



Extracting Randomness from Extractor-Dependent Sources

Yevgeniy Dodis¹, Vinod Vaikuntanathan^{2(✉)}, and Daniel Wichs^{3,4(✉)}

¹ NYU, New York, USA

² MIT, Cambridge, USA

vinod.nathan@gmail.com

³ Northeastern University, Boston, USA

danwichs@gmail.com

⁴ NTT Research Inc., East Palo Alto, USA

Abstract. We revisit the well-studied problem of extracting nearly uniform randomness from an arbitrary source of sufficient min-entropy. Strong seeded extractors solve this problem by relying on a public random seed, which is unknown to the source. Here, we consider a setting where the seed is reused over time and *the source may depend on prior calls to the extractor with the same seed*. Can we still extract nearly uniform randomness?

In more detail, we assume the seed is chosen randomly, but the source can make arbitrary oracle queries to the extractor with the given seed before outputting a sample. We require that the sample has entropy and differs from any of the previously queried values. The extracted output should look uniform even to a distinguisher that gets the seed. We consider two variants of the problem, depending on whether the source only outputs the sample, or whether it can also output some correlated public *auxiliary information* that preserves the sample's entropy. Our results are: **Without Auxiliary Information:** We show that every *pseudo-random function* (PRF) with a sufficiently high security level is a good extractor in this setting, even if the distinguisher is computationally unbounded. We further show that the source necessarily needs to be computationally bounded and that such extractors imply one-way functions.

With Auxiliary Information: We construct secure extractors in this setting, as long as both the source and the distinguisher are computationally bounded. We give several constructions based on different intermediate primitives, yielding instantiations based on the DDH, DLIN, LWE or DCR assumptions. On the negative side, we show that one cannot prove security against computationally unbounded distinguishers in this setting under any standard assumption via a black-box reduction. Furthermore, even when restricting to computationally bounded distinguishers, we show that there exist PRFs that are insecure as extractors in this setting and that a large class of constructions cannot be proven secure via a black-box reduction from standard assumptions.

1 Introduction

EXTRACTING RANDOMNESS. Randomness is an important ingredient in many algorithmic tasks, and is especially crucial in cryptography. Indeed, much of

© International Association for Cryptologic Research 2020

A. Canteaut and Y. Ishai (Eds.): EUROCRYPT 2020, LNCS 12105, pp. 313–342, 2020.

https://doi.org/10.1007/978-3-030-45721-1_12

cryptography relies on the assumption that parties can sample uniformly random bits. However, most natural sources of randomness are imperfect and not uniformly random. This motivates the study of *randomness extraction*, whose goal is to extract (nearly) uniform randomness from imperfect sources.

Ideally, we would have a deterministic function Ext that converts an imperfect source of randomness X into a (nearly) uniformly random output $\text{Ext}(X)$. Furthermore, such an extractor should work for all sources of randomness X having a sufficiently large amount of (min-)entropy. Unfortunately, this is easily seen to be impossible, even if we only want to output 1 bit [CG85]: for every extractor function Ext , there is a source X that has almost full min-entropy yet the output of $\text{Ext}(X)$ is completely fixed.

There have been two broad lines of work to get around this. The first line of work designs extractors for restricted types of sources X that satisfy additional requirements beyond just having entropy (see e.g., [von51, CGH+85, Blu86, LLS89, CG85, TV00, BST03, BIW04, CZ16]). While this is an important research direction, we often know very little about natural sources of randomness and they may fail to satisfy the imposed requirements. The second line of work considers (*strong*) *seeded* extractors [NZ93, NZ96], where the extractor is given a public uniformly random seed S , which is independent of the source X , and we require that the extracted output $\text{Ext}(X; S)$ is close to uniform even given the seed S .

EXTRACTOR-DEPENDENT SOURCES. In this work, we consider a seeded extractor and envision a scenario where a single uniformly random seed S is chosen once and then is reused over time by many different users and/or applications to extract randomness from various “natural” sources of entropy. For example, the seed S could be a part of a system *random number generator* (RNG) that extracts randomness from physical sources of entropy, such as the timing of interrupts etc. If the sources are truly independent of the seed S , then standard (strong) seeded extractors suffice to guarantee that the extracted outputs are nearly uniform. However, since the seed S is continuously reused, past outputs of the extractor will make their way back into “nature” and may affect the sources in the future. For example, interrupts may depend on processes that themselves rely on previous outputs of the extractor. Furthermore, since we cannot assume that all users/applications use the extractor securely, we have to allow for the possibility that some of the prior calls to the extractor were made on arbitrary samples that may not have any entropy. Unfortunately, if the source can depend on prior calls to the extractor with the same seed S , we violate the condition that the source is independent of the seed and can no longer rely on the security of standard seeded extractors. We emphasize that, although the seed S is public, the sources are *not* fully adversarial and *not* arbitrarily dependent on S . (A restriction of this sort is of course necessary to circumvent the obvious impossibility result.) Instead, we assume that the sources can only depend on prior calls to the extractor with the given seed S , but are otherwise independent of S . We call such sources “extractor-dependent”. Can we design *extractors*

for extractor-dependent sources (*ED-Extractors*) that manage to extract nearly uniform randomness in this setting?

DEFINING THE PROBLEM. We now specify the problem in more detail. Our goal is to design a seeded extractor EExt that extracts randomness from extractor-dependent sources. We consider a setting where a seed S is chosen uniformly at random. A source $\mathcal{S}^{\text{EExt}(\cdot, S)}$ gets oracle access to the extractor with the seed S and outputs a sample X along with some public auxiliary information AUX . We say that such a source \mathcal{S} is a *legal* extractor-dependent source of entropy α if two conditions hold: (1) the (conditional min-entropy) of X given S, AUX is at least α , and (2) the source never queries the oracle on the value X that it outputs. An α -ED-Extractor needs to ensure that for all legal extractor-dependent sources of entropy α , the output $\text{EExt}(X, S)$ is indistinguishable from uniform, even given the seed S and the auxiliary information AUX .

DISCUSSION ON THE LEGALITY CONDITIONS. We motivate the reason behind the two legality conditions imposed by the definition.

Firstly, just like for standard (seeded) extractors, we need to assume that X has a sufficient level of entropy even conditioned on AUX in order to extract randomness from it. In our case, the source also has access to the oracle $\text{EExt}(\cdot, S)$ with a random seed S , but we want the entropy to come from the internal randomness of the source rather than from the seed S since the latter is public and known to the distinguisher. Therefore, it is natural to also condition on S .

The second condition is clearly necessary: without it we could define a source that queries the oracle on random values and outputs the first such value on which the extracted output starts with a 0. Such a source would have almost full entropy, yet the extracted output would be easily distinguishable from uniform. Moreover, this condition is also reasonable when modeling our intended scenario since the sample should have entropy even given all the prior extractor calls that influenced nature, and therefore it should differ from all of them.

In particular, the two legality conditions include the following simpler subclass of sources, which already intuitively captures our intended scenario. Consider sources $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ that consists of two components. The first component $\mathcal{S}_1^{\text{EExt}(\cdot, S)}$ makes arbitrary oracle calls to the extractor and models the influence that these calls have on nature; it outputs some value state . The second component $\mathcal{S}_2(\text{state})$ then outputs X, AUX without making any further oracle queries and captures the entropic process that produces the sample. The only condition we impose is that, for every possible fixed value of state , the entropy of X conditioned on AUX when they are sampled according to $\mathcal{S}_2(\text{state})$ should be at least α . If α is large enough then \mathcal{S} satisfies both of the previous legality conditions. In particular, \mathcal{S}_1 could not have queried the oracle on X since the entropy of X comes only from the random coins of \mathcal{S}_2 that are unknown to \mathcal{S}_1 .

DISCUSSION ON AUXILIARY INFO. Our default definition allows the source to output some public auxiliary info AUX that can be correlated with the sample X as long as it preserves its (average conditional min-)entropy. It is natural that some such information may be public (e.g., the source X denotes the timing of

interrupts, but the adversary can learn some auxiliary info AUX denoting the high-order bits of such timings by interacting with the system). We also consider a weaker setting *without auxiliary info*, where we don't have AUX . In the case of standard seeded extractors, it turns out that there is not much difference between a setting with auxiliary info and without [DORS08]. However, as we will see, there is a significant difference between the two settings when it comes to ED-Extractors.

PRIOR WORK. The work of Coretti et al. [CDKT19] initiates the study of extracting from extractor-dependent sources in the special case where the extractor is a *random oracle*. While their definition is specifically tailored to the random-oracle model, our definition can be seen as the natural extension of it to the standard model. In particular, they consider the setting where $\mathcal{O}(\cdot) = \text{EExt}(\cdot, S)$ is a truly random function. They show that this is an α -ED-Extractor for any super-logarithmic entropy α , as long as the source only makes polynomially many queries, but even if the distinguisher is computationally unbounded and can see the entire truth table of the oracle. This gives us heuristic evidence that a “good” cryptographic hash function is an ED-Extractor in the standard model even against computationally unbounded distinguishers (as long as the source is computationally bounded). The main open question is therefore whether we can construct ED-Extractors under standard computational assumptions.

1.1 Our Results

We give positive and negative results for ED-Extractors with and without auxiliary info.

WITHOUT AUXILIARY INFO. On the positive side, we show that any pseudo-random function (PRF) with a sufficiently high security level is a good ED-Extractor without auxiliary info. In particular, assuming the existence of sub-exponentially secure one-way functions, there exist α -ED-Extractors with any output size m for entropy $\alpha = m + \omega(\log \lambda)$, where λ is the security parameter. Furthermore, such extractors achieve security even against computationally unbounded distinguishers, as long as the source runs in polynomial time. If we only want security against polynomial-time distinguishers, we can allow the output size to grow to an arbitrary polynomial m while only requiring entropy $\alpha = \lambda^{\Omega(1)}$.

On the negative side, we show that ED-Extractors imply one-way functions and therefore cannot be constructed unconditionally. This holds even without auxiliary info, even if we require that the source has almost full entropy, and even if the extractor outputs only 1 bit. Furthermore, we show that such ED-Extractors cannot exist for computationally unbounded sources.

WITH AUXILIARY INFO. We construct ED-Extractors in the setting with auxiliary info under standard assumptions. In particular, we give three constructions.

- The first construction relies on (adaptively secure) *constrained PRFs* [BGI14, KPTZ13, BW13] for NC1 constraints. These can be instantiated under the

- sub-exponential security of either the *learning with errors (LWE)* [BV15] or the *Decisional Diffie-Hellman Inversion (DDHI)* assumption in arbitrary prime-order groups (without requiring pairings) [AMN+18].¹
- The second construction relies on *shift-hiding shiftable functions* [PS18], which can be seen as a type of constraint-hiding constrained PRFs, and can be instantiated under LWE without requiring sub-exponential security.
 - The third construction relies on lossy functions and can be instantiated under any of: decisional Diffie-Hellman (DDH), decisional-linear (DLIN), LWE, or decisional composite residuosity (DCR) assumptions.

In all cases, we prove security against polynomial-time sources and distinguishers. Our α -ED-Extractors can have arbitrarily large polynomial input size n and output size m , and require entropy $\alpha = \lambda^{\Omega(1)}$.

Note that, in the setting without auxiliary info, we achieved security even against computationally unbounded distinguishers. Furthermore, the random-oracle based result of [CDKT19] heuristically suggests that good cryptographic hash functions achieve security against computationally unbounded distinguishers even in the auxiliary info setting. However, our constructions in the auxiliary info setting from standard assumptions only achieve security against polynomial-time distinguishers. Unfortunately, we show that this is inherent. In particular, we show that in the auxiliary info setting, one cannot prove the security of any ED-Extractor against computationally unbounded distinguishers under any standard assumption via a black-box reduction.

Furthermore, our instantiations in the auxiliary info setting rely on “cryptomania” assumptions (known to imply public-key encryption) rather than one-way functions, and we ask whether this is necessary. While we do not resolve this question, we give some evidence that the two settings necessitate substantially different constructions. Firstly, one may be tempted to conjecture that every PRF is also a good ED-Extractor even in the auxiliary info setting. We show that this is not the case: there exist PRFs that are insecure as ED-Extractors in the auxiliary info setting even for very high levels/rates of entropy α . Moreover, we show that a large class of natural PRFs (e.g., the Naor-Reingold PRF) cannot be proven to be secure ED-Extractors in the setting of auxiliary info via a black-box reduction from any standard assumption.

1.2 Our Techniques

ED-EXTRACTORS WITHOUT AUXILIARY INFO FROM PRFs. Our first result shows that every PRF is already a good ED-Extractor in the setting without auxiliary info. In particular, the seed of the extractor is the PRF key and the extractor just evaluates the PRF on the sample X . The main difficulty in proving ED-Extractor security is that the distinguisher gets the seed of the ED-Extractor, but PRF security only holds if the key is never revealed. Our insight

¹ The DDHI assumption in a cyclic group \mathbb{G} of order q with generator g states that, given any polynomially many values of the form $(g, g^a, g^{a^2}, \dots, g^{a^L})$ where $a \leftarrow \mathbb{Z}_q$, the value $g^{1/a}$ is computationally indistinguishable from uniform.

is to design a reduction that never calls the distinguisher – indeed, this allows us to prove security even for computationally unbounded distinguishers.

Let's start with the case where the PRF/Extractor only outputs 1 bit. If the extracted output is statistically far from uniform given the seed, it means that it is biased towards either 0 or 1, but the direction of the bias is unknown and may be different for each seed. Consider running the source \mathcal{S} twice with independent randomness, while giving it oracle access to the PRF/Extractor with the same random key/seed. Let X_0, X_1 be the samples that the two runs output respectively. Then the PRF/Extractor evaluations on those samples are more likely to agree than disagree, since they are biased in the same direction. But the legality conditions ensure that X_0, X_1 were never queried during either of the two runs and are different from each other (since each run cannot query its own output and the output of the other run should have enough entropy to be unpredictable). So, given oracle access to the PRF, we can use the source \mathcal{S} to find two values X_0, X_1 that we haven't yet queried, but if we then proceed to query the PRF on them, the outputs are noticeably more likely to agree than disagree. This cannot be the case given oracle access to a random function, and therefore allows us to distinguish the two and break PRF security. The analysis extends to a larger output size m , but the advantage of the reduction shrinks by a factor of 2^{-m} . Therefore, we need very secure PRFs that cannot be distinguished from random functions with advantage better than $\text{negl}(\lambda)2^{-m}$, which requires sub-exponential security assumptions.

Note that the above argument completely breaks down in the setting with auxiliary info. The problem is that now the direction of the bias can be different for each choice of the key/seed *and* the auxiliary info. But the two independent runs of the source \mathcal{S} are unlikely to produce the same auxiliary info and hence we cannot argue that the bias would go in the same direction. Indeed, we show that there are PRFs that are completely insecure as ED-Extractors in the setting with auxiliary info.

ED-EXTRACTORS IMPLY ONE-WAY FUNCTIONS. We show that ED-Extractors cannot exist if the source is allowed to be computationally unbounded. This holds even in the setting without auxiliary info, even if we only consider polynomial-time distinguishers, even if we require that the source has almost full entropy, and even if the extractor outputs only 1 bit. The high level idea is that a computationally unbounded source \mathcal{S} with oracle access to the function $\text{EExt}(\cdot, S)$ can learn the function sufficiently well to predict its output on a random value with high probability. It can then sample a random X subject to predicting that $\text{EExt}(X, S) = 0$, without querying the extractor on X . This is a legal source with almost full entropy, yet the extractor output is highly biased towards 0. We extend the above argument to showing that such extractors imply one-way functions.

ED-EXTRACTORS WITH AUXILIARY INFO FROM CONSTRAINED PRFS. We construct ED-Extractors in the setting with auxiliary info, using constrained pseudorandom functions (C-PRF). A C-PRF allows us to constrain a PRF key

k on some constraint function C to yield a constrained key, denoted $k\{C\}$. The constrained key allows us to evaluate the PRF on all points x such that $C(x) = 0$. However, given the constrained key $k\{C\}$, the PRF outputs at all points x for which $C(x) = 1$ look random. We need to rely on *adaptively secure* constrained PRFs, where the adversary can choose the constraint C after seeing some PRF outputs.

Our construction of ED-Extractors uses a constrained PRF and a standard (seeded) randomness extractors Ext . The seed of the ED-Extractor is a constrained PRF key $k\{C_{S,U}\}$, with the constraint $C_{S,U}(X)$ that outputs 1 (i.e., prevents evaluation) on all points X such that $\text{Ext}(X; S) = U$, where S, U are chosen randomly. We choose the output size of the extractor to be $\ell = \omega(\log \lambda)$ and therefore the key is constrained on a negligible fraction of points. On input X , the ED-Extractor checks if $C_{S,U}(X) = 1$, in which case it outputs some fixed dummy value, and otherwise it uses the seed $k\{C_{S,U}\}$ to evaluate the PRF on X .

To argue ED-Extractor security, we consider a source $\mathcal{S}^{\text{EDExt}(\cdot, k\{C_{S,U}\})}$ that gets oracle access to the ED-Extractor with a random seed $k\{C_{S,U}\}$ and outputs X, AUX . A distinguisher \mathcal{D} then gets the seed $k\{C_{S,U}\}$ together with AUX and the extracted output $R = \text{EDExt}(X, k\{C_{S,U}\})$. We first argue that this is statistically indistinguishable from giving the source \mathcal{S} oracle access to the unconstrained PRF and setting R to be the output of the PRF with the unconstrained key on X (since the probability that any of the queries of \mathcal{S} or its output lie in the constrained set is negligible). Now, instead of giving the distinguisher \mathcal{D} the constrained key $k\{C_{S,U}\}$ where U is uniform, we give it the key $k\{C_{S, \text{Ext}(X; S)}\}$ which is constrained on X . This is statistically indistinguishable since X has entropy even conditioned on AUX and is sampled independently of S ; therefore $\text{Ext}(X; S)$ is close to uniform even given AUX . But now we can switch R from the output of the PRF on X to uniform, and this is computationally indistinguishable even given the constrained PRF key $k\{C_{S, \text{Ext}(X; S)}\}$ since it is constrained on X (and we know that the source didn't query the oracle on X). This shows that the extracted output is indistinguishable from uniform even given the ED-Extractor seed and the auxiliary info. (The above proof outline conveys the intuition but is slightly oversimplified and ignores some subtleties; see the full proof for details).

Since standard extractors can be evaluated in NC1, we only need constrained PRFs for NC1 circuits. Fortunately, we have such constructions from the LWE and DDHI assumptions [BV15, AMN+18]. However, they only achieve selective security, where the constrained circuit needs to be chosen ahead of time before any PRF outputs are given out, while we need adaptive security. We can get this via standard complexity leveraging at the cost of having to assume the sub-exponential security of the LWE and DDHI assumptions.

ADDITIONAL CONSTRUCTIONS (IN THE FULL VERSION). We give two alternate constructions of ED-Extractors in the setting with auxiliary info. The first uses *shift-hiding shiftable functions* [PS18], which can be instantiated from standard LWE, without needing complexity leveraging. The construction and proof

of security differ substantially from the one above. The second one uses *lossy functions*, which are essentially equivalent to lossy trapdoor functions (LTDFs) [PW08] without requiring a trapdoor. The construction can be instantiated from several different assumptions (DDH, DLIN, LWE, DCR). Both constructions are omitted from this proceedings version due to lack of space; please see the full version [DVW19].

NOT ALL PRFs ARE ED-EXTRACTORS WITH AUX INFO. We construct PRFs, which fail to be good extractors in the setting of auxiliary info. For example, consider a PRF which first hashes the input x into a small digest using a collision-resistant hash function and then applies another PRF on the output. Consider a source that chooses a random x and sets the auxiliary info to be the hash of x . Since the hash is small, this does not reduce the entropy of x by much. However, if the distinguisher is given the PRF key (which is the ED-Extractor seed) and the auxiliary info, it can compute the PRF on x and therefore easily distinguish it from uniform. In this example, the auxiliary info reduces the entropy of x by some small super-logarithmic amount. We give an even more dramatic example of this type using fully-homomorphic encryption (FHE) where the auxiliary info reduces the entropy of x by only 1 bit.

BLACK-BOX SEPARATION RESULTS. Lastly, we give two black-box separation results showing that, in the auxiliary info setting, one cannot prove security (via a black-box reduction under a standard assumption) against computationally unbounded distinguishers or for certain natural classes of constructions. Our results rely on the framework of [Wic13] and rely on the fact that the ED-Extractor definition is expressed as a two-stage game where the attacker consists of two components (the source and the distinguisher) that cannot communicate. This allows us to give black-box separations showing that, in certain cases, we cannot prove security under any standard assumption which is in the form of a single-stage game between a challenger and a single stateful adversary.

1.3 Additional Related Work

RNGS. Our scenario is partially motivated by the problem of extracting randomness from physical sources as part of a system Random Number Generator (RNG). We note that extracting randomness is only one component of a good RNG; see e.g., [BH05, DPR+13, DSSW14, GT16, Hut16, CDKT19] for works that formally deal with the broader problem of RNG design.

UNIVERSAL COMPUTATIONAL EXTRACTORS (UCE). The notion of universal computational extractors (UCE) [BHK13, ST17] was originally proposed as a way of capturing “random-oracle like” security properties of hash functions via a standard-model definition. While the format of the UCE definition is also given in terms of an extraction game with a source and a distinguisher, there are major differences between the UCE definition and that of ED-Extractors, both in terms of their syntactic structure, but also more conceptually in terms of what they aim to capture. The key such difference is that the notion of legal source is

defined in the “ideal model”, and permits sources which only have *computational* unpredictability in the “real” model (say, conditioned on the auxiliary information).² In contrast, this work only aimed to capture a smaller class of sources that have entropy even in the “real model”, but could depend of the previous extractor output.

Unfortunately, it is known that even the weakest form of UCE security cannot be achieved under standard assumptions (via black-box reductions; this indirectly follows from [Wic13]), while our work shows that ED-Extractors can. It remains an interesting open problem whether ED-Extractors can be used in place of UCEs to get any broader cryptographic applications beyond the immediate ones of extracting randomness.

LOW-COMPLEXITY SAMPLERS. Introduced by Trevisan and Vadhan [TV00] and later extended by [KRVZ11], these seedless extractors assume that the entropy source producing input X is unable to run the extractor even once. In contrast, our sampler can be much slower than the extractor, but we use a seed and give the sampler oracle access to the extractor, before releasing the seed to the distinguisher.

SEED-DEPENDENT CONDENSERS. This approach, formalized by Dodis, Ristenpart and Vadhan [DRV12], relaxes the security guarantees of the randomness extractor to only ensure that the output of the condenser is almost full entropy, but not necessarily close to uniform. In this sense it is weaker than ED-Extractors. However, the sampler is given the actual seed, which is stronger than our setting. Interestingly, the availability of auxiliary information also played a crucial role in the constructions of seed-dependent condensers from standard assumptions.

2 Preliminaries

When X is a distribution, or a random variable following this distribution, we let $x \leftarrow X$ denote the process of sampling x according to the distribution X . If X is a set, we let $x \leftarrow X$ denote sampling x uniformly at random from X .

Let X, Y be random variables with supports S_X, S_Y , respectively. We define their *statistical difference* as

$$\text{SD}(X, Y) = \frac{1}{2} \sum_{u \in S_X \cup S_Y} |\Pr[X = u] - \Pr[Y = u]|.$$

The *min-entropy* of a random variable X is $H_\infty(X) = -\log(\max_x \Pr[X = x])$. Following Dodis et al. [DORS08], we define the (average) conditional min-entropy of X given Y as: $H_\infty(X|Y) = -\log(\mathbb{E}_{y \leftarrow Y} [2^{-H_\infty(X|Y=y)}])$. Note that $H_\infty(X|Y) = k$ iff the optimal strategy for guessing X given Y succeeds with probability 2^{-k} .

² Somewhat confusingly, this is true even for so called “UCEs for statistically unpredictable sources”.

Lemma 1. *For any random variables X, Y where Y is supported over a set of size T we have $H_\infty(X|Y) \leq H_\infty(X) - \log T$.*

Definition 1 ((Strong, Average-Case) Seeded Extractor [NZ96]). *We say that an efficient function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ is an (α, ε) -extractor if for all random variables (X, Z) such that X is supported over $\{0, 1\}^n$ and $H_\infty(X|Z) \geq \alpha$ we have $\text{SD}((Z, S, \text{Ext}(X; S)), (Z, S, U_\ell)) \leq \varepsilon$ where S, U_ℓ are uniformly random and independent bit-strings of length d, ℓ respectively.*

Theorem 1 ([ILL89]). *There exist (α, ε) -extractors with input length n and output length ℓ as long as $\alpha \geq \ell + 2 \log(1/\varepsilon)$.*

Definition 2 ((Strong, Average-Case) Two-Source Extractor [CG88]). *We say that an efficient function $2\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is an (e_1, e_2, δ) -strong 2-source extractor if for all random variables (X_1, X_2, Z) such that X_1, X_2 are independent conditioned on Z and $H_\infty(X_1|Z) \geq e_1, H_\infty(X_2|Z) \geq e_2$ we have $\text{SD}((Z, X_2, 2\text{Ext}(X_1; X_2)), (Z, X_2, U_m)) \leq \delta$ where U_m is a uniformly string of length m .*

Theorem 2 ([Raz05]). *For any polynomial input length $n = \text{poly}(\lambda)$, any $e_1 = \lambda^{\Omega(1)}$ and any $e_2 = (1/2 + \Omega(1))n$, there exist (e_1, e_2, δ) -extractor with input length n , output length $m = \lambda^{\Omega(1)}$ and error $\delta = 2^{-\lambda^{\Omega(1)}}$.*

Definition 3. *The collision probability of a random variable A is defined as $\text{Col}(A) = \Pr[a = a' : a \leftarrow A, a' \leftarrow A]$. The conditional collision probability of A given B is defined as $\text{Col}(A|B) = \Pr[a = a' : b \leftarrow B, a \leftarrow (A|B = b), a' \leftarrow (A|B = b)]$.*

Claim 1 (Statistical Distance vs Collision Probability [IZ89]). *Let A be a random variable supported over $\{0, 1\}^m$ such that $\text{SD}(A, U_m) \geq \varepsilon$, where U_m is uniform over $\{0, 1\}^m$. Then $\text{Col}(A) \geq \frac{1}{2^m}(1 + 4\varepsilon^2)$.*

Furthermore, let A, B be correlated random variables, where A is supported over $\{0, 1\}^m$ and

$$\text{SD}((A, B), (U_m, B)) \geq \varepsilon.$$

Then $\text{Col}(A|B) \geq \frac{1}{2^m}(1 + 4\varepsilon^2)$.

Learning with Errors. The (n, m, q, χ) LWE assumption states that $(A, sA + e)$ is computationally indistinguishable from (A, u) where $A \leftarrow \mathbb{Z}_q^{n \times m}, s \leftarrow \mathbb{Z}_q^n, e \leftarrow \chi^m$ and $u \leftarrow \mathbb{Z}_q^m$. Throughout this work, the LWE assumption (without qualification), refers to assuming that there exists some $n = \text{poly}(\lambda)$, some $q \geq 2^{\lambda^{\Omega(1)}}$ and some distribution χ over \mathbb{Z} which is $\text{poly}(\lambda)$ bounded such that the (n, m, q, χ) assumption holds for all $m = \text{poly}(\lambda)$. This is implied by the hardness of worst-case lattice problems with sub-exponential approximation factors.

Definition 4 (Pseudorandom Function (PRF) [GGM84]). *A polynomial-time function $F : \{0, 1\}^\ell \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ with key length $\ell = \ell(\lambda)$, input*

length $n = n(\lambda)$ and output length $m = m(\lambda)$ is a PRF if for any polynomial-time attacker \mathcal{A} there exists some negligible function $\mu(\lambda) = \text{negl}(\lambda)$ such that

$$| \Pr[\mathcal{A}^{F(k,\cdot)}(1^\lambda) = 1] - \Pr[\mathcal{A}^{\mathcal{O}(\cdot)}(1^\lambda) = 1] | \leq \mu(\lambda).$$

where we choose $k \leftarrow \{0, 1\}^\ell$ and $\mathcal{O} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a uniformly random function. We say that the PRF has security level $\sigma = \sigma(\lambda)$ if $\mu(\lambda) \leq 1/\sigma(\lambda)$.

Definition 5 (Constrained PRFs (CPRF) [BGI14,KPTZ13,BW13]). A CPRF for a class of constraints $\mathcal{C} = \{C_\lambda\}$ consists of two polynomial-time algorithms (F , Constrain) where:

- $y = F(k, x)$ is a deterministic polynomial-time function that takes as input a key k (either constrained or unconstrained) and a value $x \in \{0, 1\}^n$ and outputs $y \in \{0, 1\}^m$ for some polynomial length parameters $n = n(\lambda), m = m(\lambda)$ in the security parameter λ .
- $k\{C\} \leftarrow \text{Constrain}(k, C)$ takes as input a key $k \in \{0, 1\}^\lambda$ and a constraint $C : \{0, 1\}^n \rightarrow \{0, 1\}$ with $C \in \mathcal{C}_\lambda$. It outputs a constrained key, denoted $k\{C\}$.

We require that the scheme satisfies a correctness and a security property defined below:

Correctness: We require that no adversary can find an input which is not constrained, yet the constrained key disagrees with the original key. More concretely, consider the following game between a stateful adversary \mathcal{A} and a challenger:

- The adversary $\mathcal{A}(1^\lambda)$ chooses $C \in \mathcal{C}_\lambda$.
- The challenger chooses $k \in \{0, 1\}^\lambda$ and $k\{C\} \leftarrow \text{Constrain}(k, C)$.
- The adversary $\mathcal{A}^{F(k,\cdot)}(k\{C\})$ gets the constrained key $k\{C\}$ and oracle access to $F(k, \cdot)$. It outputs a value $x \in \{0, 1\}^n$.

We require that, in the context of the above experiment, we have $\Pr[C(x) = 0 \wedge F(k, x) \neq F(k\{C\}, x)] \leq \text{negl}(\lambda)$.

(Adaptive) Security: Consider the following distinguishing game between an adversary \mathcal{A} and a challenger:

- Challenger chooses $k \leftarrow \{0, 1\}^\lambda$ and a bit $b \leftarrow \{0, 1\}$.
- Adversary $\mathcal{A}^{F(k,\cdot)}(1^\lambda)$ gets oracle access to $F(k, \cdot)$ and outputs a constraint $C \in \mathcal{C}_\lambda$ and a values x such that $C(x) = 1$ and x was never queried to the oracle.
- If $b = 0$, the challenger sets $r = F(k, x)$ and else it chooses $r \leftarrow \{0, 1\}^m$. The challenger also computes $k\{C\} \leftarrow \text{Constrain}(k, C)$.
- The adversary \mathcal{A} is given $k\{C\}$ and r . It outputs a bit b' .

We require that for all polynomial-time adversaries \mathcal{A} , we have $|\Pr[b = b'] - \frac{1}{2}| = \text{negl}(\lambda)$.

We also consider several variants of the definition. Firstly, we define the notion of no-constrained-evaluation security, where we restrict the adversary to never querying the oracle $F(k, \cdot)$ on a point x for which $C(x) = 1$. Secondly, we

consider selective security where the adversary chooses $C \in \mathcal{C}_\lambda$ at the beginning of the game before getting oracle access to $F(k, c)$. Lastly, we consider no-evaluation security where the adversary does not get oracle access to $F(k, \cdot)$ at all.

Note that, via a simple guessing argument where we guess the adversary's choice of C , selective security with a sufficiently high security level $\sigma(\lambda) = |\mathcal{C}_\lambda|^{-1} \omega(\log \lambda)$ implies adaptive security. Furthermore by the same argument, no-evaluation security (which is inherently selective) with a sufficiently high security level $\sigma(\lambda) = |\mathcal{C}_\lambda|^{-1} \omega(\log \lambda)$ implies no-constrained-evaluation security. This is because, if we guess the adversary's choice of C ahead of time and gets $k\{C\}$, we can answer queries on unconstrained points using $k\{C\}$ rather than calling the PRF oracle.

3 Defining ED-Extractors

In this section, we give a formal definition of extractors for extractor-dependent sources (ED-Extractors) and provide some discussion on the various aspects of the definition.

Definition 6 (Extractor-Dependent Extraction). *An extractor for α -entropy extractor-dependent sources (α -ED-Extractor) consists of two polynomial-time algorithms (SeedGen, EDExt) with the following syntax:*

- seed \leftarrow SeedGen(1^λ) is a randomized algorithm that generates seed.
- EDExt(x , seed) is a deterministic algorithm that takes a sample $x \in \{0, 1\}^n$, together with seed and outputs a value $y \in \{0, 1\}^m$ for some polynomial length parameters $n = n(\lambda)$, $m = m(\lambda)$.

Consider an adversarial source/distinguisher pair $(\mathcal{S}, \mathcal{D})$ and define the following extraction experiment $\text{EDGame}^{\mathcal{S}, \mathcal{D}}(1^\lambda)$:

- Sample a random bit $b \leftarrow \{0, 1\}$ and a random seed \leftarrow SeedGen(1^λ).
- Run $(x, \text{aux}) \leftarrow \mathcal{S}^{\text{EDExt}(\cdot, \text{seed})}(1^\lambda)$.
- If $b = 0$ set $r = \text{EDExt}(x, \text{seed})$ else if $b = 1$ set $r \leftarrow \{0, 1\}^m$.
- Let $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$.

We say that \mathcal{S} is an α -legal extractor-dependent source if the following conditions hold:

1. The probability that \mathcal{S} queries its oracle on the value x that it outputs is negligible.
2. $H_\infty(X|\text{AUX}, \text{SEED}) \geq \alpha(\lambda)$, where $X, \text{SEED}, \text{AUX}$ denotes the joint distribution of the values $x, \text{seed}, \text{aux}$ in the above experiment.

An α -ED-Extractor is secure if for all α -legal polynomial-time sources \mathcal{S} and all polynomial-time distinguishers \mathcal{D} , the above experiment satisfies

$$\left| \Pr[b = b'] - \frac{1}{2} \right| = \text{negl}(\lambda).$$

We can also define a weaker notion without auxiliary info by restricting aux to be empty. We can also strengthen security to computationally unbounded sources or distinguishers by removing the restriction that the source or the distinguisher runs in polynomial time.

Remark on the Legality Conditions. As we discussed in the introduction, the legality conditions above may not seem entirely intuitive on first look. For example, it may be unclear why we prohibit the source from querying the extractor on the value it outputs. Another undesirable aspect of definition is that the legality conditions are construction-dependent: in other words, a source may be legal for some constructions of the ED-Extractor but illegal for others since the entropy of the output may depend on the oracle queries. Ideally, the legality of the source could be checked independently of the construction. For these reasons, we can also consider an alternate, weaker, definition, which may be more intuitively compelling and does not suffer from the above issue. We say that source \mathcal{S} is α -super-legal if:

- It can be written as $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ where $\mathcal{S}_1^{\text{Ext}(\cdot, \text{seed})}(1^\lambda)$ gets oracle access to the extractor and outputs some value $\text{state} \in \{0, 1\}^{p(\lambda)}$ for some polynomial p , and $\mathcal{S}_2(\text{state})$ outputs x, aux without getting any further access to the extractor.
- For all choices of $\text{state} \in \{0, 1\}^{p(\lambda)}$ it holds that $H_\infty(X|\text{AUX}) \geq \alpha(\lambda)$, where (X, AUX) are random variables for the output of $\mathcal{S}_2(\text{state})$.

Note that “super-legality” is only a condition of \mathcal{S}_2 which does not have oracle access to the extractor, and is therefore construction-independent.

We claim that for any $\alpha(\lambda) = \omega(\log \lambda)$, every α -super-legal source \mathcal{S} is also α -legal. Firstly, if \mathcal{S}_1 only makes polynomially many queries and has a non-negligible probability of querying the oracle on the value x that \mathcal{S}_2 outputs then there must be some value of state for which we can predict the value x that $\mathcal{S}_2(\text{state})$ outputs with non-negligible probability. But this contradicts $H_\infty(X) \geq H_\infty(X|\text{AUX}) \geq \omega(\log \lambda)$. Therefore \mathcal{S} satisfies the first legality condition. Secondly, let STATE be a random variable for the value $\text{state} \leftarrow \mathcal{S}_1^{\text{Ext}(\cdot, \text{seed})}(1^\lambda)$. Then SEED is independent of (X, AUX) if we condition on STATE . Therefore, $H_\infty(X|\text{AUX}, \text{SEED}) \geq H_\infty(X|\text{AUX}, \text{STATE}) \geq \min_{\text{state}} H_\infty(X_{\text{state}}|\text{AUX}_{\text{state}}) \geq \alpha(\lambda)$ where $X_{\text{state}}, \text{AUX}_{\text{state}}$ is the conditional distribution of X, AUX conditioned on $\text{STATE} = \text{state}$, which is just the distribution of the outputs of $\mathcal{S}_2(\text{state})$. Therefore \mathcal{S} satisfies the second legality condition.

As discussed in the introduction, the super-legality condition can also be interpreted very intuitively: we think of \mathcal{S}_1 as capturing all of the influence that prior extractor call can have on nature and \mathcal{S}_2 as modeling the entropic process that’s responsible for generating x, aux . We chose to use “legality” rather than “super-legality” in our default definition since it makes the definition stronger and thus gives stronger positive results. We mention that (by simple inspection) all of our negative results also hold for the weaker definition using super-legality.

Remark about Conditioning on the Seed. Our legality condition in the formal definition requires that the entropy $H_\infty(X|AUX, SEED) \geq \alpha(\lambda)$, where we condition on SEED. Note that we could remove this conditioning and have an alternate, stronger, definition where we only require $H_\infty(X|AUX) \geq \alpha(\lambda)$. We observe that, assuming one-way functions, any α -ED-Extractor according to our definition can be converted into an $(\alpha' = \alpha + \lambda^\varepsilon)$ -ED-Extractor according to the stronger definition for any constant $\varepsilon > 0$. The idea is that we can modify the SeedGen algorithm to only use λ^ε random bits by expanding them out using a PRG to get as many pseudorandom bits as needed by the original algorithm. By the security of the PRG, this change cannot harm ED-Extractor security. But now SEED comes from a domain of size only 2^{λ^ε} and therefore $H_\infty(X|AUX, SEED) \geq H_\infty(X|AUX) - \lambda^\varepsilon \geq \alpha' - \lambda^\varepsilon \geq \alpha$. Hence the new construction is an α' -ED-Extractor according to the stronger definition. The takeaway is that (as long as we're only considering polynomial-time distinguishers) it does not make much difference whether or not we condition on the seed in the definition.

Remark on Output Size. Note that if we have an α -ED-Extractor with output size $m = \lambda^\varepsilon$ for some constant $\varepsilon > 0$ then, assuming one-way functions, we can also construct an α -ED-Extractor for arbitrarily large output size $m = \lambda^c$ for any constant c just by using a pseudorandom generator (PRG) to expand the output. This holds as long as we're only considering polynomial-time distinguishers.

4 Security Without Auxiliary Info

4.1 Construction from Any PRF

We first show that every pseudorandom function (PRF) with a sufficiently high level of security is a good ED-Extractor in the setting without auxiliary info.

Theorem 3. *Let $F(\cdot, \cdot) : \{0, 1\}^\ell \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a pseudorandom function (PRF) with key-length $\ell = \ell(\lambda)$, input length $n = n(\lambda)$ and output length $m = m(\lambda)$, having security level $\sigma(\lambda) = 2^{m(\lambda)}\omega(\log \lambda)$. Define (SeedGen, EDExt) where SeedGen(1^λ) outputs $\text{seed} \leftarrow \{0, 1\}^{\ell(\lambda)}$ and EDExt(x, seed) = $F(\text{seed}, x)$. Then (SeedGen, EDExt) is an α -ED Extractor without auxiliary info for any $\alpha \geq m + \omega(\log \lambda)$. Furthermore, it has security for unbounded distinguishers.*

Proof. Assume that $(\mathcal{S}, \mathcal{D})$ is some α -legal source and distinguisher pair with advantage $\varepsilon = \varepsilon(\lambda)$ in the ED-Extractor security game. Assume that \mathcal{S} runs in polynomial time, but \mathcal{D} can be unbounded. We define a polynomial-time adversary \mathcal{A} that has $(\varepsilon^2 - \text{negl}(\lambda))/2^m$ advantage in the PRF game. In particular, $\mathcal{A}^{\mathcal{O}(\cdot)}$ is given access to an oracle \mathcal{O} and runs $\mathcal{S}^{\mathcal{O}(\cdot)}$ twice with independent randomness to derive two values x, x' . Then $\mathcal{A}^{\mathcal{O}(\cdot)}$ computes $r = \mathcal{O}(x), r' = \mathcal{O}(x')$. If $r = r'$, it outputs 1 else 0.

Firstly, consider the experiment where we sample $k \leftarrow \{0, 1\}^\ell, x \leftarrow \mathcal{S}^{F(k, \cdot)}$, $r = F(k, x)$ and let K, R be the random variables for the values k, r respectively.

Then the statistical distance $\text{SD}(((K, R), (K, U_m)) \geq \varepsilon$ since \mathcal{D} distinguishes the two distributions with probability ε . Therefore, by Claim 1, we have $\text{Col}(R|K) \geq \frac{1}{2^m}(1 + 4\varepsilon^2)$ where Col denotes the collision probability (Definition 3). It's easy to see that, by the definition of \mathcal{A} , we have $\Pr[\mathcal{A}^{F(k,\cdot)} = 1 : k \leftarrow \{0, 1\}^\ell] = \text{Col}(R|K) \geq \frac{1}{2^m}(1 + 4\varepsilon^2)$.

Secondly, consider the experiment where we sample $k \leftarrow \{0, 1\}^\ell$ and then sample $x \leftarrow \mathcal{S}^{F(k,\cdot)}, x' \leftarrow \mathcal{S}'^{F(k,\cdot)}$ by running \mathcal{S} twice with independent randomness and let K, X, X' be the random variables for the value k, x, x' in the experiment. Since \mathcal{S} is an α -legal source we know that:

- The probability that \mathcal{S} queried the oracle on x during the first run or on x' during the second run is negligible.
- Since $H_\infty(X|K) = H_\infty(X'|K) \geq \alpha \geq m + \omega(\log \lambda)$, the probability that either (1) \mathcal{S} queried the oracle on x' during the first run or (2) \mathcal{S} queried the oracle on x during the second run or (3) $x = x'$ is bounded by $\text{negl}(\lambda)/2^m$.

To summarize, in the above experiment, if we define the ‘bad event’ that $x = x'$ or that the oracle is queried on one of x, x' during the course of the experiment, then the probability of the bad event is at most $\text{negl}(\lambda)/2^m$. Now, consider the modified experiment where we sample $x \leftarrow \mathcal{S}^{U(\cdot)}, x' \leftarrow \mathcal{S}'^{U(\cdot)}$ and U is a truly random function. By $\sigma(\lambda)$ security of the PRF, the probability of the bad event occurring in the modified experiment is still bounded by $\text{negl}(\lambda)/2^m$. If the bad event does not occur, then $r = U(x), r' = U(x')$ are random and independent values and therefore $\Pr[r = r'] = \frac{1}{2^m}$. This shows that $\Pr[\mathcal{A}^{U(\cdot)} = 1] \leq (1 + \text{negl}(\lambda))2^m$.

This shows that the advantage of \mathcal{A} in the PRF security game is $(4\varepsilon^2(\lambda) - \text{negl}(\lambda))/2^m$ which must be $\leq 1/\sigma(\lambda) \leq \text{negl}(\lambda)/2^m$, by the $\sigma(\lambda)$ security of the PRF. Therefore $\varepsilon(\lambda) = \text{negl}(\lambda)$, which concludes the proof of the ED-Extractor security.

Corollary 1. *Assuming the existence of sub-exponentially secure one-way functions, for any polynomial input size $n = n(\lambda)$ the following holds:*

- *For any polynomial output size $m = m(\lambda)$, there exists an α -ED Extractor in the setting without auxiliary info and with security for unbounded distinguishers as long as $\alpha \geq m + \omega(\log \lambda)$.*
- *For any constant $\varepsilon > 0$ and any polynomial output size $m = m(\lambda)$, there exists an α -ED Extractor in the setting without auxiliary info and security for polynomial-time distinguishers as long as $\alpha \geq \lambda^\varepsilon$.*

Proof. We note that sub-exponentially secure one-way functions imply the existence of PRFs with security level $2^{p(\lambda)}$ for any polynomial p (by making the key sufficiency large). Therefore the first part of the corollary follows directly from the preceding Theorem. The second part follows by using a pseudorandom generator (PRG) to expand the output-size of the ED-Extractor as discussed in the Remark on Output Size in Sect. 3.

4.2 Necessity of One-Way Functions

Theorem 4. *For any input length $n = n(\lambda)$, the existence of an $(\alpha = n - 1)$ -ED-Extractor, even without auxiliary info and even with output length $m = 1$, implies the existence of one-way functions. Furthermore, such extractors cannot be secure for computationally unbounded sources, even if we restrict to polynomial-time distinguishers.*

Proof. Let $(\text{SeedGen}, \text{EExt})$ be an ED Extractor as in the theorem statement. Assume $\text{SeedGen}(1^\lambda)$ uses at most $\ell = \ell(\lambda)$ bits of randomness and let $q = 7\ell + \lambda$. Define the function $f(r, x_1, \dots, x_q) = (x_1, \dots, x_q, y_1, \dots, y_q)$, which takes as input a uniformly random $r \in \{0, 1\}^\ell$ and $x_i \in \{0, 1\}^n$ and computes $\text{seed} = \text{SeedGen}(1^\lambda; r)$ and $y_i = \text{EExt}(\text{seed}, x_i)$ for $i \in [q]$. Then we claim that f is a one-way function.

Assume by way of contradiction that a polynomial-size adversary \mathcal{A} breaks the one-wayness of f with non-negligible probability. We define a source $\mathcal{S}^{\text{EExt}(\text{seed}, \cdot)}$ as follows:

1. Choose x_1, \dots, x_q uniformly at random from $\{0, 1\}^n$. Query the oracle to learn $y_i = \text{EExt}(\text{seed}, x_i)$ for each $i \in [q]$.
2. Run $\mathcal{A}(x_q, \dots, x_q, y_1, \dots, y_q)$ and get some value (r', x'_1, \dots, x'_q) .
3. Test if $f(r', x_q, \dots, x'_q) = (x_1, \dots, x_q, y_1, \dots, y_q)$. If not, output a uniformly random $x_0^* \leftarrow \{0, 1\}^n$ and halt. Else continue.
4. Compute $\text{seed}' = \text{SeedGen}(1^\lambda; r')$. Choose a random $x_1^* \leftarrow \{0, 1\}^n$ and if $\text{EExt}(\text{seed}', x_1^*) = 0$ output x_1^* and halt. Else continue.
5. Choose a random $x_2^* \leftarrow \{0, 1\}^n$ and output it.

We define a corresponding distinguisher $D(\text{seed}, r)$, which outputs r . We claim that the pair (\mathcal{S}, D) breaks the $(\alpha = n - 1)$ -ED-Extractor security.

Firstly, we claim that \mathcal{S} is an $(\alpha = n - 1)$ -legal source. It is easy to see that the probability of it outputting a value x that it previously queried is negligible since it outputs one of x_0^*, x_1^*, x_2^* each of which is individually uniformly random and independent of the prior queries. To analyze entropy, let us fix any choice of the values of $\text{seed}, x_1, \dots, x_q$ and randomness of \mathcal{A} and let X be the random variable for the output of \mathcal{S} in the above experiment. We argue that, even for any choice of the fixed values, it holds that $H_\infty(X) \geq n - 1$, which proves the claim. The fixed values determine whether the test in line 3 passes or fails. If it fails, then X is uniformly random and so $H_\infty(X) = n$. If it passes, then let us define the variable V which is 0 if x is output in line 4 and 1 if it is output in line 5. Let us define the value $P_0 = |\{x : \text{EExt}(x, \text{seed}) = 0\}|$. Then we have

$$\begin{aligned} \max_x \Pr[X = x] &= \max_x (\Pr[X = x|V = 0] \Pr[V = 0] + \Pr[X = x|V = 1] \Pr[V = 1]) \\ &\leq \frac{1}{P_0} \cdot \frac{P_0}{2^n} + \frac{1}{2^n} \left(1 - \frac{1}{P_0}\right) \\ &\leq 2^{-(n-1)} \end{aligned}$$

and therefore $H_\infty(X) \geq n - 1$.

Next, we analyze the success probability of the pair $(\mathcal{S}, \mathcal{D})$ in the ED-Extractor security game. If the challenger chooses the challenge bit $b = 1$ then, since r is uniformly random, we have $\Pr[b' = 1] = \frac{1}{2}$. On the other hand, let's analyze the security game when the challenge chooses the bit $b = 0$. Assume that the adversary \mathcal{A} breaks the security of the one-way function f with some non-negligible probability $\varepsilon = \varepsilon(\lambda)$. Then $\varepsilon(\lambda) \geq 1/p(\lambda)$ for some polynomial p and for infinitely many values of λ . We define several events in the context of the ED-Extractor security game with the particular sampler defined above:

FAIR: Let's call a seed *biased* if $\Pr_{x \leftarrow \{0,1\}^n} [\text{EExt}(\text{seed}, x) = 0] \leq \frac{1}{2} - \delta$ where we set $\delta := \frac{1}{20p}$. Let's define the event **FAIR** to occur if the seed is not biased. Since we assumed that the ED-Extractor is secure, it must be the case that probability that a random seed is biased is negligible (otherwise the sampler that outputs a random x and the distinguisher that tests if the seed is biased and if so outputs r else outputs random would break security). Therefore, $\Pr[\text{FAIR}] = 1 - \text{negl}(\lambda)$.

INV: Let this be the event that the test in line 3 of the execution of \mathcal{S} succeeds, meaning that \mathcal{A} succeeded to invert correctly. By definition $\Pr[\text{INV}] = \varepsilon$.

CLOSE: Let this be the event that for the value seed' computed in line 4, it holds that

$$\Pr_{x \leftarrow \{0,1\}^n} [\text{EExt}(x, \text{seed}) = \text{EExt}(x, \text{seed}')] \geq .9$$

where, if the process terminates before line 4, we define $\text{seed}' = \text{seed}$. If **CLOSE** does not occur, it means that there exists some seed' for which $\Pr_{x \leftarrow \{0,1\}^n} [\text{EExt}(x, \text{seed}) = \text{EExt}(x, \text{seed}')] < .9$ but for all $i \in [q]$ it holds that $\text{EExt}(x_i, \text{seed}) = \text{EExt}(x_i, \text{seed}')$. The probability of this happening for any fixed seed' is $.9^q \leq .9^{7\ell + \lambda} \leq 2^{-\ell} \text{negl}(\lambda)$. By taking a union bound over all 2^ℓ values of seed' the probability that some such seed' exists is negligible and therefore $\Pr[\text{CLOSE}] \geq 1 - \text{negl}(\lambda)$.

For simplicity, we also define the event $\text{IFC} = \text{INV} \wedge \text{FAIR} \wedge \text{CLOSE}$. When $b = 0$ we therefore have:

$$\begin{aligned} \Pr[b' = 0] &\geq \Pr[b' = 0 \wedge \text{INV}] + \Pr[b' = 0 \wedge \neg \text{INV}] \\ &\geq \Pr[b' = 0 \wedge \text{INV} \wedge \text{FAIR} \wedge \text{CLOSE}] + \Pr[b' = 0 \wedge \neg \text{INV} \wedge \text{FAIR}] \\ &\geq \Pr[b' = 0 | \text{IFC}] (\Pr[\text{INV}] - \Pr[\neg \text{FAIR}] - \Pr[\neg \text{CLOSE}]) \\ &\quad + \Pr[b' = 0 | \neg \text{INV} \wedge \text{FAIR}] (\Pr[\neg \text{INV}] - \Pr[\neg \text{FAIR}]) \\ &\geq \Pr[b' = 0 | \text{IFC}] (\varepsilon - \text{negl}(\lambda)) + \Pr[b' = 0 | \neg \text{INV} \wedge \text{FAIR}] (1 - \varepsilon - \text{negl}(\lambda)) \\ &\geq \Pr[b' = 0 | \text{IFC}] (\varepsilon - \text{negl}(\lambda)) + \left(\frac{1}{2} - \delta \right) (1 - \varepsilon - \text{negl}(\lambda)) \end{aligned}$$

To analyze $\Pr[b' = 0 | \text{IFC}]$ let us fix all randomness z of the experiment except for the choice of x_1^*, x_2^* , such that this fixing makes the event **IFC** occurs. Let IFC_z be the event that the randomness takes on this value. For any such choice, let E_1 be

the event that $\text{EDExt}(x_1^*, \text{seed}) = 0$, let E'_1 be the event that $\text{EDExt}(x_1^*, \text{seed}') = 0$, let A be the event that $\text{EDExt}(x_1^*, \text{seed}) = \text{EDExt}(x_1^*, \text{seed}')$ and let E_2 be the event that $\text{EDExt}(\text{seed}, x_2^*) = 0$, where the randomness is only over the choice of x_1^*, x_2^* . Since we conditioned on **CLOSE** we have $\Pr[A] \geq .9$. Since we conditioned on **FAIR** we have $\Pr[E_1] \geq (1/2 - \delta), \Pr[E_2] \geq (1/2 - \delta)$. Therefore, for any such choice of randomness z we have:

$$\begin{aligned} \Pr[b' = 0 | \text{IFC}_z] &= \Pr[E_1 \wedge E'_1] + \Pr[E_2 \wedge \neg E'_1] \\ &= \Pr[A \wedge E'_1] + \Pr[E_2] (1 - \Pr[E'_1]) \\ &\geq \Pr[E'_1] - \Pr[\neg A] + \left(\frac{1}{2} - \delta\right) (1 - \Pr[E'_1]) \\ &\geq \frac{1}{2} - \delta - .1 + \frac{1}{2} \Pr[E'_1] \\ &\geq \frac{1}{2} - \delta - .1 + \frac{1}{2} (\Pr[E_1] - \Pr[\neg A]) \\ &\geq \frac{1}{2} - \delta - .1 + \frac{1}{2} (\frac{1}{2} - \delta - .1) \\ &\geq .6 - \frac{3}{2} \delta \end{aligned}$$

which also implies that $\Pr[b' = 0 | \text{IFC}] \geq .6 - \frac{3}{2} \delta$. Combining, we have:

$$\begin{aligned} \Pr[b' = 0] &\geq \left(.6 - \frac{3}{2} \delta\right) (\varepsilon - \text{negl}(\lambda)) + \left(\frac{1}{2} - \delta\right) (1 - \varepsilon - \text{negl}(\lambda)) \\ &\geq \frac{1}{2} - \delta + \varepsilon(.1 - \delta/2) - \text{negl}(\lambda) \\ &\geq \frac{1}{2} + \varepsilon/10 - (3/2)\delta - \text{negl}(\lambda) \\ &\geq \frac{1}{2} + \frac{1}{10p(\lambda)} - \frac{3}{40p(\lambda)} - \text{negl}(\lambda) \\ &\geq \frac{1}{2} + \frac{1}{40p(\lambda)} - \text{negl}(\lambda) \end{aligned}$$

for infinitely many values of λ . Therefore $\Pr[b' = b] - \frac{1}{2}$ is non-negligible, which leads to a contradiction and hence f must be one-way.

For the second part of the theorem, note that we showed how to convert an inverter for f into a source \mathcal{S} together with an efficient distinguisher \mathcal{D} that break ED-Extractor security. Since an inefficient inverter for f always exists, it means that there exists an inefficient source \mathcal{S} and an efficient distinguisher \mathcal{D} that break the security of the ED-Extractor.

5 Security with Auxiliary Info

5.1 Construction via Constrained PRFs

We now show how to construct an ED-Extractor in the setting with auxiliary info, using constrained PRFs (Definition 5) and standard seeded extractors (Definition 1).

Construction. Let $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^\ell$ be an (α', ε) -seeded extractor for some lengths $n = n(\lambda), d = d(\lambda), \ell = \ell(\lambda)$ and some $\alpha' = \alpha'(\lambda), \varepsilon = \varepsilon(\lambda)$. Further let Ext also be a universal hash function. Let $(F, \text{Constrain})$ be a constrained PRF with input length n and output length $m = m(\lambda)$ for the class of constraints $\mathcal{C} = \{C_{s,u}\}_{s \in \{0,1\}^d, u \in \{0,1\}^\ell}$ where $C_{s,u}(x) = 1$ iff $\text{Ext}(x; s) = u$. We construct an ED-Extractor ($\text{SeedGen}, \text{EExt}$) as follows:

- $\text{SeedGen}(1^\lambda)$: Choose a random $k \leftarrow \{0, 1\}^\lambda$. Choose a random $s \leftarrow \{0, 1\}^d, u \leftarrow \{0, 1\}^\ell$ and let $C_{s,u} \in \mathcal{C}$ be the corresponding constraint. Let $k\{C_{s,u}\} \leftarrow \text{Constrain}(k, C_{s,u})$. Output $\text{seed} = k\{C_{s,u}\}$.
- $\text{EExt}(x, \text{seed})$: Output $F(k\{C_{s,u}\}, x)$.

Note that F always outputs some value, even if x is in the constrained set. Without loss of generality, we can assume that the constrained key $k\{C_{s,u}\}$ reveals s, u in the clear and that, $F(k\{C_{s,u}\}, x)$ outputs 0^m whenever $C_{s,u}(x) = 1$.

Theorem 5. *Assuming the constrained PRF has no-constrained-evaluation security, the construction above is an α -entropy secure ED-Extractor for $\alpha = \alpha' + m$, as long as the parameters satisfy $\ell(\lambda) = \omega(\log \lambda)$, and $\varepsilon(\lambda) = \text{negl}(\lambda)$.*

Proof. Our proof of security follows by a sequence of hybrid games:

Hybrid 0: This is the game $\text{EDGame}^{\mathcal{S}, \mathcal{D}}(1^\lambda)$ with a source \mathcal{S} and a distinguisher \mathcal{D} as in Definition 6. The game proceeds as follows:

- Sample a random bit $b \leftarrow \{0, 1\}$ and a random $\text{seed} \leftarrow \text{SeedGen}(1^\lambda)$. The latter consists of sampling $k \leftarrow \{0, 1\}^\lambda, s \leftarrow \{0, 1\}^d, u \leftarrow \{0, 1\}^\ell, k\{C_{s,u}\} \leftarrow \text{Constrain}(k, C_{s,u})$ and setting $\text{seed} = k\{C_{s,u}\}$.
- Run $(x, \text{aux}) \leftarrow \mathcal{S}^{\text{EExt}(\cdot, \text{seed})}(1^\lambda)$.
- If $b = 0$ set $r = \text{EExt}(x, \text{seed})$ else if $b = 1$ set $r \leftarrow \{0, 1\}^m$.
- Let $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$.

Hybrid 1: In this game, instead of giving the source $\mathcal{S}^{\text{EExt}(\cdot, \text{seed})}$ access to the oracle $\text{EExt}(\cdot, \text{seed}) = F(k\{C_{s,u}\}, \cdot)$, we replace it with the oracle $F(k, \cdot)$ using the unconstrained key k . Furthermore, if $b = 0$, instead of setting $r = \text{EExt}(x, \text{seed}) = F(k\{C_{s,u}\}, x)$, we now set $r = F(k, x)$. In detail, the hybrid is defined as follows:

1. Sample a random bit $b \leftarrow \{0, 1\}$ and a random $k \leftarrow \{0, 1\}^\lambda$.
2. Run $(x, \text{aux}) \leftarrow \mathcal{S}^{F(k, \cdot)}(1^\lambda)$.
3. If $b = 0$ set $r = F(k, x)$ else if $b = 1$ set $r \leftarrow \{0, 1\}^m$. Choose $s \leftarrow \{0, 1\}^d, u \leftarrow \{0, 1\}^\ell$ and $\text{seed} \leftarrow \text{Constrain}(k, C_{s,u})$.
4. Let $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$.

Hybrids 0 and 1 are indistinguishable. The only time Hybrid 0 differs from Hybrid 1 is if in Hybrid 0 either: (A) some oracle query or the final output x produced by \mathcal{S} satisfy $\text{Ext}(x; s) = u$, or (B) some oracle query or the final output x produced by \mathcal{S} satisfy $C_{s,u}(x) = 0 \wedge F(k, x) \neq F(k\{C_{s,u}\}, x)$. Since u is uniformly random, the probability of (A) happening when \mathcal{S} makes q queries is at most $(q + 1)/2^\ell$ which is negligible. By the correctness of the constrained PRF, the probability of (B) happening is also negligible.

Hybrid 2: This is the same as Hybrid 1, except that we give the source access to an oracle $\text{EExt}(\cdot; \text{seed}')$ where $\text{seed}' = k\{C_{s',u'}\} \leftarrow \text{Constrain}(k, C_{s',u'})$ is a constrained PRF key for random and independent values s', u' . In detail, the hybrid is defined as follows:

1. Sample a random bit $b \leftarrow \{0, 1\}$ and a random $k \leftarrow \{0, 1\}^\lambda$. Choose $s' \leftarrow \{0, 1\}^d, u' \leftarrow \{0, 1\}^\ell$ and $\text{seed}' \leftarrow \text{Constrain}(k, C_{s',u'})$.
2. Run $(x, \text{aux}) \leftarrow \mathcal{S}^{\text{EExt}(\cdot, \text{seed}')}(1^\lambda)$.
3. If $b = 0$ set $r = F(k, x)$ else if $b = 1$ set $r \leftarrow \{0, 1\}^m$. Choose $s \leftarrow \{0, 1\}^d, u \leftarrow \{0, 1\}^\ell$ and $\text{seed} \leftarrow \text{Constrain}(k, C_{s,u})$.
4. Let $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$.

Hybrids 1 and 2 are statistically close. The only time Hybrid 1 differs from Hybrid 2 is if in Hybrid 2 either: (A) some oracle query x_i produced by \mathcal{S} satisfies $\text{Ext}(x_i; s') = u'$, or (B) some oracle query x_i produced by \mathcal{S} satisfy $C_{s',u'}(x) = 0 \wedge F(k, x) \neq F(k\{C_{s',u'}\}, x)$. Since u' is uniformly random, the probability of (A) happening when \mathcal{S} makes q queries is at most $q/2^\ell$ which is negligible. By the correctness of the constrained PRF, the probability of (B) happening is also negligible.

Hybrid 3: This is the same as Hybrid 2, except that in step 3, instead of choosing $u \leftarrow \{0, 1\}^\ell$ we now set $u = \text{Ext}(x; s)$. In detail, the hybrid is defined as follows:

1. Sample a random bit $b \leftarrow \{0, 1\}$ and a random $k \leftarrow \{0, 1\}^\lambda$. Choose $s' \leftarrow \{0, 1\}^d, u' \leftarrow \{0, 1\}^\ell$ and $\text{seed}' \leftarrow \text{Constrain}(k, C_{s',u'})$.
2. Run $(x, \text{aux}) \leftarrow \mathcal{S}^{\text{EExt}(\cdot, \text{seed}')}(1^\lambda)$.
3. If $b = 0$ set $r = F(k, x)$ else if $b = 1$ set $r \leftarrow \{0, 1\}^m$. Choose $s \leftarrow \{0, 1\}^d, u = \text{Ext}(x; s)$ and $\text{seed} \leftarrow \text{Constrain}(k, C_{s,u})$.
4. Let $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$.

Hybrids 2 and 3 are statistically close if Ext is an (α, ε) -extractor. To argue this, let us use capital letters to denote random variables for the corresponding values in the experiment. Firstly, note that the view of the source \mathcal{S} in hybrid 2 is identically distributed to that of hybrid 0.³ Therefore, we can rely on the legality to \mathcal{S} (which is defined relative to the distribution of hybrid 0) to argue that $H_\infty(X|\text{AUX}, \text{SEED}') \geq \alpha$. By Lemma 1, we also have $H_\infty(X|\text{AUX}, \text{SEED}', R) \geq \alpha - m \geq \alpha'$. Lastly since K is independent of X when conditioned on SEED', R , we also have $H_\infty(X|\text{AUX}, K, R) \geq \alpha'$. Therefore, by the security of the extractor, $U = \text{Ext}(X; S)$ is statistically close to a uniformly random and independent U even given AUX, K, R, S . Lastly, since the view of \mathcal{D} in hybrids 2 and 3 is a function of AUX, K, R, S, U where

³ This was the reason for introducing hybrid 2 rather than directly going from 1 to 3.

$U = \text{Ext}(X; S)$ in hybrid 3 and U is uniform/independent in hybrid 2, the two hybrids are statistically close.

Hybrid 4: This is the same as Hybrid 3, except that we switch back from giving S oracle access to $\text{EDExt}(\cdot, \text{seed}')$ to giving it access to the unconstrained PRF $F(k, \cdot)$. In detail, the hybrid is defined as follows:

1. Sample a random bit $b \leftarrow \{0, 1\}$ and a random $k \leftarrow \{0, 1\}^\lambda$.
2. Run $(x, \text{aux}) \leftarrow \mathcal{S}^{F(k, \cdot)}(1^\lambda)$.
3. If $b = 0$ set $r = F(k, x)$ else if $b = 1$ set $r \leftarrow \{0, 1\}^m$. Choose $s \leftarrow \{0, 1\}^d$ $u = \text{Ext}(x; s)$ and $\text{seed} \leftarrow \text{Constrain}(k, C_{s,u})$.
4. Let $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$.

Hybrids 3 and 4 are indistinguishable by the same argument as the indistinguishability of hybrid 1 and 2.

Advantage in Hybrid 4: We now claim that in Hybrid 4, the advantage $|\Pr[b = b'] - \frac{1}{2}|$ is negligible by the no-constrained-evaluation security of the constrained PRF. In particular, we define a reduction that runs $(x, \text{aux}) \leftarrow \mathcal{S}^{F(k, \cdot)}(1^\lambda)$ by making queries to its PRF oracle. The reduction chooses $s \leftarrow \{0, 1\}^d$, sets $u = \text{Ext}(x; s)$ and gives the constraint $C_{s,u}$ together with the value x to its challenger. Since S is a legal source, x was never queried by the oracle and, by the definition of the constraint, we have $C_{s,u}(x) = 1$. Secondly, since $\text{Ext}(\cdot; s)$ is a universal hash function, the probability that of any of the previous queries x_i made by S satisfy $\text{Ext}(x_i; s) = \text{Ext}(x; s)$ is also negligible. Therefore, our reduction makes no constrained-evaluation queries to the PRF.

So, the reduction is a legal attacker in the no-constrained-evaluation security game of constrained PRF. The reduction receives a value r , which is either $F(k, x)$ or uniform, along with a constrained key $k\{C_{s,u}\}$. It sets $\text{seed} = k\{C_{s,u}\}$ and outputs the bit $b' = \mathcal{D}(1^\lambda, \text{seed}, \text{aux}, r)$. The advantage of the reduction in the constrained PRF security game is exactly the same as that of the adversary in hybrid 3, and therefore the latter is negligible.

Since the advantage in hybrid 3 is negligible and hybrid 3 is indistinguishable from hybrid 0, the advantage in hybrid 0 must be negligible as well. This proves the theorem.

Corollary 2. *Under the sub-exponential security of either the LWE assumption or the DDHI assumption in an arbitrary prime-order group, there exists an ED-Extractor for α -entropy sources with auxiliary info, for any $\alpha = \lambda^{\Omega(1)}$ and with any polynomial input length n and output length m . Security holds against polynomial-time sources and distinguishers.*

Proof. The work of [BV15] construct selectively secure constrained PRFs for all circuits from LWE. We can then use complexity leveraging to get adaptive security by assuming sub-exponential LWE. The results of [AMN+18] constructs no-evaluation secure PRFs for NC1 from the DDHI assumption in arbitrary prime-order groups (the also construct selectively secure PRFs from the DDHI assumption in specific groups). We then use complexity leveraging to get no-constrained-evaluation security under sub-exponential DDHI, as discussed in the remarks after Definition 5.

We use an extractor with output length $\alpha/4$ which is secure for entropy $\alpha' = \alpha/2$ with $\varepsilon = 2^{-(\alpha/8)} = \text{negl}(\lambda)$. We combine this with a constrained PRF with output length $m = \alpha/2$ which ensures $\alpha \geq \alpha' + m$. This gives us an ED-Extractor with output length $\alpha/2 = \lambda^{\Omega(1)}$. We can then use a PRG to then get arbitrarily large polynomial output size as discussed in the Remark on Output Size in Sect. 3.

5.2 Negative Results for ED Extractors with Auxiliary Info

Our constructions of ED-Extractors in the auxiliary info setting have several disadvantages compared to our construction in the setting without auxiliary info. Firstly, in the auxiliary info setting we needed complex constructions based on “cryptomania” assumptions (LWE and DDHI), whereas in the setting without auxiliary info, we showed that any sufficiently secure PRF is a good ED-Extractor. Secondly, in the auxiliary info setting we only achieved security for polynomial-time distinguishers while in the setting without auxiliary info we got security even for computationally unbounded distinguishers. In this section, we give some evidence that the two settings are substantially different and that we indeed need to work harder and cannot hope for as much in the setting with auxiliary info.

Not All PRFs Are ED-Extractors with Aux Info. Firstly, we show that *not* every PRF is a good ED-Extractor in the setting with auxiliary info. We give two variants of this result. The first is based on collision-resistant hash functions (CRHFs) and gives a PRF that is not an α -ED-Extractor for entropy $\alpha = n - \lambda^\varepsilon$. The second one is based on fully homomorphic encryption and gives a PRF that is not an α -ED Extractor even for entropy $\alpha = n - 1$. In both cases, the result holds even if the PRF/ED-Extractor only outputs 1 bit.

CRHF-based Construction. Let $F' : \{0, 1\}^\ell \times \{0, 1\}^{n'} \rightarrow \{0, 1\}$ be a PRF with key-length $\ell = \ell(\lambda)$, input length $n' = n'(\lambda)$ and output length 1. Let $H : \{0, 1\}^d \times \{0, 1\}^n \rightarrow \{0, 1\}^{n'}$ be a collision-resistant hash function (CRHF) with seed length $d = d(\lambda)$, input length $n = n(\lambda)$ and output length $n' = n'(\lambda)$. We define a PRF $F : \{0, 1\}^{\ell+d} \times \{0, 1\}^n \rightarrow \{0, 1\}$ as follows. Parse the key $k = (k', s)$ with $k' \in \{0, 1\}^\ell, s \in \{0, 1\}^d$. Define $F(k, x)$:

- If $x \leq d$ output $s[x]$, where we interpret x as an integer in the range $[2^n]$ and $s[x]$ denotes the x 'th bit of s .
- Else output $F(k', H(s, x))$.

It is easy to see that F is a PRF if F' is a PRF and H is a CRHF. On the other hand it is not an $\alpha = (n - n')$ -ED-Extractor. In particular, consider the source that queries the oracle on values $1, \dots, d$ to learn the CRHF seed s . It then chooses a random $x \leftarrow \{0, 1\}^n$ and outputs $x, \text{aux} = H(s, x)$. It is clearly an α legal source. Yet we can define a distinguisher D that gets $k = (k', s), \text{aux}, r$ and outputs 1 iff $r = F(k', \text{aux})$. Then D always outputs 1 if r is the outputs of the

ED-Extractor on x but only outputs 1 with probability $1/2$ if r is truly random, giving it a non-negligible advantage of $1/2$. For parameters, we note that the existence of CRHFs implies the existence of a CRHF with arbitrary polynomial input size $n = n(\lambda)$ and output size λ^ϵ for any constant $\epsilon > 0$. Therefore, we get a PRF with arbitrary polynomial input size $n = n(\lambda)$ and output size $m = 1$, which is not an α -ED-Extractor for $\alpha = n - \lambda^\epsilon$.

Theorem 6. *Assuming the existence of collision-resistant hash functions, for every polynomial $n = n(\lambda)$ and every constant $\epsilon > 0$ there exists a PRF with n -bit input and 1-bit output which is not a secure α -ED-Extractor with auxiliary input for $\alpha = n - \lambda^\epsilon$.*

FHE-based Construction. Let $F' : \{0, 1\}^\ell \times \{0, 1\}^{n'} \rightarrow \{0, 1\}$ be a PRF with key-length $\ell = \ell(\lambda)$, input length $n' = n'(\lambda)$ and output length 1. Let $(\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$ be an FHE scheme capable of evaluating the PRF F' . Furthermore assume that the ciphertexts are pseudorandom and that the Eval procedure is statistically circuit private. Assume that the key-generation algorithm and the encryption algorithm each use at most $d = d(\lambda)$ bits of randomness, and that the encryption of an ℓ -bit message produces an ℓ' -bit ciphertext. Define the PRF $F : \{0, 1\}^{\ell+2d} \times \{0, 1\}^n \rightarrow \{0, 1\}$ as follows. Parse the key $k = (k', s_1, s_2)$ with $k' \in \{0, 1\}^\ell, s_1, s_2 \in \{0, 1\}^d$. Define $F(k, x)$:

- Check if $x \leq \ell'$ (where we interpret x as an integer in the range $[2^n]$). If so let $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda; s_1), \text{ct} \leftarrow \text{Enc}(\text{pk}, k; s_2)$. Output the x 'th bit of ct denoted by $\text{ct}[x]$.
- Else output $F(k, x)$.

It is easy to see that F is a secure PRF: by the security of the FHE with pseudorandom ciphertexts, we can replace ct by a uniformly random value independent of k , and by the security of the PRF F' the above is then a good PRF. On the other hand it is not an $\alpha = (n - 1)$ -ED-Extractor. In particular, consider the source that queries the oracle on values $1, \dots, \ell'$ to learn the the ciphertext ct . It then chooses a random $x \leftarrow \{0, 1\}^n$ and outputs $x, \text{aux} = \text{Eval}(F'(\cdot, x), \text{ct})$ so that aux is an FHE encryption of $F'(k, x)$. Since Eval is circuit private aux does not reveal anything about x beyond $F(k, x)$ and therefore is an $\alpha = n - 1$ legal source. Yet we can define a distinguisher D that gets $k = (k', s_1, s_2), \text{aux}, r$ and outputs 1 iff $\text{Dec}(\text{sk}, \text{aux}) = r$ where $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^\lambda; s_1)$. Then D outputs 1 with probability 1) if r is the outputs of the ED-Extractor on x , but only outputs 1 with probability $1/2$ if r is truly random, giving it a non-negligible advantage of $1/2 - \text{negl}(\lambda)$. Therefore, we get a PRF with arbitrary polynomial input size $n = n(\lambda)$ and output size $m = 1$, which is not an α -ED-Extractor for $\alpha = n - 1$.

Theorem 7. *Assuming the existence of Fully Homomorphic Encryption (FHE) with statistical circuit privacy and pseudorandom ciphertexts, for every polynomial $n = n(\lambda)$ there exists a PRF with n -bit input and 1-bit output which is not a secure α -ED-Extractor with auxiliary input for $\alpha = n - 1$.*

Black-Box Separations. We now show give two black-box separation results, showing that certain types of ED-Extractors cannot be proven secure via a black-box reduction from virtually any “standard” computational assumption (e.g., including general assumptions such as the existence of one-way functions or public-key encryption, as well as specific assumptions such as DDH, LWE, RSA, etc., even if we assume (sub-)exponential security). In particular, we show two results of this type. Firstly, we show that one cannot prove the security of any ED-Extractor in the auxiliary info setting against computationally unbounded distinguishers (and polynomial-time sources) under such assumptions. This is contrast to the setting without auxiliary info, where we were able to do so. Secondly, we show that one cannot prove security in the auxiliary input setting (even for polynomial-time sources and distinguisher) of any ED-Extractor that has a certain type of *seed-committing* property: if you query the extractor EDExt on some polynomial set of values x_1, \dots, x_q then the output uniquely fixes a single possible seed that could have produced it. This is true for many natural constructions, such as the Naor-Reingold PRF or most block-cipher and hash-function based constructions. (But is crucially not true for our constructions based on constrained PRFs.) We view this as partial evidence that more complex constructions are necessary in the setting with auxiliary info.

Note that these results do not show that ED-Extractors with such properties cannot be constructed; in fact the work of Coretti et al. [CDKT19] in the random-oracle model can be interpreted as showing that “good” hash functions are heuristically likely to be good ED-Extractors in the auxiliary info setting with security even against computationally unbounded distinguishers, *and* they are also likely to be seed-committing. However, our results show that we cannot prove security under standard assumptions.

Our results are of the same flavor as the work of Wichs [Wic13]. They define the class of (single-stage) cryptographic game assumptions, which are modeled via a game between a challenger and a stateful adversary. They require that any polynomial-time (or sub-exponential time) attacker has at most a negligible (or inverse sub-exponential) success probability in winning the game. This captures essentially all standard assumptions used in cryptography. However, the security definition of ED-Extractors is *not* a single-stage game since it involves two separate entities (the source and the distinguisher) who cannot share state.

We use the “simulatable attacker” paradigm (also called a meta-reduction) to prove our black box separations. This paradigm is formalized in [Wic13] and we give a high-level overview. To prove a separation, we design a class of inefficient attackers \mathcal{A}_h indexed by some h that break the security property but otherwise satisfy any structural/legality conditions (e.g., being multi-stage, entropy conditions etc.). However we also design an efficient simulator \mathcal{A}' that may not satisfy such conditions, such that one cannot distinguish between black-box access to \mathcal{A}_h for a random h versus \mathcal{A}' . Therefore if some reduction can break an assumption given black-box access to every \mathcal{A}_h it would also be able to do so given access to \mathcal{A}' . If for any polynomial ℓ we can further show such a simulatable attack

which is $2^{-\ell(\lambda)}$ indistinguishable, then we also rule out black-box reductions under sub-exponential or even exponential assumptions.

Unbounded Distinguishers. We first give a black-box reduction for ED-Extractors in the auxiliary info setting with security against unbounded distinguishers. Since the distinguisher can be computationally unbounded, a black-box reduction cannot call it. Therefore it suffices to construct a class of simulatable inefficient sources \mathcal{A}_h that satisfy the legality conditions and ensure that for the output (x, aux) it holds that $\text{seed}, \text{aux}, \text{EDExt}(x, \text{seed})$, is statistically far from $\text{seed}, \text{aux}, u$ where u is uniform. Our a high level, the source \mathcal{A}_h that we construct makes oracle queries and inefficiently learns the function $\text{EDExt}(\cdot, \text{seed})$ sufficiently well to predict $\text{EDExt}(x, \text{seed})$ for a random x with high accuracy without querying it. It chooses such random x and sets aux to be a “statistically binding commitment” of its prediction for $\text{EDExt}(x, \text{seed})$. This ensures that the distribution of $(\text{seed}, \text{aux}, \text{EDExt}(x, \text{seed}))$ is statistically far from $(\text{seed}, \text{aux}, \text{uniform})$. The commitment is generated using an exponentially large random function h and can therefore be simultaneously statistically hiding and binding. Therefore this attack is simulatable by an efficient simulator that chooses a random x and outputs a commitment to a random value.

Theorem 8. *For any candidate ED-Extractor $(\text{SeedGen}, \text{EDExt})$ with $n(\lambda)$ -bit input and 1 bit output and for any polynomial $\ell = \ell(\lambda)$ there exists a $2^{-\ell(\lambda)}$ -simulatable attack against the $\alpha = (n-1)$ -ED-Extractor security of the candidate in the setting with auxiliary info and unbounded distinguishers.*

In particular, if there is a black-box reduction showing this type of security for the candidate based on the security of some cryptographic game \mathcal{G} , then \mathcal{G} is not secure. If the reduction is based on the $2^{\ell(\lambda)}$ -security of the game \mathcal{G} then \mathcal{G} is not $2^{\ell(\lambda)}$ secure.

Proof. Assume that the length of $\text{seed} \leftarrow \text{SeedGen}(1^\lambda)$ is bounded by $|\text{seed}| \leq p(\lambda)$ for some polynomial p . Let $q = q(\lambda) = 3p(\lambda) + \lambda$. Let H_λ be the set of all functions from $\{0, 1\}^{\ell(\lambda)}$ to $\{0, 1\}$. For any $h \in H_\lambda$, consider the inefficient source $\mathcal{S}_{\lambda, h}$ that chooses x_1, \dots, x_q uniformly at random and queries its oracle on them, gets back y_1, \dots, y_q , and finds the (lexicographically first) value seed' such that $\text{EDExt}(x_i, \text{seed}') = y_i$ for all $i \in [q]$. It chooses a random x , computes $z' = \text{EDExt}(x, \text{seed}')$ and sets $\text{aux} = (r, h(r) \oplus z')$ where $r \leftarrow \{0, 1\}^\ell$.

First we claim that for any $h \in H_\lambda$, the above source $\mathcal{S}_{\lambda, h}$ breaks the security of the ED-Extractor with auxiliary info and an unbounded distinguisher. It’s easy to see that $\mathcal{S}_{\lambda, h}$ is a legal source with entropy $n - 1$ since x is uniformly random and aux can reveal at most 1-bit of information z' about x . Secondly, we claim that if $\mathcal{S}_{\lambda, h}$ has oracle access to $\text{EDExt}(\cdot, \text{seed})$, then with overwhelming probability the value seed' that it finds must agree with seed on at least $3/4$ of all inputs. Otherwise there exists some seed' that agrees with seed on $< 3/4$ inputs yet agrees with it on x_1, \dots, x_q which occurs with probability at most $2^p(3/4)^q = \text{negl}(\lambda)$. This also implies that if we let $z' = \text{EDExt}(x, \text{seed}')$, $z = \text{EDExt}(x, \text{seed})$ in the experiment, then $z' = z$ with probability $3/4 - \text{negl}(\lambda)$. But

this shows that the distribution $(\text{seed}, \text{aux}, u = \text{EExt}(\text{seed}, x))$ is statistically far from $(\text{seed}, \text{aux}, u \leftarrow \{0, 1\})$ since in the first case, if we let $\text{aux} = (r, v)$ then $h(r) \oplus v = u$ with probability at least $3/4 - \text{negl}(\lambda)$ while in the second case this happens with probability at most $1/2$.

Secondly, we claim that for a random $h \leftarrow H_\lambda$, the above source $\mathcal{S}_{\lambda, h}$ can be simulated by an efficient \mathcal{S}'_λ that runs in time $\text{poly}(\lambda)$. We define \mathcal{S}'_λ which chooses x_1, \dots, x_q uniformly at random and queries its oracle on them, gets back y_1, \dots, y_q , and outputs a uniformly random $(r, v) \leftarrow \{0, 1\}^\ell \times \{0, 1\}$.

The only way that $\mathcal{S}_{\lambda, h}$ for a random h can be distinguished from \mathcal{S}'_λ using black-box access is if two different executions of \mathcal{S} use the same randomness r . Given Q queries to \mathcal{S} , this happens with probability at most $\text{poly}(Q)2^\ell$.

Seed-Committing Extractors. We show that one cannot prove security in the auxiliary input setting (even for polynomial-time sources and distinguisher) of any ED-Extractor that has a certain type of *seed-committing* property.

Definition 7. *An ED-Extractor is seed-committing if there exist some polynomial $q = q(\lambda)$ and some inputs $x_1, \dots, x_q \in \{0, 1\}^{n(\lambda)}$ such that for any seed, seed' for which $\text{EExt}(x_i, \text{seed}) = \text{EExt}(x_i, \text{seed}')$ for all $i \in [q]$ it must hold that for all x^* we have $\text{EExt}(x^*, \text{seed}) = \text{EExt}(x^*, \text{seed}')$.*

For example, if we use the Naor-Reingold PRF [NR97] as an ED-Extractor then it is seed-committing. Moreover, we believe that ED-Extractor constructions using standard hash-functions and block-cipher will be seed-committing.

Theorem 9. *For any candidate seed-committing ED-Extractor (SeedGen, EExt) with $n(\lambda)$ -bit input and $m(\lambda)$ bit output and for any polynomial $\ell = \ell(\lambda)$ there exists a $2^{-\ell(\lambda)}$ -simulatable attack against the $\alpha = (n - 1)$ -ED-Extractor security of the candidate in the setting with auxiliary info.*

In particular, if there is a black-box reduction showing this type of security for the candidate based on the security of some cryptographic game \mathcal{G} , then \mathcal{G} is not secure. If the reduction is based on the $2^{\ell(\lambda)}$ -security of the game \mathcal{G} then \mathcal{G} is not $2^{\ell(\lambda)}$ secure.

Proof. Let H_λ be the set of all pairs of functions $h_1 : \{0, 1\}^\ell \rightarrow \{0, 1\}^{q\ell+1}$, $h_2 : \{0, 1\}^{q\ell+1} \rightarrow \{0, 1\}^\ell$. First we define $(\text{Enc}_{h_1, h_2}, \text{Dec}_{h_1, h_2})$ to be an information-theoretic authenticated encryption scheme whose key is h_1, h_2 . In particular, $\text{Enc}_{h_1, h_2}(m) = (r, h_1(r) \oplus m, h_2(r, h_1(r) \oplus m))$ where $r \leftarrow \{0, 1\}^\ell$ is uniformly random and $\text{Dec}_{h_1, h_2}(r, c, \sigma) = h_1(r) \oplus c$ if $h_2(r, c) = \sigma$ and \perp otherwise.

For any $h = (h_1, h_2) \in H_\lambda$, consider an inefficient source/distinguisher pair $\mathcal{A}_{\lambda, h} = (\mathcal{S}_{\lambda, h}, \mathcal{D}_{\lambda, h})$ defined as follow. The source $\mathcal{S}_{\text{sec}, h}$ chooses x_1, \dots, x_q as given by the seed-committing definition and queries its oracle on them, gets back y_1, \dots, y_q , and finds the (lexicographically first) seed' such that $\text{EExt}(x_i, \text{seed}') = y_i$ for all $i \in [q]$. It chooses a random x , computes z' to be the first bit of $\text{EExt}(x, \text{seed}')$ and sets $\text{aux} \leftarrow \text{Enc}_h(y_1, \dots, y_q, z')$. The distinguisher $\mathcal{D}_{\lambda, h}$ gets $(\text{seed}, \text{aux}, u)$, it computes z to be the first bit of u . It sets

$\text{Dec}_h(\text{aux}) = (y_1, \dots, y_q, z')$. If $\text{EExt}(\text{seed}, x_i) = y_i$ for all $i \in [q]$ and $z' = z$ it outputs 0 else 1.

It is easy to see that, for any h , the adversary $\mathcal{A}_{\lambda,h}$ is an $\alpha = (n - 1)$ -legal adversary and breaks ED-Extractor security with advantage $1/4$: If the challenge bit is $b = 0$, the distinguisher always outputs 0 and if the challenge bit is $b = 1$ the distinguisher only outputs 1 with probability $> 1/2$.

Secondly, for a random $h = (h_1, h_2)$ the adversary $\mathcal{A}_{\lambda,h}$ can be efficiently simulated by a stateful adversary $\mathcal{A}' = (\mathcal{S}', \mathcal{D}')$ that acts as both the source and the distinguisher but allows them to share state. On input y_1, \dots, y_q to \mathcal{S}' , it chooses a random x , aux and remembers the tuple $(\text{aux}, y_1, \dots, y_q, x)$. On input $(\text{seed}, \text{aux}, u)$ to \mathcal{D}' it checks if it stores a tuple of the form $(\text{aux}, y_1, \dots, y_q, x)$. If it does store such a tuple and $\text{EExt}(\text{seed}, x_i) = y_i$ for all $i \in [q]$ and u is equal to the first bit of $\text{EExt}(x, \text{seed})$ it outputs 0 else 1.

To show that one cannot distinguish between black-box access to \mathcal{A} vs \mathcal{A}' we define an intermediate \mathcal{A}^* which is inefficient but also stateful. In particular, $\mathcal{A}^* = (\mathcal{S}^*, \mathcal{D}^*)$ acts just like \mathcal{A} , but instead of encrypting, the source \mathcal{S} sets aux to be uniformly random and stores the tuple $(\text{aux}, y_1, \dots, y_q, z')$ and instead of decrypting \mathcal{D}^* retrieves the tuple indexed by aux to uses the corresponding (y_1, \dots, y_q, z') .

Firstly, we claim that \mathcal{A} and \mathcal{A}^* are indistinguishable by any (comp. unbounded) distinguisher that makes Q queries with probability better than $\text{poly}(Q) \cdot 2^{-\ell}$. This essentially follows by the authenticated-encryption security of the encryption scheme.

Secondly, we claim that \mathcal{A}^* and \mathcal{A}' are perfectly indistinguishable. The only difference between them is that \mathcal{A}^* compares u against the first bit of $\text{EExt}(\text{seed}', x)$ while \mathcal{A}' compares it against $\text{EExt}(\text{seed}, x)$. But since $\text{seed}, \text{seed}'$ agree on x_1, \dots, x_q , the seed-committing property ensures that $\text{EExt}(\text{seed}', x) = \text{EExt}(\text{seed}, x)$.

Acknowledgements. YD was partially supported by gifts from VMware Labs, Facebook and Google, and NSF grants 1314568, 1619158, 1815546. VV was supported in part by NSF Grants CNS-1350619 and CNS-1414119, an NSF-BSF grant CNS-1718161, the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236, an IBM-MIT grant and a Microsoft Trustworthy and Robust AI grant. DW was supported by NSF grants CNS-1314722, CNS-1413964, CNS-1750795 and the Alfred P. Sloan Research Fellowship.

References

[AMN+18] Attrapadung, N., Matsuda, T., Nishimaki, R., Yamada, S., Yamakawa, T.: Constrained PRFs for NC^1 in traditional groups. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 543–574. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_19

- [BGI14] Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudo-random functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54631-0_29
- [BH05] Barak, B., Halevi, S.: A model and architecture for pseudo-random generation with applications to /dev/random. In: Atluri, V., Meadows, C., Juels, A. (eds.) ACM CCS 2005: 12th Conference on Computer and Communications Security, pp. 203–212. ACM Press, November 2005
- [BHK13] Bellare, M., Hoang, V.T., Keelveedhi, S.: Instantiating random oracles via UCEs. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 398–415. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_23
- [BIW04] Barak, B., Impagliazzo, R., Wigderson, A.: Extracting randomness using few independent sources. In: 45th Annual Symposium on Foundations of Computer Science, pp. 384–393. IEEE Computer Society Press, October 2004
- [Blu86] Blum, M.: Independent unbiased coin flips from a correlated biased source—a finite state Markov chain. *Combinatorica* **6**(2), 97–108 (1986)
- [BST03] Barak, B., Shaltiel, R., Tromer, E.: True random number generators secure in a changing environment. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) CHES 2003. LNCS, vol. 2779, pp. 166–180. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45238-6_14
- [BV15] Brakerski, Z., Vaikuntanathan, V.: Constrained key-homomorphic PRFs from standard lattice assumptions. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 1–30. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46497-7_1
- [BW13] Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-42045-0_15
- [CDKT19] Coretti, S., Dodis, Y., Karthikeyan, H., Tessaro, S.: Seedless fruit is the sweetest: random number generation, revisited. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 205–234. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_8
- [CG85] Chor, B., Goldreich, O.: Unbiased bits from sources of weak randomness and probabilistic communication complexity (extended abstract). In: 26th Annual Symposium on Foundations of Computer Science, pp. 429–442. IEEE Computer Society Press, October 1985
- [CG88] Chor, B., Goldreich, O.: Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.* **17**(2), 230–261 (1988)
- [CGH+85] Chor, B., Goldreich, O., Håstad, J., Friedman, J., Rudich, S., Smolensky, R.: The bit extraction problem of t -resilient functions (preliminary version). In: 26th Annual Symposium on Foundations of Computer Science, pp. 396–407. IEEE Computer Society Press, October 1985
- [CZ16] Chattopadhyay, E., Zuckerman, D.: Explicit two-source extractors and resilient functions. In: Wichs, D., Mansour, Y. (eds.) 48th Annual ACM Symposium on Theory of Computing, pp. 670–683. ACM Press, June 2016

- [DORS08] Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.D.: Fuzzy extractors: how to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.* **38**(1), 97–139 (2008)
- [DPR+13] Dodis, Y., Pointcheval, D., Ruhault, S., Vergnaud, D., Wichs, D.: Security analysis of pseudo-random number generators with input: /dev/random is not robust. In: Sadeghi, A.-R., Gligor, V.D., Yung, M. (eds.) *ACM CCS 2013: 20th Conference on Computer and Communications Security*, pp. 647–658. ACM Press, November 2013
- [DRV12] Dodis, Y., Ristenpart, T., Vadhan, S.: Randomness condensers for efficiently samplable, seed-dependent sources. In: Cramer, R. (ed.) *TCC 2012. LNCS*, vol. 7194, pp. 618–635. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28914-9_35
- [DSSW14] Dodis, Y., Shamir, A., Stephens-Davidowitz, N., Wichs, D.: How to eat your entropy and have it too – optimal recovery strategies for compromised RNGs. In: Garay, J.A., Gennaro, R. (eds.) *CRYPTO 2014, Part II. LNCS*, vol. 8617, pp. 37–54. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44381-1_3
- [DVW19] Dodis, Y., Vaikuntanathan, V., Wichs, D.: Extracting randomness from extractor-dependent sources. *Cryptology ePrint Archive*, Report 2019/1339 (2019). <https://eprint.iacr.org/2019/1339>
- [GGM84] Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: *25th Annual Symposium on Foundations of Computer Science*, pp. 464–479. IEEE Computer Society Press, October 1984
- [GT16] Gazi, P., Tessaro, S.: Provably robust sponge-based PRNGs and KDFs. In: Fischlin, M., Coron, J.-S. (eds.) *EUROCRYPT 2016, Part I. LNCS*, vol. 9665, pp. 87–116. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_4
- [Hut16] Hutchinson, D.: A robust and sponge-like PRNG with improved efficiency. *Cryptology ePrint Archive*, Report 2016/886 (2016). <http://eprint.iacr.org/2016/886>
- [ILL89] Impagliazzo, R., Levin, L.A., Luby, M.: Pseudo-random generation from one-way functions (extended abstracts). In: Johnson, D.S. (ed.) *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, Seattle, Washington, USA, 14–17 May 1989, pp. 12–24. ACM (1989)
- [IZ89] Impagliazzo, R., Zuckerman, D.: How to recycle random bits. In: *30th Annual Symposium on Foundations of Computer Science*, pp. 248–253. IEEE Computer Society Press, October/November 1989
- [KPTZ13] Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: Sadeghi, A.-R., Gligor, V.D., Yung, M. (eds.) *ACM CCS 2013: 20th Conference on Computer and Communications Security*, pp. 669–684. ACM Press, November 2013
- [KRVZ11] Kamp, J., Rao, A., Vadhan, S.P., Zuckerman, D.: Deterministic extractors for small-space sources. *J. Comput. Syst. Sci.* **77**(1), 191–220 (2011)
- [LLS89] Lichtenstein, D., Linial, N., Saks, M.E.: Some extremal problems arising from discrete control processes. *Combinatorica* **9**(3), 269–287 (1989)
- [NR97] Naor, M., Reingold, O.: Number-theoretic constructions of efficient pseudo-random functions. In: *38th Annual Symposium on Foundations of Computer Science*, pp. 458–467. IEEE Computer Society Press, October 1997

- [NZ93] Nisan, N., Zuckerman, D.: More deterministic simulation in logspace. In: 25th Annual ACM Symposium on Theory of Computing, pp. 235–244. ACM Press, May 1993
- [NZ96] Nisan, N., Zuckerman, D.: Randomness is linear in space. *J. Comput. Syst. Sci.* **52**(1), 43–52 (1996)
- [PS18] Peikert, C., Shiehian, S.: Privately constraining and programming PRFs, the LWE way. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 675–701. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76581-5_23
- [PW08] Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C. (eds.) 40th Annual ACM Symposium on Theory of Computing, pp. 187–196. ACM Press, May 2008
- [Raz05] Raz, R.: Extractors with weak random seeds. In: Gabow, H.N., Fagin, R. (eds.) 37th Annual ACM Symposium on Theory of Computing, pp. 11–20. ACM Press, May 2005
- [ST17] Soni, P., Tessaro, S.: Public-seed pseudorandom permutations. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part II. LNCS, vol. 10211, pp. 412–441. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_14
- [TV00] Trevisan, L., Vadhan, S.P.: Extracting randomness from samplable distributions. In: 41st Annual Symposium on Foundations of Computer Science, pp. 32–42. IEEE Computer Society Press, November 2000
- [von51] von Neumann, J.: Various techniques used in connection with random digits. In: Householder, A.S., Forsythe, G.E., Germond, H.H. (eds.) Monte Carlo Method. National Bureau of Standards Applied Mathematics Series, vol. 12, pp. 36–38. U.S. Government Printing Office, Washington, D.C. (1951)
- [Wic13] Wichs, D.: Barriers in cryptography with weak, correlated and leaky sources. In: Kleinberg, R.D. (ed.) ITCS 2013: 4th Innovations in Theoretical Computer Science, pp. 111–126. Association for Computing Machinery, January 2013