013, Part I*, pages 531–550, 2013. (Cited on page 3.)
2. B. Barak, G. Kindler, R. Shaltiel, B. Sudakov, and A. Wigderson. Simulating independence: new constructions of condensers, ramsey graphs, dispersers, and extractors. In *37th ACM STOC*, pages 1–10, 2005. (Cited on page 31.)
3. B. Bauer, P. Farshim, and S. Mazaheri. Combiners for backdoored random oracles. In *CRYPTO 2018, Part II*, pages 272–302, 2018. (Cited on pages 2, 4, 6, and 7.)
4. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. On the indistinguishability of the sponge construction. In *EUROCRYPT 2008*, pages 181–197, 2008. (Cited on page 3.)
5. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145, 2001. (Cited on page 3.)
6. S. Coretti, Y. Dodis, and S. Guo. Non-uniform bounds in the random-permutation, ideal-cipher, and generic-group models. In *CRYPTO 2018, Part I*, pages 693–721, 2018. (Cited on pages 2 and 4.)
7. S. Coretti, Y. Dodis, S. Guo, and J. P. Steinberger. Random oracles and non-uniformity. In *EUROCRYPT 2018, Part I*, pages 227–258, 2018. (Cited on pages 2, 4, 9, 12, 15, and 36.)
8. J.-S. Coron, Y. Dodis, C. Malinaud, and P. Puniya. Merkle-Damgård revisited: How to construct a hash function. In *CRYPTO 2005*, pages 430–448, 2005. (Cited on page 3.)
9. J.-S. Coron, J. Patarin, and Y. Seurin. The random oracle model and the ideal cipher model are equivalent. In *CRYPTO 2008*, pages 1–20, 2008. (Cited on page 3.)
10. Y. Dodis, S. Guo, and J. Katz. Fixing cracks in the concrete: Random oracles with auxiliary input, revisited. In *EUROCRYPT 2017, Part II*, pages 473–495, 2017. (Cited on pages 2 and 4.)
11. Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *EUROCRYPT 2004*, pages 523–540, 2004. (Cited on page 12.)
12. Y. Dodis, M. Stam, J. P. Steinberger, and T. Liu. Indistinguishability of confusion-diffusion networks. In *EUROCRYPT 2016, Part II*, pages 679–704, 2016. (Cited on page 3.)
13. M. Fischlin, C. Janson, and S. Mazaheri. Backdoored hash functions: Immunizing HMAC and HKDF. In *CSF 2018*, pages 105–118, 2018. (Cited on page 2.)

14. M. Göös, S. Lovett, R. Meka, T. Watson, and D. Zuckerman. Rectangles are nonnegative juntas. In *47th ACM STOC*, pages 257–266, 2015. (Cited on pages 4, 8, and 10.)
15. J. J. Hoch and A. Shamir. On the strength of the concatenated hash combiner when all the hash functions are weak. In *ICALP 2008, Part II*, pages 616–630, 2008. (Cited on page 2.)
16. T. Holenstein, R. Künzler, and S. Tessaro. The equivalence of the random oracle model and the ideal cipher model, revisited. In *43rd ACM STOC*, pages 89–98, 2011. (Cited on page 3.)
17. A. Joux. Multicollisions in iterated hash functions. Application to cascaded constructions. In *CRYPTO 2004*, pages 306–316, 2004. (Cited on page 34.)
18. P. K. Kothari, R. Meka, and P. Raghavendra. Approximating rectangles by juntas and weakly-exponential lower bounds for LP relaxations of CSPs. In *49th ACM STOC*, pages 590–603, 2017. (Cited on pages 9 and 10.)
19. C. Lee, C. Lu, S. Tsai, and W. Tzeng. Extracting randomness from multiple independent sources. *IEEE Trans. Information Theory*, 51(6):2224–2227, 2005. (Cited on page 31.)
20. G. Leurent and T. Peyrin. From collisions to chosen-prefix collisions application to full SHA-1. In *EUROCRYPT 2019, Part III*, pages 527–555, 2019. (Cited on page 2.)
21. X. Li. Three-source extractors for polylogarithmic min-entropy. In *56th FOCS*, pages 863–882, 2015. (Cited on page 31.)
22. M. Liskov. Constructing an ideal hash function from weak ideal compression functions. In *SAC 2006*, pages 358–375, 2007. (Cited on page 2.)
23. U. M. Maurer, R. Renner, and C. Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *TCC 2004*, pages 21–39, 2004. (Cited on pages 3 and 7.)
24. B. Pfitzmann and M. Waidner. A model for asynchronous reactive systems and its application to secure message transmission. In *2001 IEEE Symposium on Security and Privacy*, pages 184–200, 2001. (Cited on page 3.)
25. R. Raz. Extractors with weak random seeds. In *37th ACM STOC*, pages 11–20, 2005. (Cited on page 31.)
26. M. Stevens, E. Bursztein, P. Karpman, A. Albertini, and Y. Markov. The first collision for full SHA-1. In *CRYPTO 2017, Part I*, pages 570–596, 2017. (Cited on page 2.)
27. D. Unruh. Random oracles and auxiliary input. In *CRYPTO 2007*, pages 205–223, 2007. (Cited on pages 2, 4, and 9.)