Electronic Colloquium on Computational Complexity, Report No. 75 (2025)

Leakage-Resilient Extractors against Number-on-Forehead Protocols

Eshan Chattopadhyay*
Cornell University
eshan@cs.cornell.edu

Jesse Goodman†
The University of Texas at Austin
jpmgoodman@utexas.edu

Abstract

Given a sequence of N independent sources $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_N \sim \{0,1\}^n$, how many of them must be good (i.e., contain some min-entropy) in order to extract a uniformly random string? This question was first raised by Chattopadhyay, Goodman, Goyal and Li (STOC '20), motivated by applications in cryptography, distributed computing, and the unreliable nature of real-world sources of randomness. In their paper, they showed how to construct explicit low-error extractors for just $K \geq N^{1/2}$ good sources of polylogarithmic min-entropy. In a follow-up, Chattopadhyay and Goodman improved the number of good sources required to just $K \geq N^{0.01}$ (FOCS '21). In this paper, we finally achieve $K = 3$.

Our key ingredient is a near-optimal explicit construction of a new pseudorandom primitive, called a leakage-resilient extractor (LRE) against number-on-forehead (NOF) protocols. Our LRE can be viewed as a significantly more robust version of Li's low-error three-source extractor (FOCS '15), and resolves an open question put forth by Kumar, Meka, and Sahai (FOCS '19) and Chattopadhyay, Goodman, Goyal, Kumar, Li, Meka, and Zuckerman (FOCS '20). Our LRE construction is based on a simple new connection we discover between multiparty communication complexity and non-malleable extractors, which shows that such extractors exhibit strong average-case lower bounds against NOF protocols.

1 Introduction

Randomness is a surprisingly useful tool, with important applications in algorithm design, combinatorics, cryptography, distributed computing, and more [Vad12]. This raises the natural question:

Where does this randomness come from?

In the real world, random bits are harvested from a variety of natural phenomena, including atmospheric noise, thermal noise, radioactive decay, and other quantum phenomena [HG17]. Unfortunately, most applications of randomness require access to uniformly random bits, while the bits produced in nature can have various correlations and biases. This motivates the need for a device that can purify these weak sources of randomness into perfectly random bits. A *randomness extractor* is exactly such a device.

Definition 0 (Randomness extractor)**.** *Let $\mathcal{X}$ be a family of distributions over $\{0,1\}^n$. A function $\mathsf{Ext} : \{0,1\}^n \to \{0,1\}^m$ is called an* extractor for $\mathcal{X}$ with error ε *if for every $\mathbf{X} \in \mathcal{X}$,*

$$|\mathsf{Ext}(\mathbf{X}) - \mathbf{U}_m| \leq \varepsilon,$$

where $\mathbf{U}_m$ denotes the uniform distribution over $\{0,1\}^m$ and $|\cdot|$ denotes statistical distance.

*Supported by a Sloan Research Fellowship and NSF CAREER Award 2045576.
†Supported by a Simons Investigator Award (#409864, David Zuckerman).

ISSN 1433-8092

Ever since the classical work of von Neumann [vN51], extractors have become a central object in complexity theory, finding applications in cryptography [BBR88], combinatorics [Li23], coding theory [Gur04], and more. Along the way, a beautiful theory has grown around extractors, resulting in a rich literature spanning four decades and hundreds of papers. The fundamental question underlying much of this work is,

Which family of distributions $\mathcal{X}$ should we try to extract from?

Of course, we would like the family $\mathcal{X}$ to be as general as possible, but some assumptions are necessary in order for extraction to actually *be* possible. First, each source $\mathbf{X} \in \mathcal{X}$ must clearly have *some* randomness in it. In the extractor literature, it is common to formalize this using the notion of *min-entropy*, defined as

$$H_\infty(\mathbf{X}) := \min_{x \in \text{support}(\mathbf{X})} \log\left(1/\Pr[\mathbf{X} = x]\right).$$

However, even this assumption is not enough, and a classic impossibility result [CG88] says that it is impossible to extract even from the family of distributions with nearly the *maximum* amount of min-entropy $k = n - 1$. As such, researchers have looked towards additional assumptions to make on the family $\mathcal{X}$.

Independent sources

Perhaps the oldest and most well-studied assumption to make on the family $\mathcal{X}$ is that each source $\mathbf{X} \in \mathcal{X}$ actually consists of *several* independent sources $\mathbf{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_N)$, each over n bits, and each guaranteed to have some min-entropy, k.[1] A systematic study of this model began with the seminal works of Vazirani [Vaz87] and Chor and Goldreich [CG88], and has continued until the present day. Today, after a long line of work stretching over three decades (e.g., [Vaz87, CG88, BIW06, Bou05, Rao09, Li15, CZ19, Li23]), we now have near-optimal extractors for the independent source model.[2]

Adversarial sources

While the study of extractors for independent sources has seen great success, it is not always clear if this is actually a realistic model. In the real world, natural sources of randomness can be unreliable, and may sometimes (at unknown times) produce samples with no entropy at all. This motivates the study of a more general model, where each source need not have a min-entropy guarantee. Such a model was introduced by Chattopadhyay, Goodman, Goyal and Li [CGGL20], and are known as *adversarial sources*.[3]

Definition 1 (Adversarial source)**.** *A source $\mathbf{X} \sim (\{0,1\}^n)^N$ is called an* (N, K, n, k)-adversarial source *if it is of the form $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_N)$, where each $\mathbf{X}_i$ is an independent source over n bits, and at least K of them are* good*: i.e., there is some set $S \subseteq [N]$ of size $|S| \geq K$ such that $H_\infty(\mathbf{X}_i) \geq k$ for all $i \in S$.*

If one could construct explicit extractors for adversarial sources, one could harvest uniform bits in a much more robust way, which doesn't completely break down whenever a source fails to output min-entropy. Moreover, as discussed in [CGGL20], extractors for adversarial sources have applications in cryptography

[1] In this paper, N and n are completely unrelated variables (unlike many papers in the field that take N to mean 2^n).

[2] In more detail, we now have explicit three-source extractors for $k = \text{polylog}(n)$ entropy and exponentially small error [Li15], explicit two-source extractors for $k = \text{polylog}(n)$ entropy and polynomially small error [CZ19], and explicit two-source extractors for $k = O(\log n)$ entropy and constant error [Li23].

[3] In the original paper [CGGL20], the authors also consider a more general notion of adversarial sources, where each bad source can depend on up to d good sources. However, the majority of the work [CGGL20], and all of the follow-up work [CG21], focuses on the fundamental setting of $d = 0$ (which, as we will see, already enjoys many applications). We take the same focus, here.

(such as in generating *common random strings* (CRS) in the presence of adversaries [GK08,GGJS11,GO14]) and distributed computing (i.e., in the execution of *collective coin flipping protocols* [BOL85,CFG+85]).[4] In these applications, it is often crucial to construct extractors that have negligible error $\varepsilon = n^{-\omega(1)}$ [DOPS04]. This leads us to our main motivating question:

Can we construct explicit low-error extractors for adversarial sources?

Prior constructions

While the adversarial source model was first introduced in [CGGL20], it was implicitly studied in several prior works [KM04, KM05, GSV05, LLTT05, CL16]. There, it was shown that a two-source extractor can be applied in a black-box manner to construct extractors for adversarial sources with $K = 2$ good sources. Plugging in the best-known low-error two-source extractors [Bou05, Lew19], this gives low-error extractors for adversarial sources with $K = 2$ good sources of min-entropy roughly $k \geq \frac{4}{9}n$. In related work [BGK06, KRVZ11], it was shown how to extract from $K = O(1)$ good sources of min-entropy $k = 0.01n$.

Following these early papers, all subsequent work has focused on reducing the entropy requirement of the good sources (while using many more of them). In particular, Kamp, Rao, Vadhan and Zuckerman [KRVZ11] showed how to construct low-error extractors for adversarial sources with roughly $K = \sqrt{N}$ good sources of min-entropy $k \geq n^{0.99}$. Then, in [CGGL20], it was shown how to reduce the entropy requirement to just $k \geq \text{polylog}(n)$, while still using roughly $K = \sqrt{N}$ good sources. Finally, a follow-up work [CG21] showed how to extract from just $K = N^{0.01}$ good sources of min-entropy $k \geq \text{polylog}(n)$. To date, these remained the best-known low-error extractors for adversarial sources.

1.1 Our results

Extractors for adversarial sources

In this paper, we construct explicit low-error extractors for adversarial sources with just $K = 3$ good sources, dramatically improving on the previous best requirement of $K = N^{0.01}$. We prove the following.

Theorem 1 (Extractors for adversarial sources)**.** *There exist universal constants $C, \gamma > 0$ such that the following holds. There exists an explicit extractor $\mathsf{Ext} : (\{0,1\}^n)^N \to \{0,1\}^m$ for (N, K, n, k)-adversarial sources with $K = 3$ good sources of min-entropy $k \geq \log^C n$, which has output length $m = \lceil k^\gamma \rceil$ and error $\varepsilon = 2^{-\lceil k^\gamma \rceil}$.*

To summarize, the original paper on adversarial sources required $K = \sqrt{N}$ good sources [CGGL20], while the follow-up work [CG21] improved this requirement to $K = N^{0.01}$. In this work, we reduce the number of good sources required to just $K = 3$, no matter how many sources, N, there are in total (including super-constant and beyond). This essentially "finishes off" the problem of extracting from adversarial sources, as any further improvement is equivalent to constructing improved low-error two-source extractors – one of the foremost remaining challenges in extractor theory.[5]

[4] CRS's are useful, because they enable the construction of otherwise impossible crytographic primitives, such as *non-interactive zero-knowledge (NIZK) proofs* [BFM88] and *universally composable (UC) commitment schemes* [CF01]. And the ability to generate a CRS *in the presence of adversaries* has natural applications in blockchains, where leader election protocols must succeed even if some players are malicious.

[5] As alluded to in the footnote preceding Definition 1, there is still a more general notion of adversarial sources (where the bad sources can *depend* on the good sources) for which much exciting work remains to be done.

Theorem 1 has some interesting interpretations in the context of both extractors and cryptography. In the context of cryptography, this result gives a way for any number N of players to agree on a uniformly random string, provided that just 3 of them are honest. In the context of extractors, this result gives a way to extract from just 3 sources of weak randomness, hidden among any number of bad sources. In this way, it can be thought of as a significantly more robust version of Li's classical three-source extractor [Li15].

Finally, our construction is significantly simpler than prior constructions of extractors for adversarial sources. Indeed, Theorem 1 follows immediately from a powerful new object that we construct, called a *leakage-resilient extractor (LRE) against number-on-forehead (NOF) protocols*. Such objects lie directly at the intersection of extractor theory and communication complexity, and we discuss them next.

Leakage-resilient extractors against number-on-forehead protocols

LREs against NOFs are a brand new type of pseudorandom primitive, first introduced in [KMS19] (under the name *cylinder intersection extractors*). There are two equivalent ways to define LREs against NOFs: either as (1) a robust extractor for independent sources, or (2) a strong average-case lower bound against communication protocols. Let us take a moment to motivate and present definition (2). To do so, we must briefly recall a rich subfield of complexity theory, called *(multiparty) communication complexity* [Yao79, DF92, CFL83].

In multiparty communication complexity, the goal is to understand the amount of communication necessary to compute a function $f : (\{0,1\}^n)^N \to \{0,1\}^m$ when its input is split up among N parties. In the most well-studied model, the parties must determine the value of f via a *number-on-forehead* (NOF) protocol [CFL83] (see Definition 4). Here, each party $i \in [N]$ receives one piece of the input $x_i \in \{0,1\}^n$, which is metaphorically written on their forehead. Then, the parties take turns writing a bit on a public chalkboard, using all of the inputs *except* the one written on their forehead. The protocol continues until everyone knows the value of f, and the NOF communication complexity of f is the number of bits written on the chalkboard (in the worst-case over all possible inputs x, for the best possible protocol).

As it turns out, NOF protocols are an incredibly powerful computational model, capable of simulating many other well-studied models. This means that explicit lower bounds against NOF protocols immediately yield explicit lower bounds against important models such as ACC^0 circuits [Yao90, RW93, BT94] and low-degree polynomials over $\mathbb{F}_2$ [Vio09]. Because of this, researchers have spent a significant amount of effort trying to construct *explicit NOF lower bounds* (i.e., explicit functions that have high NOF communication complexity) [BNS92], in the hopes of achieving bounds that are strong enough to yield breakthrough lower bounds for these other computational models.

The most basic form of an explicit NOF lower bound that one might hope for is a *worst-case bound*. This is an explicit function f such that for every NOF protocol Π computing f, there is *some* input x for which Π must write many bits on the blackboard in order to compute $f(x)$.

An even stronger lower bound one might pursue is an *average-case bound*. This is an explicit function f such that for every NOF protocol Π computing f, it holds that for *many* inputs x, Π must write many bits μ on the blackboard in order to compute $f(x)$. Equivalently, for every protocol Π' that writes $< \mu$ bits on the blackboard (on every input), it holds that $\Pi'(\mathbf{X})$ is highly uncorrelated with $f(\mathbf{X})$, where $\mathbf{X} \sim (\{0,1\}^n)^N$ is a uniform random variable.

This leads us to the strongest of all lower bounds, which we might call a *$\mathcal{X}$-average-case lower bound*. This is an explicit function $f : (\{0,1\}^n)^N \to \{0,1\}^m$ such that for all protocols Π' that write $< \mu$ bits on the blackboard, and all sources $\mathbf{X} \in \mathcal{X}$, it holds that $\Pi'(\mathbf{X})$ is highly uncorrelated with $f(\mathbf{X})$.[6] Given that

[6] The average-case lower bound definition corresponds to the case where $\mathcal{X}$ consists of a single source, namely the uniform one.

each $\mathbf{X} \sim (\{0,1\}^n)^N$ consists of N parts $\mathbf{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_N)$, it is natural to take $\mathcal{X}$ to be the family of independent sources, where each part $\mathbf{X}_i \sim \{0,1\}^n$ has min-entropy k.

As it turns out, this is exactly the definition of an LRE against NOF protocols (see Definition 5 for a formal definition). Indeed, in order for f to be highly uncorrelated with Π', a trivial requirement is that f looks uniform on any $\mathbf{X} \in \mathcal{X}$ (for otherwise it is correlated with a trivial protocol that always outputs a constant). Moreover, it is not too difficult to see that the output $f(\mathbf{X})$ must look uniform *and remain uniform*, even conditioned on fixing the "NOF leakage" $\Pi'(\mathbf{X})$. In this way, since $\mathcal{X}$ denotes the family of independent sources, f can be thought of as a more robust type of extractor for independent sources – definition (1).

LREs against NOFs were first introduced (under the name *cylinder intersection extractors*) in the work [KMS19], and systematically studied in [CGG+20]. It has long been known [CGGL20] that if one could construct low-error LREs against NOFs for polylogarithmic min-entropy (and three parties), then one would immediately obtain Theorem 1. However, such an object seemed challenging to construct, as it must not only be as strong as Li's classical three-source extractor [Li15], but also offer powerful robustness guarantees. In this paper, we finally construct them.

Theorem 2 (LREs against NOFs)**.** *There exist universal constants $C, \gamma > 0$ such that the following holds. There exists an explicit three-source extractor $\mathsf{Ext} : (\{0,1\}^n)^3 \to \{0,1\}^m$ for min-entropy $k \geq \log^C n$, which is leakage-resilient against number-on-forehead protocols with $\mu = \lceil k^\gamma \rceil$ bits of communication, and which has output length $m = \lceil k^\gamma \rceil$ and error $\varepsilon = 2^{-\lceil k^\gamma \rceil}$.*

Prior to our work, all LREs against NOFs required high min-entropy, k. In the paper that introduced these objects [KMS19], the authors showed how to construct such LREs for min-entropy $k = 0.99n$, by observing that existing average-case NOF lower bounds can handle a little missing entropy in the inputs for free. In the follow-up work [CGG+20], the authors improve this entropy requirement to $k = 0.3n$, via intricate exponential sum estimates. In this paper, we reduce the entropy requirement to $k = \mathrm{polylog}(n)$, via a simple new reduction to *non-malleable extractors* (in fact, a very weak version of this object - see Definition 6).

Non-malleable extractors (NMEs) are a different, much more well-studied flavor of robust extractor, and were first introduced in [DW09, CG17]. We find the fact that there exists a reduction from LREs to NMEs to be quite surprising, for two reasons:

- First, prior to this work, LREs and NMEs were thought to be incomparable. Indeed, even though they are both robust flavors of randomness extractors, they are robust in seemingly incomparable ways. In particular, while an LRE must defend against *arbitrary* leaks, these leaks can only act on $N-1$ of the inputs. An NME, on the other hand, must only defend against "self-leaks" (where the leaks are the NME itself, called with slightly modified inputs), but these leaks can depend on all N of the inputs. This work shows that NMEs are, in fact, the stronger object.

- Second, since LREs against NOFs are simply a strong type of lower bound against NOF protocols (as discussed above), and we show that every non-malleable extractor is such an LRE, our work also shows that *non-malleable extractors witness strong NOF lower bounds*. In fact, the parameters we obtain show that NMEs achieve essentially the best-known lower bounds against NOF protocols [BNS92], and are even able to do so in the (stronger) low-entropy setting.[7] This provides a brand new

[7] Indeed, while Theorem 2 is stated for just three parties, our full result (Theorem 5) works for any number N of parties with min-entropy k, and achieves a lower bound of the form $k^{\Omega(1)}/2^N$ for any $k \geq \mathrm{polylog}(n)$. Since the current best lower bounds against NOF protocols are of the form $\Omega(n/2^N)$, our result comes close to matching these bounds, and does so in the stronger

family of black-box functions that are hard, and gives further support for a recent blossoming line of work on extractor-based lower bounds. In particular, prior to this work, it was known that basic extractors witness the best-known explicit lower bounds against general circuits [LY22], DNFs of parities [CS16], linear ROBPs [GPT22, CL23], and any model that shrinks under random restrictions [GG25]. In this paper, we show that essentially the best-known *NOF lower bounds* are also directly witnessed by basic extractors - namely, non-malleable extractors.

Finally, we note that the reduction from LREs against NOFs to NMEs is simple. In particular, it just involves viewing (and slightly rephrasing) standard Cauchy-Schwarz-based proofs of NOF lower bounds in the language of extractors.[8] Thus, beyond this insight, all of the work in constructing our low-entropy LREs against NOFs (Theorem 2) goes into constructing low-entropy NMEs, which is done via a composition of basic tools from extractor theory, combined with a slightly clever fixing argument.

Organization

The rest of the paper is organized as follows. In Section 2, we start with some basic preliminaries. Then, in Section 3, we give the reduction from LREs to NMEs. Then, we construct our explicit NMEs in Section 4. Finally, in Section 5, we instantiate the reduction with our explicit NMEs to obtain our explicit LREs (Theorem 2), and recall the standard reduction from extractors for adversarial sources to LREs to immediately obtain Theorem 1.

2 Preliminaries

Notation First, let us introduce some helpful notation. To start, for an integer $n \in \mathbb{N}$, we let $[n] := \{1, 2, \ldots, n\}$. Given a string $x \in \{0,1\}^n$, we let x_i denote the value it holds at its i^{th} coordinate, and for any subset $S \subseteq [n]$, we let $(x_i)_{i \in S}$ denote the concatenation of all bits x_i in increasing order of i. Speaking of concatenation, we sometimes represent it using the symbol $\circ$, in the sense that $x \circ y := (x, y)$. In particular, throughout this paper, $\circ$ does *not* denote function composition. Another perhaps unconventional notation we use throughout is that N, n will always denote *completely unrelated* integers. Namely, we *do not* take N to mean 2^n. (The same goes for K and k.)

Now, given a subset $S \subseteq [n]$, we write x_S as shorthand for $(x_i)_{i \in S}$, and for any $m \in [n]$, we write $x_{\leq m}$ as shorthand for $x_{[m]}$ (i.e., the first m bits of x), and define $x_{<m}$ and $x_{>m}$ and $x_{\geq m}$ in the analogous way. Next, for any $i \in [n]$ we define x_{-i} to mean $x_{[n] \setminus \{i\}}$. All of this notation extends naturally to larger alphabets. For example, for any alphabet Σ and $x \in \Sigma^n$, we let x_i denote the i^{th} symbol in x. If $\Sigma = \{0,1\}^m$, then x_i denotes the i^{th} consecutive chunk of m bits. Finally, for any positive integer t, we let $\mathbb{F}_{2^t}$ denote the finite field over 2^t elements, and whenever we write $x \cdot y$ for $x, y \in \{0,1\}^t$, we intend this to mean their product over $\mathbb{F}_{2^t}$. As usual, all logs will be base 2.

2.1 Probability

Throughout, we use bold font, such as $\mathbf{X}$, to denote random variables. The *support* of $\mathbf{X}$, denoted $\text{support}(\mathbf{X})$, contains all elements x such that $\Pr[\mathbf{X} = x] > 0$, and we write $\mathbf{X} \sim V$ if $\text{support}(\mathbf{X}) \subseteq V$.

"low-entropy" setting. (In the full min-entropy setting $k = n$, one can obtain "NME-based" NOF lower bounds that essentially match the state-of-the-art – for more details, see the full version.)

[8] In more detail, we show that NMEs (for independent sources) have small *cube norm*, which implies that they have large NOF communication complexity. A similar idea can be used to show that NMEs for affine sources (in fact, even just *directional affine extractors* [GPT22]) have small *Gowers norm*, and thus exhibit correlation bounds against low-degree $\mathbb{F}_2$-polynomials.

As per tradition in the extractor literature, we will often refer to random variables as *sources*. Furthermore, for any set V and $S \subseteq V$, we let $\mathbf{U}_S$ denote a uniform random variable over S, and use the shorthand $\mathbf{U}_m$ when $S = \{0,1\}^m$. That is, $\mathbf{U}_m$ is a uniform random variable over bitstrings of length m.

Next, when talking about extractors, the canonical distance that is used is called the *statistical distance*:

Definition 2 (Statistical distance)**.** *The* statistical distance *between two random variables $\mathbf{X}, \mathbf{Y} \sim V$ is*

$$|\mathbf{X} - \mathbf{Y}| := \max_{S \subseteq V} |\Pr[\mathbf{X} \in S] - \Pr[\mathbf{Y} \in S]| = \frac{1}{2} \sum_{v \in V} |\Pr[\mathbf{X} = v] - \Pr[\mathbf{Y} = v]|.$$

If $|\mathbf{X} - \mathbf{Y}| \leq \varepsilon$, we write $\mathbf{X} \approx_\varepsilon \mathbf{Y}$ and say $\mathbf{X}, \mathbf{Y}$ are ε-close*. Otherwise, we say they are* ε-far*. Moreover, if $|\mathbf{X} - \mathbf{Y}| = 0$, we write $\mathbf{X} \equiv \mathbf{Y}$.*

Next, we introduce three extremely useful facts about statistical distance. First, when dealing with extractors, the classical goal is to show that its output is close to uniform, $\mathbf{U}_m$. Towards this end, the following tool is invaluable. It says that two random variables cannot get further apart if you process them with the same deterministic function.

Fact 1 (Data-processing inequality)**.** *For any random variables $\mathbf{X}, \mathbf{Y} \sim V$, and any function $f : V \to W$,*

$$|\mathbf{X} - \mathbf{Y}| \geq |f(\mathbf{X}) - f(\mathbf{Y})|.$$

Second, it is well-known that the *triangle inequality* also holds for statistical distance:

Fact 2 (Triangle inequality)**.** *For any random variables $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$,*

$$|\mathbf{X} - \mathbf{Y}| \leq |\mathbf{X} - \mathbf{Z}| + |\mathbf{Z} - \mathbf{Y}|.$$

Third, the following fact is very useful when thinking about "fixing" random variables.

Fact 3 (Averaging principle)**.** *For any random variables $\mathbf{X}, \mathbf{Y} \sim V$ and $\mathbf{Z}, \mathbf{Z}' \sim Z$ such that $\mathbf{Z} \equiv \mathbf{Z}'$,*

$$|\mathbf{X} \circ \mathbf{Z} - \mathbf{Y} \circ \mathbf{Z}'| = \mathbb{E}_{z \sim \mathbf{Z}}[|(\mathbf{X} \mid \mathbf{Z} = z) - (\mathbf{Y} \mid \mathbf{Z}' = z)|].$$

Speaking of fixing random variables, the following standard lemma will be of great help whenever we want to argue that fixing one random variable doesn't cause another to lose too much min-entropy.

Lemma 1 (Min-entropy chain rule [MW97])**.** *For any random variables $\mathbf{X} \sim X, \mathbf{Y} \sim Y$ and any $\varepsilon > 0$,*

$$\Pr_{y \sim \mathbf{Y}}[H_\infty(\mathbf{X} \mid \mathbf{Y} = y) \geq H_\infty(\mathbf{X}) - \log|Y| - \log(1/\varepsilon)] \geq 1 - \varepsilon.$$

Finally, we record one last lemma, which (to the best of our knowledge) is new. The goal of this lemma is to formally "capture" (or describe) the randomness that is lost after applying a deterministic function f to a random variable $\mathbf{X}$. It strengthens and simplifies [CGG+20, Lemma 7].

Lemma 2 (Dependency reversal)**.** *For any random variable $\mathbf{X} \sim X$ and deterministic function $f : X \to Y$, there exists an independent random variable $\mathbf{A} \sim A$ and deterministic function $g : Y \times A \to X$ such that*

$$g(f(\mathbf{X}), \mathbf{A}) \equiv \mathbf{X}.$$

Proof. We define an independent random variable $\mathbf{A} \sim A := X^Y$ as a sequence of independent random variables $\mathbf{A} = (\mathbf{A}_y)_{y \in Y}$ where each $\mathbf{A}_y \sim X$ is defined as

$$\Pr[\mathbf{A}_y = x] := \Pr[\mathbf{X} = x \mid f(\mathbf{X}) = y] \text{ for all } x \in X.$$

If we then define the deterministic function $g : Y \times A \to V$ as

$$g(y, a) := a_y,$$

it directly follows that $g(f(\mathbf{X}), \mathbf{A}) \equiv \mathbf{X}$ via the law of total probability, as desired. □

2.2 Randomness condensers

We now recall the notion of *condensers*, which are slightly weaker versions of extractors. While an extractor guarantees that its output is close to *uniform* $\mathbf{U}_m$, a condenser just guarantees that its output is close to *having high min-entropy*. Formally, (two-source) condensers are defined as follows.

Definition 3 (Two-source condensers)**.** *A function* $\mathsf{2Cond} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ *is called an* $(n,k) \to_\varepsilon (m,\ell)$ two-source condenser[9] *if for any two independent sources* $\mathbf{X}, \mathbf{Y} \sim \{0,1\}^n$ *each with min-entropy at least* k*,* $\mathsf{2Cond}(\mathbf{X},\mathbf{Y})$ *is* ε*-close to some source* $\mathbf{Z} \sim \{0,1\}^m$ *with min-entropy at least* ℓ*. We say* 2Cond *is* strong *if*

$$\Pr_{y\sim\mathbf{Y}}[\mathsf{2Cond}(\mathbf{X},y) \text{ is not } \varepsilon\text{-close to a source with min-entropy at least } \ell] \le \varepsilon$$

and

$$\Pr_{x\sim\mathbf{X}}[\mathsf{2Cond}(x,\mathbf{Y}) \text{ is not } \varepsilon\text{-close to a source with min-entropy at least } \ell] \le \varepsilon.$$

As it turns out, every two-source condenser is also a *strong* two-source condenser, provided that we feed it two sources with a little more entropy. In particular, the following is true.

Fact 4 (Every two-source condenser is strong)**.** *For any* $n \ge k$ *and* $m \ge \ell$*, if* $\mathsf{2Cond} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ *is an* $(n,k) \to_\varepsilon (m,\ell)$ *two-source condenser, then it is a* strong $(n, k+m+\log(1/\varepsilon)) \to_\varepsilon (m,\ell)$ *two-source condenser.*

The proof of this fact is almost identical to a proof, due to Barak, that every two-source *extractor* is also a strong two-source extractor (with slightly weaker parameters) [Rao07]. For completeness, we include its proof below. Beyond the ideas in Barak's analogous proof for two-source extractors, we need the following fact, which gives a nice characterization of sources that are close to having high min-entropy. (See [GLZ24, Lemma 2] for a proof.)

Fact 5 (An alternative characterization for being close to high min-entropy [Zuc07, Lemma 2.2])**.** *For any* $n \ge k$*, a source* $\mathbf{X} \sim \{0,1\}^n$ *is* ε*-close to a source of min-entropy at least* k *iff for every* $S \subseteq \{0,1\}^n$*,*

$$\Pr[\mathbf{X} \in S] \le |S| \cdot 2^{-k} + \varepsilon.$$

With this characterization in hand, it is now straightforward to extend Barak's argument about two-source extractors [Rao07] to also work for two-source condensers, and thereby prove Fact 4.

Proof of Fact 4. We prove that for any independent sources $\mathbf{X}, \mathbf{Y} \sim \{0,1\}^n$ with min-entropy at least k',

$$\Pr_{y\sim\mathbf{Y}}[\mathsf{2Cond}(\mathbf{X},y) \text{ is not } \varepsilon\text{-close to a source with min-entropy at least } \ell] \le 2^{k+m-k'}$$

and

$$\Pr_{x\sim\mathbf{X}}[\mathsf{2Cond}(x,\mathbf{Y}) \text{ is not } \varepsilon\text{-close to a source with min-entropy at least } \ell] \le 2^{k+m-k'},$$

[9] Or equivalently, a two-source condenser with input entropy k, output entropy ℓ, output length m, and error ε.

and the result will follow by setting $k' = k + m + \log(1/\varepsilon)$. Actually, we will just prove the first claim, as the proof of the second claim is identical. Towards this end, for any $S \subseteq \{0,1\}^m$, define

$$\mathsf{Bad}_S := \{y \in \{0,1\}^m : \Pr[\mathsf{2Cond}(\mathbf{X}, y) \in S] > |S| \cdot 2^{-\ell} + \varepsilon\},$$

and notice that Fact 5 tells us that $\mathsf{2Cond}(\mathbf{X}, \mathbf{U}_{\mathsf{Bad}_S})$ is not ε-close to a source with min-entropy at least ℓ. We must therefore have $|\mathsf{Bad}_S| \leq 2^k$, or else 2Cond is not the condenser it claimed to be. Thus, if we define $\mathsf{Bad} := \bigcup_S \mathsf{Bad}_S$, we know that $|\mathsf{Bad}| \leq 2^{k+m}$. Now, again using Fact 5, we have

$$\begin{aligned}
&\Pr_{y \sim \mathbf{Y}}[\mathsf{2Cond}(\mathbf{X}, y) \text{ is not } \varepsilon\text{-close to an } (m, \ell) \text{ source}] \\
&= \Pr_{y \sim \mathbf{Y}}[\exists S \subseteq \{0,1\}^m : \Pr[\mathsf{2Cond}(\mathbf{X}, y) \in S] > |S| \cdot 2^{-\ell} + \varepsilon] \\
&= \Pr_{y \sim \mathbf{Y}}[y \in \mathsf{Bad}] \leq |\mathsf{Bad}| \cdot 2^{-k'} = 2^{k+m-k'},
\end{aligned}$$

as desired. □

2.3 Communication complexity

Finally, we provide a formal definition of *number-on-forehead* (NOF) multiparty communication protocols.

Definition 4 (Number-on-forehead protocol)**.** *A function $\Pi : (\{0,1\}^n)^N \to \{0,1\}^\mu$ is called a* number-on-forehead protocol *if for every $i \in [\mu]$ there exist functions $S_i : \{0,1\}^{i-1} \to [N]$ and $g_i : \{0,1\}^{i-1} \times (\{0,1\}^n)^{N-1} \to \{0,1\}$ such that*

$$\Pi_i(x) = g_i(\Pi_1(x), \ldots, \Pi_{i-1}(x), x_{-S_i(\Pi_1(x), \ldots, \Pi_{i-1}(x))}).$$

We say that Π is non-adaptive *if each g_i does not depend on its first $i-1$ inputs and each S_i is a constant.*

We say that a function $f : (\{0,1\}^n)^N \to \{0,1\}^m$ has ε-average-case NOF communication complexity $> \mu$ if for any NOF protocol Π with output length μ, it holds that $f(\mathbf{X}) \circ \Pi(\mathbf{X}) \approx_\varepsilon \mathbf{U} \circ \Pi(\mathbf{X})$, where $\mathbf{X}, \mathbf{U}$ are independent uniform random variables. Equivalently, we say that f is *ε-hard* for such protocols.

To conclude this section, we record the following lemma, which says that an average-case NOF-hard function is still hard even if the input distributions are missing a little bit of min-entropy.

Lemma 3 (NOF-hard functions can tolerate a little missing min-entropy [CGG+20, Lemma 3])**.** *Let $\mathsf{NOF} : (\{0,1\}^n)^N \to \{0,1\}^m$ be a function with ε-average-case NOF communication complexity $> \mu$. Then for any independent sources $\mathbf{X}_1, \ldots, \mathbf{X}_N \sim \{0,1\}^n$ each with min-entropy at least k, and any number-on-forehead protocol $\Pi : (\{0,1\}^n)^N \to \{0,1\}^{\mu-2}$, it holds that*

$$|\mathsf{NOF}(\mathbf{X}_1, \ldots, \mathbf{X}_N) \circ \Pi(\mathbf{X}_1, \ldots, \mathbf{X}_N) - \mathbf{U}_m \circ \Pi(\mathbf{X}_1, \ldots, \mathbf{X}_N)| \leq \varepsilon \cdot 2^{N(n-k)}.$$

As we will soon see (via Definition 5), another way to summarize this lemma is that any (average-case) NOF-hard function is automatically a basic LRE against NOF protocols, which can handle slightly less than full min-entropy.

3 A reduction from leakage-resilient extractors to non-malleable extractors

In order to construct our extractors for adversarial sources with $K = 3$ good sources (Theorem 1), we will build a three-source leakage-resilient extractor (LRE) against number-on-forehead protocols (Theorem 2), and the adversarial source extractor will easily follow (simply by calling the LRE on all triples of sources and XORing the results). To get started, we formally introduce LREs against NOFs.

Definition 5 (LREs against NOFs)**.** *A function* $\mathsf{LRE} : (\{0,1\}^n)^N \to \{0,1\}^m$ *is called a* leakage-resilient extractor *for min-entropy* k *and error* ε *against number-on-forehead protocols with* μ *bits of communication if the following holds. For any independent sources* $\mathbf{X}_1, \ldots, \mathbf{X}_N \sim \{0,1\}^n$ *each with min-entropy at least* k*, and any number-on-forehead protocol* $\Pi : (\{0,1\}^n)^N \to \{0,1\}^\mu$*,*

$$\mathsf{LRE}(\mathbf{X}_1, \ldots, \mathbf{X}_N) \circ \Pi(\mathbf{X}_1, \ldots, \mathbf{X}_N) \approx_\varepsilon \mathbf{U}_m \circ \Pi(\mathbf{X}_1, \ldots, \mathbf{X}_N).$$

LREs against NOFs were first introduced by Kumar, Meka and Sahai (under the name *cylinder intersection extractors*) [KMS19], and systematically studied for the first time in [CGG+20] (where the more general notion of LREs was introduced). Looking at Definition 5, it is easy to see that these are powerful objects, which simultaneously generalize both (1) extractors for independent sources and (2) average-case lower bounds against NOF protocols.[10] In this section, we take our first step towards constructing them.

In order to build our LREs against NOFs, we provide a simple reduction from these objects to a new, weak type of non-malleable extractor, which we define below. Then, we will show how to build one.

Definition 6 (Weak NMEs)**.** *A function* $\mathsf{nmExt} : (\{0,1\}^n)^N \to \{0,1\}^m$ *is called a* weak non-malleable extractor *for min-entropy* k *and error* ε *if the following holds. For any independent sources* $\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0, \mathbf{X}_1^1, \ldots, \mathbf{X}_N^1 \sim \{0,1\}^n$*, each with min-entropy at least* k*, and* $\mathbf{X}_i^0 \equiv \mathbf{X}_i^1$ *for all* $i \in [N]$*,*

$$\mathsf{nmExt}(\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0) \circ \left(\mathsf{nmExt}(\mathbf{X}_1^{b_1}, \ldots, \mathbf{X}_N^{b_N})\right)_{b \neq \vec{0} \in \{0,1\}^N} \approx_\varepsilon \mathbf{U}_m \circ \left(\mathsf{nmExt}(\mathbf{X}_1^{b_1}, \ldots, \mathbf{X}_N^{b_N})\right)_{b \neq \vec{0} \in \{0,1\}^N}.$$

With these definitions in hand, we proceed with the first main lemma of this paper, which shows that weak NMEs are, in fact, LREs against NOFs.

Lemma 4 (Weak NMEs $\implies$ LREs against NOFs)**.** *Let* $\mathsf{nmExt} : (\{0,1\}^n)^N \to \{0,1\}^m$ *be a weak non-malleable extractor for min-entropy* k *and error* ε*. Then* nmExt *is a leakage-resilient extractor against number-on-forehead protocols with* μ *bits of communication for min-entropy* k *and error* $2^{m+\mu} \cdot (2\varepsilon)^{1/2^N}$*.*

Proof. The proof, in fact, is not too difficult. It follows from observing that the standard technique for obtaining NOF lower bounds (i.e., via the cube norm) both (1) automatically work for low entropy, and (2) can be interpreted as seeking out the non-malleability property put forth by Definition 6. Indeed, the novelty of the proof is just a shift in perspective.

To make things more formal, let $\mathsf{nmExt} : (\{0,1\}^n)^N \to \{0,1\}^m$ be the weak non-malleable extractor from the lemma statement. We want to upper bound the statistical distance

$$|\mathsf{nmExt}(\mathbf{X}) \circ \Pi(\mathbf{X}) - \mathbf{U}_m \circ \Pi(\mathbf{X})|,$$

where Π is an NOF protocol of length μ, and $\mathbf{X}$ consists of N independent sources $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_N)$, each over n bits and each with min-entropy at least k. The first issue is that lower-bounding the above expression amounts to achieving "multi-bit" hardness, whereas standard communication complexity seems to

[10] Indeed, setting $\mu = 0$ recovers object (1), while setting $k = n$ recovers object (2).

focus on proving hardness for Boolean-valued functions. However, there are standard tools from cryptography and extractor theory that can be exploited to get such a result. Here, we will use a "non-uniform" XOR lemma of Dodis, Li, Wooley, and Zuckerman [DLWZ14, Lemma 3.8], which says that we can upper bound the above statistical distance by $2^{m+\mu}\alpha$, provided we can show that

$$\left| \left(\bigoplus_{i \in S} \mathsf{nmExt}_i(\mathbf{X}) \oplus \bigoplus_{j \in T} \Pi_j(\mathbf{X}) \right) - \mathbf{U}_1 \right| \leq \alpha,$$

for all $S \neq \emptyset \subseteq [m], T \subseteq [\mu]$. Fix one such choice of S, T, and define $f(x) := \bigoplus_{i \in S} \mathsf{nmExt}_i(x)$ and $g(x) := \bigoplus_{j \in T} \Pi_j(x)$. By the data-processing inequality (Fact 1), the above statistical distance can be upper bounded by

$$|f(\mathbf{X}) \circ g(\mathbf{X}) - \mathbf{U}_1 \circ g(\mathbf{X})|,$$

where $\mathbf{U}_1$ is independent of everything else. Applying the data-processing inequality once more, the above can further be upper bounded by

$$|f(\mathbf{X}) \circ \Pi(\mathbf{X}) - \mathbf{U}_1 \circ \Pi(\mathbf{X})|.$$

The goal is now to upper bound the above, by showing that if nmExt started out as a weak non-malleable extractor, then $f(\mathbf{X})$ has low correlation with all NOF protocols. To show this, we can simply repeat known arguments originating in [BNS92][11] through the lens of extractors.

To get started, note that

$$\begin{aligned}
2|f(\mathbf{X}) \circ \Pi(\mathbf{X}) - \mathbf{U}_1 \circ \Pi(\mathbf{X})| &= 2\mathbb{E}_{\pi \sim \Pi(\mathbf{X})} |(f(\mathbf{X}) \mid \Pi(\mathbf{X}) = \pi) - \mathbf{U}_1| \\
&= \mathbb{E}_{\pi \sim \Pi(\mathbf{X})} |\Pr[f(\mathbf{X}) = 1 \mid \Pi(\mathbf{X}) = \pi] - \Pr[f(\mathbf{X}) = 0 \mid \Pi(\mathbf{X}) = \pi]| \\
&\leq 2^{\mu} \max_{\pi} |\Pr[f(\mathbf{X}) = 1 \wedge \Pi(\mathbf{X}) = \pi] - \Pr[f(\mathbf{X}) = 0 \wedge \Pi(\mathbf{X}) = \pi]| \\
&= 2^{\mu} \cdot |\Pr[f(\mathbf{X}) = 1 \wedge \Pi(\mathbf{X}) = \pi] - \Pr[f(\mathbf{X}) = 0 \wedge \Pi(\mathbf{X}) = \pi]|,
\end{aligned}$$

where the last line just fixes the worst transcript π. Now, define the function $f' := (-1)^f$ which swaps the output domain of f from $\{0,1\}$ to $\{-1,1\}$, and define the function $C : (\{0,1\}^n)^N \to \{0,1\}$ as the indicator of $\Pi(x) = \pi$. Namely, $C(x) = 1$ iff $\Pi(x) = \pi$. Using this, the most recent equation above is

$$= 2^{\mu} \cdot |\mathbb{E}[f'(\mathbf{X}) \cdot C(\mathbf{X})]|.$$

We now focus on upper bounding $|\mathbb{E}[f'(\mathbf{X}) \cdot C(\mathbf{X})]|$. This part of the argument is now standard, and we follow the exposition of Lovett [Lov19]. Towards this end, it is well-known (and not too hard to show) that C is the indicator of a "cylinder intersection" [BNS92]. This means that there exist indicator functions $C_1, \ldots, C_N : (\{0,1\}^n)^{N-1} \to \{0,1\}$ so that $C(x) = C_1(x_{-1}) \cdot C_2(x_{-2}) \cdot \cdots \cdot C_N(x_{-N})$. Thus,

$$|\mathbb{E}[f'(\mathbf{X}) \cdot C(\mathbf{X})]| = |\mathbb{E}[f'(\mathbf{X}) \cdot C_1(\mathbf{X}_{-1}) \cdot C_2(\mathbf{X}_{-2}) \cdot \cdots \cdot C_N(\mathbf{X}_{-N})]|$$

Now, the main idea introduced in [BNS92] is to use N iterations of the Cauchy-Schwarz inequality (for random variables) to get rid of the "NOF leaks" $C_1, \ldots, C_N$, and replace them with, in our language, "weak non-malleable leaks." In more detail, recall that the Cauchy-Schwarz inequality tells us that for any

[11] Our argument will follow along with the one that appears in the excellent lecture notes of Lovett [Lov19].

real-valued random variables $\mathbf{A}, \mathbf{B}$ it holds that $|\mathbb{E}[\mathbf{A} \cdot \mathbf{B}]| \leq \left(\mathbb{E}[\mathbf{A}^2] \cdot \mathbb{E}[\mathbf{B}^2]\right)^{1/2}$. Applying this once, we can continue as follows.

$$
\begin{aligned}
&= \left| \mathbb{E}_{\mathbf{X}_{-1}} \left[C_1(\mathbf{X}_{-1}) \cdot \mathbb{E}_{\mathbf{X}_1}[f'(\mathbf{X}) \cdot C_2(\mathbf{X}_{-2}) \cdot \cdots \cdot C_N(\mathbf{X}_{-N})] \right] \right| \\
&\leq \left(\mathbb{E}_{\mathbf{X}_{-1}} \left[\left(C_1(\mathbf{X}_{-1}) \right)^2 \right] \cdot \mathbb{E}_{\mathbf{X}_{-1}} \left[\left(\mathbb{E}_{\mathbf{X}_1}[f'(\mathbf{X}) \cdot C_2(\mathbf{X}_{-2}) \cdot \cdots \cdot C_N(\mathbf{X}_{-N})] \right)^2 \right] \right)^{1/2} \\
&\leq \left(1 \cdot \mathbb{E}_{\mathbf{X}_{-1}} \left[\mathbb{E}_{\mathbf{X}_1, \mathbf{X}_1'}[f'(\mathbf{X}) \cdot C_2(\mathbf{X}_{-2}) \cdots C_N(\mathbf{X}_{-N}) \cdot f'(\mathbf{X}_{-1}, \mathbf{X}_1') \cdot C_2(\mathbf{X}_{-1,-2}, \mathbf{X}_1') \cdots C_N(\mathbf{X}_{-1,-N}, \mathbf{X}_1')] \right] \right)^{1/2} \\
&= \left(\mathbb{E}_{\mathbf{X}, \mathbf{X}_1'} \left[f'(\mathbf{X}) \cdot C_2(\mathbf{X}_{-2}) \cdots C_N(\mathbf{X}_{-N}) \cdot f'(\mathbf{X}_{-1}, \mathbf{X}_1') \cdot C_2(\mathbf{X}_{-1,-2}, \mathbf{X}_1') \cdots C_N(\mathbf{X}_{-1,-N}, \mathbf{X}_1') \right] \right)^{1/2}.
\end{aligned}
$$

The first inequality is an application of the Cauchy-Schwarz inequality, where the random variable $\mathbf{B}$ looks like an expectation. The second inequality is because the range of C_1 is $\{0, 1\}$ (which, in particular, is at most 1 when squared), and because $\mathbf{X}_1, \mathbf{X}_1'$ are independent (and identically distributed). The last equality is immediate.

Notice that by using the Cauchy-Schwarz inequality once, we were able to "remove" or "condition away" one NOF leak $C_1(\mathbf{X}_{-1})$, and effectively replace it with a "non-malleable leak" $f'(\mathbf{X}_{-1}, \mathbf{X}_1')$ (i.e., instead of an arbitrary function leaking on some of the inputs, we turn it into the specific leak f', acting on all inputs, or at least independent copies of them). To get the Cauchy-Schwarz inequality to work (to remove this NOF leak), we pulled out an expectation over the random variables on which it depends $\mathbf{X}_{-1}$, leaving just $\mathbf{X}_1$ in the internal expectation. When we execute it again, we'll once again pull out almost all random variables, this time leaving just $\mathbf{X}_2$ in the internal expectation. This will let us get rid of the next NOF leak, $C_2(\mathbf{X}_{-2})$. Thus, pulling out $\mathbf{X}_{-2}, \mathbf{X}_1'$ and leaving in $\mathbf{X}_2$, we can follow the argument above once again to obtain

$$
\begin{aligned}
\leq \Bigg(\mathbb{E}_{\mathbf{X}, \mathbf{X}_1', \mathbf{X}_2'} \Big[& f'(\mathbf{X}) \cdot C_3(\mathbf{X}_{-3}) \cdots C_N(\mathbf{X}_{-N}) \cdot \\
& f'(\mathbf{X}_{-1}, \mathbf{X}_1') \cdot C_3(\mathbf{X}_{-3,-1}, \mathbf{X}_1') \cdots C_N(\mathbf{X}_{-N,-1}, \mathbf{X}_1') \cdot \\
& f'(\mathbf{X}_{-2}, \mathbf{X}_2') \cdot C_3(\mathbf{X}_{-3,-2}, \mathbf{X}_2') \cdots C_N(\mathbf{X}_{-N,-2}, \mathbf{X}_2') \cdot \\
& f'(\mathbf{X}_{-1,-2}, \mathbf{X}_1', \mathbf{X}_2') \cdot C_3(\mathbf{X}_{-3,-2,-1}, \mathbf{X}_2', \mathbf{X}_1') \cdot \cdots C_N(\mathbf{X}_{-N,-2,-1}, \mathbf{X}_2', \mathbf{X}_1') \Big] \Bigg)^{1/2^2}.
\end{aligned}
$$

As the notation is getting cumbersome, it will be convenient to switch to something different. In particular, we define random variables $\mathbf{X}_1^0, \mathbf{X}_2^0, \ldots, \mathbf{X}_N^0, \mathbf{X}_1^1, \mathbf{X}_2^1, \ldots, \mathbf{X}_N^1$ that are all independent and such that $\mathbf{X}_i \equiv \mathbf{X}_i^0 \equiv \mathbf{X}_i^1$. Then, for a fixed string $b \in \{0,1\}^N$, we let $\mathbf{X}^b$ denote $(\mathbf{X}_1^{b_1}, \ldots, \mathbf{X}_N^{b_N})$. With this notation, the above expression is much easier to write:

$$
= \mathbb{E} \left[\prod_{\substack{b \in \{0,1\}^N \\ : b_i = 0 \forall i > 2}} f'(\mathbf{X}^b) \cdot C_3(\mathbf{X}_{-3}^b) \cdot C_4(\mathbf{X}_{-4}^b) \cdots C_N(\mathbf{X}_{-N}^b) \right]^{1/4}
$$

Finally, repeating the argument with the remaining leaks $\{C_i\}_{i \geq 3}$ immediately gives

$$
\begin{aligned}
&\leq \left(\mathbb{E}\left[\prod_{b\in\{0,1\}^N} f'(\mathbf{X}^b)\right]\right)^{1/2^N} \\
&= \left(\Pr\left[\prod_{b\in\{0,1\}^N} f'(\mathbf{X}^b) = 1\right] - \Pr\left[\prod_{b\in\{0,1\}^N} f'(\mathbf{X}^b) = -1\right]\right)^{1/2^N} \\
&= \left(\Pr\left[\bigoplus_{b\in\{0,1\}^N} f(\mathbf{X}^b) = 0\right] - \Pr\left[\bigoplus_{b\in\{0,1\}^N} f(\mathbf{X}^b) = 1\right]\right)^{1/2^N} \\
&= \left(2\left|\bigoplus_{b\in\{0,1\}^N} f(\mathbf{X}^b) - \mathbf{U}_1\right|\right)^{1/2^N} \\
&\leq \left(2\cdot\left|f(\mathbf{X})\circ(f(\mathbf{X}^b))_{b\neq 0\in\{0,1\}^N} - \mathbf{U}_1\circ(f(\mathbf{X}^b))_{b\neq 0\in\{0,1\}^N}\right|\right)^{1/2^N}
\end{aligned}
$$

where the last line is a simple application of the data-processing inequality. This almost looks like the weak non-malleable extractor definition, but recall f itself is of the form $f(x) := \oplus_{i\in S}\mathsf{nmExt}_i(x)$ for some nonempty S. Since this is just an XOR of some output bits of nmExt, it is clearly a deterministic function of it, and thus we may apply the data-processing inequality (Fact 1) once more to upper bound the above by

$$
\leq \left(2\cdot\left|\mathsf{nmExt}(\mathbf{X})\circ(\mathsf{nmExt}(\mathbf{X}^b))_{b\neq 0\in\{0,1\}^N} - \mathbf{U}_m\circ(\mathsf{nmExt}(\mathbf{X}^b))_{b\neq 0\in\{0,1\}^N}\right|\right)^{1/2^N},
$$

and since we were given that nmExt is a weak-non-malleable extractor for min-entropy k and error ε, it follows that this is

$$
\leq (2\varepsilon)^{1/2^N}.
$$

To summarize, we have finished showing that

$$
\left|\left(\bigoplus_{i\in S}\mathsf{nmExt}_i(\mathbf{X})\oplus\bigoplus_{j\in T}\Pi_j(\mathbf{X})\right) - \mathbf{U}_1\right| \leq (2\varepsilon)^{1/2^N}.
$$

As discussed at the beginning of the proof, the non-uniform XOR lemma allows us to therefore conclude

$$
|\mathsf{nmExt}(\mathbf{X})\circ\Pi(\mathbf{X}) - \mathbf{U}_m\circ\Pi(\mathbf{X})| \leq 2^{m+\mu}\cdot(2\varepsilon)^{1/2^N},
$$

as desired. □

4 An explicit construction of weak non-malleable extractors

Now that we know that weak NMEs automatically give LREs against NOFs, we turn our attention to explicitly constructing them. We will prove the following lemma.

Lemma 5 (Explicit weak NMEs)**.** *There exist universal constants $C, \gamma > 0$ such that the following holds. There exists an explicit weak non-malleable extractor* $\mathsf{nmExt} : (\{0,1\}^n)^N \to \{0,1\}^m$ *for min-entropy $k \geq \log^C n$, which has output length $m = \lceil k^\gamma / 2^N \rceil$ and error $\varepsilon = 2^{-\lceil k^\gamma / 2^N \rceil}$.*

In order to prove this lemma, we show a simple way to construct weak NMEs from basic, black-box tools. As it turns out, one of these tools is *itself* a function that is hard against number-on-forehead protocols! Thus, we will have traveled from LREs against NOFs, to weak NMEs, and back to functions that are hard against NOF protocols (but for one less party). As a result, we ultimately provide a way to convert a black-box function that is hard for NOF protocols on $N-1$ parties into a function that is hard for NOF protocols on N parties. By iterating this hardness amplification down to just 2 parties, *we obtain a reduction that can convert any function that is hard for two-party communication protocols into a function that is hard for NOF multiparty communication protocols.* We find this reduction to be surprising, and believe it may be of independent interest. In more detail, we prove the following.

Lemma 6 (A recipe for weak NMEs)**.** *Let* $\mathsf{2Cond} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^r$ *be a two-source condenser for input entropy k_0, output entropy k_1, and error ε_1. Let* $\mathsf{NOF} : (\{0,1\}^r)^{N-1} \to \{0,1\}^m$ *be a function with ε_2-average-case NOF communication complexity $> \mu_2 := 2^N m$ against non-adaptive number-on-forehead protocols. Then the function* $\mathsf{nmExt} : (\{0,1\}^n)^N \to \{0,1\}^m$ *defined as*

$$\mathsf{nmExt}(x_1, \ldots, x_N) := \mathsf{NOF}(\mathsf{2Cond}(x_1, x_N), \mathsf{2Cond}(x_2, x_N), \ldots, \mathsf{2Cond}(x_{N-1}, x_N))$$

is a weak non-malleable extractor for entropy $k = k_0 + 2r + 2\log(1/\varepsilon_1)$ and error $\varepsilon = 2N\varepsilon_1 + 2^{N(r-k_1)}\varepsilon_2$.

Now, in order to actually obtain our explicit weak NMEs, we will just instantiate the above lemma with well-known NOF lower bounds and two-source condensers. It is worth noting that the NOF-hard function is just called over $N-1$ parties. Thus, in order to construct the NMEs that we actually use for our three-source LREs, we only need an NOF-hard function over 2 parties. In other words – just a function that is hard in the standard two-party communication model (such as any two-source extractor, like the inner product function). We proceed with the proof of the recipe, before instantiating it to get our weak NMEs.

Proof of Lemma 6. Let $\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0, \mathbf{X}_1^1, \ldots, \mathbf{X}_N^1 \sim \{0,1\}^n$ be a collection of independent sources, each with min-entropy at least k, such that $\mathbf{X}_i^0 \equiv \mathbf{X}_i^1$ for all $i \in [N]$. For convenience, we refer to $\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0$ as the *good sources*, and we refer to $\mathbf{X}_1^1, \ldots, \mathbf{X}_N^1$ as the *shadows*. The goal is to upper bound

$$\Big| \mathsf{nmExt}(\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0) \circ \Big(\mathsf{nmExt}(\mathbf{X}_1^{b_1}, \ldots, \mathbf{X}_N^{b_N}) \Big)_{b \neq \vec{0} \in \{0,1\}^N} - \mathbf{U}_m \circ \Big(\mathsf{nmExt}(\mathbf{X}_1^{b_1}, \ldots, \mathbf{X}_N^{b_N}) \Big)_{b \neq \vec{0} \in \{0,1\}^N} \Big|.$$

To do so, let's start by making this expression less cumbersome. Towards this end, for each $b \in \{0,1\}^N$ we define a random variable

$$\mathbf{Z}^b := \mathsf{nmExt}(\mathbf{X}_1^{b_1}, \ldots, \mathbf{X}_N^{b_N}),$$

so that we may rewrite the original expression as

$$\Big| \mathsf{nmExt}(\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0) \circ \Big(\mathbf{Z}^b \Big)_{b \neq \vec{0} \in \{0,1\}^N} - \mathbf{U}_m \circ \Big(\mathbf{Z}^b \Big)_{b \neq \vec{0} \in \{0,1\}^N} \Big|. \tag{1}$$

Now, in order to upper bound Equation (1), our analysis will proceed in four phases, in which we gradually fix randomness until a strong bound comes for free (from the ingredients that make up nmExt).

Phase I: Fixing the shadows. The first step is to fix the shadows $\mathbf{X}_1^1, \ldots, \mathbf{X}_N^1$. As these are independent of the good sources $\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0$, this fixing does not affect the latter, and can only make our expression simpler. More formally, using standard fixing arguments, we know there exist fixed strings $x_1^1, \ldots, x_N^1 \in \{0,1\}^n$ such that if we define $\tilde{\mathbf{Z}}^b := (\mathbf{Z}^b \mid \mathbf{X}_1^1 = x_1^1, \ldots, \mathbf{X}_N^1 = x_N^1)$ for each $b \in \{0,1\}^N$, we can upper bound Equation (1) by

$$\left| \mathsf{nmExt}(\mathbf{X}_1^0, \ldots, \mathbf{X}_N^0) \circ \left(\tilde{\mathbf{Z}}^b\right)_{b \neq \vec{0} \in \{0,1\}^N} - \mathbf{U}_m \circ \left(\tilde{\mathbf{Z}}^b\right)_{b \neq \vec{0} \in \{0,1\}^N} \right|. \tag{2}$$

As a sanity check, note that $\tilde{\mathbf{Z}}^{\vec{1}}$ is now a fixed constant $z \in \{0,1\}^m$.

Phase II: Fixing the 1-local condenser calls. Now, recalling the definition of nmExt from the current lemma statement, the next step is to fix $\mathsf{2Cond}(\mathbf{X}_1^0, x_N^1), \mathsf{2Cond}(\mathbf{X}_2^0, x_N^1), \ldots, \mathsf{2Cond}(\mathbf{X}_{N-1}^0, x_N^1)$. Since each output is a deterministic function of a single good source, this will not introduce any correlations. Furthermore, since each condenser doesn't output too many bits, it will decrease the entropy of each good source by just a little. More formally, by combining standard fixing arguments with the min-entropy chain rule (Lemma 1) and a simple union bound, the following holds. There exist fixed strings $y_1, \ldots, y_{N-1} \in \{0,1\}^r$ such that if we define (for each $b \in \{0,1\}^N$ and $i \in [N]$) the random variables

$$\tilde{\tilde{\mathbf{Z}}}^b := \left(\tilde{\mathbf{Z}}^b \mid \mathsf{2Cond}(\mathbf{X}_1^0, x_N^1) = y_1, \mathsf{2Cond}(\mathbf{X}_2^0, x_N^1) = y_2, \ldots, \mathsf{2Cond}(\mathbf{X}_{N-1}^0, x_N^1) = y_{N-1}\right),$$
$$\tilde{\tilde{\mathbf{X}}}_i^0 := \left(\mathbf{X}_i^0 \mid \mathsf{2Cond}(\mathbf{X}_1^0, x_N^1) = y_1, \mathsf{2Cond}(\mathbf{X}_2^0, x_N^1) = y_2, \ldots, \mathsf{2Cond}(\mathbf{X}_{N-1}^0, x_N^1) = y_{N-1}\right),$$

then we can upper bound Equation (2) by

$$\left| \mathsf{nmExt}(\tilde{\tilde{\mathbf{X}}}_1^0, \ldots, \tilde{\tilde{\mathbf{X}}}_N^0) \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \neq \vec{0} \in \{0,1\}^N} - \mathbf{U}_m \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \neq \vec{0} \in \{0,1\}^N} \right| + N\delta, \tag{3}$$

where each $\tilde{\tilde{\mathbf{X}}}_i^0$ is independent and has min-entropy at least $k - r - \log(1/\delta)$. Now, notice that for every $b \in \{0,1\}^N$ with $b_N = 1$, it actually holds that $\tilde{\tilde{\mathbf{Z}}}^b$ has been fixed to some constant $z^b \in \{0,1\}^m$. Indeed, unwrapping the definition $\tilde{\tilde{\mathbf{Z}}}^b$ and then nmExt, observe that every such $\tilde{\tilde{\mathbf{Z}}}^b$ is actually just a deterministic function of two-source condenser calls of the form $\mathsf{2Cond}(\mathbf{X}_i^0, x_N^1)$ (which have been fixed to constants) or $\mathsf{2Cond}(x_i^1, x_N^1)$ (which are clearly constant). Thus Equation (3) can be rewritten as

$$\Bigg| \mathsf{nmExt}(\tilde{\tilde{\mathbf{X}}}_1^0, \ldots, \tilde{\tilde{\mathbf{X}}}_N^0) \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \in \{0,1\}^N : b \neq \vec{0}, b_N = 0} \circ \left(z^b\right)_{b \in \{0,1\}^N : b \neq \vec{0}, b_N = 1}$$
$$- \mathbf{U}_m \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \in \{0,1\}^N : b \neq \vec{0}, b_N = 0} \circ \left(z^b\right)_{b \in \{0,1\}^N : b \neq \vec{0}, b_N = 1} \Bigg| + N\delta.$$

Applying the data-processing inequality (Fact 1), this is at most

$$\left| \mathsf{nmExt}(\tilde{\tilde{\mathbf{X}}}_1^0, \ldots, \tilde{\tilde{\mathbf{X}}}_N^0) \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \in \{0,1\}^N : b \neq \vec{0}, b_N = 0} - \mathbf{U}_m \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \in \{0,1\}^N : b \neq \vec{0}, b_N = 0} \right| + N\delta. \tag{4}$$

Now, let's analyze the remaining $\tilde{\tilde{\mathbf{Z}}}^b$. Notice that since $b \in \{0,1\}^N$ satisfies both $b \neq \vec{0}$ and $b_N = 0$, there must be some $i \in [N-1]$ such that $b_i \neq 0$. Since we have fixed all the shadows, this means that $\tilde{\tilde{\mathbf{Z}}}^b$ is

a deterministic function of $\tilde{\tilde{\mathbf{X}}}^0_{-i}$. To make this more formal, we partition the remaining $b \in \{0,1\}^N$ into buckets according to which coordinate is nonzero. In particular, define for each $i \in [N-1]$ the set

$$B_i := \{b \in \{0,1\}^n : b_i = 1, b_N = 0\}.$$

To make this into an actual partition, simply define $B'_i := B_i \setminus (\cup_{h<i} B_h)$ to only include the strings not yet accounted for. Using this partition, we can rewrite Equation (4) as

$$\left| \mathsf{nmExt}(\tilde{\tilde{\mathbf{X}}}^0_1, \ldots, \tilde{\tilde{\mathbf{X}}}^0_N) \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \in B'_i, i \in [N-1]} - \mathbf{U}_m \circ \left(\tilde{\tilde{\mathbf{Z}}}^b\right)_{b \in B'_i, i \in [N-1]} \right| + N\delta.$$

Furthermore, recall from above that if $b \in B'_i$, then $b_i = 1$, which means that $\tilde{\tilde{\mathbf{Z}}}^b$ is a deterministic function of $\tilde{\tilde{\mathbf{X}}}^0_{-i}$. Moreover, note that each $\tilde{\tilde{\mathbf{Z}}}^b$ is over m bits, and the total number of elements across all the B'_i is at most 2^{N-1}. Thus there exist deterministic functions, $g_1, \ldots, g_{N-1}$, which output $2^{N-1}m$ bits in total, such that the above expression is exactly

$$\left| \mathsf{nmExt}(\tilde{\tilde{\mathbf{X}}}^0_1, \ldots, \tilde{\tilde{\mathbf{X}}}^0_N) \circ \left(g_i\left(\tilde{\tilde{\mathbf{X}}}^0_{-i}\right)\right)_{i \in [N-1]} - \mathbf{U}_m \circ \left(g_i\left(\tilde{\tilde{\mathbf{X}}}^0_{-i}\right)\right)_{i \in [N-1]} \right| + N\delta. \tag{5}$$

Finally, we are ready to proceed to the next phase.

Phase III: Boosting the min-entropy. In this phase, our goal is to execute the two-source condenser calls baked into nmExt in order to boost the min-entropy of the sources $\tilde{\tilde{\mathbf{X}}}^0_i$, and prepare them for the final phase, in which we call an NOF-hard function that expects to receive uniformly random input. Towards this end, define for each $i \in [N-1]$ the random variable

$$\mathbf{Y}_i := \mathsf{2Cond}(\tilde{\tilde{\mathbf{X}}}^0_i, \tilde{\tilde{\mathbf{X}}}^0_N),$$

so that we may rewrite Equation (5) as

$$\left| \mathsf{NOF}(\mathbf{Y}_1, \ldots, \mathbf{Y}_{N-1}) \circ \left(g_i\left(\tilde{\tilde{\mathbf{X}}}^0_{-i}\right)\right)_{i \in [N-1]} - \mathbf{U}_m \circ \left(g_i\left(\tilde{\tilde{\mathbf{X}}}^0_{-i}\right)\right)_{i \in [N-1]} \right| + N\delta. \tag{6}$$

Then, we recall that since $\mathsf{2Cond} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^r$ is a two-source condenser for input entropy k_0, output entropy k_1, and error ε_1, it is automatically a *strong* two-source condenser for input entropy $k_0 + r + \log(1/\varepsilon_1)$, output entropy k_1, and error ε_1 (Fact 4). Thus, by standard fixing arguments, the definition of strong condenser (Definition 3), and a simple union bound, the following holds. We may fix $\tilde{\tilde{\mathbf{X}}}^0_N$ to a string $x^0_N \in \{0,1\}^n$ such that if we define $\tilde{\mathbf{Y}}_i := \mathsf{2Cond}(\tilde{\tilde{\mathbf{X}}}^0_i, x^0_N)$ for each $i \in [N-1]$, then we may upper bound Equation (6) by

$$\left| \mathsf{NOF}(\tilde{\mathbf{Y}}_1, \ldots, \tilde{\mathbf{Y}}_{N-1}) \circ \left(g_i\left(\tilde{\tilde{\mathbf{X}}}^0_{-i}\right)\right)_{i \in [N-1]} - \mathbf{U}_m \circ \left(g_i\left(\tilde{\tilde{\mathbf{X}}}^0_{-i}\right)\right)_{i \in [N-1]} \right| + N\delta + (N-1)\varepsilon_1, \tag{7}$$

where each g_i is now a deterministic function of $\tilde{\tilde{\mathbf{X}}}^0_1, \ldots, \tilde{\tilde{\mathbf{X}}}^0_{i-1}, \tilde{\tilde{\mathbf{X}}}^0_{i+1}, \ldots, \tilde{\tilde{\mathbf{X}}}^0_{N-1}$, all $\{\tilde{\tilde{\mathbf{X}}}^0_j\}$ are still independent, each $\tilde{\mathbf{Y}}_j$ is a deterministic function of $\tilde{\tilde{\mathbf{X}}}^0_j$, and each $\tilde{\mathbf{Y}}_j$ has min-entropy at least k_1. Now, note that

since $\tilde{\mathbf{Y}}_j$ is a deterministic function of $\tilde{\tilde{\mathbf{X}}}_j^0$, we may write $\tilde{\mathbf{Y}}_j = f_j(\tilde{\tilde{\mathbf{X}}}_j^0)$ for some deterministic function f_j. On the other hand, by the dependency reversal lemma (Lemma 2), we know that we can furthermore find a deterministic function h_j and a new random variable $\mathbf{A}_j$ (which is completely independent of all random variables seen so far) such that $\tilde{\tilde{\mathbf{X}}}_j^0 = h_j(f_j(\tilde{\tilde{\mathbf{X}}}_j^0), \mathbf{A}_j) = h_j(\tilde{\mathbf{Y}}_j, \mathbf{A}_j)$. This means that any fixing of $\mathbf{A}_j$ does not affect the distribution or independence of the $\tilde{\mathbf{Y}}_j$'s, and furthermore every such fixing turns $\tilde{\tilde{\mathbf{X}}}_j^0$ into a deterministic function of $\mathbf{Y}_j$. As a result, standard fixing arguments tell us that there exists *some* fixings of $\mathbf{A}_1, \dots, \mathbf{A}_{N-1}$ such that Equation (7) can be upper bounded by

$$\left| \mathsf{NOF}(\tilde{\mathbf{Y}}_1, \dots, \tilde{\mathbf{Y}}_{N-1}) \circ \left(g_i \left(\tilde{\mathbf{Y}}_{-i} \right) \right)_{i \in [N-1]} - \mathbf{U}_m \circ \left(g_i \left(\tilde{\mathbf{Y}}_{-i} \right) \right)_{i \in [N-1]} \right| + N\delta + (N-1)\varepsilon_1, \quad (8)$$

where each g_i is now a deterministic function of $\tilde{\mathbf{Y}}_{-i}$, all the $\tilde{\mathbf{Y}}_i$'s are still independent, and all the $\tilde{\mathbf{Y}}_i$'s still have min-entropy at least k_1.

Phase IV: Executing the NOF-hard function. Finally, notice that the sequence of functions $(g_i)_i$ is exactly a (non-adaptive) number-on-forehead protocol over the remaining random variables $\tilde{\mathbf{Y}}_1, \dots, \tilde{\mathbf{Y}}_{N-1} \sim \{0,1\}^r$, and each of these remaining random variables has min-entropy at least k_1. Moreover, recall that the functions $g_1, \dots, g_{N-1}$ output $2^{N-1}m$ bits in total. Additionally, the lemma statement asserted that NOF has ε_2-average-case NOF communication complexity $> \mu_2 := 2^N m$ against number-on-forehead protocols. Since $2^{N-1}m \leq 2^N m - 2$, Lemma 3 immediately tells us that Equation (8) is at most

$$\varepsilon_2 \cdot 2^{N(r-k_1)} + N\delta + (N-1)\varepsilon_1.$$

Finally, recall that for the condenser to actually work (and be strong), we needed to feed it sources with min-entropy $k_0 + r + \log(1/\varepsilon_1)$. And furthermore, recall that the sources we fed into it had min-entropy at least $k - r - \log(1/\delta)$. Thus we need to have $k - r - \log(1/\delta) \geq k_0 + r + \log(1/\varepsilon_1)$. In other words, if we set $\delta = \varepsilon_1$ and start with min-entropy at least $k \geq k_0 + 2r + 2\log(1/\varepsilon_1)$, Equation (8) is bounded above by

$$\varepsilon_2 \cdot 2^{N(r-k_1)} + N\varepsilon_1 + (N-1)\varepsilon_1 \leq \varepsilon_2 \cdot 2^{N(r-k_1)} + 2N\varepsilon_1,$$

as desired. □

Now that we have completed the recipe, let's instantiate it to obtain our explicit weak NMEs (which, in the next section, we'll use to get our LREs). Recall that the weak NME is built from an NOF-hard function and a two-source condenser. We import below the ones that we use.

Theorem 3 (Finite field multiplication is NOF-hard [FG13])**.** *There exists a universal constant $c > 0$ such that the following holds. Let $f : (\{0,1\}^n)^N \to \{0,1\}^m$ be the function $f(x_1, \dots, x_N) := (x_1 \cdot x_2 \cdots x_N)_{\leq m}$ that multiplies its inputs over $\mathbb{F}_{2^n}$, and outputs the first m bits. As long as $m \leq cn/2^N$, it holds that f has $2^{-cn/2^N}$-average-case NOF communication complexity $> \mu := \lceil cn/2^N \rceil$.*[12]

Theorem 4 (Low-error two-source condenser [BCDT19, Theorem 31])**.** *For every small enough constant $\gamma > 0$, there exists a constant $C > 0$ such that the following holds. For every $n \geq k \geq \log^C n$, there exists an explicit two-source condenser $\mathsf{2Cond} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ for min-entropy k with output length $m \geq k - 6k^\gamma$, output entropy $\ell \geq m - k^\gamma$, and error $\varepsilon = 2^{-k^\gamma}$.*

[12] In the original paper [FG13], this result is only explicitly stated for $m = 1$. However, it is straightforward to extend to larger m - for a formal presentation, see [CGG+20].

Now, let's obtain our explicit weak NMEs.

Proof of Lemma 5. Simply plug the two-source condenser from Theorem 4 and the NOF-hard function from Theorem 3 into the recipe for constructing weak NMEs (Lemma 6), and set constants C, γ appropriately. □

5 Putting everything together

Now that we have our reduction from LREs against NOFs to weak NMEs, and our explicit construction of weak NMEs, we are ready to construct our LREs against NOFs.

Theorem 5 (LREs against NOFs - Theorem 2, general version)**.** *There exist universal constants $C, \gamma > 0$ such that the following holds. There exists an explicit leakage-resilient extractor $\mathsf{Ext} : (\{0,1\}^n)^N \to \{0,1\}^m$ for min-entropy $k \geq \log^C n$ against number-on-forehead-protocols with $\mu = \lceil k^\gamma / 2^N \rceil$ bits of communication, which has output length $m = \lceil k^\gamma / 2^N \rceil$ and error $\varepsilon = 2^{-\lceil k^\gamma / 2^N \rceil}$.*

Proof. Simply combine Lemmas 4 and 5, resetting constants C, γ appropriately. □

Setting $N = 3$ in Theorem 5, we obtain the following LRE against NOF protocols.

Corollary 1 (LREs against NOFs - Theorem 2, restated)**.** *There exist universal constants $C, \gamma > 0$ such that the following holds. There exists an explicit three-source extractor $\mathsf{Ext} : (\{0,1\}^n)^3 \to \{0,1\}^m$ for min-entropy $k \geq \log^C n$, which is leakage-resilient against number-on-forehead protocols with $\mu = \lceil k^\gamma \rceil$ bits of communication, and which has output length $m = \lceil k^\gamma \rceil$ and error $\varepsilon = 2^{-\lceil k^\gamma \rceil}$.*

Given this LRE against NOFs, it is now straightforward to obtain our extractors for adversarial sources.

Theorem 6 (Extractors for adversarial sources - Theorem 1, restated)**.** *There exist universal constants $C, \gamma > 0$ such that the following holds. There exists an explicit extractor $\mathsf{Ext} : (\{0,1\}^n)^N \to \{0,1\}^m$ for (N, K, n, k)-adversarial sources with $K = 3$ good sources of min-entropy $k \geq \log^C n$, which has output length $m = \lceil k^\gamma \rceil$ and error $\varepsilon = 2^{-\lceil k^\gamma \rceil}$.*

Proof. Let $\mathsf{LRE} : (\{0,1\}^n)^3 \to \{0,1\}^m$ be the LRE against NOFs from Corollary 1, and define the function $\mathsf{Ext} : (\{0,1\}^n)^N \to \{0,1\}^m$ as

$$\mathsf{Ext}(\mathbf{X}) := \bigoplus_{1 \leq a < b < c \leq N} \mathsf{LRE}(\mathbf{X}_a, \mathbf{X}_b, \mathbf{X}_c).$$

Without loss of generality, we may assume that the good sources in $\mathbf{X}$ are $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3$. Fix all $\mathbf{X}_i, i > 3$, and observe that Ext becomes

$$\mathsf{Ext}(\mathbf{X}) = \mathsf{LRE}(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3) \oplus f_1(\mathbf{X}_2, \mathbf{X}_3) \oplus f_2(\mathbf{X}_1, \mathbf{X}_3) \oplus f_3(\mathbf{X}_1, \mathbf{X}_2)$$

for some deterministic functions f_1, f_2, f_3. By combining the definition of LREs against NOFs with a standard application of the data-processing inequality, it follows that $\mathsf{Ext}(\mathbf{X})$ is close to uniform. □

References

[BBR88] Charles H Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. *SIAM journal on Computing*, 17(2):210–229, 1988. Preliminary version in CRYPTO 1985.

[BCDT19] Avraham Ben-Aroya, Gil Cohen, Dean Doron, and Amnon Ta-Shma. Two-source condensers with low error and small entropy gap via entropy-resilient functions. In *23rd International Conference on Randomization and Computation (RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.

[BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications. In *20th Annual ACM Symposium on Theory of Computing (STOC 1988)*, pages 103–112, 1988.

[BGK06] Jean Bourgain, Alexey A Glibichuk, and Sergei Vladimirovich Konyagin. Estimates for the number of sums and products and for exponential sums in fields of prime order. *Journal of the London Mathematical Society*, 73(2):380–398, 2006.

[BIW06] Boaz Barak, Russell Impagliazzo, and Avi Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118, 2006. Preliminary version in FOCS 2004.

[BNS92] László Babai, Noam Nisan, and Márió Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *Journal of Computer and System Sciences*, 45(2):204–232, 1992. Preliminary version in STOC 1989.

[BOL85] Michael Ben-Or and Nathan Linial. Collective coin flipping, robust voting schemes and minima of banzhaf values. In *26th Annual Symposium on Foundations of Computer Science (FOCS 1985)*, pages 408–416. IEEE, 1985.

[Bou05] Jean Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 1(01):1–32, 2005.

[BT94] Richard Beigel and Jun Tarui. On ACC. *Computational Complexity*, 4:350–366, 1994. Preliminary version in FOCS 1991.

[CF01] Ran Canetti and Marc Fischlin. Universally composable commitments. In *21st Annual International Cryptology Conference (CRYPTO 2001)*, pages 19–40. Springer, 2001.

[CFG+85] Benny Chor, Joel Friedman, Oded Goldreich, Johan Håstad, Steven Rudich, and Roman Smolensky. The bit extraction problem or t-resilient functions. In *26th Annual Symposium on Foundations of Computer Science (FOCS 1985)*, pages 396–407. IEEE, 1985.

[CFL83] Ashok K Chandra, Merrick L Furst, and Richard J Lipton. Multi-party protocols. In *15th Annual ACM Symposium on Theory of Computing (STOC 1983)*, pages 94–99, 1983.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988. Preliminary version in FOCS 1985.

[CG17] Mahdi Cheraghchi and Venkatesan Guruswami. Non-malleable coding against bit-wise and split-state tampering. *Journal of Cryptology*, 30:191–241, 2017.

[CG21] Eshan Chattopadhyay and Jesse Goodman. Improved extractors for small-space sources. In *62nd Annual Symposium on Foundations of Computer Science (FOCS 2021)*, pages 610–621. IEEE, 2021.

[CGG+20] Eshan Chattopadhyay, Jesse Goodman, Vipul Goyal, Ashutosh Kumar, Xin Li, Raghu Meka, and David Zuckerman. Extractors and secret sharing against bounded collusion protocols. In *61st Annual Symposium on Foundations of Computer Science (FOCS 2020)*, pages 1226–1242. IEEE, 2020.

[CGGL20] Eshan Chattopadhyay, Jesse Goodman, Vipul Goyal, and Xin Li. Extractors for adversarial sources via extremal hypergraphs. In *52nd Annual Symposium on Theory of Computing (STOC 2020)*, pages 1184–1197. ACM, 2020.

[CL16] Eshan Chattopadhyay and Xin Li. Extractors for sumset sources. In *48th Annual ACM Symposium on Theory of Computing (STOC 2016)*, pages 299–311, 2016.

[CL23] Eshan Chattopadhyay and Jyun-Jie Liao. Hardness against linear branching programs and more. In *Leibniz International Proceedings in Informatics (LIPIcs): 38th Computational Complexity Conference (CCC 2023)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023.

[CS16] Gil Cohen and Igor Shinkar. The complexity of DNF of parities. In *7th Conference on Innovations in Theoretical Computer Science (ITCS 2016)*, pages 47–58, 2016.

[CZ19] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. *Annals of Mathematics*, 189(3):653–705, 2019. Preliminary version in STOC 2016.

[DF92] Danny Dolev and Tomás Feder. Determinism vs. nondeterminism in multiparty communication complexity. *SIAM Journal on Computing*, 21(5):889–895, 1992. Preliminary version in FOCS 1989.

[DLWZ14] Yevgeniy Dodis, Xin Li, Trevor D Wooley, and David Zuckerman. Privacy amplification and nonmalleable extractors via character sums. *SIAM Journal on Computing*, 43(2):800–830, 2014.

[DOPS04] Yevgeniy Dodis, Shien Jin Ong, Manoj Prabhakaran, and Amit Sahai. On the (im)possibility of cryptography with imperfect randomness. In *45th Annual Symposium on Foundations of Computer Science (FOCS 2004)*, pages 196–205. IEEE, 2004.

[DW09] Yevgeniy Dodis and Daniel Wichs. Non-malleable extractors and symmetric key cryptography from weak secrets. In *Proceedings of the forty-first annual ACM Symposium on Theory of Computing*, pages 601–610, 2009.

[FG13] Jeff Ford and Anna Gál. Hadamard tensors and lower bounds on multiparty communication complexity. *computational complexity*, 22:595–622, 2013. Preliminary version in ICALP 2005.

[GG25] Jesse Goodman and Vipul Goyal. Two-source extractors don't shrink. Unpublished manuscript, 2025.

[GGJS11] Sanjam Garg, Vipul Goyal, Abhishek Jain, and Amit Sahai. Bringing people of different beliefs together to do UC. In *Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28-30, 2011. Proceedings 8*, pages 311–328. Springer, 2011.

[GK08] Vipul Goyal and Jonathan Katz. Universally composable multi-party computation with an unreliable common reference string. In *Theory of Cryptography Conference*, pages 142–154. Springer, 2008.

[GLZ24] Jesse Goodman, Xin Li, and David Zuckerman. Improved condensers for Chor-Goldreich sources. In *65th Annual Symposium on Foundations of Computer Science (FOCS 2024)*. IEEE, 2024.

[GO14] Jens Groth and Rafail Ostrovsky. Cryptography in the multi-string model. *Journal of cryptology*, 27(3):506–543, 2014.

[GPT22] Svyatoslav Gryaznov, Pavel Pudlák, and Navid Talebanfard. Linear branching programs and directional affine extractors. In *37th Computational Complexity Conference (CCC 2022)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2022.

[GSV05] Shafi Goldwasser, Madhu Sudan, and Vinod Vaikuntanathan. Distributed computing with imperfect randomness. In *Distributed Computing: 19th International Conference (DISC 2005)*, pages 288–302. Springer, 2005.

[Gur04] Venkatesan Guruswami. Better extractors for better codes? In *36th annual ACM Symposium on Theory of Computing (STOC 2004)*, pages 436–444, 2004.

[HG17] Miguel Herrero-Collantes and Juan Carlos Garcia-Escartin. Quantum random number generators. *Reviews of Modern Physics*, 89(1):015004, 2017.

[KM04] Robert König and Ueli Maurer. Extracting randomness from generalized symbol-fixing and markov sources. In *International Symposium on Information Theory (ISIT 2004)*, page 232. IEEE, 2004.

[KM05] Robert König and Ueli Maurer. Generalized strong extractors and deterministic privacy amplification. In *Cryptography and Coding: 10th IMA International Conference*, pages 322–339. Springer, 2005.

[KMS19] Ashutosh Kumar, Raghu Meka, and Amit Sahai. Leakage-resilient secret sharing against colluding parties. In *60th Annual Symposium on Foundations of Computer Science (FOCS 2019)*, pages 636–660. IEEE, 2019.

[KRVZ11] Jesse Kamp, Anup Rao, Salil Vadhan, and David Zuckerman. Deterministic extractors for small-space sources. *Journal of Computer and System Sciences*, 77(1):191–220, 2011. Preliminary version in STOC 2006.

[Lew19] Mark Lewko. An explicit two-source extractor with min-entropy rate near $4/9$. *Mathematika*, 65(4):950–957, 2019.

[Li15] Xin Li. Three-source extractors for polylogarithmic min-entropy. In *56th Annual Symposium on Foundations of Computer Science (FOCS 2015)*, pages 863–882. IEEE, 2015.

[Li23] Xin Li. Two source extractors for asymptotically optimal entropy, and (many) more. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS 2023)*, pages 1271–1281. IEEE, 2023.

[LLTT05] Chia-Jung Lee, Chi-Jen Lu, Shi-Chun Tsai, and Wen-Guey Tzeng. Extracting randomness from multiple independent sources. *IEEE Transactions on Information Theory*, 51(6):2224–2227, 2005.

[Lov19] Shachar Lovett. CSE 291: Communication Complexity, Winter 2019: Multi-party protocols, 2019. Lecture notes, University of California, San Diego. Version dated March 22, 2019.

[LY22] Jiatu Li and Tianqi Yang. $3.1n - o(n)$ circuit lower bounds for explicit functions. In *54th Annual ACM Symposium on Theory of Computing (STOC 2022)*, pages 1180–1193, 2022.

[MW97] Ueli Maurer and Stefan Wolf. Privacy amplification secure against active adversaries. In *Advances in Cryptology: 17th Annual International Cryptology Conference (CRYPTO 1997)*, pages 307–321. Springer, 1997.

[Rao07] Anup Rao. An exposition of Bourgain's 2-source extractor. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 14, 2007.

[Rao09] Anup Rao. Extractors for a constant number of polynomially small min-entropy independent sources. *SIAM Journal on Computing*, 39(1):168–194, 2009. Preliminary version in STOC 2006.

[RW93] Alexander Razborov and Avi Wigderson. $n^{\omega(\log n)}$ lower bounds on the size of depth-3 threshold circuits with AND gates at the bottom. *Information Processing Letters*, 45(6):303–307, 1993.

[Vad12] Salil Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

[Vaz87] Umesh V Vazirani. Strong communication complexity or generating quasi-random sequences from two communicating semi-random sources. *Combinatorica*, 7(4):375–392, 1987. Preliminary version in STOC 1985.

[Vio09] Emanuele Viola. Guest column: correlation bounds for polynomials over $\{0,1\}$. *ACM SIGACT News*, 40(1):27–44, 2009.

[vN51] John von Neumann. Various techniques used in connection with random digits. *Applied Math Series*, 12(36-38):1, 1951.

[Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *11th Annual ACM Symposium on Theory of Computing (STOC 1979)*, pages 209–213, 1979.

[Yao90] AC-C Yao. On ACC and threshold circuits. In *31st Annual Symposium on Foundations of Computer Science (FOCS 1990)*, pages 619–627. IEEE, 1990.

[Zuc07] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory OF Computing*, 3:103–128, 2007. Preliminary version in STOC 2006.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il