

Sharp lower bounds on the extractable randomness from non-uniform sources

B. Škorić, C. Obi, E. Verbitskiy, B. Schoenmakers

Abstract

Extraction of uniform randomness from (noisy) non-uniform sources is an important primitive in many security applications, e.g. (pseudo-)random number generators, privacy-preserving biometrics, and key storage based on Physical Unclonable Functions. Generic extraction methods exist, using universal hash functions. There is a trade-off between the length of the extracted bit string and the uniformity of the string. In the literature there are proven lower bounds on this length as a function of the desired uniformity. The best known bound involves a quantity known as smooth min-entropy. Unfortunately, there exist at least three definitions of smooth entropy. In this paper we compare three of these definitions, and we derive improved lower bounds on the extractable randomness.

We also investigate the use of almost universal hash functions, which are slightly worse at extracting randomness than universal hash functions, but are preferable in practice because they require far less nonvolatile memory. We show that using them has negligible effect on the extractable randomness.

1 Introduction

1.1 Randomness extraction from non-uniform noisy sources

Many cryptographic applications require input bitstrings to be uniformly distributed and exactly reproducible. Cryptographic keys, for instance, have to be uniformly random in order to prevent efficient attacks; they have to be reproducible in order to allow for decryption of encrypted data, verification of signatures, successful authentication etc. Even a single bit error in a key causes failure.

Physical sources of randomness, however, are neither uniform nor noise-free. Consider for example biometric measurements. The patterns in fingerprints and iris scans do not follow a uniform distribution, and they are never exactly re-obtained when a measurement is repeated. Measurement noise can be due to many factors, e.g. differences in lighting conditions or sensor alignment, physiological changes, difference between sensors etc. Another class of physical sources that has received a lot of attention recently are the Physical Unclonable Functions (PUFs), also known as Physical One-Way Functions, Physical Random Functions and Physically Obscured Keys. PUFs can be regarded as ‘non-biological biometrics’ in the sense that an identification/authentication string or secret key is derived from measuring an intrinsic property of an object. Many types of physical system for use as a PUF have been described in the literature, for instance three-dimensional multiple scattering of laser light [12], reflection of laser light from paper fibers [2], randomized dielectric properties in protective coatings on integrated circuits [17], radiofrequent responses from pieces of metal wire [4] or thin-film resonators [19], delay times in chip components [3] and start-up values

of SRAM cells [7]. Even typed passphrases can be regarded as a non-uniform noisy source in view of possible typing mistakes.

For security and/or privacy reasons it is often necessary to apply a one-way hash function to the measurement of the biometric/PUF, in analogy with the `/etc/passwd` file in UNIX. The storage of biometric/PUF data is assumed to be public; the hashing step hides the measurement data. However, as the measurements are noisy, it is not possible to directly apply a hash function; a single bit error in the input causes roughly 50% of the output bits to flip. Hence, an error-correction step is required first (‘information reconciliation’). This is not trivial, since the redundancy data has to be stored publicly and may reveal too much privacy-sensitive information. Similarly, if PUF data is to be used as a cryptographic key, then it should be thoroughly noise-corrected first. Here, too, it is crucial that the publicly stored redundancy data does not reveal secrets (in this case the extracted key).

After information reconciliation, the step of *privacy amplification* is applied, mapping a non-uniform random string to a shorter, almost uniform string. The requirement of uniformity is obvious in the case of key extraction. Interestingly, extracting uniform bitstrings is also desirable in biometric identification and PUF-based anti-counterfeiting, applications where the identifiers are not considered to be secret. A uniform string is the most efficient way of storing the entropy present in a measurement. Furthermore, database search speed is improved. Apart from biometrics and PUF applications, the subject of privacy amplification is highly relevant to other areas, e.g. true random number generation.

The concept of a *Fuzzy Extractor* [5, 6], also known as a *helper data scheme* [11], was introduced as a primitive that achieves both information reconciliation and privacy amplification¹. The publicly stored enrolment data (a.k.a. secure sketch or helper data) suffices to reproducibly reconstruct a string from noisy measurements, yet leaks only a negligible amount of information about the measurement or the extracted key. An overview of privacy-preserving biometrics, PUFs and fuzzy extraction is given in [18].

1.2 Lower bound on the extractable information

It is, of course, important to know which key length is safe to extract from a source. ‘Safe’ here means that the key’s distribution is sufficiently close to uniform.

Let X denote the noisy measurement, and Y the helper data sufficient to reduce the errors to a tolerable level. The quantity $\ell_{\text{ext}}^\varepsilon(X|Y)$ is defined as the maximum achievable key length that can be extracted from X , given that the adversary knows Y , such that the distribution of the key is ε -close to uniform. (This statement is made more precise in Section 2.) Similarly, a quantity $\ell_{\text{ext}}^\varepsilon(X)$ is defined in the case of a noise-free source X . In [14] Renner and Wolf provided a lower bound on $\ell_{\text{ext}}^\varepsilon(X|Y)$ in terms of the *smooth min-entropy* of X given Y . Loosely speaking, ‘smoothing’ of a probability distribution means probing a small neighborhood of it; the smooth entropy is the entropy of a point in this neighborhood. (This is made more precise in Section 2.)

The technique of smoothing has been used before in many papers (e.g. [8]) under a variety of names. In [1], a random variable is said to have “ ε -HILL-type pseudoentropy at least k ” if there exists a random variable X' with min-entropy at least k such that the (computational) distance between X and X' is at most ε . However, Renner and Wolf actually defined the smooth min-entropy of X as a uniquely determined quantity (by taking the

¹In this setting there is no additional source of randomness to aid the privacy amplification.

maximum value for k), which can then be studied in its own right. The lower bound given in [14] is

$$\ell_{\text{ext}}^\varepsilon(X|Y) \geq H_\infty^\rho(X|Y) - 2 \log \frac{1}{\varepsilon - \rho}. \quad (1)$$

Here $\rho \in [0, \varepsilon)$ is the ‘radius’ of the smoothening neighborhood, and H_∞^ρ is the smooth min-entropy. This bound is sharper than previous results, which do not have the smoothening degree of freedom.

Eq. (1) should be read as a maximization of the right-hand side of (1) over the smoothening parameter ρ . A nonzero ρ increases the term $H_\infty^\rho(X|Y)$, but gives a penalty in the logarithmic term. The optimal ρ depends on the probability distribution of X, Y .

It is important to note that this bound is not just of theoretical interest. It guarantees that well-known generic privacy amplification methods (universal hash functions) yield a key that is ε -close to uniformity as long as the key length does not exceed the r.h.s. of (1). This is relevant even when a tailor-made privacy amplification scheme should give a better result; in practice it may be difficult to find such a scheme or to implement it on a constrained device, and one has to resort to hash functions. Hence, a sharp lower bound on $\ell_{\text{ext}}^\varepsilon$ has practical implications.

While the result (1) is very useful, some unsatisfactory issues remain. Unfortunately, at least three different definitions of smooth entropy appear in the literature [14, 15, 9, 13, 10], and it is not immediately clear if they are equivalent. While it seems that they can be used interchangeably as far as chain rules and similar inequalities are concerned, the consequences for bounds such as (1) are less evident.

Furthermore, the definition of smoothening employed in [14] involves distributions that are not normalized. While that is not a problem in itself, it can lead to technical difficulties when smooth entropies are applied in a context that requires normalization.

1.3 Contributions and outline

In this paper we address the issue of the multiple definitions of smooth entropy, and we provide lower bounds on extractable randomness that are tighter than previously stated in the literature.

We introduce notation and give a list of definitions in Section 2.1. We then present a number of useful lemmas in Section 2.2. Two of these are novel, concerning the leftover hash lemma in the case of un-normalized distributions.

Our main results are stated in Section 3. We first prove a sharper lower bound on the extractable randomness in the unconditional case (Theorem 1). The improvement has two causes:

- a. We use Rényi entropy of order 2 (H_2) instead of min-entropy. In this way the bound is written in terms of the quantity that naturally appears in the leftover hash lemma, without losses from further inequalities. The use of H_2 is not new in itself, nor is the smooth variant H_2^ρ . However, bounds involving smooth entropy have always been formulated in terms of min-entropy.
- b. We make use of the leftover hash lemma for un-normalized distributions. This has the effect of reducing the penalty term $-\log(\varepsilon - \rho)^2$ to $-\log \varepsilon(\varepsilon - \rho)$.

We illustrate the improvement in graphs. For low-entropy sources, the relative improvement in the amount of extracted random bits is significant.

We prove a bound (Theorem 2) for the conditional case that is sharper than (1). It is formulated in terms of the *smooth average conditional Rényi entropy of order two*. The average conditional Rényi entropy (\tilde{H}_2) is the natural quantity appearing in the leftover hash lemma.² The improvement that we obtain here is due only to this, without the above-mentioned effect (b).

We compare three of the existing definitions of smooth entropy regarding their impact on provable bounds on randomness extraction. It turns out that they are not all equivalent in this respect. The two definitions with un-normalized distributions give identical results. However, the ‘normalized’ version of smoothening yields less sharp bounds. The underlying reason is that the normalization constraint effectively doubles the statistical distance between a distribution and its smoothened version, as compared to un-normalized smoothening.

The results above are based on the existence of *universal* families of hash functions. In Section 4 we study the case of *almost universal* families of hash functions [16]. These do not achieve the same amount of privacy amplification, but require significantly less nonvolatile storage to implement and are therefore often preferable in practice. We show that their use has a negligible effect on the extractable randomness.

2 Preliminaries

In Section 2.1 we first introduce the notation that will be used in the rest of the paper. Then we list a number of definitions in order to specify precisely which notions of distance, entropies, smoothening and extractability we are talking about.

Throughout this paper we make an effort to use notation which explicitly shows which variables are operated on by a mapping. In particular, the notation always makes it evident that an entropy is a function of a distribution (as opposed to a random variable). For instance, for a random variable X with distribution \mathbb{P} , we always write $H(\mathbb{P})$ for the entropy instead of the more usual notation $H(X)$. This will sometimes lead to awkward superscripts, but it is a price worth paying for the improved clarity. Clarity is especially needed when smooth entropies are handled. A smoothened distribution is not necessarily normalized; this makes it difficult to write statements like “variable X has distribution \mathbb{Q} ”. In order to avoid such awkwardness we concentrate on statements about distributions. We formulate all the definitions in such a way that they apply to normalized as well as un-normalized distributions.

In Section 2.2 we list some lemmas that are necessary to prove our results.

2.1 Notation and definitions

We denote stochastic variables by capitals, specific values by lowercase letters, and spaces by calligraphic symbols. Distributions are written in mathbold font. E.g. $X \in \mathcal{X}$, $X \sim \mathbb{P}$, $\text{Prob}[X = x] = \mathbb{P}(x)$. The notation $\mathcal{P}_{\mathcal{X}}$ denotes the space of non-negative functions on \mathcal{X} .

As a distance measure between distributions we use the usual definition of statistical distance, generalized to arbitrary non-negative functions.

²It was already pointed out in [6] that worst-case conditioning in the conditional min-entropy $H_{\infty}(X|Y)$ can be replaced by average-case conditioning.

Definition 1 (Statistical distance) Let $\mathbb{P}, \mathbb{Q} \in \mathcal{P}_{\mathcal{X}}$. The statistical distance between \mathbb{P} and \mathbb{Q} is defined as

$$\Delta(\mathbb{P}, \mathbb{Q}) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathbb{P}(x) - \mathbb{Q}(x)|.$$

For $(X, Y) \sim \mathbb{P}$, the marginal distribution of X is denoted as \mathbb{P}_1 , with $\text{Prob}[X = x] = \mathbb{P}_1(x) = \sum_y \mathbb{P}(x, y)$. Similarly, $\text{Prob}[Y = y] = \mathbb{P}_2(y) = \sum_x \mathbb{P}(x, y)$. The conditional probability $\text{Prob}[X = x|Y = y]$ is denoted as $\mathbb{P}_{1|2}(x, y) = \mathbb{P}(x, y)/\mathbb{P}_2(y)$.³ The advantage of this notation is that marginals and conditioning can be easily written down for un-normalized distributions.

Definition 2 (Rényi entropy) Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X}}$. For $\alpha \in (0, 1) \cup (1, \infty)$, the Rényi entropy of \mathbb{P} is defined as

$$H_{\alpha}(\mathbb{P}) = -\frac{1}{\alpha - 1} \log \sum_{x \in \mathcal{X}} [\mathbb{P}(x)]^{\alpha}.$$

For $0 \leq \alpha \leq \beta$ it holds that $H_{\beta}(\mathbb{P}) \leq H_{\alpha}(\mathbb{P})$.

Definition 3 (Min-entropy) Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X}}$. The min-entropy of \mathbb{P} is defined as

$$H_{\infty}(\mathbb{P}) = -\log \max_{x \in \mathcal{X}} \mathbb{P}(x).$$

Definition 4 (Conditional Rényi entropy) Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X} \times \mathcal{Y}}$. Let $\alpha \in (0, 1) \cup (1, \infty)$ be a constant. The conditional Rényi entropy is defined as

$$H_{\alpha(1|2)}(\mathbb{P}) = -\frac{1}{\alpha - 1} \log \max_{y \in \text{supp } \mathbb{P}_2} \sum_{x \in \mathcal{X}} [\mathbb{P}_{1|2}(x, y)]^{\alpha}.$$

Definition 5 (Conditional min-entropy) Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X} \times \mathcal{Y}}$. The conditional min-entropy of \mathbb{P} is defined as

$$H_{\infty(1|2)}(\mathbb{P}) = -\log \max_{x \in \mathcal{X}} \max_{y \in \text{supp } \mathbb{P}_2} \mathbb{P}_{1|2}(x, y).$$

For $(X, Y) \sim \mathbb{P}$ the conditional min-entropy is usually denoted as $H_{\infty}(X|Y)$.

Definition 6 (Average conditional min-entropy) (From [6]).

Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X} \times \mathcal{Y}}$. The average conditional min-entropy of \mathbb{P} , with conditioning on the second argument, is defined as

$$\tilde{H}_{\infty(1|2)}(\mathbb{P}) = -\log \sum_{y \in \mathcal{Y}} \mathbb{P}_2(y) \max_{x \in \mathcal{X}} \mathbb{P}_{1|2}(x, y) = -\log \sum_{y \in \mathcal{Y}} \max_{x \in \mathcal{X}} \mathbb{P}(x, y).$$

The usual notation is $\tilde{H}_{\infty}(X|Y)$ for $(X, Y) \sim \mathbb{P}$.

Definition 7 (Average conditional Rényi entropy) Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$, and let $(X, Y) \sim \mathbb{P}$. Let $\alpha \in (0, 1) \cup (1, \infty)$ be a constant. The average conditional Rényi entropy of X given Y is defined as

$$\tilde{H}_{\alpha(1|2)}(\mathbb{P}) = -\frac{1}{\alpha - 1} \log \sum_{y \in \mathcal{Y}} \mathbb{P}_2(y) \sum_{x \in \mathcal{X}} [\mathbb{P}_{1|2}(x, y)]^{\alpha}.$$

³We use the convention $\mathbb{P}_{1|2}(x, y) = 0$ when $\mathbb{P}_2(y) = 0$.

The following two definitions are somewhat awkward, but they allow us to generalize stochastic variables like $f(X)$ to the case of un-normalized distributions.

Definition 8 (Induced action on a distribution) *Let $f : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ be a function, with the first argument drawn uniformly random from \mathcal{R} . Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X}}$. We define the induced action of f on \mathbb{P} as a mapping $f^* : \mathcal{P}_{\mathcal{X}} \rightarrow \mathcal{P}_{\mathcal{R} \times \{0, 1\}^\ell}$, with*

$$(f^*\mathbb{P})(r, t) = \text{Prob}[R = r, f(R, X) = t] = \frac{1}{|\mathcal{R}|} \sum_{x \in \mathcal{X}: f(r, x) = t} \mathbb{P}(x).$$

Definition 9 (Induced action on a joint distribution) *Let $F : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ be a function, with the first argument drawn uniformly random from \mathcal{R} . Let $\mathbb{P} \in \mathcal{P}_{\mathcal{X} \times \mathcal{Y}}$. We define the induced action of F on \mathbb{P} as a mapping $F^* : \mathcal{P}_{\mathcal{X} \times \mathcal{Y}} \rightarrow \mathcal{P}_{\mathcal{Y} \times \mathcal{R} \times \{0, 1\}^\ell}$, with*

$$(F^*\mathbb{P})(y, r, t) = \text{Prob}[Y = y, R = r, F(R, X) = t] = \frac{1}{|\mathcal{R}|} \mathbb{P}_2(y) \sum_{x \in \mathcal{X}: F(r, x) = t} \mathbb{P}_{1|2}(x, y).$$

Next we show three different notions of smooth entropy that appear in the literature. They differ in the sort of modifications to the probability distribution that are allowed. In order to distinguish between them we introduce the terms ‘strictly bounded’, ‘loosely bounded’ and ‘normalized’.

Definition 10 (Strictly bounded vicinity of a distribution) *(From [10])*

Let \mathbb{P} be a probability measure on \mathcal{X} . Let $\rho \geq 0$. We define the strictly bounded ρ -vicinity of \mathbb{P} as

$$B_{\text{strict}}^\rho(\mathbb{P}) = \left\{ \mathbb{Q} \in \mathcal{P}_{\mathcal{X}} : \forall x \in \mathcal{X} \mathbb{Q}(x) \leq \mathbb{P}(x) \text{ and } \sum_{x \in \mathcal{X}} \mathbb{Q}(x) \geq 1 - \rho \right\}.$$

Note that $\mathbb{Q} \in B^\rho(\mathbb{P})$ with $\mathbb{Q} \neq \mathbb{P}$ is not a probability measure.

Definition 11 (Normalized vicinity of a distribution) *Let \mathbb{P} be a probability measure on \mathcal{X} . Let $\rho \geq 0$. The normalized ρ -vicinity of \mathbb{P} is defined as*

$$B_{\text{norm}}^\rho(\mathbb{P}) = \left\{ \mathbb{Q} \in \mathcal{P}_{\mathcal{X}} : \Delta(\mathbb{Q}, \mathbb{P}) \leq \rho \text{ and } \sum_{x \in \mathcal{X}} \mathbb{Q}(x) = 1 \right\}.$$

Definition 12 (Loosely bounded vicinity of a distribution) *Let \mathbb{P} be a probability measure on \mathcal{X} . Let $\rho \geq 0$. The loosely bounded ρ -vicinity of \mathbb{P} is defined as*

$$B_{\text{loose}}^\rho(\mathbb{P}) = \left\{ \mathbb{Q} \in \mathcal{P}_{\mathcal{X}} : \Delta(\mathbb{Q}, \mathbb{P}) \leq \frac{\rho}{2} \text{ and } \sum_{x \in \mathcal{X}} \mathbb{Q}(x) \leq 1 \right\}.$$

The strictly bounded version appeared in [14] (unnumbered definition). The loosely bounded version appeared in [13] (Definition 3.2.1). The normalized version was used in [9]. (Yet another definition features in [15], but we will not discuss it here.)

For all three types of vicinity, the smooth entropy is defined analogously. In the definitions below, the subscript ‘type’ denotes the type of vicinity. The result (1) from [14] was formulated in terms of the ‘strictly bounded’ vicinity.

Definition 13 (Smooth Rényi entropy) Let \mathbb{P} be a probability measure on \mathcal{X} . Let $\rho \geq 0$ be a constant. The smooth Rényi entropy of \mathbb{P} is defined as

$$H_{\alpha}^{\rho, \text{type}}(\mathbb{P}) = \max_{\mathbb{Q} \in B_{\text{type}}^{\rho}(\mathbb{P})} H_{\alpha}(\mathbb{Q}).$$

Definition 14 (Smooth conditional Rényi entropy) Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$. Let $\rho \geq 0$ and $\alpha \in (0, 1) \cup (1, \infty)$ be constants. The ρ -smooth conditional Rényi entropy of order α of \mathbb{P} is defined as

$$H_{\alpha(1|2)}^{\rho, \text{type}}(\mathbb{P}) = \max_{\mathbb{Q} \in B_{\text{type}}^{\rho}(\mathbb{P})} -\frac{1}{\alpha - 1} \log \max_{y \in \text{supp } \mathbb{P}_2} \sum_{x \in \mathcal{X}} \left[\frac{\mathbb{Q}(x, y)}{\mathbb{P}_2(y)} \right]^{\alpha}.$$

Definition 15 (Smooth conditional min-entropy) Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$. Let $\rho \geq 0$ be a constant. The smooth conditional min-entropy is defined as

$$H_{\infty(1|2)}^{\rho, \text{type}}(\mathbb{P}) = \max_{\mathbb{Q} \in B_{\text{type}}^{\rho}(\mathbb{P})} -\log \max_{y \in \text{supp } \mathbb{P}_2} \max_{x \in \mathcal{X}} \frac{\mathbb{Q}(x, y)}{\mathbb{P}_2(y)}.$$

The usual notation is $H_{\infty}^{\rho}(X|Y)$ for $(X, Y) \sim \mathbb{P}$. Note that the smooth entropies in Definitions 14 and 15 cannot be written as a maximization $\max_{\mathbb{Q}}$ of some expression in \mathbb{Q} only. It will turn out that we need such a property later on; hence we introduce the following definition.

Definition 16 (Smooth average conditional Rényi entropy) Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$. Let $\rho \geq 0$ and $\alpha \in (0, 1) \cup (1, \infty)$ be constants. The ρ -smooth average conditional Rényi entropy of order α of \mathbb{P} is defined as

$$\tilde{H}_{\alpha(1|2)}^{\rho, \text{type}}(\mathbb{P}) = \max_{\mathbb{Q} \in B_{\text{type}}^{\rho}(\mathbb{P})} \tilde{H}_{\alpha(1|2)}(\mathbb{Q}).$$

The smooth average conditional min-entropy follows by setting $\alpha = \infty$. Next we give the formal definition of extractable randomness. This quantity plays a central role in Section 3.

Definition 17 (Extractable randomness) (Modification of Def.3 in [14]).

Let $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ be random variables, with $(X, Y) \sim \mathbb{P}$. For any $\varepsilon > 0$ we say that a finite set \mathcal{Z} is ε -allowed if there exists a function $F : \mathcal{X} \rightarrow \mathcal{Z}$ such that

$$\Delta(YF(X), YZ) \leq \varepsilon,$$

where Z is a random variable uniformly distributed on \mathcal{Z} , independent of X and Y . The ε -extractable randomness of X conditioned on Y is defined as

$$\ell_{\text{ext}}^{\varepsilon}(X|Y) = \max \{ \log |\mathcal{Z}| : \mathcal{Z} \text{ is } \varepsilon\text{-allowed} \}.$$

Note that the definition in [14] contains an extra random variable. The definition of extractable randomness in the unconditional case follows from Def. 17 by omitting Y .

Lower bounds on the extractable randomness are proven using *universal hash functions*.

Definition 18 (Universal family of hash functions) Let \mathcal{R} , \mathcal{X} and \mathcal{T} be finite sets. Let $\{\Phi_r\}_{r \in \mathcal{R}}$ be a family of hash functions from \mathcal{X} to \mathcal{T} . The family $\{\Phi_r\}_{r \in \mathcal{R}}$ is called universal iff, for R drawn uniformly from \mathcal{R} , it holds that

$$\text{Prob}[\Phi_R(x) = \Phi_R(x')] \leq 1/|\mathcal{T}|$$

for all $x, x' \in \mathcal{X}$ with $x' \neq x$.

Following Shoup [16] we give the definition of a more general class of hash functions.

Definition 19 (Almost universal family of hash functions) Let $\eta \geq 0$ be a constant. Let \mathcal{R} , \mathcal{X} and \mathcal{T} be finite sets. Let $\{\Phi_r\}_{r \in \mathcal{R}}$ be a family of hash functions from \mathcal{X} to \mathcal{T} . The family $\{\Phi_r\}_{r \in \mathcal{R}}$ is called η -almost universal iff, for R drawn uniformly from \mathcal{R} , it holds that

$$\text{Prob}[\Phi_R(x) = \Phi_R(x')] \leq \eta$$

for all $x, x' \in \mathcal{X}$ with $x' \neq x$.

Note that a $1/|\mathcal{T}|$ -almost universal family of hash functions is universal. The above defined classes of hash functions play a central role in this paper: They are used in Lemmas 4 and 5 in Section 2.2, on which our main results are based.

2.2 Lemmas

Lemma 1 (Jensen's inequality for concave functions) Let φ be a real concave function. Let n be a positive integer. Let a_1, \dots, a_n be positive weights and x_1, \dots, x_n be real numbers. Then

$$\varphi\left(\frac{\sum_{i=1}^n a_i x_i}{\sum_{i=1}^n a_i}\right) \geq \frac{\sum_{i=1}^n a_i \varphi(x_i)}{\sum_{i=1}^n a_i}.$$

Lemma 2 (Exercise 8.36 in [16]) Let q_1, \dots, q_m be real numbers satisfying $\sum_{s=1}^m q_s = 1$. Then it holds that

$$\sum_{s=1}^m q_s^2 \geq \frac{1}{m}.$$

Proof: $0 \leq \sum_s (q_s - \frac{1}{m})^2 = \sum_s q_s^2 - \frac{2}{m} \sum_s q_s + \sum_s 1/m^2 = \sum_s q_s^2 - \frac{1}{m}$. \square

Lemma 3 (Generalization of Theorem 8.36 in [16] to un-normalized distributions.) Let $\mathbb{Q} \in \mathcal{P}_{\mathcal{Z}}$ such that $\sum_{z \in \mathcal{Z}} \mathbb{Q}(z) = 1 - \alpha$. Then the statistical distance between \mathbb{Q} and the uniform distribution U can be bounded as

$$\Delta(\mathbb{Q}, U) \leq \frac{1}{2} \sqrt{|\mathcal{Z}| \sum_{z \in \mathcal{Z}} \mathbb{Q}^2(z) - 1 + 2\alpha}. \quad (2)$$

Proof: Define $D = \Delta(\mathbb{Q}, U)$ and $q_z := |\mathbb{Q}(z) - \frac{1}{|\mathcal{Z}|}|/2D$ satisfying $\sum_z q_z = 1$. Application of Lemma 2 yields $\frac{1}{|\mathcal{Z}|} \leq \frac{1}{4D^2} [\sum_z \mathbb{Q}^2(z) - \frac{2}{|\mathcal{Z}|} \sum_z \mathbb{Q}(z) + \sum_z |\mathcal{Z}|^{-2}] = \frac{1}{4D^2} [\sum_z \mathbb{Q}^2(z) - \frac{2}{|\mathcal{Z}|} (1 - \alpha) + \frac{1}{|\mathcal{Z}|}]$. Rearranging the inequality gives (2). \square

Lemma 4 (Leftover hash lemma in the unconditional case) (Generalization of Theorem 8.37 in [16] to un-normalized distributions.)

Let $\alpha \in [0, 1)$ be a constant. Let \mathcal{X} be a finite set. Let $\mathbb{Q} \in \mathcal{P}_{\mathcal{X}}$ satisfy $\sum_{x \in \mathcal{X}} \mathbb{Q}(x) = 1 - \alpha$. Let $f : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ be an η -almost universal hash function (see Def. 19), with $\eta = (1 + \delta)2^{-\ell}$, and with r drawn uniformly random from \mathcal{R} . Let $f^*\mathbb{Q}$ be the induced action of f on \mathbb{Q} , as given in Def. 8. Then

$$\Delta(f^*\mathbb{Q}, U_{\mathcal{R}} \otimes U_\ell) \leq \frac{1}{2} \sqrt{\alpha^2 + \delta + 2^{\ell - H_2(\mathbb{Q})}}, \quad (3)$$

where $U_{\mathcal{R}}$ stands for the uniform distribution on \mathcal{R} and U_ℓ stands for the uniform distribution on $\{0, 1\}^\ell$.

A proof is given in Appendix A.

Lemma 5 (Leftover hash lemma in the conditional case) (Generalization of exercise 8.71 in [16] to un-normalized distributions.)

Let $\beta \in [0, 1)$ be a constant. Let $\mathbb{Q} \in \mathcal{P}_{\mathcal{X} \times \mathcal{Y}}$ such that $\sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \mathbb{Q}(x, y) = 1 - \beta$. Let $F : \mathcal{R} \times \mathcal{X} \rightarrow \{0, 1\}^\ell$ be an η -almost universal hash function, with $\eta = (1 + \delta)2^{-\ell}$, and with r drawn randomly from \mathcal{R} . Let $F^*\mathbb{Q}$ be the induced action of F on \mathbb{Q} as given in Def. 9. Then

$$\Delta(F^*\mathbb{Q}, \mathbb{Q}_2 \otimes U_{\mathcal{R}} \otimes U_\ell) \leq \frac{1}{2} \sqrt{1 - \beta} \sqrt{\delta + 2^{\ell - \tilde{H}_{2(1|2)}(\mathbb{Q})}}.$$

We provide a proof in Appendix B. Remark: Lemmas 4 and 5 can also be formulated with H_∞ instead of H_2 , since $H_2(\mathbb{Q}) \geq H_\infty(\mathbb{Q})$.

Lemma 6 Let \mathbb{P} be a probability measure on \mathcal{X} . Let $\rho \geq 0$. Then for any $\mathbb{Q} \in B_{\text{strict}}^\rho(\mathbb{P})$

$$\Delta(\mathbb{P}, \mathbb{Q}) \leq \frac{1}{2}\rho.$$

Proof: $\Delta(\mathbb{P}, \mathbb{Q}) = \frac{1}{2} \sum_{x \in \mathcal{X}} |\mathbb{P}(x) - \mathbb{Q}(x)| = \frac{1}{2} [1 - \sum_{x \in \mathcal{X}} \mathbb{Q}(x)] \leq \frac{1}{2}\rho.$ \square

Lemma 7 Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$. Let $\rho \geq 0$. Then for any $\mathbb{Q} \in B_{\text{strict}}^\rho(\mathbb{P})$

$$\Delta(\mathbb{P}_2, \mathbb{Q}_2) \leq \frac{1}{2}\rho.$$

Proof: $\Delta(\mathbb{P}_2, \mathbb{Q}_2) = \frac{1}{2} \sum_y |\sum_x \mathbb{P}(x, y) - \sum_x \mathbb{Q}(x, y)| = \frac{1}{2} \{1 - \sum_{xy} \mathbb{Q}(x, y)\} \leq \frac{1}{2}\rho.$ \square

Lemma 8 Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$. Let ρ be a constant. Then it holds that

$$\tilde{H}_{2(1|2)}^{\rho, \text{type}}(\mathbb{P}) \geq H_{\infty(1|2)}^{\rho, \text{type}}(\mathbb{P}).$$

Proof: We use that $\tilde{H}_{2(1|2)}^{\rho, \text{type}}(\mathbb{P}) \geq \tilde{H}_{\infty(1|2)}^{\rho, \text{type}}(\mathbb{P})$. Moreover,

$$\begin{aligned} \tilde{H}_{\infty(1|2)}^{\rho, \text{type}}(\mathbb{P}) &= \max_{\mathbb{Q} \in B_{\text{type}}^\rho(\mathbb{P})} \tilde{H}_{\infty(1|2)}(\mathbb{Q}) = \max_{\mathbb{Q} \in B_{\text{type}}^\rho(\mathbb{P})} -\log \sum_{y \in \mathcal{Y}} \mathbb{Q}_2(y) \max_{x \in \mathcal{X}} \mathbb{Q}_{1|2}(x, y) \\ &= \max_{\mathbb{Q} \in B_{\text{type}}^\rho(\mathbb{P})} -\log \sum_{y \in \mathcal{Y}} \max_{x \in \mathcal{X}} \mathbb{Q}(x, y) \\ &= \max_{\mathbb{Q} \in B_{\text{type}}^\rho(\mathbb{P})} -\log \sum_{y \in \text{supp } \mathbb{P}_2} \mathbb{P}_2(y) \max_{x \in \mathcal{X}} \frac{\mathbb{Q}(x, y)}{\mathbb{P}_2(y)} \\ &\geq \max_{\mathbb{Q} \in B_{\text{type}}^\rho(\mathbb{P})} -\log \max_{y \in \text{supp } \mathbb{P}_2} \max_{x \in \mathcal{X}} \frac{\mathbb{Q}(x, y)}{\mathbb{P}_2(y)} = H_{\infty(1|2)}^{\rho, \text{type}}(\mathbb{P}). \quad \square \end{aligned}$$

3 Main result

We prove two theorems regarding the lower bound on extractable randomness (see Def. 17). Both provide an improved bound in terms of the ‘strictly bounded’ type of smooth entropy. The first theorem addresses the unconditional case. The second theorem addresses the conditional case. The proofs are based on the existence of *universal* families of hash functions (see Def. 18). The more general (and more easy to implement) case of *almost universal* hash functions is treated in Section 4.

Theorem 1 *Let $X \sim \mathbb{P}$ be a random variable. Let $\varepsilon > 0$ be a constant. Then*

$$\ell_{\text{ext}}^\varepsilon(X) \geq \max_{\rho \in [0, \varepsilon]} \left[\mathsf{H}_2^{\rho, \text{strict}}(\mathbb{P}) + 2 - \log \frac{1}{\varepsilon(\varepsilon - \rho)} \right].$$

Theorem 2 *Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$, and $(X, Y) \sim \mathbb{P}$. Let $\varepsilon > 0$. Then*

$$\ell_{\text{ext}}^\varepsilon(X|Y) \geq \max_{\rho \in [0, \varepsilon]} \left[\tilde{\mathsf{H}}_{2(1|2)}^{\rho, \text{strict}}(\mathbb{P}) + 2 - 2 \log \frac{\sqrt{1 - \rho}}{\varepsilon - \rho} \right].$$

Proof of Theorem 1: Consider $\hat{\mathbb{Q}} \in B_{\text{strict}}^\rho(\mathbb{P})$ achieving $\mathsf{H}_2(\hat{\mathbb{Q}}) = \mathsf{H}_2^{\rho, \text{strict}}(\mathbb{P})$. We explicitly use the fact that $\sum_{x \in \mathcal{X}} \hat{\mathbb{Q}}(x) = 1 - \rho$. Consider a strong extractor $f : \mathcal{X} \times \mathcal{R} \rightarrow \{0, 1\}^\ell$, with $\ell = \mathsf{H}_2^{\rho, \text{strict}}(X) + 2 - \log \frac{1}{\varepsilon(\varepsilon - \rho)}$ and distance from uniformity $\varepsilon - \rho/2$. The existence of such an extractor is guaranteed by Lemma 4 (taking $\delta = 0$ since we consider universal hash functions, and substituting $\alpha = \rho$). It is readily verified that substitution of ℓ into (3) yields statistical distance $\varepsilon - \rho/2$ from uniformity.

When f is applied to $\hat{\mathbb{Q}}$, the result is a string of length ℓ , with uniformity $\varepsilon - \rho/2$. What happens when f is applied to \mathbb{P} ? Then the uniformity is bounded as follows,

$$\Delta(f^*\mathbb{P}, U_{\mathcal{R}} \otimes U_\ell) \leq \Delta(f^*\mathbb{P}, f^*\hat{\mathbb{Q}}) + \Delta(f^*\hat{\mathbb{Q}}, U_{\mathcal{R}} \otimes U_\ell) \quad (4)$$

$$\leq \Delta(U_{\mathcal{R}} \otimes \mathbb{P}, U_{\mathcal{R}} \otimes \hat{\mathbb{Q}}) + (\varepsilon - \rho/2) \quad (5)$$

$$\leq \rho/2 + (\varepsilon - \rho/2) = \varepsilon. \quad (6)$$

In (4) we used the triangle inequality. In (5) we used the uniformity property of the extractor f . Finally in (6) we applied Lemma 6. \square

Proof of Theorem 2: Consider $\hat{\mathbb{Q}} \in B_{\text{strict}}^\rho(\mathbb{P})$ achieving $\tilde{\mathsf{H}}_{2(1|2)}(\hat{\mathbb{Q}}) = \tilde{\mathsf{H}}_{2(1|2)}^{\rho, \text{strict}}(\mathbb{P})$. The existence of such a $\hat{\mathbb{Q}}$ follows directly from Def. 16. Next we explicitly use that $\sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \hat{\mathbb{Q}}(x, y) = 1 - \rho$. Consider a strong extractor $F : \mathcal{X} \times \mathcal{R} \rightarrow \{0, 1\}^\ell$, acting on the source $\hat{\mathbb{Q}}$, extracting a string of length $\ell = \tilde{\mathsf{H}}_{2(1|2)}^{\rho, \text{strict}}(\mathbb{P}) + 2 - 2 \log \frac{\sqrt{1 - \rho}}{\varepsilon - \rho}$ that has uniformity $\varepsilon - \rho$. The existence of such a strong extractor F is guaranteed by Lemma 5 (again setting $\delta = 0$, and substituting $\beta = \rho$). What happens when F is applied to \mathbb{P} instead of $\hat{\mathbb{Q}}$? We first apply the triangle inequality,

$$\Delta(F^*\mathbb{P}, \mathbb{P}_2 \otimes U_{\mathcal{R}} \otimes U_\ell) \leq \Delta(F^*\mathbb{P}, F^*\hat{\mathbb{Q}}) + \Delta(F^*\hat{\mathbb{Q}}, \mathbb{P}_2 \otimes U_{\mathcal{R}} \otimes U_\ell). \quad (7)$$

The first term in (7) is bounded by using the fact that conditioning reduces the statistical distance,

$$\Delta(F^*\mathbb{P}, F^*\hat{\mathbb{Q}}) \leq \Delta(U_{\mathcal{R}} \otimes \mathbb{P}, U_{\mathcal{R}} \otimes \hat{\mathbb{Q}}) = \Delta(\mathbb{P}, \hat{\mathbb{Q}}) \leq \frac{1}{2}\rho. \quad (8)$$

The last term in (7) is bounded by using one more triangle inequality.

$$\begin{aligned} \Delta(F^*\hat{\mathbb{Q}}, \mathbb{P}_2 \otimes U_{\mathcal{R}} \otimes U_{\ell}) &\leq \Delta(F^*\hat{\mathbb{Q}}, \hat{\mathbb{Q}}_2 \otimes U_{\mathcal{R}} \otimes U_{\ell}) \\ &\quad + \Delta(\hat{\mathbb{Q}}_2 \otimes U_{\mathcal{R}} \otimes U_{\ell}, \mathbb{P}_2 \otimes U_{\mathcal{R}} \otimes U_{\ell}) \end{aligned} \quad (9)$$

$$\leq (\varepsilon - \rho) + \Delta(\hat{\mathbb{Q}}_2, \mathbb{P}_2) \quad (10)$$

$$\leq \varepsilon - \frac{1}{2}\rho. \quad (11)$$

In (10) we used the uniformity property of F . In (11) we applied Lemma 7. It follows that

$$\Delta(F^*\mathbb{P}, \mathbb{P}_2 \otimes U_{\mathcal{R}} \otimes U_{\ell}) \leq \varepsilon. \quad \square$$

Compared to the proof of Theorem 1 there is one extra ingredient: the extra triangle inequality required to handle the distance between the marginals \mathbb{P}_2 and $\hat{\mathbb{Q}}_2$.

3.1 Extractable randomness in the unconditional case

3.1.1 Comparison of Theorem 1 with prior results

- The smooth Rényi entropy of order 2 ($H_2^{\rho, \text{strict}}$) is larger than the smooth min-entropy ($H_{\infty}^{\rho, \text{strict}}$) that features in (1). Fig. 1 shows for a simple example distribution that the difference can be several bits.
- The smoothening ‘penalty’ paid in the logarithm term is less severe than the $-2\log(\varepsilon - \rho)$ penalty in (1), as is shown in Fig. 2. The difference is $-2\log(\varepsilon - \rho) + \log \varepsilon(\varepsilon - \rho) = -\log(1 - \rho/\varepsilon)$, which can amount to several bits in practice.

For low-entropy sources, the relative improvement due to these two differences can be substantial.

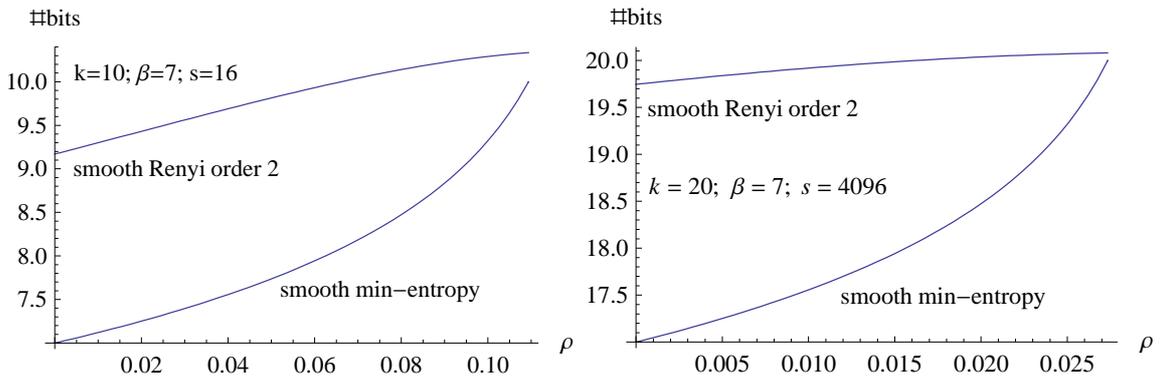


Figure 1: Comparison of $H_{\infty}^{\rho, \text{strict}}(X)$ and $H_2^{\rho, \text{strict}}(X)$ in a toy example: $X \in \{0, 1\}^k$, s elements have probability $p_1 = 2^{-k}(1 + \beta)$, and the remaining $2^k - s$ elements all have probability $p_2 = (1 - sp_1)/(2^k - s)$.

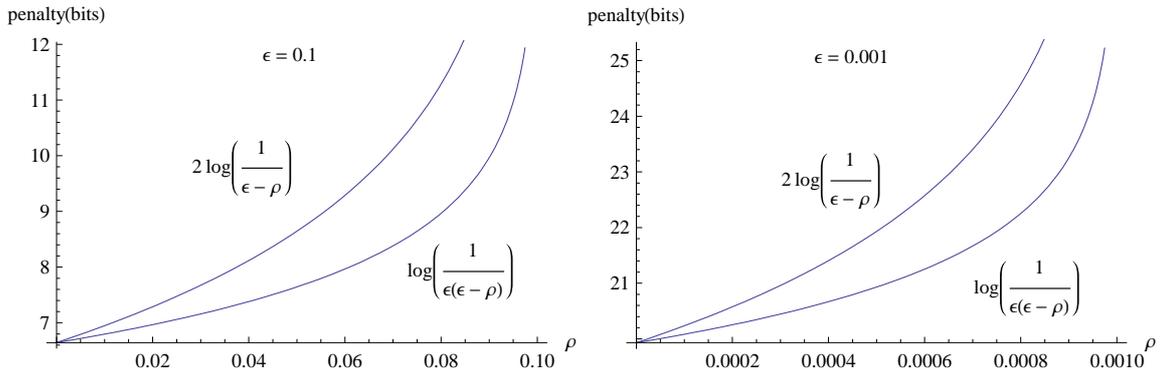


Figure 2: Comparison of the penalty terms in the unconditional case.

3.1.2 Effect of different smoothening definitions

What changes when another version of smooth entropy is considered? The two smoothening-related parts of the proof are the use of the un-normalized version of the leftover hash lemma (Lemma 4), and the final step (6), where the statistical distance between \mathbb{Q} and \mathbb{P} enters. Let us first consider the loosely bounded $H_2^{\rho, \text{loose}}(\mathbb{P})$. Here, just as in the ‘strict’ case, the norm of $\hat{\mathbb{Q}}$ is $1 - \rho$, and $\Delta(\mathbb{P}, \hat{\mathbb{Q}})$ is $\frac{1}{2}\rho$. Hence, in terms of $H^{\rho, \text{loose}}$ we get exactly the same statements as in Theorem 1.

Next we consider the ‘normalized’ $H_2^{\rho, \text{norm}}$. Now we have to use $\alpha = 0$ in Lemma 4 since $\hat{\mathbb{Q}}$ is normalized. Furthermore we have $\Delta(\mathbb{P}, \hat{\mathbb{Q}}) \leq \rho$ from definition 11 as opposed to $\frac{1}{2}\rho$ in the other two cases. Following all the steps of the above proof yields the following result.

$$\ell_{\text{ext}}^{\varepsilon}(X) \geq \max_{\rho \in [0, \varepsilon]} \left[H_2^{\rho, \text{norm}}(\mathbb{P}) + 2 - 2 \log \frac{1}{\varepsilon - \rho} \right]. \quad (12)$$

How does this compare to Theorem 1? For most⁴ distributions \mathbb{P} , the smooth entropies $H_2^{\rho, \text{norm}}(\mathbb{P})$ and $H_2^{\rho, \text{strict}}(\mathbb{P})$ will be roughly the same, since both definitions of smoothening allow one to remove an amount ρ from the highest probabilities $\mathbb{P}(x)$. The logarithmic penalty term of (12), however, is more severe: $\log \frac{1}{(\varepsilon - \rho)^2} \geq \log \frac{1}{\varepsilon(\varepsilon - \rho)}$. Hence Theorem 1 gives a sharper bound.

3.2 Extractable randomness in the conditional case

3.2.1 Comparison of Theorem 2 with prior results

- The expression $\tilde{H}_{2(1|2)}^{\rho, \text{strict}}(\mathbb{P})$ is larger than $H_{\infty(1|2)}^{\rho, \text{strict}}(\mathbb{P})$ (see Lemma 8) for two reasons. (a) Rényi entropy of order 2 is larger than min-entropy (see Fig. 1); (b) taking the average over y is more favorable than taking the worst-case y .
- The factor $\sqrt{1 - \rho}$ gives a slight improvement over (1) by reducing the penalty term. In fact, this factor was already present in an expression for the distance between mixed quantum states in [13] (Corollary 5.6.1 on page 88), but the consequences for extractable

⁴As long as \mathbb{P} is not too close to uniform.

randomness were never explicitly stated. In practice ρ is so small that $\sqrt{1-\rho} \approx 1$, i.e. the improvement is negligible.

3.2.2 Effect of different smoothening definitions

We look what changes occur in the proof of Theorem 2 when the ‘strict’ smoothening is replaced by another type. In the ‘loosely’ bounded version nothing changes, since $\hat{\mathbb{Q}}$ will be the same (i.e. an amount ρ removed from the highest peaks), and $\Delta(\mathbb{P}_2, \hat{\mathbb{Q}}_2)$ as well as $\Delta(\mathbb{P}, \hat{\mathbb{Q}})$ remain unchanged.

In the ‘normalized’ version, Lemma 5 must be used with $\beta = 0$, which is a slight disadvantage. The distances $\Delta(\mathbb{P}_2, \hat{\mathbb{Q}}_2)$ and $\Delta(\mathbb{P}, \hat{\mathbb{Q}})$ are doubled from $\frac{1}{2}\rho$ to ρ , which is a strong disadvantage. This yields the following result.

Theorem 3 *Let \mathbb{P} be a probability measure on $\mathcal{X} \times \mathcal{Y}$, and $(X, Y) \sim \mathbb{P}$. Let $\varepsilon > 0$. Then*

$$\ell_{\text{ext}}^\varepsilon(X|Y) \geq \max_{\rho \in [0, \varepsilon]} \left[\tilde{\mathbb{H}}_{2(1|2)}^{\rho, \text{norm}}(\mathbb{P}) + 2 - 2 \log \frac{1}{\varepsilon - 2\rho} \right].$$

How does this compare to Theorem 2? As in the unconditional case, for most distributions \mathbb{P} it will hold that $\tilde{\mathbb{H}}_{2(1|2)}^{\rho, \text{norm}}(\mathbb{P}) \approx \tilde{\mathbb{H}}_{2(1|2)}^{\rho, \text{strict}}(\mathbb{P})$. The penalty term is larger, since $\frac{1}{\varepsilon - 2\rho} \geq \frac{\sqrt{1-\rho}}{\varepsilon - \rho}$. Hence Theorem 2 gives a sharper bound.

It must be remarked, however, that Eq. 12 and Theorem 3 can be useful in contexts where normalization must be strictly adhered to. For instance, it regularly occurs that a distribution \mathbb{Q} is not exactly known, but that is guaranteed to be at most ρ -distant from some distribution \mathbb{P} with known properties. Bridging the gap from \mathbb{Q} to \mathbb{P} can be done by ‘smoothening’, but only if the normalized version of Def. 11 is applied.

4 Reduction of storage requirements by using almost universal instead of universal hash functions

An important practical consideration for implementing a randomness extractor is the available nonvolatile memory in a device. As mentioned in Section 1.1, the information reconciliation step requires storing the helper data. Furthermore, the privacy amplification step also needs storage space. This may not be immediately obvious from the previous sections. In this step the device needs an implementation of an appropriate hash function. The proofs of Theorems 1 and 2 require the use of a *universal* family of hash functions mapping an input $x \in \mathcal{X}$ to a smaller target space \mathcal{T} . As can be seen in Def. 18 (page 7), these hash functions are labeled by a random variable $R \in \mathcal{R}$; the value of R needs to be stored in the enrolment phase. Both the helper data and the R value are ‘unique’ in the sense that they are different for each instance of a source and for each enrolment. Hence these pieces of data are preferably stored in flash or eeprom memory. In resource-constrained devices this type of memory is expensive. Typical constructions of a universal family of hash functions require that $\log |\mathcal{R}|$ is (almost) as large as $\log |\mathcal{X}|$. For instance, the construction of Example 8.39 in [16] requires

$$\#\text{bits} = \log |\mathcal{R}| = \log |\mathcal{X}| - \log |\mathcal{T}|. \quad (13)$$

For highly non-uniform sources X this is prohibitive.

It is possible to save on nonvolatile memory by relaxing the constraints on the hash function: By allowing small deviations from perfect universality, it becomes possible to have a tradeoff between the quality of the privacy amplification and the space needed to store R . Def. 19 introduces η -almost universal families of hash functions.⁵ There are constructions [16] of such hash functions, with $\eta = (1 + \delta)/|\mathcal{T}|$, requiring only

$$\log |\mathcal{R}| = \mathcal{O} \left(\log |\mathcal{T}| + \log \frac{\log |\mathcal{X}|}{\log |\mathcal{T}|} + \log \frac{1}{\delta} \right). \quad (14)$$

Comparing (14) to (13) we see that the dependence on $|\mathcal{X}|$ has changed from $\log \log |\mathcal{X}|$ to $\log \log |\mathcal{X}|$, which is much smaller. The incurred penalty $\log(1/\delta)$ is manageable.

Example: Consider a case where $\mathcal{X} = \{0, 1\}^{1024}$ and $\mathcal{T} = \{0, 1\}^{128}$. For universality we need $\log |\mathcal{R}| = 896$ bits of storage, while for almost universality with $\delta = 2^{-16}$ we only need approximately 147 bits.

Finally we show that the reduced quality of the privacy amplification due to a finite δ does not lead to significant worsening of the extractable randomness.

Theorem 4 *Let $\varepsilon \geq 0$ be a constant. Let $X \sim \mathbb{P}$ be a random variable on \mathcal{X} . Let $\{\Phi_r\}_{r \in \mathcal{R}}$ be an η -almost universal family of hash functions from \mathcal{X} to $\{0, 1\}^\ell$, with $\eta = (1 + \delta)2^{-\ell}$. Then the amount of randomness that can be extracted from X using this family of hash functions is bounded from below by*

$$\ell_{\text{ext}}^\varepsilon(X) \geq \max_{\rho \in [0, \varepsilon - \delta/4]} \left[\mathbf{H}_2^{\rho, \text{strict}}(\mathbb{P}) + 2 - \log \frac{1}{\varepsilon(\varepsilon - \rho) - \delta/4} \right].$$

Proof: Same as for Theorem 1, but without setting $\delta = 0$.

Theorem 5 *Let $\varepsilon \geq 0$ be a constant. Let $(X, Y) \sim \mathbb{P}$ be random variables on \mathcal{X} and \mathcal{Y} respectively. Let $\{\Phi_r\}_{r \in \mathcal{R}}$ be an η -almost universal family of hash functions from \mathcal{X} to $\{0, 1\}^\ell$, with $\eta = (1 + \delta)2^{-\ell}$. Then the amount of randomness that can be extracted from X , conditioned on Y , using this family of hash functions is bounded from below by*

$$\ell_{\text{ext}}^\varepsilon(X|Y) \geq \max_{\rho \in [0, \varepsilon - \sqrt{\delta}/2]} \left[\tilde{\mathbf{H}}_{2(1|2)}^{\rho, \text{strict}}(\mathbb{P}) + 2 - \log \frac{1}{(\varepsilon - \rho)^2 - \delta/4} \right].$$

Proof: Same as for Theorem 2, but without setting $\delta = 0$.

The effect of nonzero δ on the extractable randomness is negligible as long as $\delta \ll \varepsilon^2$.

5 Summary

We have improved the known lower bound on the extractable entropy from non-uniform sources in the unconditional case (Theorem 1) and in the conditional case (Theorem 2). In the unconditional case, the improvement stems from using \mathbf{H}_2 instead of \mathbf{H}_∞ and from using an un-normalized version of the leftover hash lemma. In the unconditional case, the improvement comes from using the *smooth average conditional Rényi entropy* $\tilde{\mathbf{H}}_{2(1|2)}^{\rho, \text{strict}}$, which is the best match to the entities appearing in the leftover hash lemma.

⁵Exact universality occurs at $\eta = 1/|\mathcal{T}|$.

We have investigated three of the definitions of smoothening with respect to their impact on the extractable entropy bounds. It turns out that the two un-normalized versions are equivalent in this respect, and yield a sharper bound than the normalized version.

We have studied the amount of nonvolatile memory that is needed to store the random label (R) of the hash functions employed in privacy amplification. A significant reduction is known to be achieved ($\log |\mathcal{X}|$ to $\log \log |\mathcal{X}|$) by switching from a universal family of hash functions to an η -almost universal family. We have shown that the penalty paid in the extractable randomness can be made vanishingly small.

References

- [1] B. Barak, R. Shaltiel, and A. Wigderson. Computational analogues of entropy. In *APPROX 2003 + RANDOM 2003*, volume 2764 of *Lecture Notes in Computer Science*, pages 200–215, Berlin, 2003. Springer-Verlag.
- [2] J. D. R. Buchanan, R. P. Cowburn, A. Jausovec, D. Petit, P. Seem, G. Xiong, D. Atkinson, K. Fenton, D.A. Allwood, and M. T. Bryan. Forgery: ‘fingerprinting’ documents and packaging. *Nature, Brief Communications*, 436:475, July 2005.
- [3] D. Clarke, B. Gassend, M. van Dijk, and S. Devadas. Secure hardware processors using silicon physical one-way functions. In R. Sandu, editor, *ACM CCS '02*, 2002.
- [4] G. DeJean and D. Kirovski. Radio frequency certificates of authenticity. In *IEEE Antenna and Propagation Symposium – URSI*, 2006.
- [5] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. Cryptology ePrint Archive, Report 2003/235, 2003.
- [6] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM Journal on Computing*, 38(1):97–139, 2008.
- [7] Jorge Guajardo, Sandeep S. Kumar, Geert Jan Schrijen, and Pim Tuyls. FPGA intrinsic PUFs and their use for IP protection. In Pascal Paillier and Ingrid Verbauwhede, editors, *CHES*, volume 4727 of *Lecture Notes in Computer Science*, pages 63–80. Springer, 2007.
- [8] J. Hastad, R. Impagliazzo, L.A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [9] Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In Victor Shoup, editor, *Advances in Cryptology - CRYPTO '05*, Lecture Notes in Computer Science, pages 478–493. Springer-Verlag, August 2005.
- [10] Thomas Holenstein and Renato Renner. On the randomness of independent experiments, August 2006. available at <http://arxiv.org/abs/cs.IT/0608007>.
- [11] Jean-Paul M. G. Linnartz and Pim Tuyls. New shielding functions to enhance privacy and prevent misuse of biometric templates. In J. Kittler and M. Nixon, editors, *Conference*

on *Audio and Video Based Person Authentication*, volume 2688 of *LNCS*, pages 238–250. Springer-Verlag, 2003.

- [12] R. Pappu, B. Recht, J. Taylor, and N. Gershenfeld. Physical one-way functions. *Science*, 297:2026–2030, Sept. 2002.
- [13] Renato Renner. *Security of Quantum Key Distribution*. PhD thesis, Swiss Federal Institute of Technology (ETH) Zurich, September 2005. available at <http://arxiv.org/abs/quant-ph/0512258>.
- [14] Renato Renner and Stefan Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In Bimal Roy, editor, *Advances in Cryptology - ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 199–216. Springer-Verlag, December 2005.
- [15] Renato Renner, Stefan Wolf, and Jürg Wullschleger. The single-serving channel capacity. In *Proceedings of the International Symposium on Information Theory (ISIT)*. IEEE, July 2006. available at <http://arxiv.org/abs/cs.IT/0608018>.
- [16] Victor Shoup. *A Computational Introduction to Number Theory and Algebra*. Cambridge University Press, 2008. version 2; available at <http://www.shoup.net/ntb/ntb-v2.pdf>.
- [17] P. Tuyls, G.J. Schrijen, B. Škorić, J. van Geloven, R. Verhaegh, and R. Wolters. Read-proof hardware from protective coatings. In L. Goubin and M. Matsui, editors, *Cryptographic Hardware and Embedded Systems - CHES 2006*, volume 4249 of *LNCS*, pages 369–383. Springer-Verlag, 2006.
- [18] P. Tuyls, B. Škorić, and T. Kevenaar. *Security with Noisy Data: Private Biometrics, Secure Key Storage and Anti-Counterfeiting*. Springer, 2007.
- [19] B. Škorić, T. Bel, A.H.M. Blom, B.R. de Jong, H. Kretschman, and A.J.M. Nellissen. Randomized resonators as uniquely identifiable anti-counterfeiting tags, March 2008. Secure Component and System Identification (SECSI) Workshop.

A Proof of Lemma 4

This proof follows the approach in [16], but for an un-normalized distribution \mathbb{Q} . We use Lemma 3, replacing \mathbb{Q} in the lemma by $f^*\mathbb{Q}$ and \mathcal{Z} by $\mathcal{R} \times \{0, 1\}^\ell$. We have to compute the collision probability c ,

$$c = \sum_{r \in \mathcal{R}} \sum_{t \in \{0, 1\}^\ell} [(f^*\mathbb{Q})(r, t)]^2 = \sum_{r \in \mathcal{R}} \sum_{t \in \{0, 1\}^\ell} \left[\frac{1}{|\mathcal{R}|} \sum_{x \in \mathcal{X}: f(r, x) = t} \mathbb{Q}(x) \right]^2. \quad (15)$$

We introduce the notation $\langle A \rangle$, for some predicate A , to denote 1 if A is true and 0 if A is false. This allows us to write

$$\begin{aligned} c &= |\mathcal{R}|^{-2} \sum_{x, x' \in \mathcal{X}} \mathbb{Q}(x)\mathbb{Q}(x') \sum_{r \in \mathcal{R}} \sum_{t \in \{0, 1\}^\ell} \langle f(r, x) = t \rangle \langle f(r, x') = t \rangle \\ &= \frac{1}{|\mathcal{R}|} \sum_{x, x' \in \mathcal{X}} \mathbb{Q}(x)\mathbb{Q}(x') \text{Prob} [f(R, x) = f(R, x')]. \end{aligned} \quad (16)$$

In the last expression the probability is with respect to the random variable R , while x and x' are fixed. Next we split the x, x' summation into a part where $x' = x$ and a part where $x' \neq x$, and we apply the defining property of the almost universal hash function f .

$$\begin{aligned}
c &= \frac{1}{|\mathcal{R}|} \left(\sum_{x \in \mathcal{X}} \mathbb{Q}^2(x) + \sum_{x, x' \in \mathcal{X}: x' \neq x} \mathbb{Q}(x)\mathbb{Q}(x') \text{Prob} [f(R, x) = f(R, x')] \right) \\
&\leq \frac{1}{|\mathcal{R}|} \left(\sum_{x \in \mathcal{X}} \mathbb{Q}^2(x) + \sum_{x, x' \in \mathcal{X}: x' \neq x} \mathbb{Q}(x)\mathbb{Q}(x') \frac{1 + \delta}{2^\ell} \right) \\
&= \frac{1}{|\mathcal{R} \times \{0, 1\}^\ell|} \left(2^\ell \sum_{x \in \mathcal{X}} \mathbb{Q}^2(x) + (1 + \delta) \left[(1 - \alpha)^2 - \sum_{x \in \mathcal{X}} \mathbb{Q}^2(x) \right] \right) \\
&\leq \frac{1}{|\mathcal{R} \times \{0, 1\}^\ell|} \left(2^\ell \sum_{x \in \mathcal{X}} \mathbb{Q}^2(x) + (1 - \alpha)^2 + \delta \right). \tag{17}
\end{aligned}$$

Substitution of (17) into Lemma 3 yields Lemma 4. \square

B Proof of Lemma 5

We begin by introducing the notation $\mathbb{Q}_{1|y} \in \mathcal{P}_{\mathcal{X}}$, where $\mathbb{Q}_{1|y}(x) := \mathbb{Q}_{1|2}(x, y)$ and $\sum_{x \in \mathcal{X}} \mathbb{Q}_{1|y}(x) = 1$ for all y . The statistical distance in (5) can be expressed as

$$\Delta(F^* \mathbb{Q}, \mathbb{Q}_2 \otimes U_{\mathcal{R}} \otimes U_\ell) = \sum_{y \in \mathcal{Y}} \mathbb{Q}_2(y) \Delta(F^* \mathbb{Q}_{1|y}, U_{\mathcal{R}} \otimes U_\ell). \tag{18}$$

We apply Lemma 4 (with $\alpha = 0$, since $\mathbb{Q}_{1|y}$ is normalized) to get

$$\Delta(F^* \mathbb{Q}_{1|y}, U_{\mathcal{R}} \otimes U_\ell) \leq \frac{1}{2} \sqrt{\delta + 2^{\ell - \text{H}_2(\mathbb{Q}_{1|y})}}. \tag{19}$$

We substitute this into (18) and then apply Jensen's inequality for concave functions (Lemma 1). This yields

$$\begin{aligned}
\Delta(F^* \mathbb{Q}, \mathbb{Q}_2 \otimes U_{\mathcal{R}} \otimes U_\ell) &\leq \frac{1}{2} \sqrt{\sum_{y \in \mathcal{Y}} \mathbb{Q}_2(y)} \sqrt{\sum_{y \in \mathcal{Y}} \mathbb{Q}_2(y) [\delta + 2^{\ell - \text{H}_2(\mathbb{Q}_{1|y})}]} \\
&= \frac{1}{2} \sqrt{1 - \beta} \sqrt{\delta(1 - \beta) + 2^{\ell - \tilde{\text{H}}_2(1|2)(\mathbb{Q})}}. \tag{20}
\end{aligned}$$

Lemma 5 follows. \square