How to Fake Auxiliary Input

Dimitar Jetchev and Krzysztof Pietrzak*

EPFL Switzerland and IST Austria

Abstract. Consider a joint distribution (X, A) on a set $\mathcal{X} \times \{0, 1\}^\ell$. We show that for any family $\mathcal{F}$ of distinguishers $f: \mathcal{X} \times \{0, 1\}^\ell \rightarrow \{0, 1\}$, there exists a simulator $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ such that

1. no function in $\mathcal{F}$ can distinguish (X, A) from $(X, h(X))$ with advantage ε ,
2. h is only $O(2^{3\ell}\varepsilon^{-2})$ times less efficient than the functions in $\mathcal{F}$.

For the most interesting settings of the parameters (in particular, the cryptographic case where X has superlogarithmic min-entropy, $\varepsilon > 0$ is negligible and $\mathcal{F}$ consists of circuits of polynomial size), we can make the simulator h *deterministic*.

As an illustrative application of our theorem, we give a new security proof for the leakage-resilient stream-cipher from Eurocrypt'09. Our proof is simpler and quantitatively much better than the original proof using the dense model theorem, giving meaningful security guarantees if instantiated with a standard blockcipher like AES.

Subsequent to this work, Chung, Lui and Pass gave an interactive variant of our main theorem, and used it to investigate weak notions of Zero-Knowledge. Vadhan and Zheng give a more constructive version of our theorem using their new uniform min-max theorem.

1 Introduction

Let $\mathcal{X}$ be a set and let $\ell > 0$ be an integer. We show that for any joint distribution (X, A) over $\mathcal{X} \times \{0, 1\}^\ell$ (where we think of A as a short ℓ -bit auxiliary input to X), any family $\mathcal{F}$ of functions $\mathcal{X} \times \{0, 1\}^\ell \rightarrow \{0, 1\}$ (thought of as distinguishers) and any $\varepsilon > 0$, there exists an efficient simulator $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ for the auxiliary input that fools every distinguisher in $\mathcal{F}$, i.e.,

$$\forall f \in \mathcal{F}: |\mathbb{E}[f(X, A)] - \mathbb{E}[f(X, h(X))]| < \varepsilon.$$

Here, “efficient” means that the simulator h is $\tilde{O}(2^{3\ell}\varepsilon^{-2})$ times more complex than the functions from $\mathcal{F}$ (we will formally define “more complex” in Definition 6). Without loss of generality, we can model the joint distribution (X, A) as $(X, g(X))$, where g is some arbitrarily complex and possibly probabilistic function (where $\mathbb{P}[g(x) = a] = \mathbb{P}[A = a | X = x]$ for all $(x, a) \in \mathcal{X} \times \{0, 1\}^\ell$). Let us stress that, as g can be arbitrarily complex, one cannot hope to get an efficient simulator h where $(X, g(X))$ and $(X, h(X))$ are statistically close. Yet, one can still fool all functions in $\mathcal{F}$ in the sense that no function from $\mathcal{F}$ can distinguish the distribution (X, A) from $(X, g(X))$.

* Supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC Starting Grant (259668-PSPC).

Relation to [25]. Trevisan, Tulsiani and Vadhan [25, Thm. 3.1] prove a conceptually similar result, stating that if $\mathcal{Z}$ is a set then for any distribution Z over $\mathcal{Z}$, any family $\tilde{\mathcal{F}}$ of functions $\mathcal{Z} \rightarrow [0, 1]$ and any function $\tilde{g}: \mathcal{Z} \rightarrow [0, 1]$, there exists a simulator $\tilde{h}: \mathcal{Z} \rightarrow [0, 1]$ whose complexity is only $\mathcal{O}(\varepsilon^{-2})$ times larger than the complexity of the functions from $\tilde{\mathcal{F}}$ such that

$$\forall \tilde{f} \in \tilde{\mathcal{F}} : |\mathbb{E}[\tilde{f}(Z)\tilde{g}(Z)] - \mathbb{E}[\tilde{f}(Z)\tilde{h}(Z)]| < \varepsilon. \quad (1)$$

In [25], this result is used to prove that every high-entropy distribution is indistinguishable from an efficiently samplable distribution of the same entropy. Moreover, it is shown that many fundamental results including the Dense Model Theorem [23,14,21,10,24], Impagliazzo's hardcore lemma [18] and a version of Szémerédi's Regularity Lemma [11] follow from this theorem. The main difference between (1) and our statement

$$\forall f \in \mathcal{F} : |\mathbb{E}[f(X, g(X))] - \mathbb{E}[f(X, h(X))]| < \varepsilon \quad (2)$$

is that our distinguisher f sees not only X , but also the real or fake auxiliary input $g(X)$ or $h(X)$, whereas in (1), the distinguisher $\tilde{f}$ only sees X . In particular, the notion of indistinguishability we achieve captures indistinguishability in the standard cryptographic sense. On the other hand, (1) is more general in the sense that the range of $\tilde{f}, \tilde{g}, \tilde{h}$ can be any real number in $[0, 1]$, whereas our f has range $\{0, 1\}$ and g, h have range $\{0, 1\}^\ell$.

Nonetheless, it is easy to derive (1) from (2): consider the case of $\ell = 1$ bit of auxiliary input, and only allow families $\mathcal{F}$ of distinguishers where each $f \in \mathcal{F}$ is of the form $f(X, b) = \hat{f}(X)b$ for some function $\hat{f}: \mathcal{X} \rightarrow [0, 1]$. For this restricted class, the absolute value in (2) becomes

$$|\mathbb{E}[f(X, g(X))] - \mathbb{E}[f(X, h(X))]| = |\mathbb{E}[\hat{f}(X)g(X)] - \mathbb{E}[\hat{f}(X)h(X)]| \quad (3)$$

As $\hat{f}$ is arbitrary, this restricted class almost captures the distinguishers considered in (1). The only difference is that the function $\tilde{g}$ has range $[0, 1]$ whereas our g has range $\{0, 1\}$. Yet, note that in (1), we can replace $\tilde{g}$ having range $[0, 1]$ by a (probabilistic) g with range $\{0, 1\}$ defined as $\mathbb{P}[g(x) = 1] = \tilde{g}(x)$, thus, leaving the expectation $\mathbb{E}[\tilde{f}(X)\tilde{g}(X)] = \mathbb{E}[\tilde{f}(X)g(X)]$ unchanged.¹

In [25], two different proofs for (1) are given. The first proof uses duality of linear programming in the form of the min-max theorem for two-player zero-sum games. This proof yields a simulator of complexity $\mathcal{O}(\varepsilon^{-4} \log^2(1/\varepsilon))$ times the complexity of the functions in $\mathcal{F}$. The second elegant proof uses boosting and gives a quantitatively much better $\mathcal{O}(\varepsilon^{-2})$ complexity.

¹ The simulator $\tilde{h}$ from [25] satisfies the additional property $|\mathbb{E}[\tilde{h}(X)] - \mathbb{E}[\tilde{g}(X)]| = 0$. If this property is needed, we can get it by requiring that the function $f(X, b) = b$ is in $\mathcal{F}$. Then (2) for this f implies $|\mathbb{E}[g(X)] - \mathbb{E}[h(X)]| < \varepsilon$. One can make this term exactly zero by slightly biasing h towards 0 if $\mathbb{E}[h(X)] > \mathbb{E}[g(X)]$ or 1 otherwise, slightly increasing the advantage from ε to at most 2ε .

Proof outline. As it was just explained, (1) follows from (2). We do not know if one can directly prove an implication in the other direction, so we prove (2) from scratch. Similarly to [25], the core of our proof uses boosting with the same energy function as the one used in [25].

As a first step, we transform the statement (2) into a “product form” like (1) where $\mathcal{Z} = \mathcal{X} \times \{0, 1\}^\ell$ (this results in a loss of a factor of 2^ℓ in the advantage ε ; in addition, our distinguishers $\hat{f}$ will have range $[-1, 1]$ instead of $[0, 1]$). We then prove that (1) holds for some simulator $\tilde{h}: \mathcal{Z} \rightarrow [0, 1]$ of complexity ε^{-2} relative to $\mathcal{F}$. Unfortunately, we cannot use the result from [25] in a black-box way at this point as we need the simulator $\tilde{h}: \mathcal{Z} \rightarrow [0, 1]$ to define a probability distribution in the sense that $\tilde{h}(x, b) \geq 0$ for all (x, b) and $\sum_{b \in \{0, 1\}^\ell} \tilde{h}(x, b) = 1$ for

all x . Ensuring these conditions is the most delicate part of the proof. Finally, we show that the simulator h defined via $\mathbb{P}[h(x) = b] = \tilde{h}(x, b)$ satisfies (2). Note that for h to be well defined, we need $\tilde{h}$ to specify a probability distribution as outlined above.

Efficiency of h . Our simulator h is efficient in the sense that it is only $\mathcal{O}(2^{3\ell}\varepsilon^{-2})$ times more complex than the functions in $\mathcal{F}$. We do not know how tight this bounds is, but one can prove a lower bound of $\max\{2^\ell, \varepsilon^{-1}\}$ under plausible assumptions. The dependency on 2^ℓ is necessary under exponential hardness assumptions for one-way functions.² A dependency on ε^{-1} is also necessary. Indeed, Trevisan et al. [25, Rem. 1.6] show that such a dependency is necessary for the simulator $\tilde{h}$ in (1). Since (1) is implied by (2) with h and $\tilde{h}$ having exactly the same complexity, the ε^{-1} lower bound also applies to our h .

1.1 Subsequent Work

The original motivation for this work was to give simpler and quantitatively better proofs for leakage-resilient cryptosystems as we will discuss in Section 4. Our main theorem has subsequently been derived via two different routes.

First, Chung, Lui and Pass [4] investigate weak notions of zero-knowledge. On route, they derive an “interactive” version of our main theorem. In Section 4, we will show how to establish one of their results (with better quantitative bounds), showing that every interactive proof system satisfies a weak notion of zero-knowledge.

Second, Vadhan and Zheng [26, Thm.3.1-3.2] recently proved a version of von Neumann’s min-max theorem for two-player zero sum games that does not only guarantee existence of an optimal strategy for the second player, but also constructs a nearly optimal strategy assuming knowledge of several best responses of the second player to strategies of the first player, and provide many applications of this theorem. Their argument is based on relative entropy KL projections

² More precisely, assume there exists a one-way function where inverting becomes 2^ℓ times easier given ℓ bits of leakage. It is e.g. believed that the AES block-cipher gives such a function as $(K, X) \rightarrow (\text{AES}(K, X), X)$.

and a learning technique known as weight updates and resembles the the proof of the Uniform Hardcore Lemma by Barak, Hardt and Kale [2] (see also [16] for the original application of this method). They derive our main theorem [26, Thm.6.8], but with incomparable bounds. Concretely, to fool circuits of size t , their simulator runs in time $\tilde{O}(t \cdot 2^\ell / \varepsilon^2 + 2^\ell / \varepsilon^4)$ compared to ours whose run-time is $\tilde{O}(t \cdot 2^{3\ell} / \varepsilon^2)$. In particular, their bounds are better whenever $1/\varepsilon^2 \leq t \cdot 2^{2\ell}$. The additive $2^\ell / \varepsilon^4$ term in their running time appears due to the sophisticated iterative “weight update” procedure, whereas our simulator simply consists of a weighted sum of the evaluation of $\tilde{O}(2^{3\ell} / \varepsilon^2)$ circuits from the family we want to fool (here, circuits of size t).

1.2 More Applications

Apart from reproving one of [4]’s results on weak zero-knowledge mentioned above, we give two more applications of our main theorem in Section 4:

Chain Rules for Computational Entropy. Gentry and Wichs [13] show that black-box reductions cannot be used to prove the security of any succinct non-interactive argument from any falsifiable cryptographic assumption. A key technical lemma used in their proof ([13, Lem. 3.1]) states that if two distributions X and $\tilde{X}$ over $\mathcal{X}$ are computationally indistinguishable, then for any joint distribution (X, A) over $\mathcal{X} \times \{0, 1\}^\ell$ (here, A is a short ℓ -bit auxiliary input) there exists a joint distribution $(\tilde{X}, \tilde{A})$ such that (X, A) and $(\tilde{X}, \tilde{A})$ are computationally indistinguishable. Our theorem immediately implies the stronger statement that not only such an $(\tilde{X}, \tilde{A})$ exists, but in fact, it is efficiently samplable, i.e., there exists an efficient simulator $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ such that $(\tilde{X}, h(\tilde{X}))$ is indistinguishable from $(\tilde{X}, \tilde{A})$ and thus from (X, A) . Reyzin [22, Thm.2] observed that the result of Gentry and Wichs implies a chain rule for conditional “relaxed” HILL entropy. We give a short and simple proof of this chain rule in Proposition 2 of this paper. We then show in Corollary 1 how to deduce a chain rule for (regular) HILL entropy from Proposition 2 using the simple fact (Lemma 1) that short (i.e., logarithmic in the size of the distinguishers) computationally indistinguishable random variables must already be statistically close. Chain rules for HILL entropy have found several applications in cryptography [10,21,7,12]. The chain rule that we get in Corollary 1 is the first one that does not suffer from a significant loss in the distinguishing advantage (we only lose a constant factor of 4). Unlike the case of relaxed HILL-entropy, here we only prove a chain rule for the “non-conditional” case, which is a necessary restriction given a recent counterexample to the (conditional) HILL chain rule by Krenn et al. [19]. We will provide more details on this negative result after the statement of Corollary 1.

Leakage Resilient Cryptography. The original motivation for this work is to simplify the security proofs of leakage-resilient [10,20,7] and other cryptosystems [12] whose security proofs rely on chain rules for computational entropy (as discussed in the previous paragraph). The main idea is to replace the chain rules with

simulation-based arguments. In a nutshell, instead of arguing that a variable X must have high (pseudo)entropy in the presence of a short leakage A , one could simply use the fact that the leakage can be efficiently simulated. This not only implies that X has high (pseudo)entropy given the fake leakage $h(X)$, but if X is pseudorandom, it also implies that $(X, h(X))$ is indistinguishable from $(U, h(U))$ for a uniform random variable U on the same set as X . In the security proofs, we would now replace $(X, h(X))$ with $(U, h(U))$ and will continue with a uniformly random intermediate variable U . In contrast, the approach based on chain rules only tells us that we can replace X with some random variable Y that has high min-entropy given A . This is not only much complex to work with, but it often gives weaker quantitative bounds. In particular, in Section 4.3 we revisit the security proof of the leakage-resilient stream-cipher from [20] for which we can now give a conceptually simpler and quantitatively better security proof.

2 Notation and Basic Definitions

2.1 Notation

We use calligraphic letters such as $\mathcal{X}$ to denote sets, the corresponding capital letters X to denote random variables on these sets (equivalently, probability distributions) and lower-case letters (e.g., x) for values of the corresponding random variables. Moreover, $x \leftarrow X$ means that x is sampled according to the distribution X and $x \leftarrow \mathcal{X}$ means that x is sampled uniformly at random from $\mathcal{X}$. Let U_n denote the random variable with uniform distribution on $\{0, 1\}^n$.

We denote by $\Delta(X; Y) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathbb{P}[X = x] - \mathbb{P}[Y = x]|$ the statistical distance

between X and Y . For $\varepsilon > 0, s \in \mathbb{N}$, we use $X \sim Y$ to denote that X and Y have the same distribution, $X \sim_\varepsilon Y$ to denote that their statistical distance is less than ε and $X \sim_{\varepsilon, s} Y$ to denote that no circuit of size s can distinguish X from Y with advantage greater than ε . Note that $X \sim_{\varepsilon, \infty} Y \iff X \sim_\varepsilon Y$ and $X \sim_0 Y \iff X \sim Y$.

Finally, if $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ is a probabilistic (randomized) function then we will use $[h]$ to denote the random coins used by h (a notation that will be used in various probabilities and expectations).

2.2 Entropy Measures

A random variable X has min-entropy k , if no (computationally unbounded) adversary can predict the outcome of X with probability greater than 2^{-k} .

Definition 1. (Min-Entropy H_∞) A random variable X has min-entropy k , denoted $H_\infty(X) \geq k$, if $\max_x \mathbb{P}[X = x] \leq 2^{-k}$.

Dodis et al. [8] gave a notion of average-case min-entropy defined such that X has average-case min-entropy k conditioned on Z if the probability of the best adversary in predicting X given Z is 2^{-k} .

Definition 2. (Average min-Entropy [8] $\tilde{H}_\infty$) Consider a joint distribution (X, Z) , then the average min-entropy of X conditioned on Z is

$$\tilde{H}_\infty(X|Z) = -\log\left(\mathbb{E}_{z \leftarrow Z} \left[\max_x \mathbb{P}[X = x|Z = z] \right]\right) = -\log\left(\mathbb{E}_{z \leftarrow Z} \left[2^{-H_\infty(X|Z=z)} \right]\right)$$

HILL-entropy is the computational analogue of min-entropy. A random variable X has HILL-entropy k if there exists a random variable Y having min-entropy k that is indistinguishable from X . HILL-entropy is further quantified by two parameters ε, s specifying this indistinguishability quantitatively.

Definition 3. (HILL-Entropy [15] $H_{\varepsilon,s}^{\text{HILL}}$) X has HILL entropy k , denoted by $H_{\varepsilon,s}^{\text{HILL}}(X) \geq k$, if

$$H_{\varepsilon,s}^{\text{HILL}}(X) \geq k \iff \exists Y : H_\infty(Y) \geq k \text{ and } X \sim_{\varepsilon,s} Y$$

Conditional HILL-entropy has been defined by Hsiao, Lu and Reyzin [17] as follows.

Definition 4. (Conditional HILL-Entropy [17]) X has conditional HILL entropy k (conditioned on Z), denoted $H_{\varepsilon,s}^{\text{HILL}}(X|Z) \geq k$, if

$$H_{\varepsilon,s}^{\text{HILL}}(X|Z) \geq k \iff \exists (Y, Z) : \tilde{H}_\infty(Y|Z) \geq k \text{ and } (X, Z) \sim_{\varepsilon,s} (Y, Z)$$

Note that in the definition above, the marginal distribution on the conditional part Z is the same in both the real distribution (X, Z) and the indistinguishable distribution (Y, Z) . A “relaxed” notion of conditional HILL used implicitly in [13] and made explicit in [22] drops this requirement.

Definition 5. (Relaxed Conditional HILL-Entropy [13,22]) X has relaxed conditional HILL entropy k , denoted $H_{\varepsilon,s}^{\text{rlx-HILL}}(X|Z) \geq k$, if

$$H_{\varepsilon,s}^{\text{rlx-HILL}}(X|Z) \geq k \iff \exists (Y, W) : \tilde{H}_\infty(Y|W) \geq k \text{ and } (X, Z) \sim_{\varepsilon,s} (Y, W)$$

3 The Main Theorem

Definition 6. (Complexity of a function) Let $\mathcal{A}$ and $\mathcal{B}$ be sets and let $\mathcal{G}$ be a family of functions $h: \mathcal{A} \rightarrow \mathcal{B}$. A function h has **complexity C relative to $\mathcal{G}$** if it can be computed by an oracle-aided circuit of size $\text{poly}(C \log |\mathcal{A}|)$ with C oracle gates where each oracle gate is instantiated with a function from $\mathcal{G}$.

Theorem 1. (Main) Let $\ell \in \mathbb{N}$ be fixed, let $\varepsilon > 0$ and let $\mathcal{X}$ be any set. Consider a distribution X over $\mathcal{X}$ and any (possibly probabilistic and not necessarily efficient) function $g: \mathcal{X} \rightarrow \{0, 1\}^\ell$. Let $\mathcal{F}$ be a family of deterministic (cf. remark below) distinguishers $f: \mathcal{X} \times \{0, 1\}^\ell \rightarrow \{0, 1\}$. There exists a (probabilistic) simulator $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ with complexity³

$$O(2^{3\ell} \varepsilon^{-2} \log^2(\varepsilon^{-1}))$$

³ If we model h as a Turing machine (and not a circuit) and consider the expected complexity of h , then we can get a slightly better $O(2^{3\ell} \varepsilon^{-2})$ bound (i.e. without the $\log^2(\varepsilon^{-1})$ term).

relative to $\mathcal{F}$ which ε -fools every distinguisher in $\mathcal{F}$, i.e.

$$\forall f \in \mathcal{F} : \left| \mathbb{E}_{x \leftarrow X, [g]} [f(x, g(x))] - \mathbb{E}_{x \leftarrow X, [h]} [f(x, h(x))] \right| < \varepsilon, \quad (4)$$

Moreover, if

$$H_\infty(X) > 2 + \log \log |\mathcal{F}| + 2 \log(1/\varepsilon) \quad (5)$$

then there exists a deterministic h with this property.

Remark 1 (Closed and Probabilistic $\mathcal{F}$). In the proof of Theorem 1 we assume that the class $\mathcal{F}$ of distinguishers is closed under complement, i.e., if $f \in \mathcal{F}$ then also $1 - f \in \mathcal{F}$. This is without loss of generality, as even if we are interested in the advantage of a class $\mathcal{F}$ that is not closed, we can simply apply the theorem for $\mathcal{F}' = \mathcal{F} \cup (1 - \mathcal{F})$, where $(1 - \mathcal{F}) = \{1 - f : f \in \mathcal{F}\}$. Note that if h has complexity t relative to $\mathcal{F}'$, it has the same complexity relative to $\mathcal{F}$. We also assume that all functions $f \in \mathcal{F}$ are deterministic. If we are interested in a class $\mathcal{F}$ of randomized functions, we can simply apply the theorem for the larger class of deterministic functions $\mathcal{F}''$ consisting of all pairs (f, r) where $f \in \mathcal{F}$ and r is a choice of randomness for f . This is almost without loss of generality, except that the requirement in eq.(5) on the min-entropy of X becomes slightly stronger as $\log \log |\mathcal{F}''| = \log \log (|\mathcal{F}| 2^\rho)$ where ρ is an upper bound on the number of random coins used by any $f \in \mathcal{F}$.

4 Applications

4.1 Zero-Knowledge

Chung, Lui and Pass [4] consider the following relaxed notion of zero-knowledge

Definition 7 (distributional (T, t, ε) -zero-knowledge). Let $(\mathcal{P}, \mathcal{V})$ be an interactive proof system for a language L . We say that $(\mathcal{P}, \mathcal{V})$ is distributional (T, t, ε) -zero-knowledge (where T, t, ε are all functions of n) if for every $n \in \mathbb{N}$, every joint distributions (X_n, Y_n, Z_n) over $(L \cap \{0, 1\}^n) \times \{0, 1\}^* \times \{0, 1\}^*$, and every t -size adversary $\mathcal{V}^*$, there exists a T -size simulator S such that

$$(X_n, Z_n, \text{out}_{\mathcal{V}^*}[\mathcal{P}(X_n, Y_n) \leftrightarrow \mathcal{V}^*(X_n, Z_n)]) \sim_{\varepsilon, t} (X_n, Z_n, S(X_n, Z_n))$$

where $\text{out}_{\mathcal{V}^*}[\mathcal{P}(X_n, Y_n) \leftrightarrow \mathcal{V}^*(X_n, Z_n)]$ denotes the output of $\mathcal{V}^*(X_n, Z_n)$ after interacting with $\mathcal{P}(X_n, Y_n)$.

If L is an NP language, then in the definition above, Y would be a witness for $X \in L$. As a corollary of their main theorem, [4] show that every proof system satisfies this relaxed notion of zero-knowledge where the running time T of the simulator is polynomial in t, ε and 2^ℓ . We can derive their Corollary from Theorem 1 with better quantitative bounds for most ranges of parameters than [4]: we get $\tilde{O}(t2^{3\ell}\varepsilon^{-2})$ vs. $\tilde{O}(t^32^\ell\varepsilon^{-6})$, which is better whenever $t/\varepsilon^2 \geq 2^\ell$.

Proposition 1. *Let $(\mathcal{P}, \mathcal{V})$ be an interactive proof system for a language L , and suppose that the total length of the messages sent by $\mathcal{P}$ is $\ell = \ell(n)$ (on common inputs X of length n). Then for any $t = t(n) \geq \Omega(n)$ and $\varepsilon = \varepsilon(n)$, $(\mathcal{P}, \mathcal{V})$ is distributional (T, t, ε) -zero-knowledge, where*

$$T = O(t2^{3\ell}\varepsilon^{-2}\log^2(\varepsilon^{-1}))$$

Proof. Let $M \in \{0, 1\}^\ell$ denote the messages sent by $\mathcal{P}(X_n, Y_n)$ when talking to $\mathcal{V}^*(X_n, Z_n)$. By Theorem 1 (identifying $\mathcal{F}$ from the theorem with circuits of size t) there exists a simulator h of size $O(t \cdot 2^{3\ell}\varepsilon^{-2}\log^2(\varepsilon^{-1}))$ s.t.

$$(X_n, Z_n, M) \sim_{\varepsilon, 2t} (X_n, Z_n, h(X_n, Z_n)) \quad (6)$$

Let $S(X_n, Z_n)$ be defined as follows, first compute $M' = h(X_n, Z_n)$ (with h as above), and then compute $\text{out}_{\mathcal{V}}^*[M' \leftrightarrow \mathcal{V}^*(X_n, Z_n)]$. We claim that

$$(X_n, Z_n, \text{out}_{\mathcal{V}^*}[\mathcal{P}(X_n, Y_n) \leftrightarrow \mathcal{V}^*(X_n, Z_n)]) \sim_{\varepsilon, t} (X_n, Z_n, S(X_n, Z_n)) \quad (7)$$

To see this, note that from any distinguisher D of size t that distinguishes the distributions in (7) with advantage $\delta > \varepsilon$, we get a distinguisher D' of size $2t$ that distinguishes the distributions in (6) with the same advantage by defining D' as $D'(X_n, Z_n, \tilde{M}) = D(X_n, Z_n, \text{out}_{\mathcal{V}^*}[\tilde{M} \leftrightarrow \mathcal{V}^*(X_n, Z_n)])$. $\square$

4.2 Chain Rules for (Conditional) Pseudoentropy

The following proposition is a chain rule for relaxed conditional HILL entropy. Such a chain rule for the non-conditional case is implicit in the work of Gentry and Wichs [13], and made explicit and generalized to the conditional case by Reyzin [22].

Proposition 2. ([13,22]) *Any joint distribution $(X, Y, A) \in \mathcal{X} \times \mathcal{Y} \times \{0, 1\}^\ell$ satisfies⁴*

$$\text{H}_{\varepsilon, s}^{\text{rlx-HILL}}(X|Y) \geq k \Rightarrow \text{H}_{2\varepsilon, \hat{s}}^{\text{rlx-HILL}}(X|Y, A) \geq k - \ell \text{ where } \hat{s} = \Omega\left(s \cdot \frac{\varepsilon^2}{2^{3\ell} \log^2(1/\varepsilon)}\right)$$

Proof. $\text{H}_{\varepsilon, s}^{\text{rlx-HILL}}(X|Y) \geq k$ means that there exists a random variable (Z, W) such that $\text{H}_\infty(Z|W) \geq k$ and $(X, Y) \sim_{\varepsilon, s} (Z, W)$. For any $\hat{\varepsilon}, \hat{s}$, by Theorem 1, there exists a simulator h of size $s_h = O\left(\hat{s} \cdot \frac{2^{3\ell} \log^2(1/\hat{\varepsilon})}{\hat{\varepsilon}^2}\right)$ such that (we explain the second step below)

$$(X, Y, A) \sim_{\hat{\varepsilon}, \hat{s}} (X, Y, h(X, Y)) \sim_{\varepsilon, s-s_h} (Z, W, h(Z, W))$$

⁴ Using the recent bound from [26] discussed in Section 1.1, we can get $\hat{s} = \Omega\left(s \cdot \frac{\varepsilon^2 \ell}{2^\ell} + \frac{\ell^2 \log^2(1/\varepsilon)}{\varepsilon^4}\right)$

The second step follows from $(X, Y) \sim_{\varepsilon, s} (Z, W)$ and the fact that h has complexity s_h . Using the triangle inequality for computational indistinguishability⁵ we get

$$(X, Y, A) \sim_{\hat{\varepsilon} + \varepsilon, \min(\hat{s}, s - s_h)} (Z, W, h(Z, W))$$

To simplify this expression, we set $\hat{\varepsilon} := \varepsilon$ and $\hat{s} := \Theta(s\varepsilon^2/2^{3\ell} \log^2(1/\varepsilon))$, then $s_h = O(s)$, and choosing the hidden constant in the Θ such that $s_h \leq s/2$ (and thus $\hat{s} \leq s - s_h = s/2$), the above equation becomes

$$(X, Y, A) \sim_{2\varepsilon, \hat{s}} (Z, W, h(Z, W)) \quad (8)$$

Using the chain rule for average case min-entropy in the first, and $H_\infty(Z|W) \geq k$ in the second step below we get

$$\tilde{H}_\infty(Z|W, h(Z, W)) \geq \tilde{H}_\infty(Z|W) - H_0(h(Z, W)) \geq k - \ell. \quad (9)$$

Now equations (8) and (9) imply $H_{2\varepsilon, \hat{s}}^{\text{rlx-HILL}}(X|Y, A) = k - \ell$ as claimed. $\square$

By the following lemma, conditional relaxed HILL implies conditional HILL if the conditional part is short (at most logarithmic in the size of the distinguishers considered.)

Lemma 1. *For a joint random variable (X, A) over $\mathcal{X} \times \{0, 1\}^\ell$ and $s = \Omega(\ell 2^\ell)$ (more concretely, s should be large enough to implement a lookup table for a function $\{0, 1\}^\ell \rightarrow \{0, 1\}$) conditional relaxed HILL implies standard HILL entropy*

$$H_{\varepsilon, s}^{\text{rlx-HILL}}(X|A) \geq k \quad \Rightarrow \quad H_{2\varepsilon, s}^{\text{HILL}}(X|A) \geq k$$

Proof. $H_{\varepsilon, s}^{\text{rlx-HILL}}(X|A) \geq k$ means that there exist (Z, W) where $\tilde{H}_\infty(Z|W) \geq k$ and

$$(X, A) \sim_{\varepsilon, s} (Z, W) \quad (10)$$

We claim that if $s = \Omega(\ell 2^\ell)$, then (10) implies that $W \sim_\varepsilon A$. To see this, assume the contrary, i.e., that W and A are not ε -close. There exists then a computationally unbounded distinguisher D where

$$|\mathbb{P}[D(W) = 1] - \mathbb{P}[D(A) = 1]| > \varepsilon.$$

Without loss of generality, we can assume that D is deterministic and thus, implement D by a circuit of size $\Theta(\ell 2^\ell)$ via a lookup table with 2^ℓ entries (where the i th entry is $D(i)$.) Clearly, D can also distinguish (X, A) from (Z, W) with advantage greater than ε by simply ignoring the first part of the input, thus, contradicting (10). As $A \sim_\varepsilon W$, we claim that there exist a distribution (Z, A) such that

$$(Z, W) \sim_\varepsilon (Z, A). \quad (11)$$

⁵ which states that for any random variables α, β, γ we have

$$\alpha \sim_{\varepsilon_1, s_1} \beta \quad \& \quad \beta \sim_{\varepsilon_2, s_2} \gamma \quad \Rightarrow \quad \alpha \sim_{\varepsilon_1 + \varepsilon_2, \min(s_1, s_2)} \gamma$$

This distribution (Z, A) can be sampled by first sampling (Z, W) and then outputting $(Z, \alpha(W))$ where α is a function that is the identity with probability at least $1 - \varepsilon$ (over the choice of W), i.e., $\alpha(w) = w$ and with probability at most ε , it changes W so that it matches A . The latter is possible since $A \sim_\varepsilon W$.

Using the triangle inequality for computational indistinguishability (cf. the proof of Proposition 2) we get with (10) and (11)

$$(X, A) \sim_{2\varepsilon, s} (Z, A) \quad (12)$$

As $\tilde{H}_\infty(Z|W) \geq k$ (for α as defined above)

$$\tilde{H}_\infty(Z|W) \geq k \Rightarrow \tilde{H}_\infty(Z|\alpha(W)) \geq k \Rightarrow \tilde{H}_\infty(Z|A) \geq k \quad (13)$$

The first implication above holds as applying a function on the conditioned part cannot decrease the min-entropy. The second holds as $(Z, A) \sim (Z, \alpha(W))$. This concludes the proof as (12) and (13) imply that $H_{2\varepsilon, s}^{\text{HILL}}(X|A) \geq k$. $\square$

As a corollary of Proposition 1 and Lemma 1, we get a chain rule for (non-conditional) HILL entropy. Such a chain rule has been shown by [10] and follows from the more general Dense Model Theorem (published at the same conference) of Reingold et al. [21].

Corollary 1. *For any distribution $(X, A) \in \mathcal{X} \times \{0, 1\}^\ell$ and $\hat{s} = \Omega\left(\frac{s \cdot \varepsilon^2}{2^{3\ell} \log^2(\ell)}\right)$*

$$H_{\varepsilon, s}^{\text{HILL}}(X) \geq k \Rightarrow H_{2\varepsilon, \hat{s}}^{\text{rx-HILL}}(X|A) \geq k - \ell \Rightarrow H_{4\varepsilon, \hat{s}}^{\text{HILL}}(X|A) \geq k - \ell$$

Note that unlike the chain rule for relaxed HILL given in Proposition 2, the chain rule for (standard) HILL given by the corollary above requires that we start with some non-conditional variable X . It would be preferable to have a chain rule for the conditional case, i.e., an expression of the form $H_{\varepsilon, s}^{\text{HILL}}(X|Y) = k \Rightarrow H_{\varepsilon', s'}^{\text{HILL}}(X|Y, A) = k - \ell$ for some $\varepsilon' = \varepsilon \cdot p(2^\ell)$, $s' = s/q(2^\ell, \varepsilon^{-1})$ (for polynomial functions $p(\cdot), q(\cdot)$), but as recently shown by Krenn et al. [19], such a chain rule does not hold (all we know is that such a rule holds if we also allow the security to degrade exponentially in the length $|Y|$ of the conditional part.)

4.3 Leakage-Resilient Cryptography

We now discuss how Theorem 1 can be used to simplify and quantitatively improve the security proofs for leakage-resilient cryptosystems. These proofs currently rely on chain rules for HILL entropy given in Corollary 1. As an illustrative example, we will reprove the security of the leakage-resilient stream-cipher based on any weak pseudorandom function from Eurocrypt'09 [20], but with much better bounds than the original proof.

For brevity, in this section we often write B^i to denote a sequence $B_1, \dots, B_i$ of values. Moreover, $A \| B \in \{0, 1\}^{a+b}$ denotes the concatenation of the strings $A \in \{0, 1\}^a$ and $B \in \{0, 1\}^b$.

A function $F: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ is an (ε, s, q) -**secure weak PRF** if its outputs on q random inputs look random to any size s distinguisher, i.e., for all D of size s

$$\left| \mathbb{P}_{K, X^q} [D(X^q, F(K, X_1), \dots, F(K, X_q)) = 1] - \mathbb{P}_{X^q, R^q} [D(X^q, R^q) = 1] \right| \leq \varepsilon,$$

where the probability is over the choice of the random $X_i \leftarrow \{0, 1\}^n$, the choice of a random key $K \leftarrow \{0, 1\}^k$ and random $R_i \leftarrow \{0, 1\}^m$ conditioned on $R_i = R_j$ if $X_i = X_j$ for some $j < i$.

A **stream-cipher SC**: $\{0, 1\}^k \rightarrow \{0, 1\}^k \times \{0, 1\}^n$ is a function that, when initialized with a secret initial state $S_0 \in \{0, 1\}^k$, produces a sequence of output blocks $X_1, X_2, \dots$ recursively computed by

$$(S_i, X_i) := \text{SC}(S_{i-1})$$

We say that **SC** is (ε, s, q) -secure if for all $1 \leq i \leq q$, no distinguisher of size s can distinguish X_i from a uniformly random $U_n \leftarrow \{0, 1\}^n$ with advantage greater than ε given $X_1, \dots, X_{i-1}$ (here, the probability is over the choice of the initial random key S_0)⁶, i.e.,

$$\left| \mathbb{P}_{S_0} [D(X^{i-1}, X_i) = 1] - \mathbb{P}_{S_0, U_n} [D(X^{i-1}, U_n) = 1] \right| \leq \varepsilon$$

A **leakage-resilient** stream-cipher is $(\varepsilon, s, q, \ell)$ -secure if it is (ε, s, q) -secure as just defined, but where the distinguisher in the j th round not only gets X_j , but also ℓ bits of arbitrary adaptively chosen leakage about the secret state accessed during this round. More precisely, before $(S_j, X_j) := \text{SC}(S_{j-1})$ is computed, the distinguisher can choose any leakage function f_j with range $\{0, 1\}^\ell$, and then not only get X_j , but also $A_j := f_j(\hat{S}_{j-1})$, where $\hat{S}_{j-1}$ denotes the part of the secret state that was modified (i.e., read and/or overwritten) in the computation $\text{SC}(S_{j-1})$.

Figure 1 illustrates the construction of a leakage-resilient stream cipher SC^F from any weak PRF $F: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^{k+n}$ from [20]. The initial state is $S_0 = \{K_0, K_1, X_0\}$. Moreover, in the i th round (starting with round 0), one computes $K_{i+2} \| X_{i+1} := F(K_i, X_i)$ and outputs X_{i+1} . The state after this round is $(K_{i+1}, K_{i+2}, X_{i+1})$.⁷ In this section we will sketch a proof of the following security bound on SC^F as a leakage-resilient stream cipher in terms of the security of F as a weak PRF.

⁶ A more standard notion would require $X_1, \dots, X_q$ to be indistinguishable from random; this notion is implied by the notion we use by a standard hybrid argument losing a multiplicative factor of q in the distinguishing advantage.

⁷ Note that X_i is not explicitly given as input to f_i even though the computation depends on X_i . The reason is that the adversary can choose f_i adaptively after seeing X_i , so X_i can be hard-coded into f_i .

Lemma 2. *If F is a $(\varepsilon_F, s_F, 2)$ -secure weak PRF then SC^F is a $(\varepsilon', s', q, \ell)$ -secure leakage resilient stream cipher where*

$$\varepsilon' = 4q\sqrt{\varepsilon_F 2^\ell} \quad s' = \Theta(1) \cdot \frac{s_F \varepsilon'^2}{2^{3\ell}}$$

The bound above is quantitatively much better than the one in [20]. Setting the leakage bound $\ell = \log \varepsilon_F^{-1}/6$ as in [20], we get (for small q) $\varepsilon' \approx \varepsilon_F^{5/12}$, which is by over a power of 5 better than the $\varepsilon_F^{1/13}$ from [20], and the bound on $s' \approx s_F \varepsilon_F^{4/3}$ improves by a factor of $\varepsilon_F^{5/6}$ (from $s_F \varepsilon_F^{13/6}$ in [20] to $s_F \varepsilon_F^{8/6}$ here). This improvement makes the bound meaningful if instantiated with a standard block-cipher like AES which has a keyspace of 256 bits, making the assumption that it provides $s_F/\varepsilon_F \approx 2^{256}$ security.⁸

Besides our main Theorem 1, we need another technical result which states that if F is a weak PRF secure against two queries, then its output on a single random query is pseudorandom, even if one is given some short auxiliary information about the uniform key K . The security of weak PRFs with non-uniform keys has first been proven in [20], but we will use a more recent and elegant bound from [1]. As a corollary of [1, Thm.3.7 in eprint version], we get that for any $(\varepsilon_F, s_F, 2)$ -secure weak PRF $F: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^m$, uniform and independent key and input $K \sim U_k, X \sim U_n$ and any (arbitrarily complex) function $g: \{0, 1\}^k \rightarrow \{0, 1\}^\ell$, one has⁹

$$(X, F(K, X), g(K)) \sim_{\varepsilon, s_F/2} (X, U_m, g(K)) \text{ where } \varepsilon = \varepsilon_F + \sqrt{\varepsilon_F 2^\ell} + 2^{-n} \approx \sqrt{\varepsilon_F 2^\ell} \quad (14)$$

Generalizing the notation of $\sim_{\varepsilon, s}$ from variables to interactive distinguishers, given two (potentially stateful) oracles G, G' , we write $G \sim_{\varepsilon, s} G'$ to denote that no oracle-aided adversary A of size s can distinguish G from G' , i.e.,

$$G \sim_{\varepsilon, s} G' \iff \forall A, |A| \leq s : |\mathbb{P}[A^G \rightarrow 1] - \mathbb{P}[A^{G'} \rightarrow 1]| \leq \varepsilon.$$

Proof (of Lemma 2 (Sketch)). We define an oracle G_0^{real} that models the standard attack on the leakage-resilient stream cipher. That is, G_0^{real} samples a random initial state S_0 . When interacting with an adversary $A^{G_0^{real}}$, the oracle G_0^{real} expects as input adaptively chosen leakage functions $f_1, f_2, \dots, f_{q-1}$. On input f_i , it computes the next output block $(X_i, K_{i+1}) := SC(K_{i-1}, X_{i-1})$ and the leakage $A_i = f_i(K_{i-1})$. It forwards X_i, A_i to A and deletes everything except the state $S_i = \{X_i, K_i, K_{i+1}\}$. After round $q-1$, G_0^{real} computes and forwards X_q (i.e., the next output block to be computed) to A . The game G_0^{rand} is defined in the same way, but the final block X_q is replaced with a uniformly random U_n .

⁸ We just need security against two random queries, so the well known non-uniform upper bounds on the security of block-ciphers of De, Trevisan and Tulsiani [6,5] do not seem to contradict such an assumption even in the non-uniform setting.

⁹ The theorem implies a stronger statement where one only requires that K has $k - \ell$ bits average-case min-entropy (which is implied by having K uniform and leaking ℓ bits), we state this weaker statement as it is sufficient for our application.

To prove that SC^F is an $(\varepsilon', s', \ell, q)$ -secure leakage-resilient stream cipher, we need to show that

$$G_0^{\text{real}} \sim_{\varepsilon', s'} G_0^{\text{rand}}, \quad (15)$$

for ε', s' as in the statement of the lemma.

Defining games G_i^{real} and G_i^{rand} for $1 \leq i \leq q - 1$. We define a series of games $G_1^{\text{real}}, \dots, G_{q-1}^{\text{real}}$ where G_{i+1}^{real} is derived from G_i^{real} by replacing X_i, K_{i+1} with uniformly random values $\tilde{X}_i, \tilde{K}_{i+1}$ and the leakage A_i with simulated fake leakage $\tilde{A}_i$ (the details are provided below). Games G_i^{rand} will be defined exactly as G_i^{real} except that (similarly to the case $i = 0$), the last block X_q is replaced with a uniformly random value.

For every i , $1 \leq i \leq q - 1$, the variables $\tilde{K}_i, \tilde{X}_i$ as defined by the oracles realizing the games G_j^{rand} and G_j^{real} where $j \geq i$ will satisfy the following properties (as the initial values (X_0, K_0, K_1) never get replaced, for notational convenience we define $(\tilde{X}_0, \tilde{K}_0, \tilde{K}_1) \stackrel{\text{def}}{=} (X_0, K_0, K_1)$)

- i. $\tilde{K}_i, \tilde{X}_i$ are uniformly random.
- ii. Right before the $(i - 1)$ th round (i.e. the round where the oracle computes $X_i \| K_{i+1} := F(\tilde{X}_{i-1}, \tilde{K}_{i-1})$), the oracle has leaked no information about $\tilde{K}_{i-1}$ except for the ℓ bits fake leakage $\tilde{A}_i$.
- iii. Right before the $(i - 1)$ th round $\tilde{K}_{i-1}$ and $\tilde{X}_{i-1}$ are independent given everything the oracle did output so far.

The first two properties above will follow from the definition of the games. The third point follows using Lemma 4 from [9], we will not discuss this here in detail, but only mention that the reason for the alternating structure of the cipher as illustrated in Figure 1, with an upper layer computing $K_0, K_2, \dots$ and the lower layer computing $K_1, K_3, \dots$, is to achieve this independence.

We now describe how the oracle G_{i+1}^{real} is derived from G_i^{real} . For concreteness, we set $i = 2$. In the third step, G_2^{real} computes $(X_3, K_4) := F(\tilde{K}_2, \tilde{X}_2)$, $A_3 = f_3(\tilde{K}_2)$ and forwards X_3, A_3 to A . The state stored after this step is $S_3 = \{X_3, \tilde{K}_3, K_4\}$. Let $V_2 \stackrel{\text{def}}{=} \{\tilde{X}^2, \tilde{A}^2\}$ be the view (i.e. all the outputs she got from her oracle) of the adversary A after the second round.

Fig. 1. Leakage resilient stream-cipher SC^F from a any weak pseudorandom function F . The regular evaluation is shown in black, the attack related part is shown in gray with dashed lines. The output of the cipher is $X_0, X_1, \dots$

Defining an intermediate oracle. We now define an oracle $G_{2/3}^{real}$ (which will be in-between G_2^{real} and G_3^{real}) derived from G_2^{real} by replacing $\Lambda_3 = f_3(\tilde{K}_2)$ with fake leakage $\tilde{\Lambda}_3$ computed as follows: let $h(\cdot)$ be a simulator for the leakage $\tilde{\Lambda}_3 := f_3(\tilde{K}_2)$ such that (for $\hat{\varepsilon}, \hat{s}$ to be defined)

$$(Z, h(Z)) \sim_{\hat{\varepsilon}, \hat{s}} (Z, \tilde{\Lambda}_3) \quad \text{where} \quad Z = \{V_2, X_3, K_4\} \quad (16)$$

By Theorem 1, there exists such a simulator of size $s_h \stackrel{\text{def}}{=} O(\hat{s}2^{3\ell}/\hat{\varepsilon}^2)$. Note that h not only gets the pseudorandom output X_3, K_4 whose computation has leaked bits, but also the view V_2 . The reason for the latter is that we need to fool an adversary who learned V_2 . Equation (16) then yields

$$G_2^{real} \sim_{\hat{\varepsilon}, \hat{s}-s_0} G_{2/3}^{real}, \quad (17)$$

where s_0 is the size of a circuit required to implement the real game G_0^{real} . The reason we loose s_0 in the circuit size here is that in a reduction where we use a distinguisher for G_2^{real} and $G_{2/3}^{real}$ to distinguish $(Z, h(Z))$ and $(Z, \tilde{\Lambda}_3)$ we must still compute the remaining $q-4$ rounds, and s_0 is an upper bound on the size of this computation.

The game G_3^{real} is derived from $G_{2/3}^{real}$ by replacing the values $X_3 \| K_4 := F(\tilde{K}_2, \tilde{X}_2)$ with uniformly random $\tilde{X}_3 \| \tilde{K}_4$ right after they have been computed (let us stress that also the fake leakage that is computed as in (16) now uses these random values, i.e., $Z = \{V_2, \tilde{X}_3, \tilde{K}_4\}$).

Proving indistinguishability. We claim that the games are indistinguishable with parameters

$$G_{2/3}^{real} \sim_{\sqrt{\varepsilon_F 2^\ell}, s_F/2 - s_h - s_0} G_3^{real} \quad (18)$$

Recall that in $G_{2/3}^{real}$, we compute $X_3 \| K_4 := F(\tilde{K}_2, \tilde{X}_2)$ where by i. $\tilde{X}_2, \tilde{K}_2$ are uniformly random, by ii. only ℓ bits of $\tilde{K}_2$ have leaked and iii. $\tilde{X}_2$ and $\tilde{K}_2$ are independent. Using these properties, equation (14) implies that the outputs are roughly $(\sqrt{\varepsilon_F 2^\ell}, s_F/2)$ pseudorandom, i.e.,

$$(\tilde{X}_2, X_3 \| K_4, \tilde{\Lambda}_1) \sim_{\sqrt{\varepsilon_F 2^\ell}, s_F/2} (\tilde{X}_2, \tilde{X}_3 \| \tilde{K}_4, \tilde{\Lambda}_1), \quad (19)$$

from which we derive (18). Note the loss of s_h in circuit size in equation (18) due to the fact that given a distinguisher for $G_{2/3}^{real}$ and G_3^{real} , we must recompute the fake leakage given only distributions as in (19).

We will assume that $s_0 \leq \hat{s}/2$, i.e., the real experiment is at most half as complex as the size of the adversaries we will consider (the setting where this is not the case is not very interesting anyway.) Then $\hat{s} - s_0 \geq \hat{s}/2$.

Up to this point, we have not yet defined what $\hat{\varepsilon}$ and $\hat{s}$ are, so we set them to

$$\hat{\varepsilon} \stackrel{\text{def}}{=} \sqrt{\varepsilon_F 2^\ell} \quad \text{and} \quad \hat{s} \stackrel{\text{def}}{=} \Theta(1) \frac{s_F \hat{\varepsilon}^2}{2^{3\ell}} \quad \text{then} \quad s_F = 8 \cdot s_h = \Theta(1) \frac{\hat{s} 2^{3\ell}}{\hat{\varepsilon}^2}.$$

With (17) and (18), we then get $G_2^{real} \sim_{2\hat{\varepsilon}, \hat{s}/2} G_3^{real}$. The same proof works for any $1 \leq i \leq q-1$, i.e., we have

$$G_i^{real} \sim_{2\hat{\varepsilon}, \hat{s}/2} G_{i+1}^{real}, \quad G_i^{rand} \sim_{2\hat{\varepsilon}, \hat{s}/2} G_{i+1}^{rand}.$$

Moreover, using i.-iii. with (14),

$$G_{q-1}^{real} \sim_{2\hat{\varepsilon}, \hat{s}/2} G_{q-1}^{rand}.$$

Using the triangle inequality $2q$ times, the two equations above yield

$$G_0^{real} \sim_{4q\hat{\varepsilon}, \hat{s}/2} G_0^{rand},$$

which which completes the proof of the lemma.

References

1. Barak, B., Dodis, Y., Krawczyk, H., Pereira, O., Pietrzak, K., Standaert, F.X., Yu, Y.: Leftover hash lemma, revisited. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 1–20. Springer, Heidelberg (2011)
2. Barak, B., Hardt, M., Kale, S.: The uniform hardcore lemma via approximate bregman projections. In: Mathieu, C. (ed.) SODA, pp. 1193–1200. SIAM (2009)
3. Bellare, M., Rompel, J.: Randomness-efficient oblivious sampling. In: FOCS, pp. 276–287 (1994)
4. Chung, K.M., Lui, E., Pass, R.: From weak to strong zero-knowledge and applications. Cryptology ePrint Archive, Report 2013/260 (2013), <http://eprint.iacr.org/>
5. De, A., Trevisan, L., Tulsiani, M.: Non-uniform attacks against one-way functions and prgs. Electronic Colloquium on Computational Complexity (ECCC) 16, 113 (2009)
6. De, A., Trevisan, L., Tulsiani, M.: Time space tradeoffs for attacks against one-way functions and PRGs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 649–665. Springer, Heidelberg (2010)
7. Dodis, Y., Pietrzak, K.: Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 21–40. Springer, Heidelberg (2010)
8. Dodis, Y., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 523–540. Springer, Heidelberg (2004)
9. Dziembowski, S., Pietrzak, K.: Intrusion-resilient secret sharing. In: 48th FOCS, pp. 227–237. IEEE Computer Society Press (October 2007)
10. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 49th FOCS, pp. 293–302. IEEE Computer Society Press (October 2008)
11. Frieze, A.M., Kannan, R.: Quick approximation to matrices and applications. Combinatorica 19(2), 175–220 (1999)
12. Fuller, B., O’Neill, A., Reyzin, L.: A unified approach to deterministic encryption: New constructions and a connection to computational entropy. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 582–599. Springer, Heidelberg (2012)
13. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Fortnow, L., Vadhan, S.P. (eds.) 43rd ACM STOC, pp. 99–108. ACM Press (June 2011)

14. Gowers, T.: Decompositions, approximate structure, transference, and the Hahn–Banach theorem. *Bull. London Math. Soc.* 42(4), 573–606 (2010)
15. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM Journal on Computing* 28(4), 1364–1396 (1999)
16. Herbster, M., Warmuth, M.K.: Tracking the best linear predictor. *Journal of Machine Learning Research* 1, 281–309 (2001)
17. Hsiao, C.Y., Lu, C.J., Reyzin, L.: Conditional computational entropy, or toward separating pseudoentropy from compressibility. In: Naor, M. (ed.) *EUROCRYPT 2007*. LNCS, vol. 4515, pp. 169–186. Springer, Heidelberg (2007)
18. Impagliazzo, R.: Hard-core distributions for somewhat hard problems. In: *FOCS*, pp. 538–545 (1995)
19. Krenn, S., Pietrzak, K., Wadia, A.: A counterexample to the chain rule for conditional HILL entropy - and what deniable encryption has to do with it. In: Sahai, A. (ed.) *TCC 2013*. LNCS, vol. 7785, pp. 23–39. Springer, Heidelberg (2013)
20. Pietrzak, K.: A leakage-resilient mode of operation. In: Joux, A. (ed.) *EUROCRYPT 2009*. LNCS, vol. 5479, pp. 462–482. Springer, Heidelberg (2009)
21. Reingold, O., Trevisan, L., Tulsiani, M., Vadhan, S.P.: Dense subsets of pseudorandom sets. In: *49th FOCS*, pp. 76–85. IEEE Computer Society Press (October 2008)
22. Reyzin, L.: Some notions of entropy for cryptography - (invited talk). In: Fehr, S. (ed.) *ICITS 2011*. LNCS, vol. 6673, pp. 138–142. Springer, Heidelberg (2011), <http://www.cs.bu.edu/~reyzin/papers/entropy-survey.pdf>
23. Tao, T., Ziegler, T.: The primes contain arbitrarily long polynomial progressions. *Acta Math.* 201, 213–305 (2008)
24. Trevisan, L.: Guest column: additive combinatorics and theoretical computer science. *SIGACT News* 40(2), 50–66 (2009)
25. Trevisan, L., Tulsiani, M., Vadhan, S.P.: Regularity, boosting, and efficiently simulating every high-entropy distribution. In: *IEEE Conference on Computational Complexity*, pp. 126–136 (2009)
26. Vadhan, S., Zheng, C.J.: A uniform min-max theorem with applications in cryptography. In: Canetti, R., Garay, J.A. (eds.) *CRYPTO 2013, Part I*. LNCS, vol. 8042, pp. 93–110. Springer, Heidelberg (2013)

A Proof of Theorem 1

We will prove Theorem 1 not for the family $\mathcal{F}$ directly, but for a family $\widehat{\mathcal{F}}$ which for every $f \in \mathcal{F}$ contains the function $\widehat{f}: \mathcal{X} \times \{0, 1\}^\ell \rightarrow [-1, 1]$ defined as

$$\widehat{f}(x, b) = f(x, b) - w_f(x) \quad \text{where} \quad w_f(x) = \mathbb{E}_{b \leftarrow \{0, 1\}^\ell} [f(x, b)] = 2^{-\ell} \sum_{b \in \{0, 1\}^\ell} f(x, b)$$

Any simulator that fools $\widehat{\mathcal{F}}$ also fools $\mathcal{F}$ with the same advantage since $\forall \widehat{f} \in \widehat{\mathcal{F}}$,

$$\begin{aligned} & \left| \mathbb{E}_{x \leftarrow X, [g]} [\widehat{f}(x, g(x))] - \mathbb{E}_{x \leftarrow X, [h]} [\widehat{f}(x, h(x))] \right| \\ &= \left| \mathbb{E}_{x \leftarrow X, [g]} [f(x, g(x)) - w_f(x)] - \mathbb{E}_{x \leftarrow X, [h]} [f(x, h(x)) - w_f(x)] \right| \\ &= \left| \mathbb{E}_{x \leftarrow X, [g]} [f(x, g(x))] - \mathbb{E}_{x \leftarrow X, [h]} [f(x, h(x))] \right| \end{aligned}$$

Evaluating $\widehat{f}$ requires 2^ℓ evaluations of f as we need to compute $w_f(x)$. We thus lose a factor of 2^ℓ in efficiency by considering $\widehat{\mathcal{F}}$ instead of $\mathcal{F}$. The reason that we prove the theorem for $\widehat{\mathcal{F}}$ instead of for $\mathcal{F}$ is because in what follows, we will need that for any x , the expectation over a uniformly random $b \in \{0, 1\}^\ell$ is 0, i.e.,

$$\forall \widehat{f} \in \widehat{\mathcal{F}}, x \in \mathcal{X}: \mathbb{E}_{b \leftarrow \{0, 1\}^\ell} [\widehat{f}(x, b)] = 0. \quad (20)$$

To prove the theorem, we must show that for any joint distribution $(X, g(X))$ over $\mathcal{X} \times \{0, 1\}^\ell$, there exists an efficient simulator $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ such that

$$\forall \widehat{f} \in \widehat{\mathcal{F}}: \left| \mathbb{E}_{x \leftarrow X} [\widehat{f}(x, g(x)) - \widehat{f}(x, h(x))] \right| < \varepsilon. \quad (21)$$

Moving to product form. We define the function $\widetilde{g}: \mathcal{X} \times \{0, 1\}^\ell \rightarrow [0, 1]$ as $\widetilde{g}(x, a) := \mathbb{P}_{[g]}[g(x) = a]$. Note that for every $x \in \mathcal{X}$, we have

$$\sum_{a \in \{0, 1\}^\ell} \widetilde{g}(x, a) = 1. \quad (22)$$

We can write the expected value of $\widehat{f}(X, g(X))$ as follows:

$$\begin{aligned} \mathbb{E}_{x \leftarrow X, [g]} [\widehat{f}(x, g(x))] &= \sum_{a \in \{0, 1\}^\ell} \mathbb{E}_{x \leftarrow X} \left[\widehat{f}(x, a) \mathbb{P}_{[g]}[g(x) = a] \right] = \\ &= \sum_{a \in \{0, 1\}^\ell} \mathbb{E}_{x \leftarrow X} \left[\widehat{f}(x, a) \widetilde{g}(x, a) \right] = \\ &= 2^\ell \mathbb{E}_{x \leftarrow X, u \leftarrow \{0, 1\}^\ell} [\widehat{f}(x, u) \widetilde{g}(x, u)]. \end{aligned} \quad (23)$$

We will construct a simulator $\widetilde{h}: \mathcal{X} \times \{0, 1\}^\ell \rightarrow [0, 1]$ such that for $\gamma > 0$ (to be defined later),

$$\forall \widehat{f} \in \widehat{\mathcal{F}}: \mathbb{E}_{x \leftarrow X, b \leftarrow \{0, 1\}^\ell} [\widehat{f}(x, b)(\widetilde{g}(x, b) - \widetilde{h}(x, b))] < \gamma. \quad (24)$$

From this $\widetilde{h}$, we can then get a simulator $h(\cdot)$ like in (21) assuming that $\widetilde{h}(x, \cdot)$ is a probability distribution for all x , i.e., $\forall x \in \mathcal{X}$,

$$\sum_{b \in \{0, 1\}^\ell} \widetilde{h}(x, b) = 1, \quad (25)$$

$$\forall b \in \{0, 1\}^\ell : \widetilde{h}(x, b) \geq 0. \quad (26)$$

We will define a sequence $h_0, h_1, \dots$ of functions where $h_0(x, b) = 2^{-\ell}$ for all x, b .¹⁰ Define the energy function

$$\Delta_t = \mathbb{E}_{x \leftarrow X, b \leftarrow \{0, 1\}^\ell} [(\widetilde{g}(x, b) - h_t(x, b))^2].$$

¹⁰ It is not relevant how exactly h_0 is defined, but we need $\sum_{b \leftarrow \{0, 1\}^\ell} [h_0(x, b)] = 1$ for all $x \in \mathcal{X}$.

Assume that after the first t steps, there exists a function $\widehat{f}_{t+1}: \mathcal{X} \times \{0, 1\}^\ell \rightarrow [-1, 1]$ such that

$$\mathbb{E}_{x \leftarrow X, b \leftarrow \{0, 1\}^\ell} [\widehat{f}_{t+1}(x, b)(g(x, b) - h_t(x, b))] \geq \gamma,$$

and define

$$h_{t+1}(x, b) = h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b) \quad (27)$$

The energy function then decreases by γ^2 , i.e.,

$$\begin{aligned} & \Delta_{t+1} \\ &= \mathbb{E}_{x \leftarrow X, b \leftarrow \{0, 1\}^\ell} [(\widetilde{g}(x, b) - h_t(x, b) - \gamma \widehat{f}_{t+1}(x, b))^2] = \\ &= \Delta_t + \underbrace{\mathbb{E}_{x \leftarrow X, b \leftarrow \{0, 1\}^\ell} [\gamma^2 \widehat{f}_{t+1}^2(x, b)]}_{\leq \gamma^2} - \underbrace{\mathbb{E}_{x \leftarrow X, b \leftarrow \{0, 1\}^\ell} [2\gamma \widehat{f}_{t+1}(x, b)(\widetilde{g}(x, b) - h_t(x, b))]}_{\geq 2\gamma^2} \\ &\leq \Delta_t - \gamma^2. \end{aligned}$$

Since $\Delta_0 \leq 1$, $\Delta_t \geq 0$ for any t (as it is a square) and $\Delta_i - \Delta_{i+1} \geq \gamma^2$, this process must terminate after at most $1/\gamma^2$ steps meaning that we have constructed $\widetilde{h} = h_t$ that satisfies (24). Note that the complexity of the constructed $\widetilde{h}$ is bounded by $2^\ell \gamma^{-2}$ times the complexity of the functions from $\mathcal{F}$ since, as mentioned earlier, computing $\widetilde{f}$ requires 2^ℓ evaluations of f . In other words, $\widetilde{h}$ has complexity $\mathcal{O}(2^\ell \gamma^{-2})$ relative to $\mathcal{F}$.

Moreover, since for all $x \in \mathcal{X}$ and $\widehat{f} \in \widehat{\mathcal{F}}$, we have $\sum_{b \in \{0, 1\}^\ell} h_0(x, b) = 1$ and

$\sum_{b \in \{0, 1\}^\ell} \widehat{f}(x, b) = 0$, condition (25) holds as well. Unfortunately, (26) does not

hold since it might be the case that $h_{t+1}(x, b) < 0$. We will explain later how to fix this problem by replacing $\widehat{f}_{t+1}$ in (27) with a similar function $\widehat{f}_{t+1}^*$ that satisfies $h_{t+1}(x, b) = h_t + \gamma \widehat{f}_{t+1}^* \geq 0$ for all x and b in addition to all of the properties just discussed. Assume for now that $\widetilde{h}$ satisfies (24)-(26).

Let $h: \mathcal{X} \rightarrow \{0, 1\}^\ell$ be a probabilistic function defined as follows: we set $h(x) = b$ with probability $\widetilde{h}(x, b)$. Equivalently, imagine that we have a biased dice with 2^ℓ faces labeled by $b \in \{0, 1\}^\ell$ such that the probability of getting the face with label b is $\widetilde{h}(x, b)$. We then define $h(x)$ by simply throwing this dice and reading off the label. It follows that $\mathbb{P}_{[h]}[h(x) = b] = \widetilde{h}(x, b)$. This probabilistic function satisfies

$$\begin{aligned} \mathbb{E}_{[h], x \leftarrow X} [\widehat{f}(x, h(x))] &= \mathbb{E}_{x \leftarrow X} \sum_{a \in \{0, 1\}^\ell} \widehat{f}(x, a) \mathbb{P}_{[h]}[h(x) = a] \\ &= \mathbb{E}_{x \leftarrow X} \sum_{a \in \{0, 1\}^\ell} \widehat{f}(x, a) h_t(x, a) \\ &= \mathbb{E}_{x \leftarrow X, u \leftarrow \{0, 1\}^\ell} 2^\ell \widehat{f}(x, u) h_t(x, u). \end{aligned} \quad (28)$$

Plugging (28) and (23) into (24), we obtain

$$\forall \widehat{f} \in \mathcal{F} : \mathbb{E}_{x \leftarrow X, [h]} \left[\frac{\widehat{f}(x, g(x))}{2^\ell} - \frac{\widehat{f}(x, h(x))}{2^\ell} \right] < \gamma.$$

Equivalently,

$$\forall \widehat{f} \in \mathcal{F} : \mathbb{E}_{x \leftarrow X, [h]} \left[\widehat{f}(x, g(x)) - \widehat{f}(x, h(x)) \right] < \gamma 2^\ell \quad (29)$$

We get (4) from the statement of the theorem by setting $\gamma := \varepsilon/2^\ell$. The simulator $\widetilde{h}$ is thus of complexity $\mathcal{O}(2^{3\ell}(1/\varepsilon)^2)$ relative to $\mathcal{F}$.

Enforcing $h_t(x, b) \geq 0$ for $\ell = 1$. We now fix the problem with the positivity of $h_t(x, b)$. Consider the case $\ell = 1$. Consider the following properties:

- i. $\sum_{b \in \{0,1\}} h_t(x, b) = 1$ for $x \in \mathcal{X}$,
- ii. $\forall b \in \{0, 1\}, h_t(x, b) \geq 0$ for $x \in \mathcal{X}$,
- iii. $\mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}_{t+1}(x, b)(g(x, b) - h_t(x, b))] \geq \gamma$ for $\gamma > 0$.

Assume that $h_t: \mathcal{X} \rightarrow \{0, 1\}$ and $\widehat{f}_{t+1}: \mathcal{X} \times \{0, 1\} \rightarrow [-1, 1]$ satisfy i) and ii) for all $x \in \mathcal{X}$ and iii) for some $\gamma > 0$. Recall that $\Delta_t = \mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}} [(\widetilde{g}(x, b) - h_t(x, b))^2]$. We have shown that $h_{t+1} = h_t + \gamma \widehat{f}_{t+1}$ satisfies

$$\Delta_{t+1} \leq \Delta_t - \gamma^2. \quad (30)$$

Moreover, for all $x \in \mathcal{X}$, h_{t+1} will still satisfy i) but not necessarily ii). We define a function $\widehat{f}_{t+1}^*$ such that setting $h_{t+1} = h_t + \gamma \widehat{f}_{t+1}^*$ will satisfy i) and ii) for all $x \in \mathcal{X}$ and an inequality similar to (30).

First, for any $x \in \mathcal{X}$ for which condition ii) is satisfied, let $f_{t+1}^* = \widehat{f}_{t+1}$. Consider now $x \in \mathcal{X}$ for which ii) fails for some $b \in \{0, 1\}$, i.e., for which $h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b) < 0$. Let $\gamma' = -h_t(x, b)/\widehat{f}_{t+1}(x, b)$. Note that $0 \leq \gamma' \leq \gamma$ and $h_t(x, b) + \gamma' \widehat{f}_{t+1}(x, b) = 0$. Let

$$\widehat{f}_{t+1}^*(x, b) = \frac{\gamma'}{\gamma} \widehat{f}_{t+1}(x, b) \quad \widehat{f}_{t+1}^*(x, 1-b) = \widehat{f}_{t+1}(x, 1-b) + \frac{1-\gamma'}{\gamma} \widehat{f}_{t+1}(x, b).$$

Let $h_{t+1}(x, \cdot) = h_t(x, \cdot) + \gamma \widehat{f}_{t+1}^*(x, \cdot)$ and note that

$$\sum_{b \in \{0,1\}} \widehat{f}_{t+1}^*(x, b) = \sum_{b \in \{0,1\}} \widehat{f}_{t+1}(x, b) = 0.$$

Condition i) is then satisfied for h_{t+1} for any $x \in \mathcal{X}$. By the definition of γ' , condition ii) is satisfied for any $x \in \mathcal{X}$ as well. Condition iii) is more delicate and in fact need not hold. Yet, we will prove the following:

Lemma 3. *If $\widehat{f}_{t+1}$ and h_t satisfy i) and ii) for every $x \in \mathcal{X}$, and iii) then*

$$\begin{aligned} & \mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}_{t+1}(x, b)(g(x, b) - h_t(x, b))] \\ & - \mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}_{t+1}^*(x, b)(g(x, b) - h_t(x, b))] \leq \frac{\gamma}{4}. \end{aligned} \quad (31)$$

Proof. To prove (31), it suffices to show that for every $x \in \mathcal{X}$,

$$\sum_{b \in \{0,1\}} \widehat{f}_{t+1}(x, b)(g(x, b) - h_t(x, b)) - \sum_{b \in \{0,1\}} \widehat{f}_{t+1}^*(x, b)(g(x, b) - h_t(x, b)) \leq \frac{\gamma}{2}. \quad (32)$$

If $x \in \mathcal{X}$ is such that ii) is satisfied for h_{t+1} then there is nothing to prove. Suppose that ii) fails for some $x \in \mathcal{X}$ and $b \in \{0,1\}$. For brevity, let $f := \widehat{f}_{t+1}(x, b)$, $g := g(x, b)$, $h = h_t(x, b)$. We have $-1 \leq f < 0$, $h + \gamma f < 0$, $0 \leq g \leq 1$ and $h = -\gamma f^*$. Using $g - h \geq -h$, the left-hand side of (32) then satisfies

$$\begin{aligned} & 2(f + h/\gamma)(g - h) \leq 2(f + h/\gamma)(-h) \\ & = \frac{2}{\gamma}(-f\gamma - h)h \leq \frac{2}{\gamma} \left(\frac{-f\gamma - h + h}{2} \right)^2 = \frac{\gamma f^2}{2} \leq \frac{\gamma}{2}, \end{aligned} \quad (33)$$

where we have used the inequality $uv \leq \left(\frac{u+v}{2}\right)^2$.

If iii) holds then Lemma 3 implies $\gamma - \mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}_{t+1}^*(x, b)(g(x, b) - h_t(x, b))] \leq \frac{\gamma}{4}$. Equivalently,

$$\mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}_{t+1}^*(x, b)(g(x, b) - h_t(x, b))] \geq \frac{3\gamma}{4}. \quad (34)$$

Defining $h_{t+1} = h_t + \gamma \widehat{f}_{t+1}^*$, we still get

$$\Delta_{t+1} \leq \Delta_t - \left(\frac{3\gamma}{4}\right)^2 = \Delta_t - \frac{9\gamma^2}{16}. \quad (35)$$

Remark 2. In this case, the slightly worse inequality (35) will increase the complexity of $\widetilde{h}$, but only by a constant factor of 16/9, i.e., $\widetilde{h}$ will still have complexity $\mathcal{O}(2^\ell \gamma^{-2})$ relative to $\mathcal{F}$.

Enforcing $h_t(x, b) \geq 0$ for general ℓ . Let $\widehat{f}_{t+1}(x, b)$ be as before and suppose that there exists $x \in \mathcal{X}$ such that $h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b) < 0$ for at least one $b \in \{0,1\}^\ell$. We will show how to replace $\widehat{f}_{t+1}$ with another function $\widehat{f}_{t+1}^*$ such that it satisfies an inequality of type (34) and such that $h_{t+1}(x, b) = h_t(x, b) + \gamma \widehat{f}_{t+1}^*(x, b) \geq 0$. Let S be the set of all elements $b \in \{0,1\}^\ell$ for which $h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b) < 0$. For $b \in S$, it follows that $\widehat{f}_{t+1}(x, b) < 0$. As before, for $b \in S$,

define $\widehat{f}_{t+1}^*(x, b) = -\frac{h_t(x, b)}{\gamma}$. Note that for each such b , we have added a positive mass $-\frac{h_t(x, b) + \gamma \widehat{f}_{t+1}^*(x, b)}{\gamma}$ to modify each $\widehat{f}_{t+1}(x, b)$. Let

$$M = \sum_{b \in S} -\left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right) \quad (36)$$

be the total mass. For $b \notin S$, define $\widehat{f}_{t+1}^*(x, b) = \widehat{f}_{t+1}(x, b) - \frac{M}{2^\ell - s}$. Clearly, $\mathbb{E}_{b \leftarrow \{0,1\}^\ell} \widehat{f}_{t+1}^*(x, b) = 0$. We will now show the following

Lemma 4. *For every $x \in \mathcal{X}$, the function $\widehat{f}_{t+1}^*$ satisfies*

$$\sum_{b \in \{0,1\}^\ell} (\widehat{f}_{t+1}(x, b) - \widehat{f}_{t+1}^*(x, b))(g(x, b) - h_t(x, b)) < 2^{\ell-1}\gamma.$$

Proof. Let $s = |S|$ and $h_S = \sum_{i=1}^s h_t(x, b_i)$. First, note that (as in the case $\ell = 1$)

$$\begin{aligned} \forall b \in S: \quad & \left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right)(g(x, b) - h_t(x, b)) \\ & \leq -\left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right)h_t(x, b). \end{aligned} \quad (37)$$

Moreover,

$$\sum_{b \notin S} g(x, b) \leq \sum_{b \in \{0,1\}^\ell} g(x, b) = 1. \quad (38)$$

The difference that we want to estimate is then

$$\begin{aligned} \Delta &= \sum_{b \in \{0,1\}^\ell} (\widehat{f}_{t+1}(x, b) - \widehat{f}_{t+1}^*(x, b))(g(x, b) - h_t(x, b)) \\ &= \sum_{b \in S} \left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right)(g(x, b) - h_t(x, b)) + \frac{M}{2^\ell - s} \sum_{b \notin S} (g(x, b) - h_t(x, b)) \\ &\stackrel{(37), (38)}{\leq} \sum_{b \in S} -\left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right)h_t(x, b) + \frac{M}{2^\ell - s} \underbrace{\left(1 - \sum_{b \notin S} h_t(x, b)\right)}_{=h_S} \\ &\stackrel{(36)}{=} \sum_{b \in S} \underbrace{-\left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right)h_t(x, b)}_{\leq \gamma/4} + \frac{h_S}{2^\ell - s} \sum_{b \in S} -\left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right) \\ &\stackrel{(33)}{\leq} \frac{s\gamma}{4} - \frac{h_S}{2^\ell - s} \sum_{b \in S} \widehat{f}_{t+1}(x, b) - \frac{h_S^2}{\gamma(2^\ell - s)} = \frac{s\gamma}{4} + \frac{h_S f_S}{2^\ell - s} - \frac{h_S^2}{\gamma(2^\ell - s)}, \end{aligned}$$

where $f_S = -\sum_{b \in S} \hat{f}_{t+1}(x, b)$. Note that $\sum_{b \in S} -\hat{f}_{t+1}(x, b) \leq s$ and (using (20)) $\sum_{b \in S} -\hat{f}_{t+1}(x, b) = \sum_{b \notin S} \hat{f}_{t+1}(x, b) \leq 2^\ell - s$, i.e., $f_S \leq \min\{s, 2^\ell - s\}$. Since

$$\begin{aligned} \frac{h_S f_S}{2^\ell - s} - \frac{h_S^2}{\gamma(2^\ell - s)} &= \frac{1}{\gamma(2^\ell - s)} h_S (\gamma f_S - h_S) \\ &\leq \frac{1}{\gamma(2^\ell - s)} \left(\frac{h_S + (\gamma f_S - h_S)}{2} \right)^2 \leq \frac{s\gamma}{4}, \end{aligned}$$

where we have used that $f_S^2 \leq s(2^\ell - s)$. Since $s < 2^\ell$, we obtain $\Delta \leq \frac{s\gamma}{2} < 2^{\ell-1}\gamma$ which proves the lemma.

To complete the proof, note that the above lemma implies that

$$\begin{aligned} &\mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}^\ell} [\hat{f}_{t+1}(x, b)(g(x, b) - h_t(x, b))] \\ &- \mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}^\ell} [\hat{f}_{t+1}^*(x, b)(g(x, b) - h_t(x, b))] < \frac{\gamma}{2}, \end{aligned}$$

and hence,

$$\mathbb{E}_{x \leftarrow X, b \leftarrow \{0,1\}^\ell} [\hat{f}_{t+1}^*(x, b)(g(x, b) - h_t(x, b))] > \frac{\gamma}{2}. \quad (39)$$

Remark 3. Similarly, the slightly worse inequality (39) will increase the complexity of $\tilde{h}$ by a constant factor of 4, i.e., $\tilde{h}$ will still have complexity $\mathcal{O}(2^\ell \gamma^{-2})$ relative to $\mathcal{F}$.

A.1 Derandomizing $\tilde{h}$

Next, we discuss how to derandomize $\tilde{h}$. We can think of the probabilistic function $\tilde{h}$ as a deterministic function $\tilde{h}'$ taking two inputs where the second input represents the random coins used by $\tilde{h}$. More precisely, for $R \leftarrow \{0,1\}^\rho$ (ρ is an upper bound on the number of random bits used by $\tilde{h}$) and for any x in the support of X , we have $\tilde{h}'(x, R) \sim \tilde{h}(x)$.

To get our derandomized $\hat{h}$, we replace the randomness R with the output of a function ϕ chosen from a family of t -wise independent functions for some large t , i.e., we set $\hat{h}(x) = \tilde{h}'(x, \phi(x))$. Recall that a family Φ of functions $\mathcal{A} \rightarrow \mathcal{B}$ is t -wise independent if for any t distinct inputs $a_1, \dots, a_t \in \mathcal{A}$ and a randomly chosen $\phi \leftarrow \Phi$, the outputs $\phi(a_1), \dots, \phi(a_t)$ are uniformly random in $\mathcal{B}^t$. In the proof, we use the following tail inequality for variables with bounded independence:

Lemma 5 (Lemma 2.2 from [3]). *Let $t \geq 6$ be an even integer and let $Z_1, \dots, Z_n$ be t -wise independent variables taking values in $[0, 1]$. Let $Z = \sum_{i=1}^n Z_i$, then for any $A > 0$*

$$\mathbb{P}[|Z - \mathbb{E}[Z]| \geq A] \leq \left(\frac{nt}{A^2} \right)^{t/2}$$

Recall that the min-entropy of X is $H_\infty(X) = -\log\left(\max_x \mathbb{P}[X = x]\right)$, or equivalently, X has min-entropy k if $\mathbb{P}[X = x] \leq 2^{-k}$ for all $x \in \mathcal{X}$.

Lemma 6. (Deterministic Simulation) *Let $\varepsilon > 0$ and assume that*

$$H_\infty(X) > 2 + \log \log |\mathcal{F}| + 2 \log(1/\varepsilon). \quad (40)$$

For any (probabilistic) $\tilde{h}: \mathcal{X} \rightarrow \{0, 1\}^\ell$, there exists a deterministic $\hat{h}$ of the same complexity relative to $\mathcal{F}$ as $\tilde{h}$ such that

$$\forall f \in \mathcal{F}: \left| \mathbb{E}_{x \leftarrow X, [\tilde{h}]} [f(x, \tilde{h}(x))] - \mathbb{E}_{x \leftarrow X} [f(x, \hat{h}(x))] \right| < \varepsilon \quad (41)$$

Remark 4. About the condition (40). A lower bound on the min-entropy of X in terms of $\log \log |\mathcal{F}|$ and $\log(1/\varepsilon)$ as in (40) is necessary. For example one can show that for $\varepsilon < 1/2$, (41) implies $H_\infty(X) \geq \log \log |\mathcal{F}|$. To see this, consider the case when X is uniform over $\{0, 1\}^m$ (so $H_\infty(X) = m$), $\mathcal{F}$ contains all 2^{2^m} functions $f: \{0, 1\}^m \times \{0, 1\} \rightarrow \{0, 1\}$ satisfying $f(x, 1 - b) = 1 - f(x, b)$ for all $x, b \in \{0, 1\}^{m+1}$, and $\tilde{h}(x) \sim U_1$ is uniformly random for all x (so it ignores its input). Now, given any deterministic $\hat{h}$, we can choose $f \in \mathcal{F}$ where $f(x, \hat{h}(x)) = 1$ for all $x \in \{0, 1\}^m$ (such an f exists by definition of $\mathcal{F}$). For this f ,

$$\left| \underbrace{\mathbb{E}_{x \leftarrow X, [\tilde{h}]} [f(x, \tilde{h}(x))]}_{=1/2} - \underbrace{\mathbb{E}_{x \leftarrow X} [f(x, \hat{h}(x))]}_{=1} \right| = 1/2.$$

In terms of $\log(1/\varepsilon)$, one can show that (41) implies $H_\infty(X) \geq \log(1/\varepsilon) - 1$ (even if $|\mathcal{F}| = 1$). For this, let $\tilde{h}$ and X be as above, $\mathcal{F} = \{f\}$ is defined as $f(x, b) = b$ if $x = 0^m$ and $f(x, b) = 0$ otherwise. For any deterministic $\hat{h}$, we get

$$\left| \underbrace{\mathbb{E}_{x \leftarrow X, [\tilde{h}]} [f(x, \tilde{h}(x))]}_{1/2^{m+1}} - \underbrace{\mathbb{E}_{x \leftarrow X} [f(x, \hat{h}(x))]}_{1/2^m \text{ or } 0} \right| = 1/2^{m+1}$$

and thus, $\varepsilon = 1/2^{m+1}$. Equivalently $H_\infty(X) = m = \log(1/\varepsilon) - 1$. The condition (40) is mild and in particular, it covers the cryptographically interesting case where $\mathcal{F}$ is the family of polynomial-size circuits (i.e., for a security parameter n and a constant c , $|\mathcal{F}| \leq 2^{n^c}$), X has superlogarithmic min-entropy $H_\infty(X) = \omega(\log n)$ and $\varepsilon > 0$ is negligible in n . Here, (40) becomes

$$\omega(\log n) > 2 + c \log n + 2 \log \varepsilon^{-1}$$

which holds for a negligible $\varepsilon = 2^{-\omega(\log n)}$.

Proof (Proof of Lemma 5). Let $m = H_\infty(X)$. We will only prove the lemma for the restricted case where X is flat, i.e., it is uniform on a subset $\mathcal{X}' \subseteq \mathcal{X}$ of size 2^m .¹¹ Consider any fixed $f \in \mathcal{F}$ and the 2^m random variables $Z_x \in \{0, 1\}$ indexed by $x \in \mathcal{X}'$ sampled as follows: first, sample $\phi \leftarrow \Phi$ from a family of t -wise independent functions $\mathcal{X} \rightarrow \{0, 1\}^\rho$ (recall that ρ is an upper bound on the number of random bits used by $\tilde{h}$). Now, Z_x is defined as

$$Z_x = f(x, \tilde{h}'(x, \phi(x))) = f(x, \hat{h}(x))$$

and $Z = \sum_{x \in \mathcal{X}'} Z_x$. Note that the same ϕ is used for all Z_x .

1. The variables Z_x for $x \in \mathcal{X}'$ are t -wise independent, i.e., for any t distinct $x_1, \dots, x_t$, the variables $Z_{x_1}, \dots, Z_{x_t}$ have the same distribution as $Z'_{x_1}, \dots, Z'_{x_t}$ sampled as $Z'_{x_i} \leftarrow f(x_i, \tilde{h}'(x_i, R))$. The reason is that the randomness $\phi(x_1), \dots, \phi(x_t)$ used to sample the $Z_{x_1}, \dots, Z_{x_t}$ is uniform in $\{0, 1\}^\rho$ as ϕ is t -wise independent.
2. $\mathbb{E}[Z_x] = \mathbb{E}_{\phi \leftarrow \Phi} [f(x, \tilde{h}'(x, \phi(x)))] = \mathbb{E}_{[\tilde{h}]} [f(x, \tilde{h}(x))]$.
3. $\mathbb{P}_{x \leftarrow X, [\tilde{h}]} [f(x, \tilde{h}(x)) = 1] = \mathbb{E}_{\phi \leftarrow \Phi} [Z/2^m]$.

Let $\mu = \mathbb{E}_{\phi \leftarrow \Phi} [Z] = \mathbb{E}_{\phi \leftarrow \Phi} \left[\sum_{x \in \mathcal{X}'} Z_x \right]$. By Lemma 5, we have

$$\mathbb{P}[|Z - \mu| \geq \varepsilon 2^m] \leq \left(\frac{t}{\varepsilon^2 2^m} \right)^{t/2}.$$

Let us call ϕ bad for f if $|Z - \mu| \geq \varepsilon 2^m$ (or equivalently, using iii)),

$$\left| \mathbb{E}_{x \leftarrow X, [\tilde{h}]} [f(x, \tilde{h}(x))] - \mathbb{E}_{x \leftarrow X} [f(x, \phi(x))] \right| \geq \varepsilon$$

We want to choose t such that the probability of ϕ being bad for any particular $f \in \mathcal{F}$ is less than $1/|\mathcal{F}|$, i.e.

$$(t/\varepsilon^2 2^m)^{t/2} < |\mathcal{F}|^{-1}. \quad (42)$$

We postpone for a second how to choose t and discussing when this is even possible. Assuming (42),

$$\mathbb{P}_{\phi \leftarrow \Phi} [|Z - \mu| \geq \varepsilon 2^m] \leq \left(\frac{t}{\varepsilon^2 2^m} \right)^{t/2} < |\mathcal{F}|^{-1},$$

¹¹ Any distribution satisfying $H_\infty(X) = m$ can be written as a convex combination of flat distributions with min-entropy m . Often, this fact is sufficient to conclude that a result proven for flat distributions with min-entropy m implies the result for any distribution with the same min-entropy. Here, this is not quite the case, because we might end up using a different ϕ for every flat distribution. But as the only property we actually require from X is $\mathbb{P}[X = x] \leq 2^{-m}$, the proof goes through for general X , but becomes somewhat more technical.

and by taking a union bound over all $f \in \mathcal{F}$, we get

$$\mathbb{P}_{\phi \leftarrow \Phi} [\exists f \in \mathcal{F} : |Z - \mu| \geq \varepsilon 2^m] < 1,$$

which implies that there exists $\phi \in \Phi$ such that

$$\forall f \in \mathcal{F} : |Z - \mu| < \varepsilon 2^m.$$

Equivalently, using how Z and μ were defined,

$$\forall f \in \mathcal{F} : \left| \sum_{x \in \mathcal{X}'} f(x, \hat{h}(x)) - \sum_{x \in \mathcal{X}'} f(x, \tilde{h}(x)) \right| < \varepsilon 2^m.$$

Finally, using that X is uniform over $\mathcal{X}'$, we get (for the above choice of ϕ) the statement of the lemma

$$\forall f \in \mathcal{F} : \left| \mathbb{E}_{x \leftarrow X} [f(x, \hat{h}(x))] - \mathbb{E}_{x \leftarrow X, [\tilde{h}]} [f(x, \tilde{h}(x))] \right| < \varepsilon.$$

We still have to determine when t can be chosen so that (42) holds. By taking logarithm and rearranging the terms, (42) becomes

$$mt/2 > \log |\mathcal{F}| + (t/2) \log(t) + t \log(1/\varepsilon),$$

i.e.,

$$m > 2 \log |\mathcal{F}|/t + \log(t) + 2 \log(1/\varepsilon).$$

Setting $t = \log |\mathcal{F}|$, we get

$$m > 2 + \log \log |\mathcal{F}| + 2 \log(1/\varepsilon).$$

which holds as it is the condition (5) we made on the min-entropy $m = H_\infty(X)$.