Electronic Colloquium on Computational Complexity, Report No. 192 (2026)

Optimal Amplification via Bias-Resilient Combiners*

Benny Applebaum† Nathan Geier‡

Abstract

Cryptographic combiners take n candidate schemes for some primitive P, and realize it securely provided that at least k candidates are secure. Intuitively, we expect that if we plug n independent instances of a δ-weak candidate for P into the combiner, assuming $\delta \ll (n-k)/n$, security should be amplified and the resulting candidate should exhibit a significantly smaller weakness. This intuition relies on the implicit "all-or-nothing" assumption that each candidate fails with probability δ and is otherwise perfectly secure, allowing us to bound the failure probability of the combiner using a simple binomial tail bound. However, this intuition often fails for standard security notions where, for example, a weak candidate may consistently leak partial information rather than exhibit a clean all-or-nothing failure.

Recently, Applebaum, Bitansky and Geier (CRYPTO 2026) showed that indistinguishability combiners inherently act as security amplifiers. However, this general result incurs a multiplicative loss of roughly 2^{n-k} in the error parameter relative to the natural all-or-nothing bound. Moreover, the combiner-is-amplifier approach is fundamentally restricted to the regime $\delta < 0.5$. Consequently, both the resulting error rate and the amplification threshold δ_0 are suboptimal.

In this work, we overcome these limitations by establishing a new specialized framework of *bias-resilient indistinguishability combiners*. This formulation allows us to achieve the optimal all-or-nothing bound without the exponential penalty. We observe that bias-resilience provides a unifying abstraction for security amplification across different primitives, neatly capturing prior ad-hoc results such as those for weak PRGs and weak NIZK.

As our main application, we use this framework to establish a generalized XOR lemma over prime fields $\mathbb{F}_p$, showing that the sum modulo p of independent weakly pseudorandom elements becomes computationally indistinguishable from uniform. This improves upon a recent work by Shimizu and Yasunaga (STOC 2026) by achieving a sample complexity that is independent of the field size.

Finally, we explore the idealized notion of an *all-or-nothing amplifier*. We establish a tight characterization of the multiplicative penalty incurred when applying such an amplifier to candidate schemes that only guarantee standard weak indistinguishability error.

*The authors are supported by Israel Science Foundation (ISF) grant no. 2805/21 and by the European Union (ERC-2022-ADG) under grant agreement no.101097959 NFITSC.

†Tel Aviv University, Israel, benny.applebaum@gmail.com

‡Tel Aviv University, Israel, nathangeier.cs@gmail.com

ISSN 1433-8092

Contents

1 Introduction

Security amplification is a fundamental goal in cryptography: given a cryptographic primitive that is weakly secure in some sense, we would like to make it strongly secure. It provides a method to guarantee constructions remain secure even if their underlying components exhibit weaknesses, for example, due to physical side-channel leakage, an unproven mathematical assumption that turns out to only partially hold, or advances in technology and the computational power of adversaries. It also allows researchers to simplify the task of constructing advanced primitives by letting them aim merely for weak security. Foundational examples of security amplification include Yao's amplification of weak one-way functions (OWFs) and Yao's XOR lemma for amplifying unpredictability [Yao82, GNW11]. This line of research has been extended to public-key encryption (PKE) [DNR04, HR05, LT13] and key-agreement (KE) [Hol05], weakly verifiable puzzles [CHS05], oblivious transfer (OT) [CK90, DKS99, Wul07], non-interactive zero knowledge (NIZK) [GJS19, BKP+24, BG24, AK25b], and many others.

The notion of *robust combiners* [HKN+05, Her05] is closely related, yet distinct. A (k, n)-robust combiner is a construction that takes n candidate schemes and guarantees security as long as at least k of them are secure. Such combiners are typically built using simple secret-sharing techniques, or more sophisticated tools such as secure multiparty computation (MPC) [HIKN08, AK25a].

To illustrate the relation between combiners and amplifiers, let us review two classical examples that will be used across the introduction.

Direct product. Given n candidate one-way functions $f_1, \ldots, f_n$, it is straightforward to show that the mapping $(x_1, \ldots, x_n) \mapsto (f_1(x_1), \ldots, f_n(x_n))$ is one-way provided that at least one of the functions is one-way. Thus, the direct product yields a robust combiner. Demonstrating that the same construction also serves as an amplifier is considerably more subtle (see Section 2.3 of [Gol01]).

The XOR lemma (generalized). Suppose we are given n weakly random elements $g_1, \ldots, g_n$ of some group G, whose statistical distance from uniform is $\delta > 0$. Trivially, applying the group operation $\prod_{i=1}^{n} g_i$ results in a robust combiner: if at least one element is uniformly distributed, then so is their product. However, analyzing the amplification properties of this combiner is nontrivial. In particular, considering any characteristic-2 finite field $\mathbb{F}_{2^m}$ and the addition operation (XOR) reveals an inherent barrier to any general combiner-is-amplifier theorem: while the combiner guarantee remains unchanged, amplification fails completely for weakness parameter $\delta \geq 0.5$; we observe this by noting that a uniformly distributed even element (LSB 0) is 0.5-close to the uniform distribution over the field, yet the sum of such elements is always even. On the other hand, by switching to a prime field $\mathbb{F}_p$ the above no longer applies, and in fact, one can show the sum acts as a valid amplifier as long as δ is bounded away from $(1 - 1/p)$.[1]

The above example underscores the distinction between the inherently "all-or-nothing" nature of combiners and the "consistent small bias" phenomenon that may arise in weak candidates, including the fact that in some settings it may be overcome.

Indistinguishability combiners. To study combiners in a general way, we must abstract away the concrete properties of any specific cryptographic primitive. However, a fully general combiner definition seems difficult to reason about, since the underlying security requirement could be arbitrary. Instead, we choose to focus on the class of indistinguishability combiners (iC) [ABG26], i.e.,

[1] Distance $1 - 1/p$ corresponds to a "deterministic" random variable (RV) whose entire probability mass is concentrated on a single field element. In this case, amplification is impossible.

combiners for which the security requirement is given by an indistinguishability game. As demonstrated by [ABG26], this class is structured enough to support general theorems, yet expressive enough to capture many cryptographic primitives of interest.

Syntactically, an indistinguishability combiner $(C_0^{(\cdot)}, C_1^{(\cdot)})$ is a pair of oracle-aided algorithms given access to n candidate constructions $\vec{S} = (S_1, \ldots, S_n)$ over some input/output space.[2] Such a vector of candidates is said to be *k-ideal* with respect to a vector of ideal systems $\vec{I} = (I_1, \ldots, I_n)$ if at least k out of the n candidates are ideal. The indistinguishability combiner property asserts that for every k-ideal vector of candidates, the combined constructions $C_0^{\vec{S}}$ and $C_1^{\vec{S}}$ are indistinguishable.

For example, to capture the XOR combiner, we can set the ideal candidates $\vec{I}$ to the uniform distribution, have $C_0^{\vec{S}}$ output the XOR of the candidates, and $C_1^{\vec{S}}$ output a uniform element irrespective of the candidates. This yields a $(1, n)$ iC since perfect indistinguishability between the two modes is guaranteed for every 1-ideal vector $\vec{S}$. For the XOR lemma, we would like to bound the distinguishing advantage between $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ when no candidate is fully ideal, but rather each candidate R_i is δ-close to ideal I_i.

A dream-version amplification theorem. Throughout the paper, let $\delta^{(k,n)}$ denote the probability of failing to achieve at least k successes in n independent attempts, where each attempt fails with probability δ. That is,

$$\delta^{(k,n)} := \Pr[\mathrm{Bin}(n, \delta) > n - k] = \Pr[\mathrm{Bin}(n, 1-\delta) < k] = \sum_{i=0}^{k-1} \binom{n}{i} \delta^{n-i} (1-\delta)^i \,.$$

A very intuitive yet false (as demonstrated above) general combiner-is-amplifier theorem is:

Idealized Theorem 1 (Combiner-is-Amplifier; Dream Version)**.** *Given a (k, n) combiner $(C_0^{(\cdot)}, C_1^{(\cdot)})$, let $\vec{R} = (R_1, \ldots, R_n)$ be such that each R_i is δ-close to ideal I_i. Then $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ are $\delta^{(k,n)}$-close.*

The $\delta^{(k,n)}$ bound naturally captures the idealized all-or-nothing setting, as illustrated by the following canonical example:

Example 1. Consider the (k, n) indistinguishability combiner over bits

$$C_b(b_1, \ldots, b_n) := b \cdot \mathbb{I}\left[\sum_{i=1}^{n} b_i > n - k\right] ,$$

with the ideal $I_i \equiv 0$. If each R_i is δ-close to I_i, then $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ are $\delta^{(k,n)}$-close.

However, the dream version fails in general because using standard measures of "closeness" such as maximal distinguishing advantage, systems may exhibit weakness in much more involved ways. For example, for a (bit-string, leakage) pair that is 0.5-close to uniform $(S, f(S)) \approx_{0.5} (U, f(S))$, it could be the case that f always leaks the least significant bit (LSB), rather than leaking all of S with probability half and otherwise leaking nothing. Because standard security definitions must accommodate this type of consistent partial leakage, any general proof attempting to bound the combiner's amplification error cannot simply rely on independent all-or-nothing failures, but must instead account for the worst-case accumulation of leakage across all n candidates.

[2] Throughout, distinct candidate systems are instantiated with mutually independent random tapes.

Prior work. Building upon an earlier line of work on neutralizing constructions [MPR07, MT09, LM20], Applebaum et al. [ABG26] introduced the notion of indistinguishability combiners, proving that they inherently act as security amplifiers. However, forcing standard combiners to act as amplifiers in the presence of arbitrary partial leakage incurs a steep cost: the penalty in their general result yields a security error bounded by

$$\delta'(k,n) := \binom{n}{k-1}(2\delta)^{n-k+1}.$$

Compared to the natural all-or-nothing bound $\delta^{(k,n)} \leq \binom{n}{k-1}\delta^{n-k+1}$, the universal bound $\delta'(k,n)$ suffers a severe exponential penalty of roughly 2^{n-k+1}. As the base of the exponent becomes 2δ, this penalty represents an inherent limitation of the standard combiner-is-amplifier approach: it is fundamentally restricted to the regime where $\delta < 0.5$, meaning the general approach fails entirely for weaker candidates where the distinguishing advantage between real and ideal is $\delta \geq 0.5$.

This work. We continue the theoretical study of the relationship between combiners and amplifiers. Concretely, we ask:

> Is there a stronger notion of combiner that yields "ideal" amplification—namely, amplification with the optimal error rate $\delta^{(k,n)}$—and that applies to every amplification threshold $\delta < 1$?
>
> More generally, what is the quantitative gap between amplification in the all-or-nothing setting (where each candidate fails with probability δ) and standard amplification (where each candidate is merely δ-close to ideal)?

We provide complete answers to both questions and use the resulting framework to derive a new arithmetic XOR lemma.

1.1 Our Contributions

Our first contribution is definitional and conceptual: we introduce a new framework of *bias-resilient indistinguishability combiners* that overcomes the limitations of the existing indistinguishability combiner notion. Roughly speaking, we strengthen the combiner requirement such that a system counts as "good" not only if it is ideal, but also if it is not too biased with respect to ideal. Specifically, for $\delta \in (0,1)$ we say a system I' is δ-biased with respect to ideal system I, if for every transcript τ we have $I'(\tau) \leq I(\tau)/(1-\delta)$, i.e., the probability (over the random coins of the system) to obtain each and every transcript increases at most by a multiplicative factor $(1-\delta)^{-1}$. We observe that this framework provides a natural abstraction for tight security amplification: several previously known ad-hoc amplification results in the literature—such as those for weak PRGs and weak NIZK—can be seen as implicitly bias-resilient.

For our first result, we show that bias-resilient indistinguishability combiners act as optimal amplifiers, achieving the all-or-nothing bound.

Informal Theorem 1 (Bias-Resilient Combiners Are Optimal Amplifiers; See Theorem 1)**.** *Given a (k,n) δ-bias-resilient indistinguishability combiner $(C_0^{(\cdot)}, C_1^{(\cdot)})$ with respect to ideal systems $\vec{I} = (I_1,\ldots,I_n)$, let $\vec{R} = (R_1,\ldots,R_n)$ be a tuple of n real systems such that $\forall i \in [n] : R_i \approx_\delta I_i$. Then $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ are $\delta^{(k,n)}$-indistinguishable.*

Apart from unifying previous results, the theorem also yields a new generalization of the XOR lemma to prime fields. Specifically, by showing that summation acts as a bias-resilient combiner with respect to ideal uniform random variables, we obtain the following theorem.

Informal Theorem 2 (Prime Field XOR Lemma; See Theorems 5 and 6)**.** *Let $p \geq 3$ be a prime, $\delta < 1 - 1/p$ and $n \gg (1-\delta-1/p)^{-3}$. Then for independent RVs $(X_1, Z_1), \ldots, (X_n, Z_n)$ over $\mathbb{F}_p \times \mathcal{Z}$ such that $\forall i \in [n] : (X_i, Z_i) \approx_\delta (U_p, Z_i)$, we have $\left(\sum_{i=1}^n X_i,\ \vec{Z}\right) \approx \left(U_p,\ \vec{Z}\right)$.*

That is, if given Z_i (the "puzzle" or, more generally, some "side-information"), the random variable X_i (the "solution" or the "hidden-value") is somewhat hidden—in the sense that it is δ-close to uniform over $\mathbb{F}_p$—then given n puzzles, the sum of the solutions $\sum_{i=1}^n X_i$ becomes statistically close to uniform. This type of arithmetic XOR lemma was recently studied by [SY26], motivated in part by applications to graph theory. Our version improves upon [SY26] in most parameter regimes (see the technical overview and Section 4.3.2 for a discussion). As another application, consider the problem of amplifying the security of a weakly secure encryption scheme whose message space is $\mathbb{F}_p$. The computational variant of the theorem (Theorem 6) can be used to show that if we additively share the message into n shares and encrypt each share separately, the resulting scheme will be secure even for a mild choice of n and δ.

Finally, we bound the multiplicative penalty when feeding a δ-weak candidate into an all-or-nothing-amplifier, i.e., an amplifier analyzed with respect to candidates exhibiting perfect security with probability p.

Informal Theorem 3 (All-or-Nothing Amplifiers as Standard Amplifiers; See Theorem 7)**.** *Given a (p, n) all-or-nothing indistinguishability amplifier $(C_0^{(\cdot)}, C_1^{(\cdot)})$ with respect to ideal systems $\vec{I} = (I_1, \ldots, I_n)$, let $\vec{R} = (R_1, \ldots, R_n)$ be a tuple of n real systems such that $\forall i \in [n] : R_i \approx_\delta I_i$. Then $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ are distinguishable with advantage at most $\left(1 + \frac{2p\delta}{1-p}\right)^n$ times larger than the amplifier's guarantee for p-perfect candidates.*

We also show that this penalty ratio is tight.

1.2 Technical Overview

Attempting to amplify weakly secure primitives using robust combiners, or more generally under a clean "all-or-nothing" failure assumption, is very natural and tempting. And yet, standard security notions such as the maximal distinguishing advantage (total variation distance) do not offer such a guarantee. As an instructive example, let us consider a random variable X distributed over some finite field $\mathbb{F}_{2^m}$ that is equal to 0 with probability 1/2 and is otherwise uniformly distributed over even elements (LSB 0). In this setting we have $X \approx_{0.5} U_{2^m}$, yet the natural combiner by addition over $\mathbb{F}_{2^m}$ of n independent instances completely fails to amplify. We observe that, in fact, the combiner faces two separate issues: (1) Bad event: with probability 2^{-n} all inputs will be zero, and the output remains a constant zero. (2) "Good" event: at least one instance was sampled from the uniform distribution over even elements. However, even under the most optimistic assumption that all random variables were sampled like that, their sum is uniformly distributed over the even elements and thus the statistical distance to uniform remains 1/2—no amplification occurred at all.

In order to bypass the inherent barrier above and obtain a general combiner-is-amplifier bound, [ABG26] restricted the statistical distance δ to be strictly less than 1/2 (in fact, even smaller depending on k and n), and proved an upper bound on the penalty resulting from feeding statistically weak candidates into any indistinguishability combiner.

Instead, we take a completely different approach: rather than trying to force general combiners to act as imperfect amplifiers, we aim to characterize a stronger definition that allows combiners to amplify ideally.

1.2.1 Bias-Resilient Indistinguishability Combiners

Motivated by the example above, we introduce the framework of *bias-resilient indistinguishability combiners*: such a combiner must be able to accept "semi-ideal" candidates that are at most δ-biased compared to ideal, counting them toward the k threshold. Specifically, we say that S is δ-biased with respect to I and denote $S \leq I/(1-\delta)$, if the probability to obtain every element τ (or transcript in the interactive setting) increases by at most a multiplicative factor of $1/(1-\delta)$, that is, $\forall \tau : S(\tau) \leq I(\tau)/(1-\delta)$. Note that every combiner is trivially 0-bias-resilient, and generally speaking, a combiner that is δ-bias-resilient is also δ'-bias-resilient for every $\delta' < \delta$.

Our main technical tool in proving that bias-resilient combiners amplify optimally is the statistical coupling lemma, or more generally, its generalization to the interactive setting by Lanzenberger and Maurer [LM20]. Focusing on the non-interactive setting for the sake of simplicity, a natural way to interpret the coupling lemma is as follows: For any pair of distributions R and I such that $R \approx_\delta I$, there exist some shared distribution S and residual failure distributions R', I', such that we can write $R \equiv (1-\delta)S + \delta R'$ and $I \equiv (1-\delta)S + \delta I'$.[3] Then the key observation is that the shared distribution S is at most δ-biased with respect to ideal I, and the real distribution R is equal to S in a clean all-or-nothing sense. This allows us to achieve the $\delta^{(k,n)}$ all-or-nothing binomial-tail error bound. While the proof is rather clean and direct, the notion itself offers a surprisingly useful shift in perspective.

1.2.2 Generalizing the XOR Lemma

Let us get back to the example of a uniform even element over a finite field. Intuitively, the reason why we failed in $\mathbb{F}_{2^m}$ is that the biased shared distribution, conditioned on the coupling event (between X and U), was trapped inside a subgroup. By switching to a prime field $\mathbb{F}_p$, we eliminate all non-trivial subgroups.

Next, we make the observation that any random variable S that is δ-biased with respect to U_p has min-entropy at least $H_\infty(S) \geq \log(p) - \log(1/(1-\delta))$, i.e., it is missing at most $\log(1/(1-\delta))$ bits of min-entropy. Hence, in order to construct an appropriate (k,n) δ-bias-resilient combiner, it is sufficient to show a sum-extraction lemma over prime fields. Specifically, we need to bound the statistical distance between the sum of any k RVs over $\mathbb{F}_p$, each with missing $\log(1/(1-\delta))$ bits of min-entropy, and the uniform distribution. In fact, our extraction theorem (Theorem 3) is somewhat more subtle. Usually, in the setting of the XOR lemma, we wish to consider random variables that are weakly indistinguishable from uniform given some side information $(X,Z) \approx_\delta (U_p, Z)$. Consequently, the min-entropy $H_\infty(X_{|Z=z})$ may vary depending on the specific leakage z, prompting us to use the notion of average min-entropy [DORS08]. Our extraction theorem is based on Fourier analysis, building upon and extending ideas introduced by Benhamouda et al. [BDIR18].

We highlight the main differences between our approach and [SY26], who recently studied an extension of the XOR lemma to arithmetic over prime fields. Most obviously, they deal with the setting where X is d-unpredictable given Z, i.e., cannot be guessed with probability better than $d \in (0,1)$ given z. In the statistical setting, this translates to $\widetilde{H}_\infty(X|Z) \geq \log(1/d)$, which they state in different terms since their main focus is on the computational setting. However, there is a general technique (discussed below) to lift statistical results to the computational setting for both unpredictability and pseudorandomness. We note that d-unpredictability translates into statistical distance at most $1 - \frac{1}{p \cdot d}$ to uniform, hence our setting captures theirs.

Conversely, attempting to capture weak indistinguishability via unpredictability fails severely

[3] A convex combination over distributions denotes the distribution resulting from sampling according to the given weights.

over large fields. Suppose that $(X, Z) \approx_{0.5} (U_p, Z)$. Working directly with indistinguishability, we can condition on the coupling event occurring with probability $1/2$, meaning we effectively sum $\approx n/2$ random variables that each carry an average min-entropy of $\log(p) - 1$. However, if we view this 0.5 indistinguishability gap strictly through the lens of unpredictability, we can only guarantee $d = 0.5$, which implies a mere $\widetilde{H}_\infty(X|Z) \geq 1$. Here lies the core of the issue: under the unpredictability framework, we end up summing n random variables with a poor average min-entropy of 1 bit, rather than focusing on the $\approx n/2$ good instances with average min-entropy $\log(p)-1$. Because of this severe entropy bottleneck, achieving ε-uniformity requires $n \gg p^2 \log(1/\varepsilon)$ copies. This p^2 dependence is catastrophic if p is exponentially large, such as when amplifying a weak encryption scheme over large prime fields. In contrast, our analysis bypasses the field size entirely, requiring only $O(\log(1/\varepsilon))$ copies. We also remark that their Fourier analysis itself is focused on the setting where the min-entropy is smaller than 1, as the main focus of their paper is the very weak unpredictability regime where d is close to 1. For a more detailed comparison of parameters, see Section 4.3.2.

Finally, let us address the computational generalization, which is done using a generic approach. Using the hardcore lemma [Imp95, Hol05], or more specifically its indistinguishability variant [MT10, BG24], we have that: If $(X, Z) \overset{\mathsf{c}}{\approx}_\delta (U_p, Z)$, then there exists some hybrid distribution (X', Z') such that $(X, Z) \overset{\mathsf{c}}{\approx} (X', Z') \approx_\delta (U_p, Z)$. Importantly, the latter transition is *statistical*. We can therefore invoke the XOR lemma on $((X'_i, Z'_i))_{i\in[n]}$ and switch back to $((X_i, Z_i))_{i\in[n]}$ using the fact that the sum modulo p is efficiently computable.

1.2.3 All-or-Nothing Amplifiers

From a broader perspective, we can view the combiner-is-amplifier approach as attempting to simplify the security analysis using an idealized all-or-nothing assumption, then paying some penalty to switch to standard security notions. In this section, we explicitly formulate and study the framework of *all-or-nothing amplifiers*, parametrized by (p, n), where n denotes the number of candidates and p the probability that each candidate is perfectly ideal. In particular, note that any (k, n) combiner induces a (p, n) all-or-nothing amplifier with additional error $(1 - p)^{(k,n)}$.

Formally, we say that a candidate S is p-ideal if the ideal system I is at most $(1 - p)$-biased with respect to S, i.e., $I \leq S/p$. In order to obtain an upper bound on the penalty when applying the combiner on real systems $\vec{R} = (R_1, \ldots, R_n)$ where $\forall i \in [n] : R_i \approx_\delta I_i$, we express the real system R_n as an affine combination of p-ideal systems (and continue recursively). Specifically, we define the following hybrids using the residual failure systems R', I' given by the coupling lemma, as presented in Section 1.2.1:

$$\begin{aligned} H_R &\equiv pI + (1-p)R\,, \\ H_{R'} &\equiv pI + (1-p)R'\,, \\ H_{I'} &\equiv pI + (1-p)I'\,, \end{aligned}$$

and observe that

$$R \equiv H_R + \frac{p\delta}{1-p} H_{R'} - \frac{p\delta}{1-p} H_{I'}\,.$$

Then, by summing the absolute values of the coefficients, the error resulting from using R as the last system is at most $(1 + \frac{2p\delta}{1-p})$ times larger than the maximal error when using any one of the three. We apply the induction hypothesis on each of them to obtain that the total error over n real systems is at most $(1 + \frac{2p\delta}{1-p})^n$ times larger than the combiner's guarantee for p-ideal candidates.

It is worth noting that we may rederive the combiner-is-amplifier bound for $k = 1$, which is $(2\delta)^n$ (corresponding to the boolean XOR lemma), via the all-or-nothing framework. Specifically, for every $p \in (0,1)$, a $(1,n)$ combiner implies a (p,n) all-or-nothing amplifier with additional error $(1-p)^n$. Multiplying this by the all-or-nothing penalty, we obtain

$$(1-p)^n \cdot \left(1 + \frac{2p\delta}{1-p}\right)^n = (1 - p + 2p\delta)^n .$$

Since this is true for every p, we may take $p \to 1$ to achieve $(2\delta)^n$.

Finally, we demonstrate that the above penalty ratio is tight. At a high level, we construct a non-interactive combiner over bits, where the real candidate is 0 and the ideal candidate is $\mathrm{Ber}(\delta)$. For every input bit b_i, the combiner samples a biased bit with bias either μ_0 or μ_1 depending on the value of the bit b_i. The indistinguishability combiner $C_b(b_1, \ldots, b_n)$ then XORs all these biased bits, and uses the result to mask its secret bit b. By carefully choosing the parameters μ_0, μ_1 and analyzing the (p,n) all-or-nothing security of the construction, we show that: for every $p \in (0,1), n \in \mathbb{N}$, target real indistinguishability $\varepsilon' \in (0,1)$ and $\delta \in [0.5, 1)$, the tight ratio may be obtained, i.e., we may set μ_0, μ_1 such that $C_0(\vec{R}) \approx_{\varepsilon'} C_1(\vec{R})$ and $C_0(\vec{S}) \approx_{\varepsilon} C_1(\vec{S})$ for every p-ideal $\vec{S}$, with $\varepsilon = \varepsilon'/(1 + \frac{2p\delta}{1-p})^n$. If $\delta \in (0, 0.5)$, we obtain a slightly weaker ratio-tightness guarantee that only holds for sufficiently small ε'. For more details, see Section 5.1.

2 Preliminaries

For $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \ldots, n\}$. For a sequence $(x_1, \ldots, x_n)$ and a subset $V \subseteq [n]$, we denote by x_V the subsequence corresponding to indices in V. For a distribution X over a finite set $\mathcal{X}$, we use $x \leftarrow X$ to denote the result of sampling according to X, and $x \leftarrow \mathcal{X}$ to denote a uniformly random sample from the set. For a random variable X, we denote by $H_\infty(X) := -\log(\max_x \Pr[X = x])$ its min-entropy. For jointly distributed random variables (X, Z), we denote by $\widetilde{H}_\infty(X|Z)$ the average min-entropy of X given Z [DORS08]:

$$\widetilde{H}_\infty(X \mid Z) = -\log\left(\mathop{\mathbb{E}}_{z \leftarrow Z}\left[\max_x \Pr[X = x \mid Z = z]\right]\right) .$$

For two random variables X, Y and $\varepsilon \in [0,1]$, we write $X \approx_\varepsilon Y$ to denote the fact that the *statistical distance* between X and Y is at most ε, and say that X and Y are ε *statistically indistinguishable*. For $s \in \mathbb{N}$, we write $X \overset{\mathsf{c}}{\approx}_{s,\varepsilon} Y$ to denote the fact that the *computational distance* between X and Y against s-sized circuits is at most ε, namely for every such Boolean circuit C we have $|\mathbb{E}[C(X) - C(Y)]| \leq \varepsilon$. In that case, we say that X and Y are ε *computationally indistinguishable against s-sized circuits*.

2.1 Finite Discrete Systems

In this paper, we focus on finite discrete systems. Our tree-based formalism is essentially equivalent to the standard characterization of random systems using sequences of conditional distributions [Mau02], and our classification is based on [MR11]. We emphasize that our framework aims to capture the external *behavior* of a system, not its implementation. A reader already comfortable with the notion of a system may skip to Section 2.2.

Formalizing Interactive Systems. We characterize external interaction using three types of systems: *resources*, *distinguishers*, and *converters*. In a nutshell:

- *Resources* abstract randomized stateful (possibly inefficient) algorithms: Given an input, the algorithm updates its internal state and samples an output based on a probability distribution that depends on the internal state.
- *Distinguishers* abstract randomized oracle-aided algorithms with a binary output. When the oracles are instantiated with a resource, the algorithm defines a probability distribution over 0/1 values.
- *Converters* abstract oracle-aided randomized stateful (possibly inefficient) algorithms. Given an input, the algorithm makes a series of calls to its oracles while updating its internal state, and generates an output. By instantiating the oracles with resources, the converter simplifies into a resource.

We denote by $\mathcal{X}, \mathcal{Y}$ the spaces for incoming and outgoing messages. We always assume these are finite sets that include the special symbol $\perp$: if a resource accepts at most $h \in \mathbb{N}$ queries any subsequent queries are assumed to return $\perp$, and similarly, any invalid queries (not in input space) are assumed to return $\perp$. See more notational conventions below.

We model systems using alternating trees, which naturally capture the back-and-forth of stateful interaction: each alternating level represents a shift between queries and responses, and every node corresponds to a partial transcript.

Definition 1 (Finite Alternating Tree)**.** For finite sets $\mathcal{X}, \mathcal{Y}$ and $h \in \mathbb{N}$, an $(h, \mathcal{X}, \mathcal{Y})$ alternating tree is a rooted tree of height $2h$ where every node at even depth has $|\mathcal{X}|$ children labeled by the elements of $\mathcal{X}$, and every node at odd depth has $|\mathcal{Y}|$ children labeled by the elements of $\mathcal{Y}$.

Definition 2 (Resource)**.** For finite sets $\mathcal{X}, \mathcal{Y}$ and $h \in \mathbb{N}$, an h-round $(\mathcal{X}, \mathcal{Y})$-resource is an $(h, \mathcal{X}, \mathcal{Y})$ alternating tree where every *odd-depth* node v is assigned a distribution $D_v : \mathcal{Y} \to [0, 1]$.

Given a sequence of inputs $(x_i)_{i \in [h]}$, the outputs $(y_i)_{i \in [h]}$ are obtained by advancing from the root v_ϵ to the x_1th child v_{x_1}, sampling $y_1 \leftarrow D_{v_{x_1}}$ and advancing according to y_1 then x_2 arriving at v_{x_1, y_1, x_2}, sampling $y_2 \leftarrow D_{v_{x_1, y_1, x_2}}$ and so on.

For example, a random variable Y over $\mathcal{Y}$ could be modeled as a 1-round resource that ignores its first query x and returns $y \leftarrow Y$ (i.e., $D_v \equiv Y$ for all depth-1 nodes).

Definition 3 (Distinguisher)**.** For finite sets $\mathcal{X}, \mathcal{Y}$ and $h \in \mathbb{N}$, an h-round $(\mathcal{X}, \mathcal{Y})$-distinguisher is an $(h, \mathcal{X}, \mathcal{Y})$ alternating tree where every *internal even-depth* node v is assigned a distribution $D_v : \mathcal{X} \to [0, 1]$, and every *leaf* v is assigned a boolean value $b_v \in \{0, 1\}$.

Given a distinguisher D and a resource S, their interaction D^S results in a boolean distribution: fill the odd-depth nodes of D using the distributions of S to obtain a tree where all internal nodes are assigned distributions and all leaves are assigned boolean values—which naturally induces a boolean distribution. We remark that due to our $\perp$ conventions, the interaction is well defined even if S is an h'-round resource for $h' < h$.

Definition 4 (Converter)**.** For finite sets $\mathcal{X}, \mathcal{Y}$ and $h, n \in \mathbb{N}$, an h-round n-oracle-aided $(\mathcal{X}, \mathcal{Y})$-converter is an h-round $(\mathcal{X}, \mathcal{Y}')$-resource for $\mathcal{Y}' := (\{0\} \cup [n]) \times \mathcal{Y}$.

A 0-oracle-aided converter corresponds to a plain $(\mathcal{X}, \mathcal{Y})$-resource. Intuitively, the augmented alphabet $\mathcal{Y}'$ routes the converter's outgoing messages either to the external distinguisher (index 0) or to one of its n underlying oracles. Given an index $i \in [n]$ and a $(\mathcal{Y}, \mathcal{X})$-resource S_i, we naturally obtain an $(n-1)$-oracle-aided converter by embedding it.

Notational Conventions.

1. Interaction: Given a distinguisher D, an n-oracle-aided converter C, and n resources $\vec{S} = (S_1, \ldots, S_n)$, we denote the execution of D and $C^{\vec{S}}$ by $D(C^{\vec{S}})$ rather than $D^{C^{\vec{S}}}$ for notational convenience.

2. Terminology: We may refer to a resource simply as a system, and to an n-oracle-aided converter simply as an n-oracle-aided system.

3. Alphabets: In the above, it is possible to consider distinct input and output spaces for every resource, or even for every individual node in the tree. However, for the sake of simplicity, we choose to consider a single universal input space $\mathcal{X}$ and output space $\mathcal{Y}$, which can be viewed as the union of all sub-spaces. Furthermore, throughout most of the paper, we do not explicitly mention the sets $(\mathcal{X}, \mathcal{Y})$ and also omit the upper bound h on the number of rounds (when irrelevant).

4. Transcripts: A transcript represents a (possibly partial) interaction path $\tau = (x_1, y_1, \ldots, x_{h'}, y_{h'})$ for some $h' \leq h$. For a system I, we denote by $I(\tau)$ the probability of generating the responses in τ given the queries in τ. This is exactly the product of the probabilities assigned by the distributions D_v along the corresponding path in the tree.

5. Convex Combinations: Given systems S_0, S_1 and $p \in [0,1]$, we denote by $(1-p)S_0 + pS_1$ the system equivalent to sampling $b \leftarrow \mathrm{Ber}(p)$ and using S_b.[4] Furthermore, for a set $\mathcal{X}$ and given an ensemble of systems $\{S_x\}_{x \in \mathcal{X}}$ with a random variable X over $\mathcal{X}$, we denote by S_X the system equivalent to sampling $x \leftarrow X$ and using S_x.

6. Implementation: For a resource S, we define $\mathrm{size}(S)$ as the size of the smallest circuit realizing it. We allow circuits to receive advice sampled from arbitrary distributions, as we deal with indistinguishability where the best such advice can later be fixed. For a converter $C^{(\cdot)}$, we define $\mathrm{size}(C^{(\cdot)})$ as the size of its oracle-aided circuit. We define the composed size $\mathrm{size}(C^{\vec{S}})$ as the size of the circuit obtained by replacing the oracle gates in $C^{(\cdot)}$ with the circuits realizing $\vec{S}$, capturing the actual cost of the combined implementation rather than the minimal circuit size of the resulting resource.

   We may sometimes abuse notation and associate an implementation with the system it realizes, feeding it to a distinguisher or a converter.

Definition 5 (Statistical Distance of Systems)**.** For $\delta \in [0,1]$, we say that two systems R and I are δ statistically close, denoted $R \approx_\delta I$, if for every (unbounded adaptive) distinguisher D we have

$$\left| \mathbb{E}\left[D^R - D^I \right] \right| \leq \delta \,.$$

If $\delta = 0$, we write $R \equiv I$.

Lemma 1 (System Decomposition Lemma, Reinterpretation of [LM20, Theorem 2])**.** *For any finite discrete systems R and I with statistical distance $\delta \in [0,1]$, there exist systems S, R', I' such that*

$$R \equiv (1-\delta)S + \delta R' \,,$$
$$I \equiv (1-\delta)S + \delta I' \,.$$

[4] Equivalently, we can view a system as a consistent weight assignment over all partial transcripts (nodes). Then we simply take the point-wise convex combination over these weights $(1-p)S_0(\tau) + pS_1(\tau)$.

In [LM20], randomized systems are modeled as probability distributions over deterministic systems. To be more accurate, since this representation is not unique, randomized systems are in fact seen as equivalence classes over these distributions. It is shown that the maximal distinguishing advantage between two random systems is equal to the minimal statistical distance between distributions in their equivalence classes. Then, standard coupling of probability distributions implies the above.

2.2 Indistinguishability Combiners

Definition 6 (Indistinguishability Combiner [ABG26])**.** For finite sets $(\mathcal{X}, \mathcal{Y})$, $\varepsilon \in [0,1]$, $k, n \in \mathbb{N}$ with $k \leq n$, an (ε, k, n) indistinguishability combiner with respect to ideal $(\mathcal{Y}, \mathcal{X})$-systems $\vec{I} = (I_1, \ldots, I_n)$ against an $(\mathcal{X}, \mathcal{Y})$-distinguisher class $\mathcal{D}$, is a pair of n-oracle-aided $(\mathcal{X}, \mathcal{Y})$-systems $(C_0^{(\cdot)}, C_1^{(\cdot)})$ such that for all $(\mathcal{Y}, \mathcal{X})$-systems $\vec{S} = (S_1, \ldots, S_n)$:

$$|\{i \in [n] \mid S_i \equiv I_i\}| \geq k \implies \forall D \in \mathcal{D} : \left| \mathbb{E}\left[D(C_0^{\vec{S}}) - D(C_1^{\vec{S}}) \right] \right| \leq \varepsilon \,.$$

Remark 1. While the above definition may seem too restrictive at first since it targets a specific choice for the ideal systems $\vec{I}$, typically the same combiner $(C_0^{(\cdot)}, C_1^{(\cdot)})$ satisfies it with respect to any ideal $\vec{I}$ in some "good class" of systems.

3 Bias-Resilient Indistinguishability Combiners

We now formalize our main abstraction: bias-resilient combiners. Rather than forcing standard combiners to act as imperfect amplifiers, we identify the "semi-ideal" condition under which a combiner amplifies optimally.

First, to establish a baseline for our parameters, we note that Example 1 naturally generalizes to the imperfect setting. In this setting, combiners do not guarantee perfect indistinguishability when k systems are ideal, but rather guarantee only ε-indistinguishability.

Example 2. Consider the (ε, k, n) indistinguishability combiner

$$C_b(b_1, \ldots, b_n) := b \cdot \left(\rho \vee \mathbb{I}\left[\sum_{i=1}^{n} b_i > n - k \right] \right) ,$$

with an independent biased bit $\rho \equiv \mathrm{Ber}(\varepsilon)$ and the ideal $I_i \equiv 0$. If each R_i is δ-close to ideal, then $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ are $\left((1-\varepsilon) \cdot \delta^{(k,n)} + \varepsilon\right)$-close.

We exactly match this bound using the following notion of a bias-resilient indistinguishability combiner.

Definition 7 (Biased System)**.** For $\delta \in [0,1]$ and two h-query $(\mathcal{Q}, \mathcal{A})$ systems I', I, we say that I' is δ-biased with respect to I if for every leaf (transcript) τ we have

$$I'(\tau) \leq I(\tau)/(1-\delta) \,.$$

In that case, we denote $I' \leq I/(1-\delta)$.

Definition 8 (Bias-Resilient Indistinguishability Combiner)**.** For finite sets $(\mathcal{Q}, \mathcal{A})$, $\varepsilon, \delta \in [0,1]$, $k, n \in \mathbb{N}$, an (ε, k, n) δ-bias-resilient indistinguishability combiner with respect to ideal $(\mathcal{A}, \mathcal{Q})$-systems $\vec{I} = (I_1, \ldots, I_n)$ against a $(\mathcal{Q}, \mathcal{A})$-distinguisher class $\mathcal{D}$, is a pair of n-oracle-aided $(\mathcal{Q}, \mathcal{A})$-systems $(C_0^{(\cdot)}, C_1^{(\cdot)})$ such that for all $(\mathcal{A}, \mathcal{Q})$-systems $\vec{S} = (S_1, \ldots, S_n)$:

$$|\{i \in [n] \mid S_i \leq I_i/(1-\delta)\}| \geq k \implies \forall D \in \mathcal{D} : \left| \mathbb{E}\left[D(C_0^{\vec{S}}) - D(C_1^{\vec{S}}) \right] \right| \leq \varepsilon \,.$$

Note that any δ-bias-resilient combiner is also δ'-bias-resilient for all $\delta' < \delta$. Returning to Example 2, because the ideal system is deterministic, the combiner is inherently bias-resilient. Hence, our upper bound in Theorem 1 below is perfectly tight.

Theorem 1 (Bias-Resilient Indistinguishability Combiners Are Optimal Amplifiers)**.** *For $\varepsilon, \delta \in [0,1]$, $k, n \in \mathbb{N}$, given an (ε, k, n) δ-bias-resilient indistinguishability combiner $(C_0^{(\cdot)}, C_1^{(\cdot)})$ with respect to ideal systems $\vec{I} = (I_1, \ldots, I_n)$ against a distinguisher class $\mathcal{D}$, let $\vec{R} = (R_1, \ldots, R_n)$ be a tuple of n real systems such that $\forall i \in [n] : R_i \approx_\delta I_i$. Then, we have*

$$\forall D \in \mathcal{D} : \left| \mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}}) \right] \right| \leq (1-\varepsilon) \cdot \delta^{(k,n)} + \varepsilon \,.$$

Proof. Let us denote by $\Delta(k,n) = \Delta_{\varepsilon,\delta}(k,n)$ the worst-case distinguishing advantage over an (ε, k, n) δ-bias-resilient combiner. That is, for every such combiner with respect to some ideal systems and distinguisher class $\mathcal{D}$, and for any real systems such that each is δ-close statistically to ideal, $\Delta(k,n)$ bounds the distinguishing advantage of any $D \in \mathcal{D}$ when the real systems are used.

Fix any distinguisher $D \in \mathcal{D}$, and let $\Delta^D C^{\vec{R}} := \mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}}) \right]$. By the system decomposition lemma (Lemma 1), for all $i \in [n]$, there exist systems S_i, R_i', I_i' such that

$$\begin{aligned} R_i &\equiv (1-\delta) S_i + \delta R_i' \,, \\ I_i &\equiv (1-\delta) S_i + \delta I_i' \,, \end{aligned}$$

and we note that S_i is δ-biased with respect to I_i, namely $S_i \leq I_i/(1-\delta)$. Focusing on the last system, we have

$$\Delta^D C^{\vec{R}} = (1-\delta) \Delta^D C^{\vec{R}_{n-1}, S_n} + \delta \Delta^D C^{\vec{R}_{n-1}, R_n'} \,.$$

Observe that fixing the last system to the δ-biased S_n yields an $(\varepsilon, k-1, n-1)$ δ-bias-resilient combiner with respect to $\vec{I}_{n-1}$. Similarly, fixing it to an arbitrary system (such as R_n') results in an $(\varepsilon, k, n-1)$ δ-bias-resilient combiner with respect to $\vec{I}_{n-1}$. This yields the recurrence relation:

$$\begin{aligned} &\Delta(k,n) \leq (1-\delta)\Delta(k-1,n-1) + \delta \Delta(k,n-1) \,, \\ &\Delta(0,n) \leq \varepsilon, \\ &\Delta(n,n) \leq (1-(1-\delta)^n) + (1-\delta)^n \varepsilon \,. \end{aligned}$$

For the base case $k = 0$, enough δ-biased systems are used, allowing us to invoke the combiner's guarantee. For $k = n$, with probability $(1-\delta)^n$ we use the shared δ-biased systems $\vec{S}$ and maintain ε-indistinguishability, and otherwise the distinguishing advantage could be 1.

Finally, we prove the bound $\Delta(k,n) \leq (1-\varepsilon) \cdot \Pr[\mathrm{Bin}(n,\delta) > n-k] + \varepsilon$ by induction. For the base cases, it is immediate. For the inductive step, observe that

$$\Pr[\mathrm{Bin}(n,\delta) > n-k] = (1-\delta) \Pr[\mathrm{Bin}(n-1,\delta) > n-k] + \delta \Pr[\mathrm{Bin}(n-1,\delta) > n-k-1] \,,$$

by conditioning on the first trial: with probability $1-\delta$ it is 0 and we still need $> n-k$ ones, and with probability δ it is 1 and we only need another $> n-k-1$. Plugging this in below, we get:

$$\begin{aligned}\Delta(k,n) &\leq (1-\delta)\Delta(k-1,n-1) + \delta\Delta(k,n-1)\\ &\leq (1-\delta)\left((1-\varepsilon)\Pr\left[\mathrm{Bin}(n-1,\delta) > n-k\right] + \varepsilon\right)\\ &\quad + \delta\left((1-\varepsilon)\Pr\left[\mathrm{Bin}(n-1,\delta) > n-k-1\right] + \varepsilon\right)\\ &= (1-\varepsilon)\Pr\left[\mathrm{Bin}(n,\delta) > n-k\right] + \varepsilon\,.\end{aligned}$$

□

Previously Known Amplifiers as Bias-Resilient. We observe that several known amplifiers from the literature are naturally bias-resilient. First, let us consider one of the most foundational examples: amplification of weak OWFs. Although the challenge is typically described as solving a puzzle, we can view it as an interactive distinguishing game. Both the real and ideal systems sample an initial state x. Then, upon receiving a query x_q, the real system returns $(f(x), \mathbb{I}[f(x) = f(x_q)])$ while the ideal system returns $(f(x), 0)$ regardless of the query. In every interaction where a valid preimage was not fed as one of the queries, there is zero distinguishing advantage.

The corresponding perfect (k,n) bias-resilient indistinguishability combiner $C_b^{\vec{S}}$ forwards queries $(x_{q,1},\ldots,x_{q,n})$ and masks its secret bit b by multiplying it by the indicator $b'_q := \mathbb{I}\left[\sum_{i=1}^{n} b_{q,i} \geq n-k+1\right]$. If at least k systems are ideal, then b'_q is always zero and b remains completely hidden. Biasing the ideal system (i.e., the initial state) does not change this fact. We remark that although our framework is statistical (where OWFs do not exist), it can still capture weak OWF amplification, or more generally privately verifiable puzzles with unlimited attempts, using the hardcore lemma for cc-stateless systems [Tes11]. For more details on computational generalizations and viewing the direct product OWF construction as an indistinguishability combiner, see [ABG26].

As another example, consider the concatenate-and-extract approach taken by Maurer and Tessaro for weak PRG amplification [MT10], which was also later used in the context of weak NIZK amplification [BG24]. Given n weakly pseudorandom strings $(X_1,\ldots,X_n)$, it outputs $\mathrm{Ext}(X_1,\ldots,X_n;r), r$ for a strong randomness extractor Ext (where r denotes the seed).[5] Here, we rely on the fact that any random variable X that is at most δ-biased with respect to the ideal uniform distribution U_m inherently has high min-entropy $H_\infty(X) \geq m - \log(1/(1-\delta))$.

Finally, as we explore in the following section, a rich class of bias-resilient combiners can be derived from locally leakage-resilient secret-sharing schemes (e.g., [DP07, BDIR18, SV19, AK25b]).

3.1 Application to Locally Leakage-Resilient Secret Sharing

In this section, we demonstrate how the security guarantees of locally leakage-resilient secret-sharing schemes can be directly leveraged to construct bias-resilient combiners. To focus purely on indistinguishability, we formalize sharing schemes simply as randomized algorithms mapping a secret to n shares, intentionally omitting reconstruction correctness as it is irrelevant to the security of the induced combiner.

Definition 9 (Leakage-Resilient Secret Sharing)**.** Let $D : \mathcal{S} \to \mathcal{X}^n$ be a randomized sharing scheme. We say it is t-out-of-n ε-private ℓ-leakage-resilient if for any two secrets $s^0, s^1 \in \mathcal{S}$, any leakage functions $f_1,\ldots,f_n : \mathcal{X} \to \{0,1\}^\ell$ and any subset $T \subseteq [n]$ of size $|T| < t$, we have:

$$\left(x_i^0\right)_{i\in T}, \left(f_i(x_i^0)\right)_{i\notin T} \approx_\varepsilon \left(x_i^1\right)_{i\in T}, \left(f_i(x_i^1)\right)_{i\notin T},$$

[5] In fact, it is possible to use deterministic extractors for independent sources.

over sharings $\forall b \in \{0,1\} : (x_1^b, \ldots, x_n^b) \leftarrow D(s^b)$.

While the standard formulation above bounds the absolute length of the leakage, an alternative approach is to bound the missing min-entropy of a randomized leakage function. Conceptually, imagine there is some fixed universal distribution sampled using an m-bit uniform randomness tape. Given a share x, the randomized leakage function $F(x)$ outputs a biased choice for this m-bit tape in order to reveal information about x. The condition $H_\infty(F(x)) \geq m - \ell$ ensures that any specific input x can only reduce the min-entropy of the uniform m-bit tape by at most ℓ bits.

Definition 10 (Bias-Resilient Secret Sharing)**.** Let $D : \mathcal{S} \to \mathcal{X}^n$ be a randomized sharing scheme. We say it is t-out-of-n ε-private ℓ-bias-resilient if for any two secrets $s^0, s^1 \in \mathcal{S}$, any m and randomized functions $F_1, \ldots, F_n : \mathcal{X} \to \{0,1\}^m$ where $H_\infty(F_i(x)) \geq m - \ell$ for all $i \in [n]$ and $x \in \mathcal{X}$, and any subset $T \subseteq [n]$ of size $|T| < t$, we have:

$$\left(x_i^0\right)_{i\in T}, \left(F_i(x_i^0)\right)_{i\notin T} \approx_\varepsilon \left(x_i^1\right)_{i\in T}, \left(F_i(x_i^1)\right)_{i\notin T},$$

over sharings $\forall b \in \{0,1\} : (x_1^b, \ldots, x_n^b) \leftarrow D(s^b)$.

Remark 2. An alternative definition requires that for some random variables $Q_1, \ldots, Q_n$ we have $D_\infty(F_i(x) \| Q_i) \leq \ell$ for all $i \in [n]$ and $x \in \mathcal{X}$, for the Infinite Rényi Divergence

$$D_\infty(P\|Q) := \max_{x\in\mathcal{X}} \log \frac{P(x)}{Q(x)}.$$

In the current definition we force Q_i to be uniform (over arbitrarily many bits m), which is equivalent as we may approximate any distribution using uniform bits.

Observation 1. Let $D : \mathcal{S} \to \mathcal{X}^n$ be a randomized sharing scheme. If D is t-out-of-n ε-private ℓ-bias-resilient, then it is also t-out-of-n ε-private ℓ-leakage-resilient.

Proof sketch. We may simply pick $m = \ell$ and let $F_i(x) \equiv f_i(x)$, which trivially has min-entropy ≥ 0. (In fact, even for larger m we may define $F_i(x) \equiv (f_i(x), U^{m-\ell})$.) □

The next lemma is based on an idea from [GP91], transforming a simulator that works well over a dense subset into one that succeeds given a small amount of auxiliary leakage.

Lemma 2. *Let $D : \mathcal{S} \to \mathcal{X}^n$ be a randomized sharing scheme. If D is t-out-of-n ε-private ℓ'-leakage-resilient then it is also t-out-of-n 3ε-private ℓ-bias-resilient, for $\ell' := \ell + \log\ln(n/\varepsilon)$.*

Proof. Let $s^0, s^1 \in \mathcal{S}$, $F_1, \ldots, F_n : \mathcal{X} \to \{0,1\}^m$ be randomized functions with $H_\infty(F_i(x)) \geq m - \ell$ for all i, x, and $T \subseteq [n]$ of size $|T| < t$ be given such that for some distinguisher A we have:

$$\Pr\left[A\left(\left(x_i^0\right)_{i\in T}, \left(F_i(x_i^0)\right)_{i\notin T}\right) = 1\right] - \Pr\left[A\left(\left(x_i^1\right)_{i\in T}, \left(F_i(x_i^1)\right)_{i\notin T}\right) = 1\right] > 3\varepsilon,$$

over sharings $\forall b \in \{0,1\} : (x_1^b, \ldots, x_n^b) \leftarrow D(s^b)$. Without loss of generality, we may focus on the setting where each $F_i(x)$ is uniformly distributed over a subset $S_{i,x}$ of size $2^{m-\ell}$, since any RV with min-entropy $\geq m - \ell$ can be expressed as a convex combination of flat distributions over exactly $2^{m-\ell}$ elements. In addition, we will initially construct *randomized* leakage functions $f_1, \ldots, f_n : \mathcal{X} \to \{0,1\}^{\ell'}$, and that is also without loss of generality as we may eventually fix the best randomness using an averaging argument.

In order to simulate $F_i(x)$ we use rejection sampling: the randomized function $f_i : \mathcal{X} \to \{0,1\}^{\ell'}$ samples $2^{\ell'}$ elements $z_1, \ldots, z_{2^{\ell'}} \leftarrow \{0,1\}^m$ independently at random and outputs the first index

$j \in [2^{\ell'}]$ such that $z_j \in S_{i,x}$, or $j = 1$ if it failed and no such index exists. For any shares $(x_1, \ldots, x_n)$, invoking the union bound, the probability that some function failed is at most

$$n \cdot \left(1 - \frac{2^{m-\ell}}{2^m}\right)^{2^{\ell'}} \leq n \cdot \exp\left(-\frac{2^{\ell'}}{2^\ell}\right) = n \cdot \exp\left(-\ln(n/\varepsilon)\right) = \varepsilon\,,$$

and conditioned on no leakage function failing the following are identically distributed

$$(x_i)_{i\in T}\,,\,(F_i(x_i))_{i\notin T} \equiv (x_i)_{i\in T}\,,\,\left(z_{f_i(x_i)}\right)_{i\notin T}\,.$$

We conclude that the following distributions are $> \varepsilon$ distinguishable

$$\begin{aligned}
&\Pr\left[A\left(\left(x_i^0\right)_{i\in T}\,,\,\left(z_{f_i(x_i^0)}\right)_{i\notin T}\right)=1\right] - \Pr\left[A\left(\left(x_i^1\right)_{i\in T}\,,\,\left(z_{f_i(x_i^1)}\right)_{i\notin T}\right)=1\right] \\
&\geq \left(\Pr\left[A\left(\left(x_i^0\right)_{i\in T}\,,\,\left(F_i(x_i^0)\right)_{i\notin T}\right)=1\right] - \varepsilon\right) \\
&- \left(\Pr\left[A\left(\left(x_i^1\right)_{i\in T}\,,\,\left(F_i(x_i^1)\right)_{i\notin T}\right)=1\right] + \varepsilon\right) > 3\varepsilon - 2\varepsilon = \varepsilon\,.
\end{aligned}$$

Finally, once we fix the best randomness of the leakage functions $f_1, \ldots, f_n : \mathcal{X} \to \{0,1\}^{\ell'}$, the element $z_{f_i(x)}$ is completely determined by $f_i(x)$ and we arrive at a contradiction with

$$\left(x_i^0\right)_{i\in T}\,,\,\left(f_i(x_i^0)\right)_{i\notin T} \approx_\varepsilon \left(x_i^1\right)_{i\in T}\,,\,\left(f_i(x_i^1)\right)_{i\notin T}\,.$$

□

Lemma 3. *Let $D : \mathcal{S} \to \mathcal{X}^n$ be a t-out-of-n ε-private ℓ-bias-resilient randomized sharing scheme. Let $s_0, s_1 \in \mathcal{S}$ be any two secrets, and consider the n-oracle-aided algorithm $C_b^{(\cdot)}$ that shares $(x_1, \ldots, x_n) \leftarrow D(s_b)$ and outputs $(F_1(x_1), \ldots, F_n(x_n))$. Then $(C_0^{(\cdot)}, C_1^{(\cdot)})$ is an $(\varepsilon,\, k = n-t+1,\, n)$ $(\delta = 1 - 2^{-\ell})$-bias-resilient statistical indistinguishability combiner with respect to any input-independent ideal functionalities $(F_1^*, \ldots, F_n^*)$ (i.e., random variables).*

Proof. Let $(F_1, \ldots, F_n)$ be any functions such that $G := \left\{i \in [n] \mid F_i \leq F_i^*/2^{-\ell}\right\}$ is of size $|G| \geq k$. Also, let $Q_i \equiv F_i^*$ (on any input, as it is input-independent). Note that $\left|\overline{G}\right| < t$ and $\forall i \in G\ \forall x \in \mathcal{X}:\ D_\infty(F_i(x) \| Q_i) \leq \ell$, hence we may use the t-out-of-n ε-privacy ℓ-bias-resilience of D to argue that

$$\left(x_i^0\right)_{i\notin G}\,,\,\left(F_i(x_i^0)\right)_{i\in G} \approx_\varepsilon \left(x_i^1\right)_{i\notin G}\,,\,\left(F_i(x_i^1)\right)_{i\in G}\,,$$

over sharings $\forall b \in \{0,1\}:\ (x_1^b, \ldots, x_n^b) \leftarrow D(s^b)$. In particular, applying F_i on x_i for $i \notin G$ does not increase the statistical distance, hence

$$C_0^{\vec{F}} \approx_\varepsilon C_1^{\vec{F}}\,.$$

□

It is well known that if $H_\infty(X) \geq n$ and $f : \mathcal{X} \to \{0,1\}^\ell$, then $\widetilde{H}_\infty(X \mid f(X)) \geq n - \ell$. It is worth noting that the same implication also holds for our bounded-bias notion:

Observation 2. Let X be a random variable over $\mathcal{X}$ with $H_\infty(X) \geq n$, and $F : \mathcal{X} \to \{0,1\}^m$ be a randomized function such that $\forall x \in \mathcal{X} : H_\infty(F(x)) \geq m - \ell$. Then $\widetilde{H}_\infty(X \mid F(X)) \geq n - \ell$.

Proof. To see this, observe that

$$\begin{aligned}
2^{-\widetilde{H}_\infty(X|F(X))} &= \mathop{\mathbb{E}}_{z\leftarrow F(X)}\left[\max_x \Pr\left[X=x \mid F(X)=z\right]\right] \\
&= \sum_{z\in\{0,1\}^m} \Pr\left[F(X)=z\right]\cdot \max_x \Pr\left[X=x \mid F(X)=z\right] \\
&= \sum_{z\in\{0,1\}^m} \max_x \Pr\left[X=x \cap F(X)=z\right] \\
&= \sum_{z\in\{0,1\}^m} \max_x \left(\Pr\left[X=x\right]\cdot \Pr\left[F(X)=z \mid X=x\right]\right) \\
&\le 2^{\ell-m}\cdot \sum_{z\in\{0,1\}^m} \max_x \Pr\left[X=x\right] \le 2^{\ell-m}\cdot 2^m\cdot 2^{-n} = 2^{\ell-n}\,.
\end{aligned}$$

□

3.2 Single-Bit Leakage-Resilient Indistinguishability Combiners

Another natural approach to formalize the all-or-nothing intuition and bypass the $\delta < 0.5$ barrier is to use leakage-resilient combiners. Specifically, we consider resilience against single-bit leakage from each interaction of the combiner with its oracle systems. Note that this notion rules out the canonical bad example of addition over characteristic-2 finite fields, since that combiner is not resilient against LSB leakage.

Let us start with an example to set the limitation and baseline for the parameters.

Example 3. Consider the $((1-\delta)^k, k, n)$ indistinguishability combiner

$$C_b(b_1,\ldots,b_n) := b\cdot\prod_{i=1}^n b_i\,,$$

with the ideal $I_i \equiv \mathrm{Ber}(1-\delta)$ and real $R_i \equiv 1$. Then $C_0^{\vec{R}}$ and $C_1^{\vec{R}}$ are fully distinguishable, while in the k-ideal setting, the probability that the product evaluates to 1 is at most $(1-\delta)^k$. Observe that even revealing all inputs entirely $(b_1,\ldots,b_n)$ does not increase the distinguishing advantage beyond what the output already provides, hence the combiner is fully resilient to leakage.

Definition 11 (Single-Bit Leakage-Resilient Indistinguishability Combiner)**.** For finite sets $(\mathcal{X},\mathcal{Y})$, $\varepsilon \in [0,1]$, $k,n,s \in \mathbb{N}$ with $k \le n$, an (ε,k,n) single-bit leakage-resilient indistinguishability combiner with respect to ideal $(\mathcal{X},\mathcal{Y})$-systems $\vec{I} = (I_1,\ldots,I_n)$ against a $(\mathcal{Y},\mathcal{X})$-distinguisher class $(\mathcal{D},s)$, is a pair of n-oracle-aided $(\mathcal{Y},\mathcal{X})$-systems $(C_0^{(\cdot)},C_1^{(\cdot)})$ such that for all $(\mathcal{X},\mathcal{Y})$-systems $\vec{S} = (S_1,\ldots,S_n)$ and all leakage functions $\vec{f} = (f_1,\ldots,f_n)$ over $(\mathcal{X}\times\mathcal{Y})^* \to \{0,1\}$:

$$|\{i\in[n] \mid S_i \equiv I_i\}| \ge k \implies \forall D\in\mathcal{D}:\ \left(D(C_0^{\vec{S}}),(f_i(\tau_i))_{i\in[n]}\right) \overset{\mathsf{c}}{\approx}_{s,\varepsilon} \left(D(C_1^{\vec{S}}),(f_i(\tau_i))_{i\in[n]}\right)\,,$$

where τ_i denotes the transcript of S_i during the respective execution.

Theorem 2 (Leakage-Resilient Indistinguishability Combiners Are Almost Optimal Amplifiers)**.** *For $\varepsilon,\delta\in[0,1]$, $k,n\in\mathbb{N}$, given an (ε,k,n) single-bit leakage-resilient indistinguishability combiner $(C_0^{(\cdot)},C_1^{(\cdot)})$ with respect to ideal systems $\vec{I} = (I_1,\ldots,I_n)$ against a distinguisher class $(\mathcal{D},2n)$, let $\vec{R} = (R_1,\ldots,R_n)$ be a tuple of n real systems such that $\forall i\in[n]: R_i \approx_\delta I_i$. Then, we have*

$$\forall D\in\mathcal{D}:\ \left|\mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}})\right]\right| \le \delta^{(k,n)} + (1-\delta^{(k,n)})\cdot\varepsilon(1-\delta)^{-k}\,.$$

Remark 3. Note that for $\varepsilon \geq (1-\delta)^k$ nothing can be guaranteed, as demonstrated in Example 3. Otherwise, the bound is guaranteed to be nontrivial, i.e., smaller than 1 for $\delta < 1$. In general, we may want to ignore the factor $(1-\delta^{(k,n)})$ for simplicity and use the looser bound $\delta^{(k,n)} + \varepsilon(1-\delta)^{-k}$.

Proof. By the system decomposition lemma (Lemma 1), for all $i \in [n]$, there exist systems S_i, R'_i, I'_i such that

$$R_i \equiv (1-\delta)S_i + \delta R'_i\,,$$
$$I_i \equiv (1-\delta)S_i + \delta I'_i\,.$$

Let $\vec{B} = (B_1, \ldots, B_n) \in \{0,1\}^n$ be independent random variables with $\forall i \in [n]:\ B_i \equiv \mathrm{Ber}(1-\delta)$. For $i \in [n], b \in \{0,1\}$, define the system $T_{i,b} \equiv S_i$ if $b = 1$, and otherwise R'_i. Then $R_i \equiv T_{i,B_i}$ and $\vec{R} \equiv \vec{T}_{\vec{B}}$.

Assume toward contradiction that for some distinguisher $D \in \mathcal{D}$, letting $\varepsilon' := \varepsilon(1-\delta)^{-k}$, we have

$$\left|\mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}})\right]\right| > \delta^{(k,n)} + (1-\delta^{(k,n)}) \cdot \varepsilon'\,.$$

Observe that

$$\begin{aligned}
\left|\mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}})\right]\right| &= \left|\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{B}}}) - D(C_1^{\vec{T}_{\vec{B}}})\right]\right| \\
&= \left|\mathop{\mathbb{E}}_{\vec{b}\leftarrow\vec{B}}\left[\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}}}) - D(C_1^{\vec{T}_{\vec{b}}})\right]\right]\right| \\
&\leq \mathop{\mathbb{E}}_{\vec{b}\leftarrow\vec{B}}\left[\left|\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}}}) - D(C_1^{\vec{T}_{\vec{b}}})\right]\right|\right] \\
&= \Pr\left[\left|\vec{b}\right| < k \mid \vec{b} \leftarrow \vec{B}\right] \cdot \mathop{\mathbb{E}}_{\vec{b}\leftarrow\vec{B}_{<k}}\left[\left|\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}}}) - D(C_1^{\vec{T}_{\vec{b}}})\right]\right|\right] \\
&+ \Pr\left[\left|\vec{b}\right| \geq k \mid \vec{b} \leftarrow \vec{B}\right] \cdot \mathop{\mathbb{E}}_{\vec{b}\leftarrow\vec{B}_{\geq k}}\left[\left|\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}}}) - D(C_1^{\vec{T}_{\vec{b}}})\right]\right|\right] \\
&\leq \delta^{(k,n)} \cdot 1 + (1-\delta^{(k,n)}) \cdot \mathop{\mathbb{E}}_{\vec{b}\leftarrow\vec{B}_{\geq k}}\left[\left|\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}}}) - D(C_1^{\vec{T}_{\vec{b}}})\right]\right|\right]\,.
\end{aligned}$$

Thus, according to our assumption, there must exist some fixed $\vec{b}^* = (b_1^*, \ldots, b_n^*)$ with $\sum_{i=1}^n b_i^* \geq k$, such that

$$\left|\mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}^*}}) - D(C_1^{\vec{T}_{\vec{b}^*}})\right]\right| > \varepsilon'\,.$$

In particular, there exists some subset $V \subseteq [n]$ of size $|V| = k$, such that $\forall i \in V:\ b_i^* = 1$.

Next, for $i \in V, b \in \{0,1\}$, define the system $T'_{i,b} \equiv S_i$ if $b = 1$, and otherwise I'_i. If $i \notin V$ then $T'_{i,b} \equiv T_{i,b}$. Let $\vec{B}' = (B'_1, \ldots, B'_n) \in \{0,1\}^n$ be independent random variables with $\forall i \in V:\ B'_i \equiv \mathrm{Ber}(1-\delta)$, and $B'_i \equiv b_i^*$ otherwise. Then, $\forall i \in V:\ T'_{i,B'_i} \equiv I_i$, and hence the vector $\vec{T}'_{\vec{B}'}$ is k-ideal. Furthermore, we observe that $\vec{T}'_{\vec{b}^*} \equiv \vec{T}_{\vec{b}^*}$ since $\forall i \in V:\ b_i^* = 1$ and $T'_{i,1} \equiv S_i \equiv T_{i,1}$.

We construct a distinguisher against $\vec{T}'_{\vec{B}'}$ by focusing on the event that $\vec{B}' = \vec{b}^*$, occurring with probability $(1-\delta)^k$, in which case we answer according to D. In order to maintain the advantage $(1-\delta)^k \cdot \varepsilon' = \varepsilon$ and not reduce it by guessing incorrectly, on the complement event we aim to ignore D and output 0 (or any independent bit). This is where we use the single-bit leakage, to identify

whether or not we should use D: for $i \in V$, define $f_i(\tau_i)$ according to the conditional distribution $\left(B'_i \mid \mathrm{tr}(T'_{i,B'_i}) = \tau_i\right) \equiv \mathrm{Ber}\left((1-\delta)\cdot S_i(\tau_i)/I_i(\tau_i)\right)$.[6] For $i \notin V$ we may simply define $f_i \equiv b^*_i$. Then, for any n-oracle-aided distinguisher and in particular for $D(C_b^{(\cdot)})$ with any $b \in \{0,1\}$, we have

$$\left(D(C_b^{\vec{T}'_{\vec{B}'}}), (f_i(\tau_i))_{i\in[n]}\right) \equiv \left(D(C_b^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right) .$$

Finally, we construct a distinguisher D' between

$$\left(D(C_0^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right) \quad \text{and} \quad \left(D(C_1^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right) .$$

Denote its input by $(b_D, b'_1, \ldots, b'_n) \in \{0,1\}^{n+1}$, then we output b_D only if $\vec{b'} = \vec{b}^*$ and otherwise 0. More explicitly, we output $b_D \cdot \prod_{i=1}^{n}(1 \oplus b'_i \oplus b^*_i)$, requiring a size $2n$ circuit of depth $\lceil \log(n+1) \rceil + 1$. The resulting advantage is

$$\begin{aligned}
&\mathbb{E}\left[D'\left(D(C_0^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right) - D'\left(D(C_1^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right)\right] \\
&= \Pr\left[\vec{B}' \neq \vec{b}^*\right] \cdot \mathop{\mathbb{E}}_{\vec{b'} \leftarrow \vec{B}' \setminus \vec{b}^*}\left[\mathbb{E}\left[D'\left(D(C_0^{\vec{T}'_{\vec{b'}}}), \vec{b'}\right) - D'\left(D(C_1^{\vec{T}'_{\vec{b'}}}), \vec{b'}\right)\right]\right] \\
&+ \Pr\left[\vec{B}' = \vec{b}^*\right] \cdot \mathbb{E}\left[D'\left(D(C_0^{\vec{T}'_{\vec{b}^*}}), \vec{b}^*\right) - D'\left(D(C_1^{\vec{T}'_{\vec{b}^*}}), \vec{b}^*\right)\right] \\
&= 0 + (1-\delta)^k \cdot \mathbb{E}\left[D(C_0^{\vec{T}_{\vec{b}^*}}) - D(C_1^{\vec{T}_{\vec{b}^*}})\right] ,
\end{aligned}$$

yielding the desired contradiction

$$\left|\mathbb{E}\left[D'\left(D(C_0^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right) - D'\left(D(C_1^{\vec{T}'_{\vec{B}'}}), \vec{B}'\right)\right]\right| > (1-\delta)^k \cdot \varepsilon' = \varepsilon .$$

□

4 Prime Fields XOR Lemma

In this section, we prove a generalized XOR lemma over prime fields. As discussed in the introduction, addition over characteristic-2 fields inherently fails to amplify distance to uniform, if the initial weak indistinguishability is at least 1/2. We overcome this by showing that summation over a prime field $\mathbb{F}_p$ naturally acts as a bias-resilient combiner, allowing us to ensure the sum becomes indistinguishable from uniform even if the initial indistinguishability is very weak, as long as it is non-trivial.

We begin in Section 4.1 with a brief overview of Fourier analysis, which is our main tool to show a statistical sum extraction lemma over prime fields later in that section. In Section 4.2, we show that the sum-extraction theorem naturally induces a bias-resilient combiner. Finally, in Section 4.3, we invoke our main amplification theorem to derive the statistical generalized XOR lemma over prime fields, and then prove that it generalizes to the computational setting.

[6] Using randomized boolean leakage functions is without loss of generality, as we may later fix the best choices.

4.1 The Statistical Sum Extraction Theorem

4.1.1 Fourier Analysis over Prime Fields

We begin with a brief overview of Fourier analysis over prime fields. For a more detailed exposition of these techniques in the context of additive combinatorics, we refer the reader to [TV07, Gre07]. Our analysis builds on and extends [BDIR18].

Let p be a prime number. We consider random variables distributed over the prime field $\mathbb{F}_p$. Let $\omega := \exp(2\pi i/p)$ denote the primitive p-th root of unity. For a random variable X over $\mathbb{F}_p$ with probability distribution $P(x) := \Pr[X = x]$, the Fourier coefficient at frequency $\alpha \in \mathbb{F}_p$ is defined as the expected value of the complex character

$$\hat{X}(\alpha) := \mathbb{E}\left[\omega^{\alpha X}\right] = \sum_{x \in \mathbb{F}_p} P(x)\omega^{\alpha x}\,.$$

Because $P(x)$ is a valid probability distribution, the coefficient at the trivial frequency $\alpha = 0$ is always $\hat{X}(0) = \sum_{x\in\mathbb{F}_p} P(x) = 1$. The magnitudes of the non-trivial Fourier coefficients ($\alpha \neq 0$) measure how far the distribution is from uniform. These definitions extend naturally to conditional distributions: for a joint distribution (X, Z), we denote the Fourier coefficient of X conditioned on a fixed leakage $Z = z$ using $\hat{X}_{|Z=z}(\alpha)$.

Our analysis relies on three standard properties of Fourier transforms over finite abelian groups:

Convolution of Independent Variables. The Fourier transform maps the addition of independent random variables to point-wise multiplication. If $X_1, \ldots, X_n$ are mutually independent random variables, then the Fourier coefficients of their sum $S := \sum_{i=1}^{n} X_i \pmod p$ satisfy:

$$\hat{S}(\alpha) = \prod_{i=1}^{n} \hat{X}_i(\alpha)\,.$$

Parseval's Identity. The total squared magnitude of the Fourier coefficients is directly proportional to the collision probability of the distribution, denoted $\mathrm{CP}(X) := \sum_{x\in\mathbb{F}_p} P(x)^2$. Parseval's identity states that $\frac{1}{p}\sum_{\alpha\in\mathbb{F}_p} \left|\hat{X}(\alpha)\right|^2 = \mathrm{CP}(X)$, because

$$\begin{aligned}\sum_{\alpha\in\mathbb{F}_p} \left|\hat{X}(\alpha)\right|^2 &= \sum_{\alpha\in\mathbb{F}_p} \hat{X}(\alpha)\overline{\hat{X}(\alpha)} = \sum_{\alpha\in\mathbb{F}_p}\sum_{x_0\in\mathbb{F}_p}\sum_{x_1\in\mathbb{F}_p} P(x_0)P(x_1)\omega^{\alpha(x_0-x_1)} \\ &= \sum_{x_0\in\mathbb{F}_p}\sum_{x_1\in\mathbb{F}_p} P(x_0)P(x_1)\sum_{\alpha\in\mathbb{F}_p}\omega^{\alpha(x_0-x_1)} = \sum_{x\in\mathbb{F}_p} P(x)^2\cdot p\,.\end{aligned}$$

By separating the trivial frequency $\hat{X}(0) = 1$, the sum of the non-trivial squared magnitudes is exactly

$$\sum_{\alpha\neq 0}\left|\hat{X}(\alpha)\right|^2 = p\cdot\mathrm{CP}(X) - 1\,.$$

We remark that $\mathrm{CP}(X) \leq 2^{-H_\infty(X)}$ since $\sum_{x\in\mathbb{F}_p} P(x)^2 \leq \max_x P(x)\sum_{x\in\mathbb{F}_p} P(x) = 2^{-H_\infty(X)}$.

Distance from Uniform. The statistical distance between a distribution X and the uniform distribution U_p over $\mathbb{F}_p$ can be bounded by its non-trivial Fourier coefficients. Applying the Cauchy-Schwarz inequality ($L_1 \le \sqrt{p \cdot L_2^2}$) to the definition of statistical distance, and substituting Parseval's identity, yields the following standard bound:

$$\begin{aligned}\mathrm{SD}(X, U_p) &= \frac{1}{2}\sum_{x\in\mathbb{F}_p}\left|P(x) - \frac{1}{p}\right| \le \frac{1}{2}\sqrt{p\sum_{x\in\mathbb{F}_p}\left(P(x) - \frac{1}{p}\right)^2} \\ &= \frac{1}{2}\sqrt{p\sum_{x\in\mathbb{F}_p}\left(P(x)^2 - \frac{2P(x)}{p} + \frac{1}{p^2}\right)} \\ &= \frac{1}{2}\sqrt{p\,\mathrm{CP}(X) - 2 + 1} = \frac{1}{2}\sqrt{\sum_{\alpha\neq 0}\left|\hat{X}(\alpha)\right|^2}\,.\end{aligned}$$

4.1.2 Useful Lemmas

We defer to Section A the proof of the following propositions:

Proposition 1. Let $f(y) := \left(\frac{y\sin(\pi/y)}{p\sin(\pi/p)}\right)^2$ for some real number $p \ge 2$. Then for any random variable Y taking values in $[1, p]$, we have:

$$\mathbb{E}\left[f(Y)\right] \le f\big(\max(2, \mathbb{E}\left[Y\right])\big)\,.$$

Proposition 2. For real numbers $\ell, p \in \mathbb{R}^+$ such that $2^\ell < p$ we have

$$c_\ell := \frac{2^\ell \sin(\pi 2^{-\ell})}{p\sin(\pi/p)} < \exp\left(-\left(\frac{1}{2^{2\ell}} - \frac{1}{p^2}\right)\right)\,.$$

We now bound the non-trivial Fourier coefficients of random variables over prime fields, depending on their average min-entropy.

Lemma 4 (Generalizing [BDIR18, Lemma 3.11])**.** *Let (X, Z) be a pair of random variables distributed over $\mathbb{F}_p \times \mathcal{Z}$ for a prime p and finite set $\mathcal{Z}$. If $\widetilde{H}_\infty(X|Z) \ge \log(p) - \ell$ for some real number $\ell \ge 1$, then for every non-trivial frequency $\alpha \in \mathbb{F}_p \setminus \{0\}$, we have*

$$\mathop{\mathbb{E}}_{z\leftarrow Z}\left[\left|\hat{X}_{|Z=z}(\alpha)\right|^2\right] \le c_\ell^2\,,$$

where $c_\ell := \frac{2^\ell \sin(\pi 2^{-\ell})}{p\sin(\pi/p)}$.

Proof. For a fixed leakage observation z, let $P_z(x) := \Pr_{X|Z=z}\left[X = x\right]$, and let $q_z := \max_x P_z(x)$ denote the maximum guessing probability. Because the map $x \mapsto \alpha x \pmod p$ is a permutation for any $\alpha \neq 0$ and the constraint $P_z(x) \le q_z$ is permutation-invariant, we may assume without loss of generality that $\alpha = 1$.

The problem reduces to bounding the following optimization problem: maximize $\left|\sum_{x\in\mathbb{F}_p} P_z(x)\omega^x\right|$ subject to the constraints $\sum_x P_z(x) = 1$ and $0 \le P_z(x) \le q_z$ for all x.

By a standard extremal argument (we refer the reader to [BDIR18, Lemma 3.11] for further elaboration), the magnitude is maximized when the probability mass is greedily packed into a contiguous interval. The geometric intuition is that, in order to maximize the magnitude of a sum

of complex vectors, the vectors must point as closely to the same direction as possible. Because our vectors are of the form ω^x, minimizing the angle between them corresponds to selecting a sequence of adjacent exponents.

Because global phase shifts do not alter the magnitude of the sum, we may assume without loss of generality that this contiguous interval starts at 0. Let $K = \lfloor 1/q_z \rfloor$ be the number of elements assigned full mass q_z, and let $\delta = 1/q_z - K \in [0, 1)$ be the fractional remainder. Assigning weight q_z to $0, \ldots, K-1$ and δq_z to K, and using the identity $|1 - \omega^x| = 2\sin(\pi x/p)$, we obtain:

$$\begin{aligned}
\left|\hat{X}_{|Z=z}(1)\right| &\leq \left| q_z \sum_{x=0}^{K-1} \omega^x + \delta q_z \omega^K \right| = \left| q_z \frac{1-\omega^K}{1-\omega} + \delta q_z \omega^K \right| \\
&= \frac{q_z}{|1-\omega|} \left| 1 - \omega^K + \delta(1-\omega)\omega^K \right| \\
&= \frac{q_z}{2\sin(\pi/p)} \left| (1-\delta)(1-\omega^K) + \delta(1-\omega^{K+1}) \right| \\
&\leq \frac{q_z}{2\sin(\pi/p)} \left((1-\delta)\left|1-\omega^K\right| + \delta\left|1-\omega^{K+1}\right| \right) \\
&= \frac{q_z}{\sin(\pi/p)} \left((1-\delta)\sin\left(\frac{\pi K}{p}\right) + \delta \sin\left(\frac{\pi(K+1)}{p}\right) \right) \\
&\leq \frac{q_z}{\sin(\pi/p)} \sin\left(\frac{\pi(K+\delta)}{p}\right),
\end{aligned}$$

where in the last inequality we used that the function $g(x) = \sin(\pi x/p)$ is concave on the interval $[0, p]$. Substituting the definition $K + \delta = 1/q_z$ and multiplying the numerator and denominator by p, we conclude that

$$\left|\hat{X}_{|Z=z}(1)\right| \leq \frac{q_z p \sin(\pi/(q_z p))}{p\sin(\pi/p)}.$$

We now bound the expected squared magnitude $\mathbb{E}_{z\leftarrow Z}\left[\left|\hat{X}_{|Z=z}(1)\right|^2\right]$. Let $y_z := q_z p$ be the normalized maximum guessing probability, observing that $y_z \in [1, p]$ since $1/p \leq q_z \leq 1$. Note that our assumption $\widetilde{H}_\infty(X|Z) \geq \log(p) - \ell$ is equivalent to $\mathbb{E}_{z\leftarrow Z}[y_z] \leq 2^\ell$ by the definition of average min-entropy. By the bound derived above, $\left|\hat{X}_{|Z=z}(1)\right|^2 \leq f(y_z)$, where $f(y) := \left(\frac{y\sin(\pi/y)}{p\sin(\pi/p)}\right)^2$. By Proposition 1 (extending Jensen's inequality), and noting that $f(y)$ is increasing for $y \geq 2$,[7] we obtain:

$$\mathop{\mathbb{E}}_{z\leftarrow Z}\left[\left|\hat{X}_{|Z=z}(1)\right|^2\right] \leq \mathop{\mathbb{E}}_{z\leftarrow Z}[f(y_z)] \leq f\big(\max(2, \mathop{\mathbb{E}}_{z\leftarrow Z}[y_z])\big) \leq f\big(\max(2, 2^\ell)\big).$$

Because our assumption $\ell \geq 1$ guarantees $2^\ell \geq 2$, this simplifies to $f(2^\ell)$, yielding:

$$\mathop{\mathbb{E}}_{z\leftarrow Z}\left[\left|\hat{X}_{|Z=z}(1)\right|^2\right] \leq f(2^\ell) = \left(\frac{2^\ell \sin(\pi 2^{-\ell})}{p\sin(\pi/p)}\right)^2 = c_\ell^2,$$

which concludes the proof. □

4.1.3 The Extraction Theorem

Next, we show that the sum modulo p of independent random variables with weak average min-entropy becomes statistically close to uniform.

[7] See the proof of Proposition 1 in Section A for the full analysis of f.

Theorem 3 (Sum Extraction over Prime Fields)**.** *Let* $(X_1, Z_1), \ldots, (X_n, Z_n)$ *be mutually independent pairs of random variables distributed over* $\mathbb{F}_p \times \mathcal{Z}$ *for a prime* p *and finite set* $\mathcal{Z}$*. Let* $S := \sum_{i=1}^{n} X_i \pmod p$ *and* $\vec{Z} := (Z_1, \ldots, Z_n)$*. If for some real number* $\ell \geq 1$ *we have* $\forall i \in [n] : \widetilde{H}_\infty(X_i|Z_i) \geq \log(p) - \ell$*, then*

$$(S, \vec{Z}) \approx_\varepsilon (U_p, \vec{Z}) \,,$$

for

$$\varepsilon := \frac{1}{2} 2^{\ell/2} c_\ell^{n-1} \quad \textit{where} \quad c_\ell := \frac{2^\ell \sin(\pi 2^{-\ell})}{p \sin(\pi/p)} \,.$$

Proof. Because the uniform distribution U_p is independent of the leakage $\vec{Z}$, we express the statistical distance given the joint leakage as an expectation over $\vec{z} \leftarrow \vec{Z}$ of the conditional distance to uniform. The following sequence of inequalities, justified below, establishes the core bound:

$$\begin{aligned}
\mathrm{SD}((S, \vec{Z}), (U_p, \vec{Z})) &= \mathop{\mathbb{E}}_{\vec{z} \leftarrow \vec{Z}} \left[\mathrm{SD}(S_{|\vec{Z}=\vec{z}}, U_p) \right] && (1) \\
&\leq \mathop{\mathbb{E}}_{\vec{z} \leftarrow \vec{Z}} \left[\frac{1}{2} \sqrt{\sum_{\alpha \neq 0} \left| \hat{S}_{|\vec{Z}=\vec{z}}(\alpha) \right|^2} \right] && (2) \\
&\leq \frac{1}{2} \sqrt{\sum_{\alpha \neq 0} \mathop{\mathbb{E}}_{\vec{z} \leftarrow \vec{Z}} \left[\left| \hat{S}_{|\vec{Z}=\vec{z}}(\alpha) \right|^2 \right]} && (3) \\
&= \frac{1}{2} \sqrt{\sum_{\alpha \neq 0} \mathop{\mathbb{E}}_{\vec{z} \leftarrow \vec{Z}} \left[\prod_{i=1}^{n} \left| \hat{X}_{i|Z_i=z_i}(\alpha) \right|^2 \right]} && (4) \\
&= \frac{1}{2} \sqrt{\sum_{\alpha \neq 0} \prod_{i=1}^{n} \mathop{\mathbb{E}}_{z_i \leftarrow Z_i} \left[\left| \hat{X}_{i|Z_i=z_i}(\alpha) \right|^2 \right]} && (5) \\
&\leq \frac{1}{2} \sqrt{(c_\ell^2)^{n-1} \sum_{\alpha \neq 0} \mathop{\mathbb{E}}_{z_n \leftarrow Z_n} \left[\left| \hat{X}_{n|Z_n=z_n}(\alpha) \right|^2 \right]} && (6) \\
&= \frac{1}{2} c_\ell^{n-1} \sqrt{\sum_{\alpha \neq 0} \mathop{\mathbb{E}}_{z_n \leftarrow Z_n} \left[\left| \hat{X}_{n|Z_n=z_n}(\alpha) \right|^2 \right]} \,, && (7)
\end{aligned}$$

where Eq. (2) follows from Cauchy-Schwarz and Parseval's identity, Eq. (3) applies Jensen's inequality to the concave square root function, Eq. (4) applies the standard convolution property of the Fourier transform for independent variables, Eq. (5) relies on the mutual independence of the pairs (X_i, Z_i), and Eq. (6) follows by globally bounding the first $n-1$ variables using Lemma 4.

For the final variable X_n, applying Parseval's identity for any fixed leakage z_n yields

$$\sum_{\alpha \neq 0} \left| \hat{X}_{n|Z_n=z_n}(\alpha) \right|^2 = p \sum_{x \in \mathbb{F}_p} P_{z_n}(x)^2 - 1 < p \cdot q_{z_n} \sum_{x \in \mathbb{F}_p} P_{z_n}(x) = y_{z_n} \,,$$

where $P_{z_n}(x) := \Pr_{X_n|Z_n=z_n}[X_n = x]$, $q_{z_n} := \max_x P_{z_n}(x)$ is the maximum guessing probability, and $y_{z_n} := p \cdot q_{z_n}$ is the normalized guessing probability. Taking the expectation over Z_n, the average min-entropy condition $\widetilde{H}_\infty(X_n|Z_n) \geq \log(p) - \ell$ is equivalent to $\mathbb{E}_{z_n \leftarrow Z_n}[y_{z_n}] \leq 2^\ell$. Thus,

$$\mathop{\mathbb{E}}_{z_n \leftarrow Z_n} \left[\sum_{\alpha \neq 0} \left| \hat{X}_{n|Z_n=z_n}(\alpha) \right|^2 \right] < \mathop{\mathbb{E}}_{z_n \leftarrow Z_n} [y_{z_n}] \leq 2^\ell \,.$$

Substituting this bound into Eq. (7) bounds the statistical distance by

$$\varepsilon := \frac{1}{2} c_\ell^{n-1} \sqrt{2^\ell} = \frac{1}{2} 2^{\ell/2} c_\ell^{n-1} ,$$

concluding the proof. □

4.2 The Induced Bias-Resilient Combiner

The sum extraction theorem naturally allows us to obtain a bias-resilient combiner, by invoking it on the good random variables.

Theorem 4 (Prime Field Sum is Bias-Resilient)**.** *Let p be a prime, and let $a_0, a_1 \in \mathbb{F}_p$ be two arbitrary elements. Let $k, n \in \mathbb{N}$ and $\delta \in [1/2, 1 - 1/p)$ be given. Let $Z_1, \ldots, Z_n$ be arbitrary random variables over some finite set $\mathcal{Z}$. Consider the following n-input non-interactive combiner over $(\mathbb{F}_p \times \mathcal{Z})^n$:*

$$C_b((x_1, z_1), \ldots, (x_n, z_n)) := (a_b + s,\ \vec{z}) ,$$

where $s := \sum_{i=1}^n x_i \pmod p$ and $\vec{z} := (z_1, \ldots, z_n)$. Then (C_0, C_1) is an (ε, k, n) δ-bias-resilient indistinguishability combiner with respect to the ideal random variables $I_i \equiv (U_p, Z_i)$, with error

$$\varepsilon := \frac{1}{\sqrt{1-\delta}} \cdot \exp\left(- (k-1) \left((1-\delta)^2 - \frac{1}{p^2} \right) \right) .$$

Remark 4. Note that δ-bias-resilience implies δ'-bias-resilience for any $\delta' < \delta$, hence the theorem could be used even for $\delta' < 0.5$, but without improving the error ε. That is, $\varepsilon(\delta', k) = \varepsilon(1/2, k)$ for every $\delta' < 1/2$.

Proof. Let $(X_1, Z_1'), \ldots, (X_n, Z_n')$ be n mutually independent pairs (in each pair X_i and Z_i' are jointly distributed), such that for some "good" subset $G \subseteq [n]$ of size $|G| = k$, we have $\forall i \in G : (X_i, Z_i') \leq I_i/(1-\delta)$. Observe that for each such good pair

$$\begin{aligned}
2^{-\widetilde{H}_\infty(X_i|Z_i')} &= \mathop{\mathbb{E}}_{z \leftarrow Z_i'} \left[\max_x \Pr\left[X_i = x \mid Z_i' = z\right] \right] \\
&= \sum_{z \in \mathcal{Z}} \Pr\left[Z_i' = z\right] \cdot \max_x \Pr\left[X_i = x \mid Z_i' = z\right] \\
&= \sum_{z \in \mathcal{Z}} \max_x \Pr\left[X_i = x \cap Z_i' = z\right] \\
&\leq \sum_{z \in \mathcal{Z}} \max_x \left(\Pr\left[U_p = x \cap Z_i = z\right] / (1-\delta) \right) \\
&= \sum_{z \in \mathcal{Z}} 1/p \cdot \Pr\left[Z_i = z\right] \cdot 1/(1-\delta) \\
&= \frac{1}{p(1-\delta)} ,
\end{aligned}$$

hence we have $\forall i \in G :\ \widetilde{H}_\infty(X_i|Z_i') \geq \log(p) - \ell$, for $\ell := \log(1/(1-\delta))$.

Our assumption $\delta \in [1/2,\ 1 - 1/p)$ implies $1 \leq \ell$ for Theorem 3 and $\ell < \log(p)$ for Proposition 2. We may now invoke them to obtain

$$(S_G, \vec{Z}_G') \approx_{\varepsilon/2} (U_p, \vec{Z}_G') ,$$

where $S_G := \sum_{i \in G} X_i \pmod p$ and $\vec{Z}'_G := (Z'_i)_{i \in G}$, using that

$$\begin{aligned}
\varepsilon/2 &= \frac{1}{2} \cdot \frac{1}{\sqrt{1-\delta}} \cdot \exp\left(-(k-1)\left((1-\delta)^2 - \frac{1}{p^2}\right)\right) \\
&= \frac{1}{2} 2^{\ell/2} \exp\left(-(k-1)\left(\frac{1}{2^{2\ell}} - \frac{1}{p^2}\right)\right) \\
&\geq \frac{1}{2} 2^{\ell/2} c_\ell^{k-1} \,.
\end{aligned}$$

Since the pairs $(X_i, Z'_i)_{i \notin G}$ outside of G are independent of G and adding to uniform remains uniform, the transformation that adds them (and their leakage) does not increase statistical distance

$$(S, \vec{Z}') \approx_{\varepsilon/2} (U_p, \vec{Z}') \,.$$

Finally, we use the fact that $a_0 + U_p \equiv U_p \equiv a_1 + U_p$ and the triangle inequality to derive

$$(a_0 + S, \vec{Z}') \approx_{\varepsilon/2} (a_0 + U_p, \vec{Z}') \equiv (a_1 + U_p, \vec{Z}') \approx_{\varepsilon/2} (a_1 + S, \vec{Z}') \,,$$

and conclude the proof with

$$(a_0 + S, \vec{Z}') \approx_{\varepsilon} (a_1 + S, \vec{Z}') \,.$$

□

4.3 The Generalized XOR Lemma

We may now obtain a generalized XOR lemma over prime fields, by combining the bias-resilience of the XOR transformation with our main amplification theorem for bias-resilient combiners.

Theorem 5 (Prime Field XOR Lemma)**.** *Let $p \geq 3$ be a prime, and let $n \in \mathbb{N}$ and $\delta \in (0,\ 1 - 1/p)$ be given. Let $(X_1, Z_1), \ldots, (X_n, Z_n)$ be mutually independent pairs of random variables over $\mathbb{F}_p \times \mathcal{Z}$ for some finite set $\mathcal{Z}$, and $(Z'_1, \ldots, Z'_n)$ be any random variables over $\mathcal{Z}$, such that $\forall i \in [n]$: $(X_i, Z_i) \approx_\delta (U_p, Z'_i)$. Let $S := \sum_{i=1}^n X_i \pmod p$ and $\vec{Z} := (Z_1, \ldots, Z_n)$. Then, we have*

$$(S, \vec{Z}) \approx_{\varepsilon^*} (U_p, \vec{Z}) \,,$$

for

$$\begin{aligned}
\varepsilon^* &:= \min_{k \in [n]} \left(\delta^{(k,n)} + \varepsilon_k \right) \,, \\
\delta^* &:= \max(\delta,\ 1/2) \,, \\
\varepsilon_k &:= \frac{1}{\sqrt{1-\delta^*}} \cdot \exp\left(-(k-1)\left((1-\delta^*)^2 - \frac{1}{p^2}\right)\right) \,.
\end{aligned}$$

Proof. Observe that $\delta^* \in [1/2, 1 - 1/p)$. Invoking Theorem 4, for any $a_0, a_1 \in \mathbb{F}_p$ and any $k \in [n]$, we have that (C_0, C_1) is an $(\varepsilon_k,\ k,\ n)$ δ^*-bias-resilient indistinguishability combiner with respect to the ideal random variables $I_i := (U_p, Z'_i)$ for

$$C_b((x_1, z_1), \ldots, (x_n, z_n)) := (a_b + s,\ \vec{z}) \,.$$

Next, applying Theorem 1 while noting that δ^*-bias-resilience implies δ-bias-resilience as $\delta \leq \delta^*$, we obtain

$$(a_0 + S, \vec{Z}) \approx_{\delta^{(k,n)} + \varepsilon_k} (a_1 + S, \vec{Z}) \,.$$

In particular, we may choose $a_0 = 0$ and a uniform $a_1 \leftarrow U_p$, using the convexity of statistical distance:

$$\begin{aligned}\mathrm{SD}((S,\vec{Z}),(U_p,\vec{Z})) &= \mathrm{SD}((0+S,\vec{Z}),(U_p+S,\vec{Z})) \\ &\leq \mathop{\mathbb{E}}_{a_1\leftarrow U_p}\left[\mathrm{SD}((0+S,\vec{Z}),(a_1+S,\vec{Z}))\right] \\ &\leq \delta^{(k,n)} + \varepsilon_k\,.\end{aligned}$$

Finally, we note that the above holds for any $k \in [n]$ and in particular for $\varepsilon^* := \min_{k\in[n]}\left(\delta^{(k,n)} + \varepsilon_k\right)$. □

Observation 3. We can instantiate the Generalized $\mathbb{F}_p$-XOR Lemma with threshold $k = \lfloor \frac{1-\delta}{2} n \rfloor$. By Hoeffding's inequality, the binomial failure rate is at most $\delta^{(k,n)} \leq \exp\left(-\frac{(1-\delta)^2}{2} n\right)$, yielding a logarithmic decay of $\ln(\delta^{(k,n)}) \leq -\frac{(1-\delta)^2}{2} n = -k(1-\delta)$. Note that

$$(1-\delta^*)^2 - \frac{1}{p^2} < (1-\delta^*)^2 \leq (1-\delta)^2 < (1-\delta)\,,$$

hence the total error is dominated by the Fourier term ($\varepsilon^* \approx \varepsilon_k$). Thus, in order to bound the total error by a given threshold μ, we need

$$n \geq C \cdot \frac{\ln(1/\mu) + \ln(1/(1-\delta))}{(1-\delta)\left((1-\delta)^2 - \frac{1}{p^2}\right)}\,,$$

for some universal constant $C > 0$, which can be relaxed to the looser bound

$$n \geq C \cdot \frac{\ln(1/\mu) + \ln(1/(1-\delta))}{(1-\delta-\frac{1}{p})^3}\,.$$

By assuming that $1 - \delta - 1/p = \Omega(1/p)$ and $\mu \leq (1-\delta)^{\Omega(1)}$, we can further simplify to

$$n \geq C \cdot \frac{\ln(1/\mu)}{(1-\delta)^3}\,.$$

In particular, if δ is a constant, we only require $n \geq C \cdot \ln(1/\mu)$.

4.3.1 Computational Generalization

We generalize the statistical prime field XOR lemma to the computational setting using known techniques.

Lemma 5 (Hybrid Indistinguishability Lemma [BG24])**.** *Let X_0 and X_1 be random variables, and let $\delta, \gamma \in (0,1)$ and $s \in \mathbb{N}$. Define $s' := \frac{s\cdot\gamma^2}{128(\log|\mathrm{Im}(X_0)|+\log|\mathrm{Im}(X_1)|+1)}$. Assume that*

$$X_0 \overset{\mathsf{c}}{\approx}_{s,\delta} X_1.$$

Then there exists a hybrid distribution X' such that

$$X_0 \overset{\mathsf{c}}{\approx}_{s',\gamma(1-\delta)} X' \approx_\delta X_1\,.$$

Theorem 6 (Computational Prime Field XOR Lemma)**.** *Let $p \geq 3$ be a prime, and let $n, s \in \mathbb{N}$ and $\delta \in (0,\ 1-1/p)$ be given. Let $(X_1, Z_1), \ldots, (X_n, Z_n)$ be mutually independent pairs of random variables over $\mathbb{F}_p \times \{0,1\}^t$ for some $t \in \mathbb{N}$, and $(Z_1^*, \ldots, Z_n^*)$ be any random variables over $\{0,1\}^t$, such that $\forall i \in [n] :\ (X_i, Z_i) \overset{\mathsf{c}}{\approx}_{s,\delta} (U_p, Z_i^*)$. Let $S := \sum_{i=1}^n X_i \pmod p$ and $\vec{Z} := (Z_1, \ldots, Z_n)$. Then, for some universal constant $c > 1$ and every slackness parameter $\gamma \in (0,1)$, we have*

$$(S, \vec{Z}) \overset{\mathsf{c}}{\approx}_{s^*,\, \varepsilon^* + 2n\gamma(1-\delta)} (U_p, \vec{Z}) \,,$$

for

$$\begin{aligned}
\varepsilon^* &:= \min_{k \in [n]} \left(\delta^{(k,n)} + \varepsilon_k \right) \,, \\
\delta^* &:= \max(\delta,\ 1/2) \,, \\
\varepsilon_k &:= \frac{1}{\sqrt{1-\delta^*}} \cdot \exp\left(-(k-1) \left((1-\delta^*)^2 - \frac{1}{p^2} \right) \right) \,, \\
s^* &:= \frac{s \cdot \gamma^2}{256\,(\log(p) + t + 1)} - cn(\log(p) + t) \,.
\end{aligned}$$

Proof. Invoking Lemma 5, there exist hybrid pairs $(X_i', Z_i')_{i \in [n]}$ such that

$$\forall i \in [n] :\ (X_i, Z_i) \overset{\mathsf{c}}{\approx}_{s', \gamma(1-\delta)} (X_i', Z_i') \approx_\delta (U_p, Z_i^*) \,,$$

for $s' := (s \cdot \gamma^2)/(256\,(\log(p) + t + 1))$. Invoking the statistical Prime Field XOR Lemma (Theorem 5), for $S' := \sum_{i=1}^n X_i' \pmod p$ and $\vec{Z}' := (Z_1', \ldots, Z_n')$, we obtain

$$(S', \vec{Z}') \approx_{\varepsilon^*} (U_p, \vec{Z}') \,.$$

Finally, using a standard hybrid argument and the fact that the circuit size of computing the sum modulo p is of linear size in the binary representation, we obtain

$$(S, \vec{Z}) \overset{\mathsf{c}}{\approx}_{s^*,\, n\gamma(1-\delta)} (S', \vec{Z}') \approx_{\varepsilon^*} (U_p, \vec{Z}') \overset{\mathsf{c}}{\approx}_{s^*,\, n\gamma(1-\delta)} (U_p, \vec{Z}) \,.$$

By the triangle inequality, the total distinguishing advantage against size s^* circuits is at most $\varepsilon^* + 2n\gamma(1-\delta)$, to conclude the proof. □

4.3.2 Implications to Unpredictability

In a recent work, Shimizu and Yasunaga [SY26] established an $\mathbb{F}_p$-XOR lemma centered on computational unpredictability. While our indistinguishability-based framework naturally captures the unpredictability setting, the other direction is not true.

If a random variable X over $\mathbb{F}_p$ is d-unpredictable given leakage Z, then (X, Z) is weakly indistinguishable from uniform (U_p, Z) with parameter $\delta = 1 - \frac{1}{pd}$ (see Section B). Plugging this δ into our computational prime field XOR lemma (Theorem 6) yields an amplification requiring $n \geq C\left(\frac{p^3 d^3}{1-d} \log(pd/\mu)\right)$ copies to reach an error of μ. In contrast, the unpredictability amplification of [SY26, Theorem 3.4] requires $C\left(\frac{p^2}{1-d} \log(p/\mu)\right)$ copies. Intuitively, the extra factor p in our generic reduction arises because translating unpredictability into indistinguishability introduces an artificial error: rather than each RV having $\log(1/d)$ AME, we now need to condition on the good coupling event, occurring with probability $\frac{1}{pd}$, to achieve the same guarantee. This conversion overhead

can be bypassed for $d \geq 1/2$, which is the exact regime strictly focused on by [SY26], using the same generic approach taken by them: invoking the pseudo-average-min-entropy characterization of weakly unpredictable random variables (introduced by [Zhe14]; revisited in [SY26, Theorem 2.11]), which can be combined directly with our statistical sum-extraction theorem (Theorem 3) to recover the same bound.

Crucially, attempting to go the other direction—using the unpredictability-based approach to amplify weak indistinguishability—fails severely over large fields. Suppose (X, Z) has a constant indistinguishability gap from uniform, e.g., $\delta = 0.5$. When working directly with indistinguishability, our amplification theorem requires only $n \geq C \cdot \ln(1/\mu)$ copies, which is completely independent of the field size p. However, if we instead attempt to employ the unpredictability amplification theorem, a 0.5 indistinguishability gap merely guarantees $d = 0.5 + 1/p$ unpredictability. Feeding this constant d into the [SY26] framework forces a sample complexity of $C(p^2 \log(p/\mu))$. If p is exponentially large, as is standard in the context of cryptography when amplifying weak schemes over large prime fields, this p^2 dependence leads to an exponential blowup in the required number of copies. By treating indistinguishability as the primary metric, our combiner completely bypasses this artificial bottleneck.

5 All-or-Nothing Indistinguishability Amplifiers

From a broader perspective, the combiner-is-amplifier approach attempts to simplify security proofs by assuming idealized all-or-nothing failures, and then paying a penalty to switch back to standard security notions. In this section, we explicitly formalize this general paradigm via all-or-nothing amplifiers. We tightly characterize the multiplicative penalty incurred when translating idealized all-or-nothing amplification bounds to standard security.

Definition 12 (All-or-Nothing Indistinguishability Amplifier)**.** For finite sets $(\mathcal{Q}, \mathcal{A})$, $\varepsilon, p \in (0, 1)$, $n \in \mathbb{N}$, an (ε, p, n) all-or-nothing indistinguishability amplifier with respect to ideal $(\mathcal{A}, \mathcal{Q})$-systems $\vec{I} = (I_1, \ldots, I_n)$ against a $(\mathcal{Q}, \mathcal{A})$-distinguisher class $\mathcal{D}$, is a pair of n-oracle-aided $(\mathcal{Q}, \mathcal{A})$-systems $(C_0^{(\cdot)}, C_1^{(\cdot)})$ such that for all $(\mathcal{A}, \mathcal{Q})$-systems $\vec{S} = (S_1, \ldots, S_n)$ with $\forall i \in [n] : I_i \leq S_i/p$, we have:

$$\forall D \in \mathcal{D} : \left| \mathbb{E}\left[D(C_0^{\vec{S}}) - D(C_1^{\vec{S}}) \right] \right| \leq \varepsilon \,.$$

Theorem 7 (All-or-Nothing Indistinguishability Amplifiers Are Standard Amplifiers)**.** *For $\varepsilon, p, \delta \in (0, 1)$, $n \in \mathbb{N}$, given an (ε, p, n) all-or-nothing indistinguishability amplifier $(C_0^{(\cdot)}, C_1^{(\cdot)})$ with respect to ideal systems $\vec{I} = (I_1, \ldots, I_n)$ against a distinguisher class $\mathcal{D}$, let $\vec{R} = (R_1, \ldots, R_n)$ be a tuple of n real systems such that $\forall i \in [n] : R_i \approx_\delta I_i$. Then, we have*

$$\forall D \in \mathcal{D} : \left| \mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}}) \right] \right| \leq \varepsilon \cdot \left(1 + \frac{2p\delta}{1-p} \right)^n .$$

Proof. The proof is by induction on n. For $n = 0$ the statement is trivial. Let us focus on the induction step, and for simplicity drop the subscript and denote $I := I_n, R := R_n$. By the system decomposition lemma (Lemma 1), there exist systems S, R', I' such that

$$R \equiv (1-\delta)S + \delta R' \,,$$
$$I \equiv (1-\delta)S + \delta I' \,.$$

We define three systems such that the ideal system I is $(1-p)$-biased with respect to each:

$$\begin{aligned} H_R &\equiv pI + (1-p)R\,, \\ H_{R'} &\equiv pI + (1-p)R'\,, \\ H_{I'} &\equiv pI + (1-p)I'\,. \end{aligned}$$

Observe that fixing the last system to each of the above results in an $(\varepsilon, p, n-1)$ all-or-nothing indistinguishability amplifier, on which we may invoke the induction hypothesis. Our goal is now to express R as an affine combination $R \equiv \alpha H_R + \beta H_{R'} + \gamma H_{I'}$, implying that

$$\begin{aligned} \mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}})\right] &= \alpha \cdot \mathbb{E}\left[D(C_0^{\vec{R}_{n-1},H_R}) - D(C_1^{\vec{R}_{n-1},H_R})\right] \\ &\quad + \beta \cdot \mathbb{E}\left[D(C_0^{\vec{R}_{n-1},H_{R'}}) - D(C_1^{\vec{R}_{n-1},H_{R'}})\right] \\ &\quad + \gamma \cdot \mathbb{E}\left[D(C_0^{\vec{R}_{n-1},H_{I'}}) - D(C_1^{\vec{R}_{n-1},H_{I'}})\right], \\ \left|\mathbb{E}\left[D(C_0^{\vec{R}}) - D(C_1^{\vec{R}})\right]\right| &\leq (|\alpha| + |\beta| + |\gamma|) \cdot \varepsilon \cdot \left(1 + \frac{2p\delta}{1-p}\right)^{n-1}. \end{aligned}$$

Hence, in order to conclude the proof, we need to find an appropriate affine combination with $|\alpha| + |\beta| + |\gamma| = 1 + 2p\delta/(1-p)$. Such a combination follows:

$$\begin{aligned} H_R + \frac{p\delta}{1-p}H_{R'} - \frac{p\delta}{1-p}H_{I'} &= 1 \cdot (pI + (1-p)R) \\ &\quad + \frac{p\delta}{1-p} \cdot \big(pI + (1-p)R'\big) \\ &\quad - \frac{p\delta}{1-p} \cdot \big(pI + (1-p)I'\big) \\ &= pI + (1-p)R + p\delta R' - p\delta I' \\ &= p\big((1-\delta)S + \delta I'\big) + (1-p)\big((1-\delta)S + \delta R'\big) + p\delta R' - p\delta I' \\ &= (p + 1 - p)(1-\delta)S + (p-p)\delta I' + (1-p+p)\delta R' \\ &= (1-\delta)S + \delta R' = R\,. \end{aligned}$$

□

5.1 Penalty Tightness

We show that the multiplicative penalty established in Theorem 7 is tight for all $\delta \geq 0.5$, and tight for $\delta < 0.5$ when the advantage is sufficiently small.

Proposition 3 (All-or-Nothing Penalty Tightness)**.** For every $\varepsilon', p \in (0,1)$, $n \in \mathbb{N}$, and $\delta \in [0.5, 1)$, let $\varepsilon := \varepsilon' / \left(1 + \frac{2p\delta}{1-p}\right)^n$. There exist ideal systems $\vec{I} = (I_1, \ldots, I_n)$, a statistical (ε, p, n) all-or-nothing indistinguishability amplifier with respect to $\vec{I}$, and real systems $\vec{R} = (R_1, \ldots, R_n)$ satisfying $R_i \approx_\delta I_i$, such that

$$\mathrm{SD}(C_0^{\vec{R}}, C_1^{\vec{R}}) = \varepsilon' = \varepsilon \cdot \left(1 + \frac{2p\delta}{1-p}\right)^n.$$

For $\delta \in (0, 0.5)$, the same statement holds provided $\varepsilon' \leq \left(\frac{1-p+2p\delta}{1+p-2p\delta}\right)^n$.

Proof. In the following example, it will be more convenient to work over $\{-1,1\}$ rather than $\{0,1\}$. For a non-interactive all-or-nothing indistinguishability amplifier $(C_z(z_1,\ldots,z_n))_{z\in\{-1,1\}}$ with inputs and output in $\{-1,1\}$, the maximal distinguishing advantage over random variables $Z_1,\ldots,Z_n$ simplifies to

$$\left|\mathop{\mathbb{E}}_{z\leftarrow\{-1,1\}}\left[z\cdot C_z(Z_1,\ldots,Z_n)\right]\right|.$$

Our all-or-nothing amplifier is defined as:

$$C_z(z_1,\ldots,z_n) := z\cdot\prod_{i=1}^{n} Z'_i\,,$$

where each independent random variable Z'_i over $\{-1,1\}$ is sampled according to the input z_i, such that its expectation is $\mu(z_i)$ for some constants $-1\le\mu_{-1}<0<\mu_1\le 1$ to be specified later. Using the fact that $z^2\equiv 1$ and that the variables $Z_1,\ldots,Z_n$ are independent, we have

$$\left|\mathop{\mathbb{E}}_{z\leftarrow\{-1,1\}}\left[z\cdot C_z(Z_1,\ldots,Z_n)\right]\right| = \left|\mathbb{E}\left[\prod_{i=1}^{n} Z'_i\right]\right| = \left|\prod_{i=1}^{n}\mathbb{E}\left[Z'_i\right]\right| = \prod_{i=1}^{n}\left|\mathbb{E}\left[\mu(Z_i)\right]\right|.$$

Let us define the real and ideal systems as $R_i\equiv 1$ and $I_i\equiv(-1)^{\mathrm{Ber}(\delta)}$, which are exactly δ-statistically close. The distinguishing advantage over the real systems $R_1,\ldots,R_n$ is thus μ_1^n. To construct a tight example, we wish to maximize the ratio μ_1^n/ε, where ε bounds the distinguishing advantage over any valid all-or-nothing random variables satisfying $I_i\le Z_i/p$ for all i. Let μ^* denote the maximal value of $|\mathbb{E}\left[\mu(Z)\right]|$ over any such random variable Z with $(-1)^{\mathrm{Ber}(\delta)}\le Z/p$; then $\varepsilon := (\mu^*)^n$ and our goal simplifies to maximizing μ_1/μ^*.

Observe that $\mathbb{E}\left[\mu(I_i)\right] = \delta\mu_{-1}+(1-\delta)\mu_1$. The expectation $\mathbb{E}\left[\mu(Z_i)\right]$ of an all-or-nothing random variable Z_i satisfying $I_i\le Z_i/p$ is minimized by assigning the remaining probability mass to -1 (i.e., $Z_i := p\cdot I_i+(1-p)\cdot(-1)$) and maximized by assigning it to 1, yielding the following bounds:

$$\begin{aligned}\mu_{\min} &= (1-p(1-\delta))\cdot\mu_{-1}+p(1-\delta)\cdot\mu_1\,,\\ \mu_{\max} &= p\delta\cdot\mu_{-1}+(1-p\delta)\cdot\mu_1\,,\\ \mu^* &= \max(-\mu_{\min},\,\mu_{\max})\,.\end{aligned}$$

Without loss of generality, we may focus on the setting where $\alpha := \max(-\mu_{-1},\mu_1)=1$, as scaling both constants by $1/\alpha$ leaves the target ratio μ_1/μ^* unchanged.

- $\mu_1 = 1$: To maximize the ratio $\mu_1/\mu^* = 1/\mu^*$, we must minimize μ^*. We continuously decrease μ_{-1} from 0 until either the intersection $-\mu_{\min}=\mu_{\max}$ is reached, or we hit the domain boundary $\mu_{-1}=-1$. Note that initially for $\mu_{-1}=0$ we have $0<\mu_{\min}<\mu_{\max}$ as $p(1-\delta)<(1-p\delta)$. Furthermore, as μ_{-1} decreases, $\mu_{\min}$ decreases faster than $\mu_{\max}$ since $1-p(1-\delta)>p\delta>0$. The intersection equality $-\mu_{\min}=\mu_{\max}$ simplifies to:

$$\begin{aligned}(1-p(1-\delta))\cdot(-\mu_{-1})-p(1-\delta)\cdot\mu_1 &= p\delta\cdot\mu_{-1}+(1-p\delta)\cdot\mu_1\,,\\ (1-p+2p\delta)\cdot(-\mu_{-1}) &= (1+p-2p\delta)\cdot\mu_1\,.\end{aligned}$$

  We arrive at a separation into two regimes based on the value of δ. (1) If $\delta<0.5$ then $(1+p-2p\delta)/(1-p+2p\delta)>1$, meaning the boundary $\mu_{-1}=-1$ is reached before the intersection condition is met. Thus $\mu^*=\mu_{\max}=1-2p\delta$, yielding the looser ratio $\mu_1^n/(\mu^*)^n =$

$1/(1-2p\delta)^n$. (2) If $\delta \geq 0.5$ then the intersection is reached before the boundary, and μ_{-1} stops exactly at $-(1+p-2p\delta)/(1-p+2p\delta)$. Substituting this into $\mu_{\max}$ yields:

$$\begin{aligned}\mu^* &= (1-p\delta) - p\delta \cdot \frac{1+p-2p\delta}{1-p+2p\delta} = 1 - p\delta\left(1 + \frac{1+p-2p\delta}{1-p+2p\delta}\right) \\ &= 1 - p\delta \cdot \frac{2}{1-p+2p\delta} = \frac{1-p+2p\delta-2p\delta}{1-p+2p\delta} = \frac{1-p}{1-p+2p\delta}\,,\end{aligned}$$

which achieves the tight ratio

$$\frac{\mu_1^n}{(\mu^*)^n} = \left(\frac{1}{\mu^*}\right)^n = \left(\frac{(1-p+2p\delta)}{1-p}\right)^n = \left(1 + \frac{2p\delta}{1-p}\right)^n.$$

- $\mu_{-1} = -1$: Symmetrically, we increase μ_1 from 0 until either $-\mu_{\min} = \mu_{\max}$ or $\mu_1 = 1$. (1) If $\delta < 0.5$ then μ_1 reaches the intersection $(1-p+2p\delta)/(1+p-2p\delta) < 1$. Subsequently,

$$\begin{aligned}\mu^* &= (1-p\delta)\frac{1-p+2p\delta}{1+p-2p\delta} - p\delta = \frac{1-p+2p\delta}{1+p-2p\delta} - p\delta \cdot \left(1 + \frac{1-p+2p\delta}{1+p-2p\delta}\right) \\ &= \frac{1-p+2p\delta}{1+p-2p\delta} - p\delta \cdot \frac{2}{1+p-2p\delta} = \frac{1-p}{1+p-2p\delta}\,,\end{aligned}$$

achieving the tight ratio $\mu_1^n/(\mu^*)^n = ((1-p+2p\delta)/(1-p))^n = (1+2p\delta/(1-p))^n$. (2) If $\delta \geq 0.5$ then μ_1 reaches 1 before the intersection is met, yielding the looser ratio $1/(1-2p(1-\delta))^n$.

We conclude that for $\delta \in [0.5, 1)$, setting $\mu_1 = (\varepsilon')^{1/n}$ and $\mu_{-1} = -\mu_1 \cdot \frac{1+p-2p\delta}{1-p+2p\delta}$ achieves the tight ratio of Theorem 7, yielding $\varepsilon = (\mu^*)^n = \varepsilon' / \left(1 + \frac{2p\delta}{1-p}\right)^n$.

For $\delta \in (0, 0.5)$, a looser ratio of $(1-2p\delta)^{-n}$ is unconditionally achievable by setting $\mu_{-1} = -\mu_1$. To achieve the tight ratio in this regime, we must satisfy the intersection condition without violating the lower bound $\mu_{-1} \geq -1$. This imposes the condition $\mu_1 \leq \frac{1-p+2p\delta}{1+p-2p\delta}$. Substituting $\mu_1 = (\varepsilon')^{1/n}$, we see this restricts the target advantage to sufficiently small values, specifically $\varepsilon' \leq \left(\frac{1-p+2p\delta}{1+p-2p\delta}\right)^n$. When this holds, setting $\mu_{-1} = -\mu_1 \cdot \frac{1+p-2p\delta}{1-p+2p\delta}$ as before successfully recovers the tight ratio. □

Acknowledgments

We would like to thank Nir Bitansky for numerous helpful discussions and his valuable input throughout the course of this research.

References

[ABG26] Benny Applebaum, Nir Bitansky, and Nathan Geier. Security amplification via robust indistinguishability combiners. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part I*, volume 16800 of *Lecture Notes in Computer Science*, pages 35–66. Springer, 2026. 3, 4, 5, 6, 12, 14

[AK25a] Benny Applebaum and Eliran Kachlon. How to share an NP statement or combiners for zero-knowledge proofs. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 480–513. Springer, Cham, August 2025. 3

[AK25b] Benny Applebaum and Eliran Kachlon. NIZK amplification via leakage-resilient secure computation. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 462–479. Springer, Cham, August 2025. 3, 14

[BDIR18] Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 531–561. Springer, Cham, August 2018. 7, 14, 20, 21

[BG24] Nir Bitansky and Nathan Geier. Amplification of non-interactive zero knowledge, revisited. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part IX*, volume 14928 of *LNCS*, pages 361–390. Springer, Cham, August 2024. 3, 8, 14, 26

[BKP+24] Nir Bitansky, Chethan Kamath, Omer Paneth, Ron D. Rothblum, and Prashant Nalini Vasudevan. Batch proofs are statistically hiding. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 435–443. ACM Press, June 2024. 3

[CHS05] Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 17–33. Springer, Berlin, Heidelberg, February 2005. 3

[CK90] Claude Crépeau and Joe Kilian. Weakening security assumptions and oblivious transfer (abstract). In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 2–7. Springer, New York, August 1990. 3

[DKS99] Ivan Damgård, Joe Kilian, and Louis Salvail. On the (im)possibility of basing oblivious transfer and bit commitment on weakened security assumptions. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 56–73. Springer, Berlin, Heidelberg, May 1999. 3

[DNR04] Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 342–360. Springer, Berlin, Heidelberg, May 2004. 3

[DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008. 7, 9

[DP07] Stefan Dziembowski and Krzysztof Pietrzak. Intrusion-resilient secret sharing. In *48th FOCS*, pages 227–237. IEEE Computer Society Press, October 2007. 14

[GJS19] Vipul Goyal, Aayush Jain, and Amit Sahai. Simultaneous amplification: The case of non-interactive zero-knowledge. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 608–637. Springer, Cham, August 2019. 3

[GNW11] Oded Goldreich, Noam Nisan, and Avi Wigderson. On yao's xor-lemma. In Oded Goldreich, editor, *Studies in Complexity and Cryptography*, volume 6650 of *Lecture Notes in Computer Science*, pages 273–301. Springer, 2011. 3

[Gol01] Oded Goldreich. *The Foundations of Cryptography - Volume 1: Basic Techniques*. Cambridge University Press, 2001. 3

[GP91] Oded Goldreich and Erez Petrank. Quantifying knowledge complexity. In *32nd FOCS*, pages 59–68. IEEE Computer Society Press, October 1991. 15

[Gre07] Ben Green. Montreal lecture notes on quadratic fourier analysis, 2007. 20

[Her05] Amir Herzberg. On tolerant cryptographic constructions. In Alfred Menezes, editor, *CT-RSA 2005*, volume 3376 of *LNCS*, pages 172–190. Springer, Berlin, Heidelberg, February 2005. 3

[HIKN08] Danny Harnik, Yuval Ishai, Eyal Kushilevitz, and Jesper Buus Nielsen. OT-combiners via secure computation. In Ran Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 393–411. Springer, Berlin, Heidelberg, March 2008. 3

[HKN+05] Danny Harnik, Joe Kilian, Moni Naor, Omer Reingold, and Alon Rosen. On robust combiners for oblivious transfer and other primitives. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 96–113. Springer, Berlin, Heidelberg, May 2005. 3

[Hol05] Thomas Holenstein. Key agreement from weak bit agreement. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 664–673. ACM Press, May 2005. 3, 8

[HR05] Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 478–493. Springer, Berlin, Heidelberg, August 2005. 3

[Imp95] Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *36th FOCS*, pages 538–545. IEEE Computer Society Press, October 1995. 8

[LM20] David Lanzenberger and Ueli Maurer. Coupling of random systems. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 207–240. Springer, Cham, November 2020. 5, 7, 11, 12

[LT13] Huijia Lin and Stefano Tessaro. Amplification of chosen-ciphertext security. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 503–519. Springer, Berlin, Heidelberg, May 2013. 3

[Mau02] Ueli M. Maurer. Indistinguishability of random systems. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 110–132. Springer, Berlin, Heidelberg, April / May 2002. 9

[MPR07] Ueli M. Maurer, Krzysztof Pietrzak, and Renato Renner. Indistinguishability amplification. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 130–149. Springer, Berlin, Heidelberg, August 2007. 5

[MR11] Ueli Maurer and Renato Renner. Abstract cryptography. In Bernard Chazelle, editor, *ICS 2011*, pages 1–21. Tsinghua University Press, January 2011. 9

[MT09] Ueli M. Maurer and Stefano Tessaro. Computational indistinguishability amplification: Tight product theorems for system composition. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 355–373. Springer, Berlin, Heidelberg, August 2009. 5

[MT10] Ueli M. Maurer and Stefano Tessaro. A hardcore lemma for computational indistinguishability: Security amplification for arbitrarily weak PRGs with optimal stretch. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 237–254. Springer, Berlin, Heidelberg, February 2010. 8, 14

[SV19] Akshayaram Srinivasan and Prashant Nalini Vasudevan. Leakage resilient secret sharing and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 480–509. Springer, Cham, August 2019. 14

[SY26] Nobutaka Shimizu and Kenji Yasunaga. Hardness amplification beyond boolean functions. In Aditya Bhaskara and Artur Czumaj, editors, *Proceedings of the 58th Annual ACM Symposium on Theory of Computing, STOC 2026, Salt Lake City, UT, USA, June 22-26, 2026*, pages 1278–1289. ACM, 2026. 6, 7, 27, 28

[Tes11] Stefano Tessaro. Security amplification for the cascade of arbitrarily weak PRPs: Tight bounds via the interactive hardcore lemma. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 37–54. Springer, Berlin, Heidelberg, March 2011. 14

[TV07] Terence Tao and Van H. Vu. *Additive combinatorics*, volume 105 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 2007. 20

[Wul07] Jürg Wullschleger. Oblivious-transfer amplification. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 555–572. Springer, Berlin, Heidelberg, May 2007. 3

[Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd FOCS*, pages 80–91. IEEE Computer Society Press, November 1982. 3

[Zhe14] Jia Zheng. *A Uniform Min-Max Theorem and Characterizations of Computational Randomness*. PhD thesis, Harvard University, Cambridge, MA, USA, 2014. 28

A Analytical Bounds

Proposition (Restating Proposition 1)**.** Let $f(y) := \left(\frac{y\sin(\pi/y)}{p\sin(\pi/p)}\right)^2$ for some real number $p \geq 2$. Then for any random variable Y taking values in $[1, p]$, we have:

$$\mathbb{E}\left[f(Y)\right] \leq f\big(\max(2, \mathbb{E}\left[Y\right])\big)\,.$$

Proof. Because multiplying by a positive constant does not affect the direction of inequalities, we show without loss of generality that the inequality holds for the unscaled function $F(y) := y^2\sin^2(\pi/y) = (p\sin(\pi/p))^2 f(y)$.

Our proof proceeds in three steps. First, we show $F(y)$ is strictly concave and strictly increasing for $y \geq 2$. Second, we patch the partially convex region on $[1, 2]$ using its tangent line at $y = 2$ to construct a globally concave, strictly increasing function $\tilde{F}(y)$. This guarantees:

$$\mathbb{E}\left[\tilde{F}(Y)\right] \leq \tilde{F}(\mathbb{E}\left[Y\right]) \leq \tilde{F}\big(\max(2, \mathbb{E}\left[Y\right])\big) = F\big(\max(2, \mathbb{E}\left[Y\right])\big)\,.$$

Finally, we prove this patched function upper-bounds $F(y)$, allowing us to conclude $\mathbb{E}\left[F(Y)\right] \leq \mathbb{E}\left[\tilde{F}(Y)\right]$, completing the proof. For a visualization, see Fig. 1.

Concavity and Monotonicity for $y \geq 2$: We compute the derivatives of $F(y)$. By the chain and product rules, the first derivative is:

$$\begin{aligned} F'(y) &= 2y\sin^2(\pi/y) + y^2 \cdot 2\sin(\pi/y)\cos(\pi/y)\left(-\frac{\pi}{y^2}\right) \\ &= 2y\sin^2(\pi/y) - \pi\sin(2\pi/y)\,. \end{aligned}$$

Taking the derivative again yields:

$$\begin{aligned} F''(y) &= \left(2\sin^2(\pi/y) + 2y \cdot 2\sin(\pi/y)\cos(\pi/y)\left(-\frac{\pi}{y^2}\right)\right) - \pi\cos(2\pi/y)\left(-\frac{2\pi}{y^2}\right) \\ &= 2\sin^2(\pi/y) - \frac{2\pi}{y}\sin(2\pi/y) + \frac{2\pi^2}{y^2}\cos(2\pi/y)\,. \end{aligned}$$

Substituting $x = \pi/y$, the condition $F''(y) < 0$ reduces to evaluating $g(x) < 0$, where:

$$g(x) := \sin^2(x) - x\sin(2x) + x^2\cos(2x)\,.$$

For $y \geq 2$, we are evaluating $x \in (0, \pi/2]$. First, observe that at the boundary $x = 0$, $g(0) = 0$. We evaluate the derivative $g'(x)$, using the identity $\sin(2x) = 2\sin(x)\cos(x)$:

$$\begin{aligned} g'(x) &= 2\sin(x)\cos(x) - (\sin(2x) + 2x\cos(2x)) + \big(2x\cos(2x) - 2x^2\sin(2x)\big) \\ &= -2x^2\sin(2x)\,. \end{aligned}$$

For any $x \in (0, \pi/2)$, we have $x^2 > 0$ and $\sin(2x) > 0$, meaning $g'(x) < 0$. Since $g(0) = 0$ and g is strictly decreasing on the closed interval $[0, \pi/2]$, it follows that $g(x) < 0$ for all $x \in (0, \pi/2]$. Consequently, $F''(y) < 0$, confirming $F(y)$ is strictly concave for $y \geq 2$. Equivalently, this guarantees that the first derivative $F'(y)$ is strictly decreasing for $y \geq 2$. Furthermore, observe the limit of the first derivative as $y \to \infty$ (or equivalently, as $x \to 0^+$):

$$\lim_{y\to\infty} F'(y) = \lim_{x\to 0^+}\left(\frac{2\pi}{x}\sin^2(x) - \pi\sin(2x)\right) = 0\,.$$

Because $F'(y)$ is strictly decreasing for $y \geq 2$ and tends to 0 as $y \to \infty$, it must be strictly positive for all $y \geq 2$.

The Patched Function: Because $F(y)$ is partially convex on $[1, 2)$, we define a patch using the tangent line to $F(y)$ at $y = 2$. Evaluating F and F' at $y = 2$ yields $F(2) = 4$ and $F'(2) = 4$. The tangent line is $T(y) = 4y - 4$. We define the patched function:

$$\tilde{F}(y) := \begin{cases} 4y - 4 & \text{if } y \in [1, 2) \\ F(y) & \text{if } y \geq 2 \end{cases}$$

Observe the behavior of its derivative, $\tilde{F}'(y)$. On $[1, 2)$, the derivative is the constant 4, which is positive. At $y = 2$, the left limit is 4 and the right limit is $F'(2) = 4$, making the derivative continuous (well defined). For $y > 2$, the derivative is $F'(y)$, which we established earlier is strictly positive and strictly decreasing. Because $\tilde{F}'(y)$ is continuous, strictly positive, and everywhere non-increasing, the patched function $\tilde{F}(y)$ is both *strictly increasing* and *globally concave*.

Proving the Upper Bound: We must show $T(y) \geq F(y)$ on $[1, 2]$. Let $\Delta(y) := F(y) - T(y)$. We evaluate the boundaries and their slopes:

- $\Delta(1) = 1^2 \sin^2(\pi) - (4-4) = 0$, with $\Delta'(1) = F'(1) - 4 = -4$.
- $\Delta(2) = 2^2 \sin^2(\pi/2) - (8-4) = 0$, with $\Delta'(2) = F'(2) - 4 = 0$.

The second derivative is $\Delta''(y) = F''(y) = 2g(x)$, where $x = \pi/y$. For $y \in (1,2)$, we have $x \in (\pi/2, \pi)$. Because $\sin(2x)$ is strictly negative in this quadrant, $g'(x) = -2x^2 \sin(2x) > 0$. Thus, $g(x)$ is strictly increasing with respect to x. Since $x = \pi/y$ is strictly decreasing with respect to y, it follows that $\Delta''(y)$ decreases as y increases from 1 to 2. Evaluating the boundaries gives $\Delta''(1) = 2g(\pi) = 2\pi^2 > 0$ and $\Delta''(2) = 2g(\pi/2) = 2(1-\pi^2/4) < 0$. Therefore, $\Delta''(y)$ starts positive and decreases to a negative value.

This shape of $\Delta''(y)$ dictates the behavior of $\Delta'(y)$: the first derivative starts at -4, increases to a positive peak, and then decreases back to 0. Because $\Delta'(y)$ transitions from negative to positive, there exists a unique root $w \in (1,2)$ such that $\Delta'(y)$ is negative on $[1,w)$ and positive on $(w,2)$. Consequently, the function $\Delta(y)$ must decrease from 0 at $y = 1$ to a global minimum at $y = w$, and then increase back to 0 at $y = 2$. Thus, $\Delta(y) \leq 0$ for all $y \in [1,2]$, and we conclude that $F(y) \leq \tilde{F}(y)$ on the entire interval $[1,p]$.

The Expectation Bound: Applying our established properties of $\tilde{F}$ (global concavity, monotonicity, and the upper bound property $F(Y) \leq \tilde{F}(Y)$), along with Jensen's inequality, we obtain:

$$\mathbb{E}\left[F(Y)\right] \leq \mathbb{E}\left[\tilde{F}(Y)\right] \leq \tilde{F}\big(\mathbb{E}\left[Y\right]\big) \leq \tilde{F}\big(\max(2, \mathbb{E}\left[Y\right])\big) = F\big(\max(2, \mathbb{E}\left[Y\right])\big).$$

The final equality holds because $\max(2, \mathbb{E}\left[Y\right]) \geq 2$, placing the evaluation in the original curve's domain where $\tilde{F}(y) = F(y)$. Scaling back to $f(y)$ preserves all inequalities, concluding the proof. □

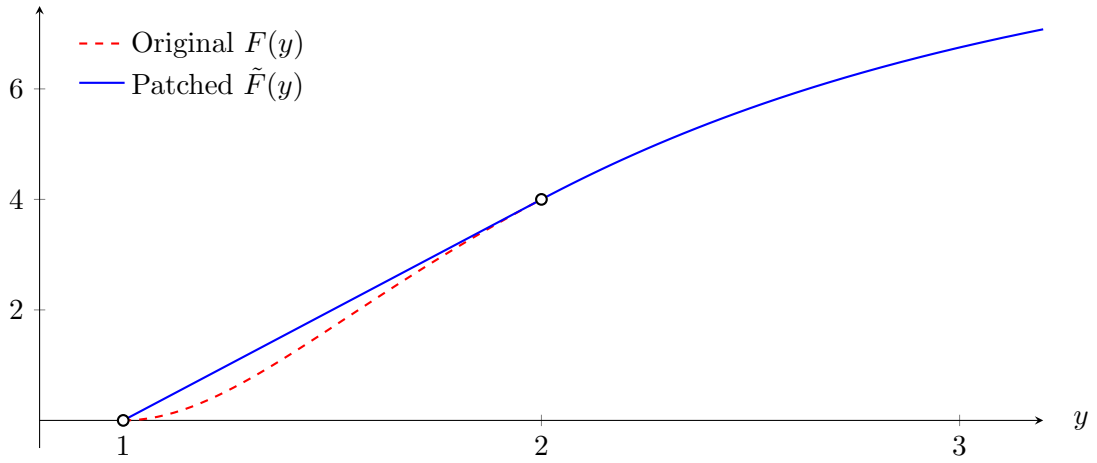

Figure 1: Visualization of the globally concave upper bound $\tilde{F}(y)$ and the partially convex region of $F(y)$ on the interval $[1,2]$.

Proposition (Restating Proposition 2)**.** For real numbers $\ell, p \in \mathbb{R}^+$ such that $2^\ell < p$ we have

$$c_\ell := \frac{2^\ell \sin(\pi 2^{-\ell})}{p \sin(\pi/p)} < \exp\left(-\left(\frac{1}{2^{2\ell}} - \frac{1}{p^2}\right)\right).$$

Proof. Let $x_0 := \pi/p$ and $x_1 := \pi 2^{-\ell}$. Since $\ell > 0$ and $p > 2^\ell$, we have $0 < x_0 < x_1 < \pi$. We can rewrite the coefficient using

$$c_\ell = \frac{\sin(x_1)/x_1}{\sin(x_0)/x_0} = \exp\left(h(x_1) - h(x_0)\right),$$

where $h(x) := \ln\left(\frac{\sin x}{x}\right) = \ln(\sin x) - \ln x$. Taking the derivative with respect to x, we get:

$$\begin{aligned} h'(x) &= (\ln(\sin x))' - (\ln x)' \\ &= \frac{\cos x}{\sin x} - \frac{1}{x} \\ &= \cot x - \frac{1}{x}. \end{aligned}$$

We invoke a known inequality for the cotangent function, derived from its Laurent series expansion, which states that for all $x \in (0, \pi)$ we have

$$\cot x < \frac{1}{x} - \frac{x}{3},$$

yielding $h'(x) < -x/3$. We can now use the fundamental theorem of calculus to bound

$$h(x_1) - h(x_0) = \int_{x_0}^{x_1} h'(x)\,dx < \int_{x_0}^{x_1} -\frac{x}{3}\,dx = -\frac{x_1^2 - x_0^2}{6},$$

which allows us to obtain

$$\ln c_\ell = h(x_1) - h(x_0) < -\frac{(\pi 2^{-\ell})^2 - (\pi/p)^2}{6} = -\frac{\pi^2}{6}\left(\frac{1}{2^{2\ell}} - \frac{1}{p^2}\right) < -\left(\frac{1}{2^{2\ell}} - \frac{1}{p^2}\right),$$

completing the proof. □

B Unpredictability to Indistinguishability

In this section, we show that if for a joint distribution (X, Z) we have that X is d-unpredictable given Z, then (X, Z) is weakly indistinguishable from uniform $(U_\mathcal{X}, Z)$. More accurately, assuming that $|\mathcal{X}|$ is polynomially bounded and $d \geq 2/|\mathcal{X}|$, we obtain indistinguishability parameter $\varepsilon := 1 - \frac{1}{d|\mathcal{X}|}$. The assumption that $|\mathcal{X}|$ is polynomially bounded is needed to efficiently implement the rejection-sampling predictor described below, and it can be removed in the statistical setting. We also remark that the restriction $d \geq 2/|\mathcal{X}|$ can be removed in the statistical setting, i.e., when considering the max-likelihood distinguisher.

Lemma 6. *Let (X, Z) be jointly distributed RVs over $\mathcal{X} \times \mathcal{Z}$ for some finite sets $\mathcal{X}, \mathcal{Z}$. Let $f : \mathcal{X} \times \mathcal{Z} \to \{0, 1\}$ be any boolean function such that*

$$\varepsilon := \mathbb{E}\left[f(X, Z) - f(U_\mathcal{X}, Z)\right] \geq \frac{1}{2}.$$

Let $P : \mathcal{Z} \to \mathcal{X}$ be the randomized rejection-sampling predictor that on input z outputs x with probability $f(x, z)/\sum_{x \in \mathcal{X}} f(x, z)$, or $1/|\mathcal{X}|$ if $\forall x :\ f(x, z) = 0$. Then $\Pr[P(Z) = X] \geq \frac{1}{|\mathcal{X}|(1-\varepsilon)}$.

Proof. First, we may assume WLOG that $\forall z \in \mathcal{Z}\, \exists x \in \mathcal{X} : f(x,z) = 1$, since switching $f(\cdot, z) \equiv 0$ to $f(\cdot, z) \equiv 1$ does not change the advantage ε or the behavior of the predictor P. Let us denote

$$\begin{aligned} p_z &:= \mathbb{E}\left[f(X,Z) \mid Z = z\right], \\ k_z &:= |\mathcal{X}| \cdot \mathbb{E}\left[f(U_{\mathcal{X}}, z)\right] = |\{x \mid f(x,z) = 1\}|, \\ \varepsilon_z &:= p_z - k_z/|\mathcal{X}|, \end{aligned}$$

then we have that $\mathbb{E}\left[\varepsilon_Z\right] = \varepsilon$. Further, let us define $\mathcal{Z}^+ := \{z \mid \varepsilon_z \geq 0\}$, $p^+ := \Pr\left[Z \in \mathcal{Z}^+\right]$, and $\varepsilon^+ := \mathbb{E}\left[\varepsilon_Z \mid Z \in \mathcal{Z}^+\right]$. Observe that $\varepsilon \leq p^+ \cdot \varepsilon^+$ and that equality is only obtained if $p^+ = 1$, since otherwise the complement term $(1-p^+) \cdot \mathbb{E}\left[\varepsilon_Z \mid Z \notin \mathcal{Z}^+\right]$ is strictly negative.

We have that

$$\begin{aligned} \Pr\left[P(Z) = X\right] &= \mathop{\mathbb{E}}_{z \leftarrow Z}\left[\frac{p_z}{k_z}\right] && (8)\\ &\geq p^+ \cdot \mathop{\mathbb{E}}_{z \leftarrow Z \in \mathcal{Z}^+}\left[\frac{p_z}{k_z}\right] && (9)\\ &= \frac{p^+}{|\mathcal{X}|} \cdot \mathop{\mathbb{E}}_{z \leftarrow Z \in \mathcal{Z}^+}\left[\frac{p_z}{p_z - \varepsilon_z}\right] && (10)\\ &\geq \frac{p^+}{|\mathcal{X}|} \cdot \mathop{\mathbb{E}}_{z \leftarrow Z \in \mathcal{Z}^+}\left[\frac{p_z}{p_z - p_z \cdot \varepsilon_z}\right] && (11)\\ &= \frac{p^+}{|\mathcal{X}|} \cdot \mathop{\mathbb{E}}_{z \leftarrow Z \in \mathcal{Z}^+}\left[\frac{1}{1 - \varepsilon_z}\right] && (12)\\ &\geq \frac{p^+}{|\mathcal{X}|} \cdot \frac{1}{\mathbb{E}_{z \leftarrow Z \in \mathcal{Z}^+}\left[1 - \varepsilon_z\right]} && (13)\\ &= \frac{p^+}{|\mathcal{X}|} \cdot \frac{1}{1 - \varepsilon^+} && (14)\\ &\geq \frac{p^+}{|\mathcal{X}|} \cdot \frac{1}{1 - \varepsilon/p^+} && (15)\\ &= \frac{1}{|\mathcal{X}|} \cdot \frac{(p^+)^2}{p^+ - \varepsilon}, && (16) \end{aligned}$$

where in Eq. (11) we used that $\forall z \in \mathcal{Z}^+ : \varepsilon_z \geq 0$ and thus $p_z \cdot \varepsilon_z \leq \varepsilon_z$, and in Eq. (13) we used Jensen's inequality for the convex function $1/x$.

Consider the function $g(x) := \frac{x^2}{x-\varepsilon}$ for $x \in (\varepsilon, 1]$. We have

$$g'(x) = \frac{2x(x-\varepsilon) - x^2}{(x-\varepsilon)^2} = \frac{x^2 - 2x\varepsilon}{(x-\varepsilon)^2} = \frac{x(x-2\varepsilon)}{(x-\varepsilon)^2},$$

which is negative for $x \in (\varepsilon, 2\varepsilon)$. Due to our assumption $\varepsilon \geq 1/2$, we have $g(x) \geq g(1) = \frac{1}{1-\varepsilon}$ for all $x \in (\varepsilon, 1]$. Hence, plugging back to Eq. (16), we obtain

$$\Pr\left[P(Z) = X\right] \geq \frac{1}{|\mathcal{X}|} \cdot \frac{1}{1-\varepsilon},$$

concluding the proof. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il