

When are Fuzzy Extractors Possible?*

Benjamin Fuller[†] Leonid Reyzin[‡] Adam Smith[§]

October 3, 2019

Abstract

Fuzzy extractors (Dodis et al., SIAM J. Computing 2008) convert repeated noisy readings of a high-entropy secret into the same uniformly distributed key. A minimum condition for the security of the key is the hardness of guessing a value that is similar to the secret, because the fuzzy extractor converts such a guess to the key. We codify this property in a new notion called *fuzzy min-entropy*. We ask: is fuzzy min-entropy sufficient to build fuzzy extractors? We provide two answers for different settings.

1. If the algorithms have precise knowledge of the probability distribution W that defines the noisy source is a sufficient condition for information-theoretic key extraction from W .
2. A more ambitious goal is to design a single extractor that works for all possible sources. This more ambitious goal is impossible: there is a family of sources with high fuzzy min-entropy for which no single fuzzy extractor is secure. This is true in three settings:
 - (a) for standard fuzzy extractors,
 - (b) for fuzzy extractors that are allowed to sometimes be wrong,
 - (c) and for secure sketches, which are the main ingredient of most fuzzy extractor constructions.

Keywords Fuzzy extractors, secure sketches, authentication, error-tolerance, key derivation, error-correcting codes, entropy

1 Introduction

Sources of reproducible secret random bits enable cryptographic applications. Often, these bits are implicit, they are obtained by repeating the same process (such as reading a biometric or a physically unclonable function [PRTG02]) that generated them the first time. However, bits obtained this way are noisy [BBR88, BS00, Dau04, EHMS00, GCVDD02, MG09, MRW02, PRTG02, SD07, TSS⁺06, ZH93]. When a secret is read multiple times readings are close (according to some metric) but not identical. To utilize such sources, it is often necessary to remove noise, in order to derive the same value in subsequent readings.

The same problem occurs in the interactive setting, in which the secret channel used for transmitting the bits between two users is noisy and/or leaky [Wyn75]. Bennett, Brassard, and Robert [BBR88] identify two fundamental tasks:

*This article is a major revision of a version published by Springer-Verlag and appearing at Asiacrypt 2016. DOI: 10.1007/978-3-662-53887-6_10

[†]Email: benjamin.fuller@uconn.edu. University of Connecticut. This work was performed in part while at Boston University and MIT Lincoln Laboratory.

[‡]Email: reyzin@cs.bu.edu. Boston University.

[§]Email: ads22@bu.edu. Boston University. Work done in part while at Pennsylvania State University.

1. Information reconciliation removes the noise with minimal leakage to an eavesdropping adversary.
2. Privacy amplification converts the high entropy secret to a uniform random value.

In this work, we consider the noninteractive version of these problems, in which these tasks are performed together with a single message.

The noninteractive setting is modeled by a primitive called a fuzzy extractor [DORS08], which consists of two algorithms. The generate algorithm (**Gen**) takes an initial reading w and produces an output key along with a nonsecret helper value p . The reproduce (**Rep**) algorithm takes the subsequent reading w' along with the helper value p to reproduce key. The correctness guarantee is that the key is reproduced precisely when the distance between w and w' is at most t .

The security requirement for fuzzy extractors is that key is uniform even to a (computationally unbounded) adversary who has observed p . This requirement is harder to satisfy as the allowed error tolerance t increases, because it becomes easier for the adversary to guess key by guessing a w' within distance t of w and running $\text{Rep}(w', p)$. This attack is enabled by the functionality of the fuzzy extractor.

Fuzzy Min-Entropy We introduce a new entropy notion that precisely measures how hard it is for the adversary to execute this attack. It measures the probability of guessing a value within distance t of the original reading w . Suppose w is sampled from a distribution W . To have the maximum chance that w' is within distance t of w , the adversary would choose the point w' that maximizes the total probability mass of W within the ball $B_t(w')$ of radius t around w' . We therefore define *fuzzy min-entropy*

$$H_{t,\infty}^{\text{fuzz}}(W) \stackrel{\text{def}}{=} -\log \max_{w'} \Pr[W \in B_t(w')].$$

A fuzzy extractor's key cannot be longer than the fuzzy min-entropy (Proposition 3.2).

However, existing constructions do not measure their security in terms of fuzzy min-entropy; instead, their security is shown to be the min-entropy of W , denoted $H_\infty(W)$, minus some loss for error-tolerance. This loss is due to an error-correction component that writes down enough information to determine which point within distance t was the original point. Since there are as many as $|B_t|$ points within radius t it takes $\log |B_t|$ bits to specify which one.¹ This value of $\log |B_t|$ is assumed to be the security loss so the residual security is $H_\infty(W) - \log |B_t|$. It is easy to show that $H_\infty(W) - \log |B_t| \leq H_{t,\infty}^{\text{fuzz}}(W)$, so it is natural to ask whether a loss of $\log |B_t|$ is necessary. This question is particularly relevant when the gap between the two sides of the inequality is high.²

As an example, in the biometric regime, entropy is estimated by comparing the distribution of distances between two different biometrics with a distribution with well understood statistical properties. For example, Daugman [Dau04] shows that the distance between two irises is distributed similarly to a binomial distribution with 249 bits of entropy. The iris is then assumed to have the same entropy as the binomial distribution. The parameter t is determined by the experimental conditions. Daugman recommends setting $t > 205$. In the Hamming metric for strings of length 2048, $|B_{205}| \approx 2^{1024}$ (see Lemma 2.2). Thus, iris scans have

$$H_\infty(W) - \log |B_t| \approx 249 - 1024 < 0,$$

¹We omit w in the notation $|B_t|$ since we study metrics where the volume of the ball $B_t(w)$ does not depend on the center w .

²For nearly uniform distributions, $H_{t,\infty}^{\text{fuzz}}(W) \approx H_\infty(W) - \log |B_t|$. In this setting, standard coding based constructions of fuzzy extractors (using almost optimal codes) yield keys of size approximately $H_{t,\infty}^{\text{fuzz}}(W)$.

Work	Metric	Dist.	Residual Entropy	
This Work/[WCD ⁺ 17]	Any discrete	Generic	$H_{t,\infty}^{\text{fuzz}}(W)$	$-2 \log 1/\epsilon - \log 1/\delta + 1$
[JW99]	Hamming	Uniform	$H_{t,\infty}^{\text{fuzz}}(W)$	$-2 \log 1/\epsilon - 2$
[LT03]/[TG04]	Any with almost perfect codes	i.i.d	$\approx H_1(W) - H_1(W W')$	$-o(n)$
[DORS08]	Hamming/Set Difference/Edit	Generic	$H_\infty(W) - \log B_t $	$-2 \log 1/\epsilon - 2$

Table 1: Comparison of Fuzzy Extractor Constructions. Here, ϵ is the statistical distance from the uniform distribution, δ is the allowed error, and t is the desired error tolerance. Note that $H_{t,\infty}^{\text{fuzz}}(W) \geq H_\infty(W) - \log |B_t|$. Here n is the dimension of the metric space. H_1 represents Shannon entropy and W' is the distribution that adds noise to the enrollment value W .

(see discussion in [BH09, Section 5]). However, iris scans for different people appear to be well-spread in the metric space [Dau06], the closest observed distance between two different irises is 548. This indicates the distribution of W has $H_{t,\infty}^{\text{fuzz}}(W)$.

The current state of fuzzy extractors is unsettling. For many distributions W with min-entropy, we have no known construction and no known impossibility result. We hope the more precise notion of fuzzy min-entropy can rectify this situation. Ideally, one could show a fuzzy extractor exists for every distribution with *enough* fuzzy min-entropy and that they are impossible with less fuzzy min-entropy. That is, we ask: *is fuzzy min-entropy sufficient for fuzzy extraction?* There is evidence that it may be sufficient when the security requirement is computational rather than information-theoretic—see Section 1.2. We provide an answer for the case of information-theoretic security in two settings.

Contribution 1: Sufficiency of $H_{t,\infty}^{\text{fuzz}}(W)$ for a Precisely Known W We first consider the case when the fuzzy extractor designer has *precise knowledge* of the probability distribution function of W . In this setting, it is possible to construct a fuzzy extractor that extracts a key almost as long as $H_{t,\infty}^{\text{fuzz}}(W)$ (Theorem 4.3).³ Our construction crucially utilizes the probability distribution function of W and, in particular, cannot necessarily be realized in polynomial time (this is similar, for example, to the interactive information-reconciliation feasibility result of [RW04]). This result shows that $H_{t,\infty}^{\text{fuzz}}(W)$ is a necessary and sufficient condition for building a fuzzy extractor for a given distribution W .

A number of previous works in the precise knowledge setting described tight bounds for specific distributions (for example, [JW99, LT03, TG04, HAD06, WRDI11, IW12]). These works are summarized in Table 1. Our characterization unifies previous work, and justifies using $H_{t,\infty}^{\text{fuzz}}(W)$ as the measure of the quality of a noisy distribution, rather than cruder measures such as $H_\infty(W) - \log |B_t|$. Our construction can be viewed as a reference to evaluate the quality of efficient constructions in the precise knowledge setting by seeing how close they get to extracting all of $H_{t,\infty}^{\text{fuzz}}(W)$.

Many works consider i.i.d sources, for example [TG04]. These works are able to derive a key qualitatively longer than fuzzy min-entropy. This is because one can characterize i.i.d sources in terms of Shannon entropy, denoted $H_1(\cdot)$, instead of min-entropy. Each symbol of W is a separate “draw” from the distribution enabling average case analysis as the dimension n of the metric space increases. Key length for i.i.d. sources asymptotically approaches $H_1(W) - H_1(W|W')$ where W' is the distribution of noisy readings around W [TG04, Theorem 2]. It is difficult to directly compare with these works as they do not specify concrete losses for a fixed source length and they consider Shannon entropy. Qualitatively, Shannon entropy can be arbitrarily higher than min-entropy. Unfortunately, many biometrics and hardware sources are not i.i.d. (see for example [Dau04]), so this analysis should be used judiciously.

Contribution 2: The Cost of Distributional Uncertainty Assuming precise knowledge of a distribution W is often unrealistic for high-entropy distributions; they can never be fully observed directly

³Woodage et al. present an improved version of this theorem [WCD⁺17]. We discuss their work below.

and must therefore be modeled. It is unrealistic to assume that the designer’s model of a distribution is as accurate as the adversary’s model. The adversary may have more resources including time to build a model when the construction is deployed. Existing fuzzy extractors are shown secure for a family of sources (for example, all sources of min-entropy at least m with at most t errors). The attacker may know more about the distribution than the designer. We call this the *distributional uncertainty* setting.

Our second contribution is a set of negative results for the distributional uncertainty setting. We provide two impossibility results for fuzzy extractors. Both demonstrate families \mathcal{W} of distributions over $\{0, 1\}^n$ such that each distribution in the family has $H_{t,\infty}^{\text{fuzz}}$ linear in n , but no fuzzy extractor can be secure for most distributions in \mathcal{W} . A fuzzy extractor designer who knows only that the distribution comes from \mathcal{W} cannot secure the family, despite the fact that fuzzy extractors can be designed for each distribution in the family individually.

The first impossibility result (Theorem 5.1) assumes that Rep is perfectly correct and rules out our fuzzy extractors for entropy rates, defined as $\mu = H_{t,\infty}^{\text{fuzz}}(W)/n$, as high as $\mu \approx 0.18$. The second impossibility result (Theorem 6.1), relying on the work of Holenstein and Renner [HR05], also rules out fuzzy extractors in which Rep is allowed to make a mistake, but applies only to distributions with entropy rates up to $\mu \approx 0.07$.

We also provide a third impossibility result (Theorem 7.2), this time for an important building block called “secure sketch.” A secure sketch is a one-round information-reconciliation component (that recovers the original w from the input w'). Secure sketches are used in most fuzzy extractor constructions. The result rules out secure sketches for a family of distributions with entropy rate up to $\mu = 0.5$, even if the secure sketches are allowed to make mistakes. We define secure sketches formally in Section 7. Most fuzzy extractor constructions are analyzed for all families with a certain amount of entropy. Thus, showing impossibility for higher entropy rates raises the lower bound on how much fuzzy min-entropy must be present in the physical distribution for security to be based on just fuzzy min-entropy. As discussed in Section 8, another alternative is to assume additional structure about the physical source.

1.1 Our Techniques

Techniques for Positive Results for a Precisely Known Distribution We now provide intuition for our positive result for a precisely known distribution W with fuzzy min-entropy. We begin with distributions in which all points in the support have the same probability (so-called “flat” distributions). Gen extracts a key from the input w using a randomness extractor [NZ93]. Consider some subsequent reading w' . To achieve correctness, the string p must permit Rep to disambiguate which point $w \in W$ within distance t of w' was given to Gen. Disambiguating multiple points can be accomplished by universal hashing, as long as the size of hash output space is slightly greater than the number of possible points. Thus, Rep includes into the public value p a “sketch” of w computed via a universal hash of w . To determine the length of that sketch, consider the heaviest (according to W) ball B^* of radius t . Because the distribution is flat, B^* is also the ball with the most points of nonzero probability. Thus, the length of the sketch needs to be slightly greater than the logarithm of the number of non-zero probability points in B^* . Since $H_{t,\infty}^{\text{fuzz}}(W)$ is determined by the weight of B^* , the number of points cannot be too high and there will be entropy left after the sketch is published. This remaining entropy suffices to extract a key.

For an arbitrary distribution, we cannot afford to disambiguate points in the ball with the greatest number of points, because there could be too many low-probability points in a single ball despite a high $H_{t,\infty}^{\text{fuzz}}(W)$. We solve this problem by splitting the arbitrary distribution into a number of nearly flat distributions we call “levels.” We then write down, as part of the sketch, the level of the original reading w and apply the above construction considering only points in that level. We call this construction *leveled*

hashing (Construction 4.2).

Techniques for Negative Results for Distributional Uncertainty We construct a family of distributions \mathcal{W} and prove impossibility for a uniformly random $W \leftarrow \mathcal{W}$. We start by observing the following asymmetry: **Gen** sees only the sample w (obtained via $W \leftarrow \mathcal{W}$ and $w \leftarrow W$), while the adversary knows W .

To exploit the asymmetry, in our first impossibility result (Theorem 5.1), we construct \mathcal{W} so that conditioning on the knowledge of W reduces the distribution to a small subspace (namely, all points on which a given universal hash function produces a given output), but conditioning on *only* w leaves the rest of the distribution uniform on a large fraction of the entire space. An adversary can exploit the knowledge of the hash value to reduce the uncertainty about key, as follows.

The nonsecret value p partitions the metric space into regions that produce a consistent value under **Rep** (preimages of each key under $\text{Rep}(\cdot, p)$). For each of these regions, the adversary knows that possible w lie at distance at least t from the boundary of the region (else, the fuzzy extractor would have a nonzero probability of error). However, in the Hamming space, the vast majority of points lie near the boundary (this result follows by combining the isoperimetric inequality [Har66], which shows that the ball has the smallest boundary, with bounds on the volume of the interior of a ball, which show that this boundary is large). This allows the adversary to rule out so many possible w that, combined with the adversarial knowledge of the hash value, many regions become empty, leaving key far from uniform.

For the second impossibility result (Theorem 6.1, which rules out even fuzzy extractors that are allowed a possibility of error), we let the adversary know some fraction of the bits of w . Holenstein and Renner [HR05] showed that if the adversary knows each bit of w with sufficient probability, and bits of w' differ from bits of w with sufficient probability, then so-called information-theoretic key agreement is impossible. Converting the impossibility of information-theoretic key agreement to impossibility of fuzzy extractors takes a bit of technical work.

1.2 Related Settings

Other settings with close readings: $H_{t,\infty}^{\text{fuzz}}$ is sufficient The security definition of fuzzy extractors can be weakened to protect only against computationally bounded adversaries [FMR13]. In this computational setting, under the assumption of semantically secure graded encoding, for most distance metrics a single fuzzy extractor can simultaneously secure all possible distributions [BCKP14, BCKP17]. This construction is secure when the adversary can rarely learn key with oracle access to the program functionality. The set of distributions with fuzzy min-entropy are exactly those where an adversary learns key with oracle access to the functionality with negligible probability. Bitansky et al.'s [BCKP17] construction requires heavy weight and disputed cryptographic tools similar to those used to construct indistinguishability obfuscation [GGH⁺13, GGH⁺16]. Their result implies that extending our negative result to the computational setting would have negative implications on the existence of certain types of obfuscation.

Furthermore, the functional definition of fuzzy extractors can be weakened to permit interaction between the party having w and the party having w' . Such a weakening is useful for secure remote authentication [BDK⁺05]. When both interaction and computational assumptions are allowed, secure two-party computation can produce a good key whenever the distribution W has fuzzy min-entropy. The two-party computation protocol needs to be secure without assuming authenticated channels; it can be built under the assumptions of collision-resistant hash functions and enhanced trapdoor permutations [BCL⁺11] or oblivious transfer and a variant of the random oracle model [DHP⁺18].

Correlated rather than close readings A different model for the problem of key derivation from noisy sources does not explicitly consider the distance between w and w' , but rather views w and w' as

realizations of a correlated pair of random variables (W, W') . This model is considered in multiple works, including [Wyn75, CK78, AC93, Mau93]; recent characterizations of when key derivation is possible in this model include [RW05] and [TW15].

Much of the work on correlated pairs considers interactive protocols (as opposed to the noninteractive agreement needed for fuzzy extractors). However, the impossibility results for that setting are directly relevant to our work, because ruling out interactive protocols also rules out noninteractive ones. Recall that our starting point is the observation that the fuzzy extractor’s output length (Proposition 3.2) is at most the fuzzy min-entropy of W . Prior work of Tyagi and Watanabe [TW15], and concurrent (with ours) work of those authors with Viswanath [TVW18], developed general upper bounds on the achievable secret key length for correlated readings (via “conditional independence testing”). We can apply these in our setting by taking W' to be a random string within distance t of W to obtain impossibility results analogous to Proposition 3.2. Our technique is less general, since it is tailored to the constraint on the distance between w and w' ; however, this specificity allows us to give a simple, direct proof.

The conditional independence testing framework also applies to what we call the “unknown distribution” case (dubbed “correlated information at the eavesdropper” in the correlated readings literature). Given a joint distribution on W, W', Z , where Z is held by the eavesdropping adversary, the conditional testing framework gives an upper bound on the achievable key length (and hence an impossibility result when that bound is small). The framework does not show how one can construct distributions for which this bound is small. In particular, it is an open question whether one can derive versions of our impossibility results for unknown distributions (Theorems 5.1 and 6.1) using the framework. Even using the specific constructions of W, W', Z that arise in our proofs, it is open whether the conditional testing framework provides a good bound on the key length.

1.3 Concurrent and subsequent work

A construction that is very similar to our positive result in the known distribution setting (Construction 4.2) was used independently in a concurrent work of Hayashi, Tyagi, and Watanabe [HTW14], who used the term “spectrum slicing” to describe it. They also extended this technique to the case of distributional uncertainty, using it in an interactive protocol, with one side telling the other to keep increasing the length of the sketch until Rep could succeed. This interactive approach was used in subsequent works, as well (e.g., [TW17, TVW18, LA18]).

Li and Anantharam [LA18] consider correlated readings in the known distribution setting. They show that maximum expected key length for interactive protocols that are allowed to output variable-lengths keys is closely related to the mutual information between W and W' .

Woodage et al. [WCD⁺17] showed a clever extension to our leveled hashing construction (Construction 4.2). They observed the level information does not have to be explicitly included as part of the sketch. The construction uses leveled hashing but with two important changes:

1. The level is not written. Denote by h_ℓ the hash with the greatest number of output bits. If the used hash, h_i , has $|h_i| < h_\ell$ then the output is extended with random bits to length $|h_\ell|$. Instead of looking for an exact match the Rep algorithm finds the close point that matches the stored string at the longest prefix. This can be seen as considering all possible levels of the original hash.
2. In the our construction, the hash output is determined by how many points are in the neighborhood of points with that probability. This may lead to some levels with short hash outputs. This is a problem for Woodage et al.’s construction, if there are multiple short levels, the longest prefix may be the wrong level with noticeable probability. To address this problem, Woodage et al. use the

(negative log of the) probability of a point to compute the length of the hash output. This ensures that all levels have length of at least the min-entropy of the distribution making collisions unlikely. Importantly, this change requires a change in the security argument, essentially arguing that all sketches are equally likely regardless of the starting level. This change requires an augmentation to the hash function called strong universality [CW79].

Fuller and Peng [FP19] extended our negative results to sources that are drawn from continuous metric spaces equipped with the Euclidean metric. There are two main differences between the Euclidean space and our setting:

1. Fewer points lie near the boundary of a ball in Euclidean space.
2. The use of continuous spaces requires volume techniques. So rather than showing that the hash value leaves few possibilities for w , they show that the “volume” of distributions is larger than the interior of parts. Thus, any choice for the interior of parts must not contain a fraction of distributions.

These changes necessitate the use of a different family that is derived from all cosets of random lattices with sufficient minimum distance (known as construction A).

2 Preliminaries

Random Variables We use uppercase letters for random variables and corresponding lowercase letters for their samples. A repeated occurrence of the same random variable signifies the same value of the random variable: for example $(W, \text{SS}(W))$ is a pair of random variables obtained by sampling w according to W and applying the algorithm SS to w . The *statistical distance* between random variables A and B with the same domain is

$$\text{SD}(A, B) = \frac{1}{2} \sum_a |\Pr[A = a] - \Pr[B = a]| = \max_S \Pr[A \in S] - \Pr[B \in S].$$

Entropy Let \log denote the base 2 logarithm. Let (X, Y) be a pair of random variables. Define *min-entropy* of X as $H_\infty(X) = -\log(\max_x \Pr[X = x])$, and the *average (conditional) min-entropy* [DORS08, Section 2.4] of X given Y as

$$\tilde{H}_\infty(X|Y) = -\log \left(\mathbb{E}_{y \leftarrow Y} \max_x \Pr[X = x | Y = y] \right).$$

Define Hartley entropy $H_0(X)$ to be the logarithm of the size of the support of X , that is $H_0(X) = \log |\{x | \Pr[X = x] > 0\}|$. Define average-case Hartley entropy by averaging the support size:

$$\tilde{H}_0(X|Y) = \log \left(\mathbb{E}_{y \leftarrow Y} |\{x | \Pr[X = x | Y = y] > 0\}| \right).$$

For $0 < a < 1$, the binary entropy is $h_2(p) = -p \log p - (1 - p) \log(1 - p)$, which corresponds to the Shannon entropy of any random variable that is 0 with probability p and 1 with probability $1 - p$.

Randomness Extractors We use randomness extractors [NZ93], as defined for the average case in [DORS08, Section 2.5].

Definition 2.1. Let \mathcal{M}, χ be finite sets. A function $\text{ext} : \mathcal{M} \times \{0, 1\}^d \rightarrow \{0, 1\}^\kappa$ a (\tilde{m}, ϵ) -average case extractor if for all pairs of random variables X, Y over \mathcal{M}, χ such that $\tilde{H}_\infty(X|Y) \geq \tilde{m}$, we have

$$\text{SD}((\text{ext}(X, U_d), U_d, Y), U_\kappa \times U_d \times Y) \leq \epsilon.$$

Metric Spaces and Balls Let \mathcal{M} be some finite space and let the function $\text{dis} : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}^+ \cup \{0\}$ be a distance metric (identity of indiscernibles, symmetric, and triangle inequality). For a metric space $(\mathcal{M}, \text{dis})$, the (closed) ball of radius t around w is the set of all points within radius t , that is, $B_t(w) = \{w' | \text{dis}(w, w') \leq t\}$. We consider the Hamming metric over vectors in \mathcal{Z}^n for some finite alphabet \mathcal{Z} , defined via $\text{dis}(w, w') = |\{i | w_i \neq w'_i\}|$. In this space, the size of a ball in a metric space does not depend on w , so we denote by $|B_t|$ the size of a ball (centered arbitrarily) of radius t . U_κ denotes the uniformly distributed random variable on $\{0, 1\}^\kappa$. We use the bounds on $|B_t|$ in $\{0, 1\}^n$, see [Ash65, Lemma 4.7.2, equation 4.7.5, p. 115] for proofs.

Lemma 2.2. Let $\tau = t/n$. The volume $|B_t|$ of the ball of radius t in the Hamming space $\{0, 1\}^n$ satisfies

$$\frac{1}{\sqrt{8n\tau(1-\tau)}} \cdot 2^{nh_2(\tau)} \leq |B_t| \leq 2^{nh_2(\tau)}.$$

We modify the definition of fuzzy extractors slightly from the work of Dodis et al. [DORS08, Sections 3.2]. First, we allow for error as discussed in [DORS08, Section 8]. Second, in the *distributional uncertainty* setting we consider a general family \mathcal{W} of distributions instead of families containing all distributions of a given min-entropy. Let \mathcal{M} be a metric space with distance function dis .

Definition 2.3. An $(\mathcal{M}, \mathcal{W}, \kappa, t, \epsilon)$ -fuzzy extractor with error δ is a pair of randomized procedures, “generate” (Gen) and “reproduce” (Rep). Gen on input $w \in \mathcal{M}$ outputs an extracted string $\text{key} \in \{0, 1\}^\kappa$ and a helper string $p \in \{0, 1\}^*$. Rep takes $w' \in \mathcal{M}$ and $p \in \{0, 1\}^*$ as inputs. (Gen, Rep) have the following properties:

1. Correctness: if $\text{dis}(w, w') \leq t$ and $(\text{key}, p) \leftarrow \text{Gen}(w)$, then $\Pr[\text{Rep}(w', p) = \text{key}] \geq 1 - \delta$.
2. Security: for any distribution $W \in \mathcal{W}$, if $(\text{Key}, P) \leftarrow \text{Gen}(W)$, then $\text{SD}((\text{Key}, P), (U_\kappa, P)) \leq \epsilon$.

In the above definition, the value of w' must be chosen before p is known in order for the correctness guarantee to hold (alternatively, w' can be sampled from a probability distribution that is independent of p).

The Case of a Precisely Known Distribution If in the above definition we take \mathcal{W} to be a one-element set containing a single distribution W , then the fuzzy extractor is said to be for a *precisely known distribution*. In this case, we need to require correctness only for w that have nonzero probability. We specify no requirement that the algorithms are compact or efficient, and so the distribution can be fully known to them.

3 New Notion: Fuzzy Min-Entropy

The fuzzy extractor helper string p allows everyone, including the adversary, to find the output of $\text{Rep}(\cdot, p)$ on any input w' . Ideally, p should not provide any useful information beyond this ability, and the outputs of Rep on inputs that are too distant from w should provide no useful information, either. In this ideal scenario, the adversary is limited to trying to guess a w' that is t -close to w . We measure the quality of a source by (the negative logarithm of) the success of this attack.

Definition 3.1. The t -fuzzy min-entropy of a distribution W in a metric space $(\mathcal{M}, \text{dis})$ is:

$$H_{t,\infty}^{\text{fuzz}}(W) = -\log \left(\max_{w'} \sum_{w \in \mathcal{M} | \text{dis}(w,w') \leq t} \Pr[W = w] \right)$$

Fuzzy min-entropy measures the functionality provided to the adversary by Rep (since p is public), and thus is a necessary condition for security. We formalize this statement in the following proposition.

Proposition 3.2. Let W be a distribution over $(\mathcal{M}, \text{dis})$ with $H_{t,\infty}^{\text{fuzz}}(W) = m$. Let (Gen, Rep) be a $(\mathcal{M}, \{W\}, \kappa, t, \epsilon)$ -fuzzy extractor with error δ . Then

$$2^{-\kappa} \geq 2^{-m} - \delta - \epsilon.$$

If $\delta = \epsilon = 2^{-\kappa}$, then κ cannot exceed $m + 2$.

Proof. Let W be a distribution where $H_{t,\infty}^{\text{fuzz}}(W) = m$. This means that there exists a point $w' \in \mathcal{M}$ such that $\sum_{w \in \mathcal{M} | \text{dis}(w,w') \leq t} \Pr[W = w] = 2^{-m}$. Consider the following function $D_{w'}$:

- Input (key, p) .
- If $\text{Rep}(w', p) = \text{key}$, output 1.
- Else output 0.

Clearly, $\Pr[D_{w'}(\text{Key}, P) = 1] \geq 2^{-m} - \delta$, while $\Pr[D_{w'}(U_\kappa, P) = 1] = 1/2^\kappa$. Thus,

$$\text{SD}((\text{Key}, P), (U_\kappa, P)) \geq \Pr[D_{w'}(\text{Key}, P) = 1] - \Pr[D_{w'}(U_\kappa, P) = 1] \geq 2^{-m} - \delta - 2^{-\kappa}.$$

□

Proposition 3.2 extends to the settings of computational security and interactive protocols if the definition gives the adversary access to the true Key. We explore properties of fuzzy min-entropy below. These properties are included to demonstrate the utility of fuzzy min-entropy and are not necessary to complete the proofs in this work. Conditioning on an event p of probability $\Pr[P = p]$ decreases fuzzy min-entropy by a factor of at most $\log 1/\Pr[P = p]$.

Lemma 3.3. $H_{t,\infty}^{\text{fuzz}}(W|P = p) \geq H_{t,\infty}^{\text{fuzz}}(W) + \log \Pr[P = p]$.

Proof.

$$\begin{aligned} H_{t,\infty}^{\text{fuzz}}(W|P = p) &= -\log \left(\max_{w'} \sum_{w | \text{dis}(w,w') \leq t} \Pr[W = w | P = p] \right) \\ &= -\log \left(\max_{w'} \sum_{w | \text{dis}(w,w') \leq t} \frac{\Pr[W = w \wedge P = p]}{\Pr[P = p]} \right) \\ &\geq -\log \left(\max_{w'} \sum_{w | \text{dis}(w,w') \leq t} \frac{\Pr[W = w]}{\Pr[P = p]} \right) \\ &= H_{t,\infty}^{\text{fuzz}}(W) + \log \Pr[P = p]. \end{aligned}$$

The second line follows from the first using Bayes rule, the third follows from the second using the monotonicity of probability. The last line follows by factoring $1/\Pr[P = p]$ from the sum, and noting the sum then represents $H_{t,\infty}^{\text{fuzz}}(W)$. □

Conditional fuzzy min-entropy Properly defined, fuzzy min-entropy obeys a chain rule. We start by defining a conditional notion of fuzzy min-entropy for a random variable P .

Definition 3.4. For distributions W, P , the t -conditional fuzzy min-entropy of $W|P$ in a metric space $(\mathcal{M}, \text{dis})$ is:

$$\tilde{H}_{t,\infty}^{\text{fuzz}}(W|P) = -\log \left(\mathbb{E}_{p \leftarrow P} \max_{w'} \sum_{w|\text{dis}(w,w') \leq t} \Pr[W = w|P = p] \right).$$

Then a chain rule analogous to average min-entropy [DORS08, Lemma 2.2b] applies:

Lemma 3.5. $\tilde{H}_{t,\infty}^{\text{fuzz}}(W|P) \geq H_{t,\infty}^{\text{fuzz}}(W) - H_0(P)$.

Proof.

$$\begin{aligned} \tilde{H}_{t,\infty}^{\text{fuzz}}(W|P) &= -\log \left(\mathbb{E}_{p \leftarrow P} \max_{w'} \sum_{w|\text{dis}(w,w') \leq t} \Pr[W = w|P = p] \right) \\ &= -\log \left(\sum_p \max_{w'} \sum_{w|\text{dis}(w,w') \leq t} \Pr[W = w|P = p] \Pr[P = p] \right) \\ &= -\log \left(\sum_p \max_{w'} \sum_{w|\text{dis}(w,w') \leq t} \Pr[W = w \wedge P = p] \right) \\ &\geq -\log \left(\sum_p \max_{w'} \sum_{w|\text{dis}(w,w') \leq t} \Pr[W = w] \right) \\ &\geq -\log \left(2^{H_0(P)} \left(\max_{w'} \sum_{w|\text{dis}(w,w') \leq t} \Pr[W = w] \right) \right) \\ &\geq H_{t,\infty}^{\text{fuzz}}(W) - H_0(P). \end{aligned}$$

Here the second line follows from the first using the definition of expectation. The third follows using Bayes rule. The fourth follows using monotonicity of expectation. By definition, there are at $2^{H_0(P)}$ possibilities for p , yielding the fifth line. The last line results by recognizing $H_{t,\infty}^{\text{fuzz}}(W)$ and converting to entropy. \square

4 $H_{t,\infty}^{\text{fuzz}}(W)$ is Sufficient in the Precise Knowledge Setting

In this section, we build fuzzy extractors that extract almost all of $H_{t,\infty}^{\text{fuzz}}(W)$ for any distribution W . These constructions assume precise knowledge of W and are not efficient. We begin with flat distributions and then turn to arbitrary distributions.

Let $\text{supp}(W) = \{w|\Pr[W = w] > 0\}$ denote the support of a distribution W . A distribution W is *flat* if all elements of $\text{supp}(W)$ have the same probability. Our construction for this case is quite simple: to produce p , Gen outputs a hash of its input point w and an extractor seed; to produce key, Gen applies

the extractor to w . Given w' , Rep looks for $w \in \text{supp}(W)$ that is near w' and has the correct hash value, and applies the extractor to this w to get key.

The specific hash function we use is *universal*. (We note that universal hashing has a long history of use for information reconciliation, for example [BBR88], [RW04], and [ST09]. This construction is not novel; rather, we present it as a stepping stone for the case of general distributions.)

Definition 4.1 ([CW79]). *Let $F : \mathcal{K} \times \mathcal{M} \rightarrow R$ be a function. We say that F is universal if for all distinct $x_1, x_2 \in \mathcal{M}$:*

$$\Pr_{K \leftarrow \mathcal{K}} [F(K, x_1) = F(K, x_2)] = \frac{1}{|R|}.$$

In our case, the hash output length needs to be sufficient to disambiguate elements of $\text{supp}(W) \cap B_t(w')$ with high probability. Observe that there are at most $2^{H_\infty(W) - H_{t,\infty}^{\text{fuzz}}(W)}$ such elements when W is flat, so output length slightly greater (by $\log 1/\delta$) than $H_\infty(W) - H_{t,\infty}^{\text{fuzz}}(W)$ will suffice. Thus, the output key length will be $H_{t,\infty}^{\text{fuzz}}(W) - \log 1/\delta - 2 \log 1/\epsilon + 2$ (by using average-case leftover hash lemma, per [DORS08, Lemma 2.2b, Lemma 2.4]). As this construction is only a warm-up, so we do not state it formally and proceed to general distributions.

4.1 Fuzzy Extractor for Arbitrary Distributions

The above hashing approach does not work for arbitrary sources. Consider a distribution W consisting of the following balls: B_t^1 is a ball with $2^{H_\infty(W)}$ points with total probability $\Pr[W \in B_t^1] = 2^{-H_\infty(W)}$, $B_t^2, \dots, B_t^{2^{-H_\infty(W)}}$ are balls with one point each with probability $\Pr[W \in B_t^i] = 2^{-H_\infty(W)}$. The above hashing algorithm writes down $H_\infty(W)$ bits to achieve correctness on B_t^1 . However, with probability $1 - 2^{-H_\infty(W)}$ the initial reading is outside of B_t^1 , and the hash completely reveals the point.

Instead, we use a layered approach: we separate the input distribution W into nearly-flat layers, write down the layer from which the input w came (i.e., the approximate probability of w) as part of p , and rely on the construction from the previous part for each layer. In other words, the hash function output is now variable-length, longer if probability of w is lower. Thus, p now reveals a bit more about w . To limit this information and the resulting security loss, we limit number of layers. As a result, we lose only $1 + \log H_0(W)$ more bits of security compared to the previous section.

The main idea is that providing the level information makes the distribution look nearly flat (the probability of points differs by at most a factor of two, which increases the entropy loss as compared to the flat case by only one bit). The level information itself increases the entropy loss by $\log H_0(W)$ bits, because there are only $H_0(W)$ levels that contain enough weight to matter. In subsequent work, Woodage et al. show that level information does not have to be leaked [WCD⁺17, Theorem 3]. We now present a formal description of our construction.

Construction 4.2. *Let W be a distribution over a metric space \mathcal{M} with $H_\infty(W) = m$.*

- *Let $\delta \leq \frac{1}{2}$ be the error parameter.*
- *Let $\ell = m + H_0(W) - 1$; round ℓ down so that $\ell - m$ is an integer (i.e., set $\ell = m + \lfloor (\ell - m) \rfloor$).*
- *For each $i = m, m + 1, \dots, \ell - 1$, let $L_i = (2^{-(i+1)}, 2^{-i}]$ and let $F_i : \mathcal{K}_i \times \mathcal{M} \rightarrow R_i$ be a family of universal hash functions with $\log |R_i| = i + 1 - H_{t,\infty}^{\text{fuzz}}(W) + \log 1/\delta$. Let $L_\ell = (0, 2^{-\ell}]$.*
- *Let **ext** be an (\tilde{m}, ϵ) -average-case extractor for $\tilde{m} = H_{t,\infty}^{\text{fuzz}}(W) - \log H_0(W) - \log 1/\delta - 1$ with output length κ .*

Define $\text{Gen}_W, \text{Rep}_W$ as:

Gen_W

1. Input: w .
2. Find i such that $\Pr[W = w] \in L_i$.
3. If $i = \ell$ then set $ss = (i, w, 0)$.
4. Else sample $K \leftarrow \mathcal{K}_i$ and set $ss = (i, F_i(K, w), K)$
5. Sample a uniform extractor seed seed
6. Output $\text{key} = \text{ext}(w, \text{seed})$, $p = (ss, \text{seed})$.

Rep_W

1. Input: $(w', p = (ss, \text{seed}))$
2. Parse ss as (i, y, K)
3. If $i = \ell$ then set $w^* = y$.
4. Else
 - (a) Let $W^* = \{w^* | \text{dis}(w^*, w') \leq t \wedge \Pr[W = w^*] \in L_i\}$.
 - (b) Find any $w^* \in W^*$ such that $F_i(K, w^*) = y$; if none exists, set $w^* = \perp$.
5. Output $\text{ext}(w^*, \text{seed})$.

We instantiate this construction with the extractor parameters given by a universal hash (namely, $\kappa = \tilde{m} - 2 \log 1/\epsilon + 2$):

Theorem 4.3. *For any metric space \mathcal{M} , distribution W over \mathcal{M} , distance t , error $\delta > 0$, and security $\epsilon > 0$, there exists a $(\mathcal{M}, \{W\}, \kappa, t, \epsilon)$ -known distribution fuzzy extractor with error δ for $\kappa = H_{t, \infty}^{\text{fuzz}}(W) - \log H_0(W) - \log 1/\delta - 2 \log 1/\epsilon + 1$.*

Proof of Theorem 4.3. We first argue **correctness**. Fix some w, w' within distance t . When $\Pr[W = w] \in L_\ell$, then Rep is always correct, so let's consider only the case when $\Pr[W = w] \notin L_\ell$. The algorithm Rep will never output \perp since at least the correct w will match the hash. Thus, an error happens when another element $w^* \in W^*$ has the same hash value $F(K_i, w^*)$ as $F(K_i, w)$. Observe that the total probability mass of W^* is greater than $|W^*| \cdot 2^{-(i+1)}$ but less than or equal to the maximum probability mass in a ball of radius t , $2^{-H_{t, \infty}^{\text{fuzz}}(W)}$. Therefore, $|W^*| \leq 2^{i+1 - H_{t, \infty}^{\text{fuzz}}(W)}$. Each element of W^* has the same hash as $F(K, w)$ with probability at most $1/|R_i|$, and thus correctness with error $|W^*|/|R_i| \leq \delta$ follows by the union bound.

Security: We now argue security of the construction. Let $W_i = \{w | \Pr[W = w] \in L_i\}$. For ease of notation, let us make the special case of $i = \ell$ as part of the general case, as follows: define $\mathcal{K}_\ell = \{0\}$, $F_\ell(0, w) = w$, and $R_\ell = W_\ell$. Also, denote by SS the randomized function that maps w to ss . First, we set up the analysis by levels:

$$\begin{aligned}
 2^{-\tilde{H}_\infty(W | \text{SS}(W))} &= \mathbb{E} \max_{ss} \Pr[W = w | \text{SS}(W) = ss] \\
 &= \sum_{ss} \max_w \Pr[W = w \wedge \text{SS}(W) = ss] \\
 &= \sum_{i=m}^{\ell} \sum_{K \in \mathcal{K}_i} \sum_{y \in R_i} \max_w \Pr[W = w \wedge \text{SS}(W) = (i, y, K)] \\
 &\leq \sum_{i=m}^{\ell} \sum_{K \in \mathcal{K}_i} \sum_{y \in R_i} \max_{w \in W_i} \Pr[W = w \wedge F_i(K, w) = y \wedge K \text{ output by Gen}].
 \end{aligned}$$

We now pay the penalty of $|R_i|$ for the presence of y (observe that removing the condition that $F_i(K, w) = y$ from the conjunction cannot reduce the probability):

$$\begin{aligned} 2^{-\tilde{H}_\infty(W|\text{SS}(W))} &\leq \sum_{i=m}^{\ell} \sum_{K \in \mathcal{K}_i} \sum_{y \in R_i} \max_{w \in W_i} \Pr[W = w \wedge K \text{ is chosen by SS}] \\ &= \sum_{i=m}^{\ell} \sum_{K \in \mathcal{K}_i} |R_i| \cdot \max_{w \in W_i} \Pr[W = w \wedge K \text{ is chosen by SS}]. \end{aligned}$$

We now get rid of the key, because it is independent:

$$\begin{aligned} 2^{-\tilde{H}_\infty(W|\text{SS}(W))} &\leq \sum_{i=m}^{\ell} \sum_{K \in \mathcal{K}_i} |R_i| \cdot \max_{w \in W_i} \Pr[W = w] \cdot \frac{1}{|\mathcal{K}_i|} \\ &= \sum_{i=m}^{\ell} |R_i| \cdot \max_{w \in W_i} \Pr[W = w] \leq |R_\ell| \cdot 2^{-\ell} + \sum_{i=m}^{\ell-1} |R_i| \cdot 2^{-i}. \end{aligned}$$

Finally, we add everything up, recalling that $|R_i|$ for $i < \ell$ is $2^{i+1 - H_{t,\infty}^{\text{fuzz}}(W) + \log 1/\delta}$.

$$\begin{aligned} 2^{-\tilde{H}_\infty(W|\text{SS}(W))} &\leq 2^{H_0(W)} \cdot 2^{-\ell} + (\ell - m) \cdot 2^{1 - H_{t,\infty}^{\text{fuzz}}(W) + \log 1/\delta} \\ &\quad (\text{next line uses } \ell > m + H_0(W) - 2) \\ &< 2^{2-m} + (\ell - m) \cdot 2^{1 - H_{t,\infty}^{\text{fuzz}}(W) + \log 1/\delta} \\ &\quad (\text{next line uses } m \geq H_{t,\infty}^{\text{fuzz}}(W) \text{ and } \log 1/\delta \geq 1) \\ &\leq (\ell - m + 1) \cdot 2^{1 - H_{t,\infty}^{\text{fuzz}}(W) + \log 1/\delta} \\ &\quad (\text{next line uses } \ell \leq m + H_0(W) - 1) \\ &\leq H_0(W) \cdot 2^{1 - H_{t,\infty}^{\text{fuzz}}(W) + \log 1/\delta}. \end{aligned}$$

Taking the negative logarithm of both sides, we obtain $\tilde{m} \stackrel{\text{def}}{=} \tilde{H}_\infty(W|\text{SS}(W)) = H_{t,\infty}^{\text{fuzz}}(W) - \log H_0(W) - \log 1/\delta - 1$. Applying the (\tilde{m}, ϵ) randomness extractor gives us the desired result. \square

5 Impossibility of Fuzzy Extractors for a Family with $H_{t,\infty}^{\text{fuzz}}$

In the previous section, we showed the sufficiency of $H_{t,\infty}^{\text{fuzz}}(W)$ for building fuzzy extractors when the distribution W is precisely known. However, it is usually infeasible to characterize a high-entropy distribution W . Traditionally, algorithms deal with this *distributional uncertainty* by providing security for a family of distributions \mathcal{W} . In this section, we show that distributional uncertainty comes at a real cost.

We demonstrate an example over the binary Hamming metric in which every $W \in \mathcal{W}$ has linear $H_{t,\infty}^{\text{fuzz}}(W)$ (which is in fact equal to $H_\infty(W)$), and yet there is some $W \in \mathcal{W}$ where even for 3-bit keys, the key distribution is far from uniform, $\epsilon = \frac{1}{4}$. In fact, we show that the adversary need not work hard: even a uniformly random choice of distribution W from \mathcal{W} will thwart the security of any (Gen, Rep). The one caveat is that, for this result, we require Rep to be always correct (i.e., $\delta = 0$). As mentioned in

the introduction, this perfect correctness requirement is removed in Sections 6 and 7 at a cost of lower entropy rate and stronger primitive, respectively.

The result is based on the following reasoning: **Gen** sees only a random sample w from a random $W \in \mathcal{W}$, but not W . The adversary sees W but not w . Because **Gen** does not know which W the input w came from, **Gen** must produce p that works for many distributions W that contain w in their support. Such p necessarily reveals a lot of information. The adversary can combine information gleaned from p with information about W to narrow down the possible choices for w and thus distinguish key from uniform.

Theorem 5.1. *There exists a family of distributions \mathcal{W} over $\{0, 1\}^n$ equipped with the Hamming metric such that for each element $W \in \mathcal{W}$, $H_{t, \infty}^{\text{fuzz}}(W) = H_{\infty}(W) \geq m$, and yet any $(\{0, 1\}^n, \mathcal{W}, \kappa, t, \epsilon)$ -fuzzy extractor with error $\delta = 0$ has $\epsilon > 1/4$.*

This holds as long as $\kappa \geq 3$ and under the following conditions on the entropy rate $\mu = m/n$, noise rate $\tau = t/n$, and n :

- any $0 \leq \tau < \frac{1}{2}$,
- any $\mu > 0$ such that $\mu < 1 - h_2(\tau)$ and $\mu < 1 - h_2(\frac{1}{2} - \tau)$, and
- any $n \geq \max\left(\frac{2}{1 - h_2(\tau) - \mu}, \frac{5}{1 - h_2(\frac{1}{2} - \tau) - \mu}\right)$.

The conditions on μ and τ imply the result applies to any entropy rate $\mu \leq .18$ as long as τ is set appropriately and n is sufficiently large (for example, the result applies to $n \geq 1275$ and $\tau = .6\sqrt{\mu}$ when $0.08 \leq \mu \leq .18$). The τ vs. μ tradeoff is depicted in Figure 1.

Here we first provide short intuition, followed by the proof. The overall goal of the proof is show a lower bound on the value of ϵ which is the quality of the output key.

- One can partition the input metric space according to what value of key is output by $\text{Rep}(w, p)$.
- The value of p reduces the set of possible w because, by correctness of **Rep**, every candidate input w to **Gen** must be such that all of its neighbors w' of distance at most t produce the same output of $\text{Rep}(w', p)$.
- The isoperimetric inequality then shows for most parts, almost all points are not in the interior (Lemma 5.2).
- The above gives a bound on the residual entropy of w conditioned on p for most values of key. The second part of the proof incorporates the adversary's knowledge of the distribution $W \in \mathcal{W}$.
- We show the theorem holds for an average member of \mathcal{W} . Let Z denote a uniform choice of W from \mathcal{W} and denote by W_z the choice specified by a particular value of z .
- Let $\{\text{Hash}_k\}_{k \in \mathcal{K}}$ be a family of hash functions with domain $\mathcal{M} = \{0, 1\}^n$ that is universal (small collision probability for any two points across the hash key), regular (large preimage size for any output value), and has preimage sets with high minimum distance. Then define $z = (k, h)$ and define $W_{(k, h)}$ as the uniform distribution over the set $\{w \mid \text{Hash}(w, k) = h\}_{(k, h)}$.
- The hash function we use is the output of a parity check matrix for a random code with high distance. Thus, each distribution W_z is a coset of some randomly chosen code C with good distance (that is not known by the construction). This family has the required properties (see Lemma 5.4).

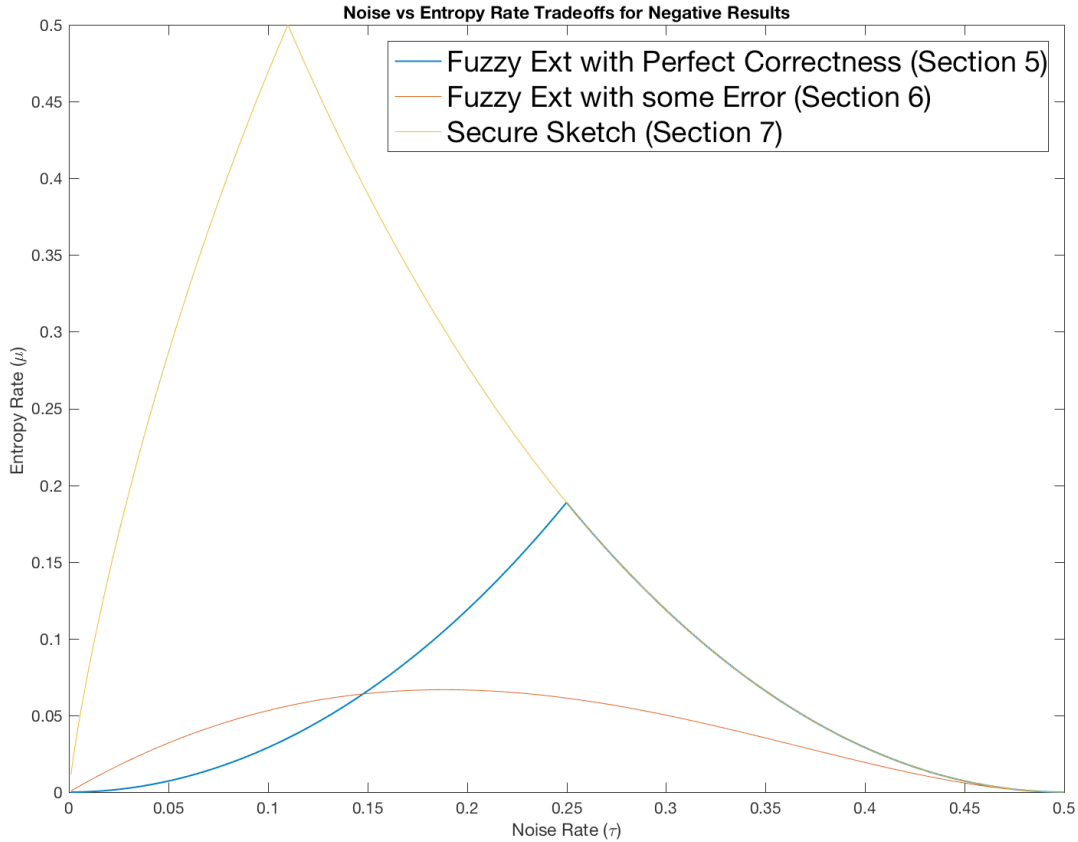


Figure 1: The region of τ (x -axis) and μ (y -axis) pairs for which the three negative results apply (Theorems 5.1, 6.1 and 7.2). The Section 5 and Section 7 curves overlap starting at $\tau = .25$.

- Since z is regular and preimage sets have minimum distance each W_z has high fuzzy min-entropy.
- The hash is universal, so learning the value of z reduces the set of possible values by another factor (Lemma 5.3).
- With this additional loss, for the average W_z , the interior of many parts contain no points from W_z . One can now build a distinguisher for a key derived from W_z from a random key. If a key comes from a part whose interior is empty the distinguisher outputs *random*, otherwise it outputs *real*.

We now proceed with the full proof.

Proof of Theorem 5.1. We show the impossibility for an average member of \mathcal{W} . We defer describing the family \mathcal{W} until after a new bound on the preimage set size of most keys for a fuzzy extractor. The following lemma shows that the knowledge of p and key reduces the entropy of w .

Lemma 5.2. *Suppose \mathcal{M} is $\{0, 1\}^n$ with the Hamming metric, $\kappa \geq 2$, $0 \leq t \leq n/2$, and $\epsilon \geq 0$. Suppose (Gen, Rep) is a $(\mathcal{M}, \mathcal{W}, \kappa, t, \epsilon)$ -fuzzy extractor with error $\delta = 0$, for some distribution family \mathcal{W} over \mathcal{M} .*

Let $\tau = t/n$. For any fixed p , there is a set $\text{GoodKey}_p \subseteq \{0, 1\}^\kappa$ of size at least $2^{\kappa-1}$ such that for every key $\in \text{GoodKey}_p$,

$$\log |\{v \in \mathcal{M} | (\text{key}, p) \in \text{supp}(\text{Gen}(v))\}| \leq n \cdot h_2\left(\frac{1}{2} - \tau\right) \leq n \cdot \left(1 - \frac{2}{\ln 2} \cdot \tau^2\right),$$

and, therefore, for any distribution $D_{\mathcal{M}}$ on \mathcal{M} ,

$$H_0(D_{\mathcal{M}} | \text{Gen}(D_{\mathcal{M}}) = (\text{key}, p)) \leq n \cdot h_2\left(\frac{1}{2} - \tau\right) \leq n \cdot \left(1 - \frac{2}{\ln 2} \cdot \tau^2\right).$$

Proof. The set GoodKey_p consists of all keys for which $H_0(\mathcal{M} | \text{Rep}(\mathcal{M}, p) = \text{key}) \leq 2^{n-\kappa+1}$.

The intuition is as follows. By perfect correctness of Rep , the input w to Gen has the following property: for all w' within distance t of w , $\text{Rep}(w', p) = \text{Rep}(w, p)$. Thus, if we partition \mathcal{M} according to the output of Rep , the true w is t away from the interior of a part. Interior sets are small, which means the set of possible of w values is small. (We note that by perfect correctness, Rep has a deterministic output even if the algorithm is randomized, so this partition is well-defined.)

To formalize this intuition, fix p and partition \mathcal{M} according to the output of $\text{Rep}(\cdot, p)$ as follows: let $Q_{p,\text{key}} = \{w' \in \mathcal{M} | \text{Rep}(w', p) = \text{key}\}$. Note that there are 2^κ keys and thus 2^κ parts $Q_{p,\text{key}}$. Let GoodKey_p be the set of keys for which these parts are not too large: $\text{key} \in \text{GoodKey}_p \Leftrightarrow |Q_{p,\text{key}}| \leq 2 \cdot \mathcal{M} / 2^\kappa = 2^{n-\kappa+1}$. Observe that GoodKey_p contains at least half the keys: $|\text{GoodKey}_p| \geq 2^{\kappa-1}$ (if not, then $\cup_{\text{key}} |Q_{p,\text{key}}| > |\mathcal{M}|$). For the remainder of the proof we focus on elements in GoodKey_p .

As explained above, if w is the input to Gen , then every point w' within distance t of w must be in the same part $Q_{p,\text{key}}$ as w , by correctness of Rep . Thus, w must come from the interior of some $Q_{p,\text{key}}$, where interior is defined as

$$\text{Inter}(Q_{p,\text{key}}) = \{w \in Q_{p,\text{key}} | \forall w' \text{ s.t. } \text{dis}(w, w') \leq t, w' \in Q_{p,\text{key}}\}.$$

We now use the isoperimetric inequality to bound the size of $\text{Inter}(Q_{p,\text{key}})$. Define a *near-ball*⁴ centered at x to be any set S that is contained in a ball of some radius η and contains the ball of radius $\eta - 1$ around x . The inequality of [FF81, Theorem 1] (the original result is due to Harper [Har66]) says that for any sets $A, B \subset \{0, 1\}^n$, there are near-balls X and Y centered at 0^n and 1^n , respectively, such that $|A| = |X|$, $|B| = |Y|$, and $\min_{a \in A, b \in B} \text{dis}(a, b) \leq \min_{x \in X, y \in Y} \text{dis}(x, y)$.

Letting A be the $\text{Inter}(Q_{p,\text{key}})$ and B be the complement of $Q_{p,\text{key}}$ and applying this inequality, we get a near-ball $S_{p,\text{key}}$ centered at 0^n and a near-ball D centered at 1^n , such that $|S_{p,\text{key}}| = |\text{Inter}(Q_{p,\text{key}})|$, $|D| = 2^n - |Q_{p,\text{key}}|$, and $\forall s \in S_{p,\text{key}}, d \in D, \text{dis}(s, d) > t$. Note that since $\text{key} \in \text{GoodKey}_p$ and $\kappa \geq 2$, we have $|Q_{p,\text{key}}| \leq 2^{n-\kappa+1}$, and therefore $|D| \geq 2^{n-1}$.

Thus, D includes all the strings of Hamming weight $\lceil n/2 \rceil$ (because it is centered at 1^n and takes up at least half the space), which means that the maximum Hamming weight of an element of $S_{p,\text{key}}$ is $\lceil n/2 \rceil - t - 1 \leq n/2 - t$ (because each element of $S_{p,\text{key}}$ is at distance more than t from D). We can now use binary entropy to bound the size of $S_{p,\text{key}}$ by Lemma 2.2:

$$|\text{Inter}(Q_{p,\text{key}})| = |S_{p,\text{key}}| \leq |\{x | \text{dis}(x, 0) \leq n/2 - t\}| \leq 2^{n \cdot h_2(\frac{1}{2} - \frac{t}{n})}.$$

The theorem statement follows by taking the logarithm of both sides and by observing (using Taylor series expansion at $\tau = 0$ and noting that the third derivative is negative) that $h_2(\frac{1}{2} - \tau) \leq 1 - \frac{2}{\ln 2} \cdot \tau^2$. \square

⁴In most statements of the isoperimetric inequality, this type of set is simply called a ball. We use the term *near-ball* for emphasis.

We now introduce the family \mathcal{W} . Let $\{\text{Hash}_k\}_{k \in \mathcal{K}}$ be a family of hash function with domain \mathcal{M} and the following properties:

- 2^{-a} -universality: for all $v_1 \neq v_2 \in \mathcal{M}$, $\Pr_{k \leftarrow \mathcal{K}}[\text{Hash}_k(v_1) = \text{Hash}_k(v_2)] \leq 2^{-a}$, where $a = n \cdot h_2(\frac{1}{2} - \tau) + 3$.
- 2^m -regularity: for each $k \in \mathcal{K}$ and h in the range of Hash_k , $|\text{Hash}_k^{-1}(h)| = 2^m$, where $m \geq \mu n$.
- preimage sets have minimum distance $t + 1$: for all $k \in \mathcal{K}$, if $v_1 \neq v_2$ but $\text{Hash}_k(v_1) = \text{Hash}_k(v_2)$, then $\text{dis}(v_1, v_2) > t$.

We demonstrate the existence of such a hash family in Lemma 5.4. Let Z be the random variable consisting of pairs (k, h) , where k is uniform in \mathcal{K} and h is uniform in the range of Hash_k . Let W_z for $z = (k, h)$ be the uniform distribution on $\text{Hash}_k^{-1}(h)$. By the 2^m -regularity and minimum distance properties of Hash , $H_\infty(W_z) = H_{t, \infty}^{\text{fuzz}}(W_z) = m$. Let $\mathcal{W} = \{W_z\}$.

We now want to show that for a random $z \leftarrow Z$, if (key, p) is the output of $\text{Gen}(W_z)$, then key can be easily distinguished from uniform in the presence of p and z .

First, view the sequence of events that we are trying to analyze as a game. The adversary chooses a uniform $k \in \mathcal{K}$ and uniform h in the range of Hash_k . A uniform w from \mathcal{M} s.t. $\text{Hash}_k(w) = h$ then gets chosen, $(\text{key}, p) = \text{Gen}(w)$ gets computed, and the adversary receives p . The output of this game is (k, h, w, p, key) .

Consider now an alternative game. A uniform w gets chosen from \mathcal{M} and uniform key k gets chosen from \mathcal{K} . $(\text{key}, p) = \text{Gen}(w)$ gets computed. The adversary receives $(k, h = \text{Hash}_k(w), p)$. The output of the game is (k, h, w, p, key) .

The distributions of the adversary's views and the outputs in the two games are identical: indeed, in both games, three random variable are uniform and independent (i.e., w is uniform in \mathcal{M} , k is uniform in \mathcal{K} , and the random coins of Gen are uniform in their domain), and the rest are determined fully by these three. However, the second game is easier to analyze, which is what we now do.

In this game, the value w is uniform on \mathcal{M} (in the absence of knowledge about w). Knowledge of p reduced the set of possible w from 2^n to $2^{n \cdot h_2(\frac{1}{2} - \tau)}$, (Lemma 5.2). We know show that knowledge of z reduces the set of possible w by another factor of 2^a . Let \mathbf{K} denote the uniform distribution on \mathcal{K} .

Lemma 5.3. *Let L be a distribution. Let $\{\text{Hash}_k\}_{k \in \mathcal{K}}$ be a family of 2^{-a} -universal hash functions on the support of L . Assume k is uniform in \mathcal{K} and independent of L . Then*

$$\tilde{H}_0(L|\mathbf{K}, \text{Hash}_{\mathcal{K}}(L)) < \log(1 + |\text{supp}(L)| \cdot 2^{-a}) \leq \max(1, 1 + H_0(L) - a).$$

Proof. Let U_L denote the uniform distribution on the support of L .

$$\begin{aligned}
2^{\tilde{H}_0(L|K, \text{Hash}_\kappa(L))} &= \mathbb{E}_{k \leftarrow K, h} |\{v \in L | \text{Hash}_k(v) = h\}| \\
&= \mathbb{E}_{k \leftarrow K} \sum_h \Pr[\text{Hash}_k(L) = h] \cdot |\{v \in L | \text{Hash}_k(v) = h\}| \\
&= \mathbb{E}_{k \leftarrow K} \sum_h |L| \cdot \Pr[\text{Hash}_k(L) = h] \cdot \Pr[\text{Hash}_k(U_L) = h] \\
&= |\text{supp}(L)| \cdot \mathbb{E}_{k \leftarrow K} \Pr_{v_1 \leftarrow L, v_2 \leftarrow U_L} [\text{Hash}_k(v_1) = \text{Hash}_k(v_2)] \\
&= |\text{supp}(L)| \cdot \Pr_{v_1 \leftarrow L, v_2 \leftarrow U_L, k \leftarrow K} [\text{Hash}_k(v_1) = \text{Hash}_k(v_2)] \\
&\leq |\text{supp}(L)| \cdot \left(\Pr_{v_1 \leftarrow L, v_2 \leftarrow U_L} [v_1 = v_2] + \Pr_{v_1 \leftarrow L, v_2 \leftarrow U_L} [v_1 \neq v_2] \cdot 2^{-a} \right) \\
&< 1 + |\text{supp}(L)| \cdot 2^{-a}.
\end{aligned}$$

This completes the proof of Lemma 5.3. \square

Let M denote the uniform distribution on \mathcal{M} . By Lemma 5.2, for any p , $H_0(M|\text{Gen}(M) = (\text{key}, p))$ such that $\text{key} \in \text{GoodKey}_p \leq n \cdot h_2\left(\frac{1}{2} - \frac{t}{n}\right) + \kappa$ (because there are most 2^κ keys in GoodKey_p). Applying Lemma 5.3 (and recalling that $\kappa \geq 3$), we get that for any p ,

$$\tilde{H}_0(M|\text{Gen}(M) = (\text{key}, p) \text{ s.t. } \text{key} \in \text{GoodKey}_p, K, \text{Hash}_\kappa(M)) < \max\left(1, 1 + n \cdot h_2\left(\frac{1}{2} - \frac{t}{n}\right) + \kappa - a\right) \leq \kappa - 2.$$

(Note carefully the somewhat confusing conditioning notation above, because we are conditioning on both events and variables. The event is $\text{key} \in \text{GoodKey}_p$ and the variables are k and $\text{Hash}_k(M)$.)

By correctness, for a fixed p , $\text{Rep}(w, p)$ can produce only one key—the same one that was produced during $\text{Gen}(w)$. Since applying a deterministic function (in this case, Rep) cannot increase H_0 , we get that for each p ,

$$\tilde{H}_0(\text{key}|\text{Gen}(M) = (\text{key}, p) \text{ s.t. } \text{key} \in \text{GoodKey}_p, K, \text{Hash}_\kappa(M)) < \kappa - 2.$$

Thus, on average over $z = (k, h)$, over half the keys in GoodKey_p (i.e., over a quarter of all possible 2^κ keys) cannot be produced. Let Implausible be the set of triples $(\text{key}, p, z = (k, h))$ such that $\Pr[\text{Gen}(W_z) = (\text{key}, p)] = 0$. Triples drawn by sampling w from W_z and computing $(p, \text{key}) = \text{Gen}(w)$ never come from this set. On other hand, random triples come Implausible at over quarter of the time. Thus, by definition of statistical distance, $\epsilon > \frac{1}{4}$. It remains to show that the hash family with the desired properties exists.

Lemma 5.4. *For any $0 \leq \tau < \frac{1}{2}$, $\mu > 0$, α , and n such that $\mu \leq 1 - h_2(\tau) - \frac{2}{n}$ and $\mu \leq 1 - \alpha - \frac{2}{n}$, there exists a family of hash functions $\{\text{Hash}_k\}_{k \in \mathcal{K}}$ on $\{0, 1\}^n$ that is 2^{-a} -universal for $a = \alpha n$, 2^m regular for $m \geq \mu n$, and whose preimage sets have minimum distance $t + 1$ for $t = \tau n$.*

Proof. Let \mathcal{C} be the set of all binary linear codes of rate μ (to be precise, dimension $m = \lceil \mu n \rceil$), length n , and minimum distance $t + 1$:

$$\mathcal{C} = \{C | C \text{ is a linear subspace of } \{0, 1\}^n, \dim(C) = m, \min_{c \in C - \{0^n\}} \text{dis}(c, 0^n) > t\}.$$

For each $C \in \mathcal{C}$, fix H_C , an $(n - m) \times n$ parity check matrix for C , such that $C = \ker H_C$. For $v \in \{0, 1\}^n$, let the syndrome $\text{syn}_C(v) = H_C \cdot v$. Let $\{\text{Hash}_k\}_{k \in \mathcal{K}} = \{\text{syn}_C\}_{C \in \mathcal{C}}$.

2^m regularity follows from the fact that for each $h \in \{0,1\}^{n-\mu n}$, $\text{Hash}_k^{-1}(h)$ is a coset of C , which has size 2^m . The minimum distance property is also easy: if $v_1 \neq v_2$ but $\text{syn}_C(v_1) = \text{syn}_C(v_2)$, then $H_C(v_1 - v_2) = 0^n$, hence $v_1 - v_2 \in C - \{0^n\}$ and hence $\text{dis}(v_1, v_2) = \text{dis}(v_1 - v_2, 0) > t$.

We show 2^{-a} -universality by first considering a slightly larger hash family. Let \mathcal{K}' be the set of all m -dimensional subspaces of $\{0,1\}^n$; for each $C' \in \mathcal{K}'$, choose a parity check matrix $H_{C'}$ such that $C' = \ker H_{C'}$, and let $\text{syn}_{C'}(v) = H_{C'} \cdot v$. Let $\{\text{Hash}'_{k'}\}_{k' \in \mathcal{K}'} = \{\text{syn}_{C'}\}_{C' \in \mathcal{K}'}$. This family is 2^{m-n} -universal: for $v_1 \neq v_2$, $\Pr_{C' \in \mathcal{K}'}[H_{C'} \cdot v_1 = H_{C'} \cdot v_2] = \Pr_{C' \in \mathcal{K}'}[v_1 - v_2 \in \ker H_{C'} = C'] = \frac{2^m}{2^n}$, because C' is a random m -dimensional subspace. Note that this family is not much bigger than our family $\{\text{Hash}_k\}_{k \in \mathcal{K}}$, because, as long as $\mu < 1 - h_2(\tau)$, almost every subspace of $\{0,1\}^n$ of dimension m has minimum distance $t + 1$ for a sufficiently large n . Formally,

$$\begin{aligned} \Pr_{C' \in \mathcal{K}'}[C' \notin \mathcal{C}] &= \Pr_{C' \in \mathcal{K}'}[\exists v_1 \neq v_2 \in C' \text{ s. t. } \text{dis}(v_1, v_2) \leq t] \\ &= \Pr_{C' \in \mathcal{K}'}[\exists v_1 \neq v_2 \in C' \text{ s. t. } \text{dis}(v_1 - v_2, 0^n) \leq t] \\ &= \Pr_{C' \in \mathcal{K}'}[\exists v \in C' - \{0^n\} \text{ s. t. } \text{dis}(v, 0^n) \leq t] \\ &\leq \sum_{v \in B_t(0^n) - \{0^n\}} \Pr_{C' \in \mathcal{K}'}[v \in C'] \leq 2^{nh_2(\tau)} \cdot \frac{2^m}{2^n} \leq \frac{1}{2} \end{aligned}$$

(the penultimate inequality follows by Lemma 2.2 and the last one from $m \leq \mu n + 1$ and $\mu \leq 1 - h_2(\tau) - \frac{2}{n}$).

Since this larger family is universal and at most factor of two bigger than our family, our family is also universal:

$$\begin{aligned} \Pr_{C \in \mathcal{C}}[\text{syn}_C(v_1) = \text{syn}_C(v_2)] &= \frac{|\{C \in \mathcal{C} | \text{syn}_C(v_1) = \text{syn}_C(v_2)\}|}{|\mathcal{C}|} \\ &\leq \frac{|\{C \in \mathcal{K}' | \text{syn}_C(v_1) = \text{syn}_C(v_2)\}|}{|\mathcal{K}'|} \cdot \frac{|\mathcal{K}'|}{|\mathcal{C}|} \leq 2^{m-n+1} \end{aligned}$$

Thus, we obtain the desired result as long as $m - n + 1 \leq -a$, which is implied by the condition $\mu \leq 1 - \alpha - \frac{2}{n}$ and the fact that $m \leq \mu n + 1$. This completes the proof of Lemma 5.4. \square

Applying Lemma 5.4 with $\alpha = h_2\left(\frac{1}{2} - \tau\right) + \frac{3}{n}$, we see that the largest possible μ is

$$\max_{\tau} \min \left(1 - h_2(\tau), 1 - h_2\left(\frac{1}{2} - \tau\right) \right) \approx 0.1887.$$

Using the quadratic approximation to $h_2\left(\frac{1}{2} - \tau\right)$ (see Lemma 5.2), we can let μ be a free variable and set $\tau = .6\sqrt{\mu}$, in which case both constraints will be satisfied for all $0 < \mu \leq .18$ and sufficiently large n , as in the theorem statement. This concludes the proof of Theorem 5.1. \square

6 Impossibility in the Case of Imperfect Correctness

The impossibility result in the previous section applies only to fuzzy extractors with perfect correctness. In this section, we build on the work of Holenstein and Renner [HR05] to show the impossibility of fuzzy extractors even when they are allowed to make mistakes a constant fraction δ (as much as 4%) of the time. However, the drawback of this result, as compared to the previous section, is that we can show impossibility only for a relatively low entropy rate of at most 7%. In Section 7, we rule out stronger primitives called secure sketches with nonzero error (which are used in most fuzzy extractor constructions), even for entropy rate as high as 50%.

Theorem 6.1. *Let \mathcal{M} denote the Hamming space $\{0, 1\}^n$. There exists a family of distributions \mathcal{W} over \mathcal{M} such that for each element $W \in \mathcal{W}$, $H_{t, \infty}^{\text{fuzz}}(W) = H_{\infty}(W) \geq m$, and yet any $(\mathcal{M}, \mathcal{W}, \kappa, t, \epsilon)$ -fuzzy extractor with error $\delta \leq \frac{1}{25}$ has $\epsilon > \frac{1}{25}$.*

This holds for any $\kappa > 0$ under the following conditions on the entropy rate $\mu = m/n$, noise rate $\tau = t/n$, and n :

- any $0 \leq \tau \leq \frac{1}{2}$ and μ such that $\mu < 4\tau(1 - \tau) \left(1 - h_2\left(\frac{1}{4-4\tau}\right)\right)$
- any sufficiently large n (as a function of τ and μ)

Note that the conditions on μ and τ imply that the result applies to any entropy rate $\mu \leq \frac{1}{15}$ as long as τ is set appropriately and n is sufficiently large. The τ vs. μ tradeoff is depicted in Figure 1.

The core structure of the proof is the same as Theorem 5.1. We construct a family \mathcal{W} where knowing the element z (specifying $W_z \in \mathcal{W}$) reveals substantial information. However, this proof uses a different family and different techniques. The outline proceeds as follows:

1. The Rep algorithm (with p specified) can be used as a decoding algorithm for a binary symmetric channel (BSC) with error probability $(1 - \alpha)/2$. To show this, we just need to argue that for a BSC with error probability $(1 - \alpha)/2$, the probability of more than t errors is small.
2. The adversary (in the auxiliary knowledge Z) will receive $\{\perp, 0, 1\}$ for each position in W . The value \perp indicates an erasure, and 0 or 1 indicates the true bit of W . So the family W_Z fixes some bits of W . These bits are known to the adversary but not the construction. Let $1 - \beta$ denote the probability of receiving a \perp . The adversary's view corresponds to a classic erasure channel.
3. As long as β is not too large, one can show that W_Z has high fuzzy min-entropy as it corresponds to the uniform distribution over a restricted set of bits. We need to cap the number of bits received by the adversary for this to be true for all elements of W_Z . However, by standard tail bounds, this removes very few distributions from the family.
4. We use a result of Holenstein and Renner [HR05, Theorem 4] that says the Shannon capacity of a β -erasure channel is greater than the capacity of a $(1 - \alpha)/2$ -binary symmetric channel.
5. From this theorem we can argue that the key has less Shannon entropy to the adversary than to Rep with a valid input.
6. The remainder of the proof is technical and converts this gap in Shannon entropy to a deficiency of the resulting key.

We now proceed with the full proof.

Proof. Similarly to the proof of Theorem 5.1, we will prove that any fuzzy extractor fails for an average element of \mathcal{W} : letting Z denote a choice of W from \mathcal{W} , we will show that $\mathbf{SD}((\text{Key}, P, Z), (U_{\kappa}, P, Z)) > \frac{1}{25}$.

We start by describing the family of distributions. In this case, Z will not be uniform but rather binomial (with tails cut off). Essentially, Z will contain each bit of w with (appropriately chosen) probability β ; given $Z = z$, the remaining bits of w will be uniform and independent.

For a string $z \in \{0, 1, \perp\}^n$, denote by $\text{info}(z)$ the number of entries in z that are not \perp : $\text{info}(z) = |\{i \text{ s.t. } z_i \neq \perp\}|$. Let W_z be the uniform distribution over all strings in $\{0, 1\}^n$ that agree with z in positions that are not \perp in z (i.e., all strings $w \in \{0, 1\}^n$ such that for $1 \leq i \leq n$, either $z_i = \perp$ or $w_i = z_i$).

Let $0 \leq \beta' \leq 1$ be a parameter (we will set it at the end of the proof). Let Z' denote the distribution on strings in $\{0, 1, \perp\}^n$ in which each symbol is, independently of other symbols, \perp with probability $1 - \beta'$, 0 with probability $\beta'/2$, and 1 with probability $\beta'/2$. Let $\beta = \beta' + \frac{1.4}{\sqrt{n}}$. Consider two distribution families: $\mathcal{W}' = \{W_z\}_{z \leftarrow Z'}$ and a smaller family $\mathcal{W} = \{W_z\}_{z \leftarrow Z}$, where $Z = Z' | \text{info}(Z') \leq \beta n$ (the second family is smaller because, although on average $\text{info}(Z') = \beta'n$, there is a small chance that $\text{info}(Z')$ is higher than even βn).

We will use \mathcal{W} to prove the theorem statement. First, we will show that every distribution $W_z \in \mathcal{W}$ has sufficient $H_{t,\infty}^{\text{fuzz}}$. Indeed, z constrains $\text{info}(z)$ coordinates out of n and leaves the rest uniform. Thus, $H_{t,\infty}^{\text{fuzz}}(W_z)$ is the same as $H_{t,\infty}^{\text{fuzz}}$ of the uniform distribution on the space $\{0, 1\}^{n - \text{info}(z)}$. Let $a = n - \text{info}(z)$. By Lemma 2.2

$$H_{t,\infty}^{\text{fuzz}}(W_z) \geq a \left(1 - h_2\left(\frac{t}{a}\right)\right) \geq n(1 - \beta) \left(1 - h_2\left(\frac{t}{n(1 - \beta)}\right)\right) = n(1 - \beta) \left(1 - h_2\left(\frac{\tau}{1 - \beta}\right)\right).$$

and therefore

$$\mu = (1 - \beta) \left(1 - h_2\left(\frac{\tau}{1 - \beta}\right)\right). \quad (1)$$

Note that smaller β gives a higher fuzzy entropy rate.

Second, we now want to show, similarly to the proof of Theorem 5.1, that $\mathbf{SD}((\text{Key}, P, Z), (U_\kappa, P, Z)) > \frac{1}{25}$. We will do so by considering the family \mathcal{W} . Observe that by triangle inequality

$$\begin{aligned} \mathbf{SD}((\text{Key}, P, Z), (U_\kappa, P, Z)) &\geq \mathbf{SD}((\text{Key}, P, Z'), (U_\kappa, P, Z')) \\ &\quad - \mathbf{SD}((\text{Key}, P, Z'), (\text{Key}, P, Z)) \\ &\quad - \mathbf{SD}((U_\kappa, P, Z), (U_\kappa, P, Z')) \\ &\geq \mathbf{SD}((\text{Key}, P, Z'), (U_\kappa, P, Z')) - 2 \cdot \mathbf{SD}(Z', Z) \\ &\geq \mathbf{SD}((\text{Key}, P, Z'), (U_\kappa, P, Z')) - \frac{1}{25}. \end{aligned}$$

The last line follows by Hoeffding's inequality [Hoe63],

$$\mathbf{SD}(Z', Z) = \Pr[\text{info}(Z') > \beta n] \leq \exp\left(-2n \left(\frac{1.4}{\sqrt{n}}\right)^2\right) < \frac{1}{50}.$$

Denote $\mathbf{SD}((\text{Key}, P, Z'), (U_\kappa, P, Z'))$ by ϵ' . To bound ϵ' , we recall a result of Holenstein and Renner [HR05, Theorem 4] (we will use the version presented in [Hol06, Lemma 4.4]). For a random variable W with a values in $\{0, 1\}^n$, let W^{noisy} denote a noisy copy of W : namely, the random variable obtained by passing W through a binary symmetric channel with error rate $\frac{1-\alpha}{2}$ (that is, $W_i^{\text{noisy}} = W_i$ with probability $\frac{1+\alpha}{2}$ and $W_i^{\text{noisy}} = 1 - W_i$ with probability $\frac{1-\alpha}{2}$, independently for each position i). Holenstein and Renner show that if $\alpha^2 \leq \beta$, then Shannon entropy of Key conditioned on P and W^{noisy} is greater than Shannon entropy of Key conditioned on Z and W^{noisy} . Intuitively, this means that the Rep , when given P and W^{noisy} , knows less about Key than the adversary (who knows P and Z).

Recall the definitions of Shannon entropy $H_1(X) \stackrel{\text{def}}{=} \mathbb{E}_{x \leftarrow X} -\log \Pr[X = x]$ and conditional Shannon entropy $H_1(X|Y) \stackrel{\text{def}}{=} \mathbb{E}_{y \leftarrow Y} H_1(X|Y = y)$.

Theorem 6.2 ([HR05, Theorem 4]; [Hol06, Lemma 4.4]). *Suppose that (P, Key) is a pair of random variables derived from W . If $\alpha^2 \leq \beta'$, then*

$$H_1(\text{Key}|P, Z') \leq H_1(\text{Key}|P, W^{\text{noisy}})$$

where H_1 denotes Shannon entropy, W^{noisy} is W passed through a binary symmetric channel with error rate $\frac{1-\alpha}{2}$, and Z' is W passed through a binary erasure channel with erasure rate $1 - \beta'$.

(For a reader interested in how our statement of Lemma 6.2 follows from [Hol06, Lemma 4.4], note that what we call $\text{Key}, P, W^{\text{noisy}}$, and Z' are called U, V, Y , and Z , respectively, in [Hol06]. Note also that we use only the part of the lemma that says that secret key rate $S_{\rightarrow} = 0$ when $\alpha^2 \leq \beta$, and the definition [Hol06, Definition 3.1] of the notion S_{\rightarrow} in terms of Shannon entropy.)

We now need to translate this bound on Shannon entropy to the language of statistical distance ϵ of the key from uniform, reliability δ of the procedure Rep , and key length κ , as used in the definition of fuzzy extractors. First, we will do this translation for the case of noisy rather than worst-case input to Rep .

Lemma 6.3. *Let $(W, W^{\text{noisy}}, Z')$ be a triple of correlated random variables such that*

- *W and W^{noisy} are uniform over $\{0, 1\}^n$,*
- *W^{noisy} is W passed through a binary symmetric channel with error rate $\frac{1-\alpha}{2}$ (that is, each bit position of W agrees with corresponding bit position of W^{noisy} with probability $\frac{1+\alpha}{2}$), and*
- *Z' is W passed through a binary erasure channel with erasure rate $1 - \beta'$ (that is, each bit position of Z' agrees with the corresponding bit position of W with probability β' and is equal to \perp otherwise).*

Suppose $\text{Gen}(W)$ produces (Key, P) with Key of length κ . Suppose $\Pr[\text{Rep}(W^{\text{noisy}}, P) = \text{Key}] = 1 - \delta'$. Suppose further that $\text{SD}((\text{Key}, P, Z'), (U_{\kappa}, P, Z')) = \epsilon'$. If $\alpha^2 \leq \beta'$, then

$$\kappa \leq \frac{h_2(\epsilon') + h_2(\delta')}{1 - \epsilon' - \delta'}.$$

In other words, if $\alpha^2 \leq \beta'$, $\epsilon' \leq \frac{1}{12}$, and $\delta' \leq \frac{1}{12}$, then even a 1-bit Key is impossible to obtain.

(We note that a similar result follows from [Hol06, Theorem 3.17] if we set the variables S_{\rightarrow} , γ , and m in that theorem to $0, \delta$, and κ , respectively. However, we could not verify the correctness of that theorem due to its informal treatment of what “ ϵ -close to uniform” means; it seems that the small correction term $-h_2(\epsilon)$, just like in our result, is needed on the right-hand side to make that theorem correct.)

Proof of Lemma 6.3. Reliability allows us to bound the entropy of the key. By Fano’s inequality [Fan61, Section 6.2, p. 187],

$$H_1(\text{Key}|P, W^{\text{noisy}}) \leq \kappa\delta' + h_2(\delta').$$

Hence, by Theorem 6.2 (and the assumption that $\alpha^2 > \beta'$), we have

$$H_1(\text{Key}|P, Z') \leq \kappa\delta' + h_2(\delta'). \tag{2}$$

We now need the following lemma, which shows that near-uniformity implies high entropy.

Lemma 6.4. *For a pair of random variables (A, B) such that the statistical distance between (A, B) and $U_{\kappa} \times B$ is ϵ , then $H_1(A|B) \geq (1 - \epsilon)\kappa - h_2(\epsilon)$.*

Proof. Let E denote a binary random variable correlated with (A, B) as follows: when $A = a$ and $B = b$, then $E = 0$ with probability

$$\max(\Pr[(A, B) = (a, b)] - \Pr[U_\kappa \times B = (a, b)], 0).$$

Similarly, let F denote a binary random variable correlated with $U_\kappa \times B$ as follows: when $U_\kappa = a$ and $B = b$, then $F = 0$ with probability

$$\max(\Pr[U_\kappa \times B = (a, b)] - \Pr[(A, B) = (a, b)], 0).$$

Note that $\Pr[E = 0] = \Pr[F = 0] = \epsilon$, by definition of statistical distance. Note also that $(A, B|E = 1)$ is the same distribution as $(U_\kappa \times B|F = 1)$. Since conditioning cannot increase Shannon entropy (by a simple argument — see, e.g., [Ash65, Theorem 1.4.4]), we get

$$\begin{aligned} H_1(A|B) &\geq H_1(A|B, E) \\ &= \Pr[E = 1]H_1(A|B, E = 1) + \Pr[E = 0]H_1(A|B, E = 0) \\ &\geq (1 - \epsilon)H_1(A|B, E = 1) = (1 - \epsilon)H_1(U_\kappa|B, F = 1). \end{aligned}$$

To bound this latter quantity, note that (the first line follows from the chain rule $H_1(X) \leq H_1(X, Y) = H_1(X|Y) + H_1(Y)$ [Ash65, Theorem 1.4.4])

$$\begin{aligned} \kappa = H_1(U_\kappa|B) &\leq H_1(U_\kappa|B, F) + H_1(F) \\ &= (1 - \epsilon)H_1(U_\kappa|B, F = 1) + \epsilon \cdot H_1(U_\kappa|B, F = 0) + h_2(\epsilon) \\ &\leq (1 - \epsilon)H_1(U_\kappa|B, F = 1) + \epsilon \cdot \kappa + h_2(\epsilon) \end{aligned}$$

Rearranging terms, we get $H_1(U_\kappa|B, F = 1) \geq \kappa - h_2(\epsilon)/(1 - \epsilon)$, and thus

$$H_1(A|B) \geq (1 - \epsilon)\kappa - h_2(\epsilon).$$

This concludes the proof of Lemma 6.4. \square

Combining (2) and Lemma 6.4 (applied to $A = \text{Key}$, $B = (P, Z')$, and $\epsilon = \epsilon'$), we get the claimed bound. This concludes the proof of Lemma 6.3. \square

Next, we translate this result from the noisy-input-case to the worst-case input case. Set $\alpha = \sqrt{\beta'}$. Suppose $t \geq n \left(\frac{1 - \sqrt{\beta'}}{2} + \frac{1.4}{\sqrt{n}} \right)$. By Hoeffding's inequality [Hoe63],

$$\Pr[\text{dis}(W, W^{\text{noisy}}) > t] \leq \exp \left(-2n \left(\frac{1.4}{\sqrt{n}} \right)^2 \right) < \frac{1}{50}.$$

Thus, a fuzzy extractor that corrects t errors with reliability δ implies that $\Pr[\text{Rep}(W^{\text{noisy}}, P) = \text{Key}] \geq 1 - \delta'$ for $\delta' = \delta + \frac{1}{50}$. Since $\delta \leq 1/25$, we have $\delta' < 1/12$ and Lemma 6.3 applies to gives us $\epsilon' > 1/12$ and $\epsilon > 1/12 - 1/25 > 1/25$ as long as $\kappa > 0$.

Finally, we work out the relationship between μ and τ and eliminate β , as follows. Recall that $\beta = \beta' + \frac{1.4}{\sqrt{n}}$; therefore $\sqrt{\beta} \leq \sqrt{\beta'} + \frac{1.2}{n^{1/4}}$, and it suffices to take $\tau \geq \frac{1 - \sqrt{\beta}}{2} + \frac{2}{\sqrt[4]{n}}$. Thus, we can set any $\tau > \frac{1 - \sqrt{\beta}}{2}$ as long as n is sufficiently large. Solving for β (that is, taking any $\beta > (1 - 2\tau)^2$) and substituting into Equation 1, we can get any $\mu < 4\tau(1 - \tau) \left(1 - h_2 \left(\frac{1}{4 - 4\tau} \right) \right)$ for a sufficiently large n . \square

7 Stronger Impossibility Result for Secure Sketches

Most fuzzy extractor constructions share the following feature with our Construction 4.2: p includes information that is needed to recover w from w' ; both Gen and Rep simply apply an extractor to w . The recovery of w from w' , known as information-reconciliation, forms the core of many fuzzy extractor constructions. The primitive that performs this information reconciliation is called *secure sketch*. In this section we show stronger impossibility results for secure sketches. First, we recall their definition from [DORS08, Section 3.1] (modified slightly, in the same way as Definition 2.3).

Definition 7.1. An $(\mathcal{M}, \mathcal{W}, \tilde{m}, t)$ -secure sketch with error δ is a pair of randomized procedures, “sketch” (SS) and “recover” (Rec). SS on input $w \in \mathcal{M}$ returns a bit string $ss \in \{0, 1\}^*$. Rec takes an element $w' \in \mathcal{M}$ and $ss \in \{0, 1\}^*$. (SS, Rec) have the following properties:

1. Correctness: $\forall w, w' \in \mathcal{M}$ if $\text{dis}(w, w') \leq t$ then $\Pr[\text{Rec}(w', \text{SS}(w)) = w] \geq 1 - \delta$.
2. Security: for any distribution $W \in \mathcal{W}$, $\tilde{H}_\infty(W | \text{SS}(W)) \geq \tilde{m}$.

Secure sketches are more demanding than fuzzy extractors (secure sketches can be converted to fuzzy extractors by using a randomness extractors like in our Construction 4.2 [DORS08, Lemma 4.1]). We prove a stronger impossibility result for them. Specifically, in the case of secure sketches, we can extend the results of Theorems 5.1 and 6.1 to cover imperfect correctness (that is, $\delta > 0$) and entropy rate μ up to $\frac{1}{2}$. We stress that most fuzzy extractor constructions rely on secure sketches.

Theorem 7.2. Let \mathcal{M} denote the Hamming space $\{0, 1\}^n$. There exists a family of distributions \mathcal{W} over \mathcal{M} such that for each element $W \in \mathcal{W}$, $H_{t, \infty}^{\text{fuzz}}(W) = H_\infty(W) \geq m$, and yet any $(\mathcal{M}, \mathcal{W}, \tilde{m}, t)$ -secure sketch with error δ has $\tilde{m} \leq 2$.

This holds under the following conditions on δ , the entropy rate $\mu = m/n$, noise rate $\tau = t/n$, and n :

- any $0 \leq \tau < \frac{1}{2}$ and $\mu > 0$ such that $\mu < h_2(\tau)$ and $\mu < 1 - h_2(\tau)$
- any $n \geq \max\left(\frac{.5 \log n + 4\delta n + 4}{h_2(\tau) - \mu}, \frac{2}{1 - h_2(\tau) - \mu}\right)$

Note that the result holds for any $\mu < 0.5$ as long as $\delta < (h_2(\tau) - \mu)/4$ and n is sufficiently large. The τ vs. μ tradeoff is depicted in Figure 1. Any fuzzy extractor that uses secure sketch (part of its output is $\text{SS}(w)$) is subject to these bounds. In addition, any fuzzy extractor where the *true* input point w can be computed from key is subject to this bound as well (called an almost injective invertible fuzzy extractor by Yasanuga and Yuzawa [YY18]).

Before starting the formal proof we note that the overall strategy is the same as Theorem 5.1. The only substantive difference is that the functionality of secure sketches allow us to prove a stronger upper bound on the number of possible $w | \text{SS}(w)$ (Lemma 7.3). The core of this proof is arguing that the set of possible $\{v | \Pr_{v' | \text{dis}(v, v') \leq t}[\text{Rec}(v', \text{SS}(v)) = v] \geq 1/2\}$ form a good error correcting code.

The family used is the same as in Theorem 5.1 with more parameter flexibility as more entropy is lost in Lemma 7.3.

Proof of Theorem 7.2. Similarly to the proof of Theorem 5.1, we will prove that any secure sketch algorithm fails for an average element of \mathcal{W} : letting Z denote a uniform choice of W from \mathcal{W} , we will show that $\tilde{H}_\infty(W_Z | \text{SS}(W_Z), Z) \leq 2$. The overall proof strategy is the same as for Theorem 5.1. We highlight only the changes here. Recall that $|B_t|$ denotes the volume of the ball of radius t in the space $\{0, 1\}^n$.

The parameters of the hash family are the same, except for universality: we require 2^{-a} -universality for $a = (n - \log |B_t| + h_2(2\delta))/(1 - 2\delta)$.

We postpone the question of the existence of such a hash family until the end of the proof. We can now state an analogue of Lemma 5.2. This result is an extension of lower bounds from [DORS08, Appendix C], which handles only the case of perfect correctness. It shows that the value of the sketch reduces the entropy of a uniform point by approximately $\log |B_t|$.

Lemma 7.3. *Let \mathcal{M} denote the Hamming space $\{0, 1\}^n$ and $|B_t|$ denote the volume of a Hamming ball of radius t in $\{0, 1\}^n$. Suppose (SS, Rec) is a $(\mathcal{M}, \mathcal{W}, \tilde{m}, t)$ secure sketch with error δ , for some distribution family \mathcal{W} over \mathcal{M} . Then for every $v \in \mathcal{M}$ there exists a set GoodSketch_v such that $\Pr[\text{SS}(v) \in \text{GoodSketch}_v] \geq 1/2$ and for any fixed ss ,*

$$\log |\{v \in \mathcal{M} | ss \in \text{GoodSketch}_v\}| \leq \frac{n - \log |B_t| + h_2(2\delta)}{1 - 2\delta},$$

and, therefore, for any distribution $D_{\mathcal{M}}$ over \mathcal{M} ,

$$H_0(D_{\mathcal{M}} | ss \in \text{GoodSketch}_{D_{\mathcal{M}}}) \leq \frac{n - \log |B_t| + h_2(2\delta)}{1 - 2\delta}.$$

Proof. For any $v \in \mathcal{M}$, define $\text{Neigh}_t(v)$ be the uniform distribution on the ball of radius t around v and let

$$\text{GoodSketch}_v = \{ss | \Pr_{v' \leftarrow \text{Neigh}_t(v)}[\text{Rec}(v', ss) \neq v] \leq 2\delta\}.$$

We prove the lemma by showing two propositions.

Proposition 7.4. *For all $v \in \mathcal{M}$, $\Pr[\text{SS}(v) \in \text{GoodSketch}_v] \geq 1/2$.*

Proof. Let the indicator variable $1_{v', ss}$ be 1 if $\text{Rec}(v', ss) = v$ and 0 otherwise. Let q_{ss} be the quality of the sketch on the ball $B_t(v)$:

$$q_{ss} = \Pr_{v' \leftarrow \text{Neigh}_t(v)}[\text{Rec}(v', ss) = v] = \mathbb{E}_{v' \leftarrow \text{Neigh}_t(v)} 1_{v', ss}.$$

By the definition of correctness for (SS, Rec) , for all $v' \in B_t(v)$,

$$\Pr_{ss \leftarrow \text{SS}(v)}[\text{Rec}(v', ss) = v] \geq 1 - \delta.$$

Hence, $\mathbb{E}_{ss \leftarrow \text{Gen}(v)} 1_{v', ss} \geq 1 - \delta$. Therefore,

$$\mathbb{E}_{ss \leftarrow \text{Gen}(v)} q_{ss} = \mathbb{E}_{ss} \mathbb{E}_{v'} 1_{v', ss} = \mathbb{E}_{v'} \mathbb{E}_{ss} 1_{v', ss} \geq \mathbb{E}_{v'} (1 - \delta) = 1 - \delta.$$

Therefore, applying Markov's inequality to $1 - q_{ss}$, we get $\Pr[q_{ss} \geq 1 - 2\delta] = \Pr[1 - q_{ss} \leq 2\delta] \leq 1/2$. \square

To finish the proof of Lemma 7.3, we will show that the set $\{v \in \mathcal{M} | ss \in \text{GoodSketch}_v\}$ forms a kind of error-correcting code, and then bound the size of the code.

Definition 7.5. *We say that a set C is an (t, δ) -Shannon code if there exists a (possibly randomized) function Decode such that for all $c \in C$,*

$$\Pr_{c' \leftarrow \text{Neigh}_t(c)}[\text{Decode}(c') \neq c] \leq \delta.$$

The set $\{v \in \mathcal{M} | ss \in \text{GoodSketch}_v\}$ forms $(t, 2\delta)$ Shannon code if we set $\text{Decode}(y) = \text{Rec}(y, ss)$. We now bound the size of such a code.

Proposition 7.6. *If $C \subseteq \{0, 1\}^n$ is a (t, δ) -Shannon code, then*

$$\log |C| \leq \frac{n - \log |B_t| + h_2(\delta)}{1 - \delta}.$$

Proof. Let the pair of random variables (X, Y) be obtained as follows: let X be a uniformly chosen element of C and Y be a uniformly chosen element of the ball of radius t around Y . By the existence of Decode and Fano's inequality [Fan61, Section 6.2, p. 187], $H_1(X|Y) \leq h_2(\delta) + \delta \log |C|$. At the same time, $H_1(X|Y) = H_1(X) - H_1(Y) + H_1(Y|X)$ (because $H_1(X, Y) = H_1(X) + H_1(Y|X) = H_1(Y) + H_1(X|Y)$), and therefore $H_1(X|Y) \geq \log |C| - n + \log |B_t|$ (because $H_1(Y) \leq n$). Therefore, $\log |C| - n + \log |B_t| \leq h_2(\delta) + \delta \log |C|$, and the lemma follows by rearranging terms. \square

Lemma 7.3 follows from Proposition 7.6. \square

We now show that entropy drops further when the adversary learns $\text{Hash}_k(w)$. Let \mathbf{M} denote the uniform distribution on \mathcal{M} and \mathbf{K} denote the uniform distribution on \mathcal{K} . Applying Lemma 5.3 to Lemma 7.3, we get that for any ss ,

$$\tilde{H}_0(\mathbf{M} | ss \in \text{GoodSketch}_{\mathbf{M}, \mathbf{K}}, \text{Hash}_{\mathbf{K}}(\mathbf{M})) < \max \left(1, 1 + \frac{n - \log |B_t| + h_2(2\delta)}{1 - 2\delta} - a \right). \quad (3)$$

To complete the proof, we will use this bound on \tilde{H}_0 as a bound on \tilde{H}_∞ , as justified by the following lemma:

Lemma 7.7. *For any random variables X and Y , $\tilde{H}_\infty(X|Y) \leq \tilde{H}_0(X|Y)$.*

Proof. Starting with the definition of \tilde{H}_∞ , recall that $-\log a = \log 1/a$, and apply Jensen's inequality to get

$$\begin{aligned} \log \frac{1}{\mathbb{E}_{y \leftarrow Y} \max_x \Pr[X = x | Y = y]} &\leq \log \mathbb{E}_{y \leftarrow Y} \frac{1}{\max_x \Pr[X = x | Y = y]} \\ &\leq \log \mathbb{E}_{y \leftarrow Y} |\{x | \Pr[X = x | Y = y] > 0\}|. \end{aligned}$$

\square

We need just one more lemma before we can complete the result, an analogue of [DORS08, Lemma 2.2b] for conditioning on a single value $Z = z$ rather than with Z on average (we view conditioning on a single value as equivalent to conditioning on an event).

Lemma 7.8. *For any pair of random variables (X, Y) and event η that is a (possibly randomized) function of (X, Y) , $\tilde{H}_\infty(X|\eta, Y) \geq \tilde{H}_\infty(X|Y) - \log 1/\Pr[\eta]$.*

Proof. The intuition is that to guess X given Y , the adversary can guess that η happened and fail if the guess is wrong. Formally,

$$\begin{aligned}
\tilde{H}_\infty(X|Y) &= -\log \mathbb{E}_{y \leftarrow Y} \max_x \Pr[X = x|Y = y] = -\log \mathbb{E}_{y \leftarrow Y} \max_x \frac{\Pr[X = x \wedge Y = y]}{\Pr[Y = y]} \\
&\leq -\log \mathbb{E}_{y \leftarrow Y} \max_x \frac{\Pr[X = x \wedge Y = y \wedge \eta]}{\Pr[Y = y]} \\
&= -\log \mathbb{E}_{y \leftarrow Y} \max_x \frac{\Pr[X = x \wedge Y = y|\eta] \Pr[\eta]}{\Pr[Y = y]} \\
&= \log \frac{1}{\Pr[\eta]} - \log \sum_{y \leftarrow Y} \max_x \Pr[X = x \wedge Y = y|\eta] \\
&= \log \frac{1}{\Pr[\eta]} - \log \mathbb{E}_{y \leftarrow Y|\eta} \max_x \frac{\Pr[X = x \wedge Y = y|\eta]}{\Pr[Y = y|\eta]} \\
&= \log \frac{1}{\Pr[\eta]} - \log \mathbb{E}_{y \leftarrow Y|\eta} \max_x \Pr[X = x|\eta \wedge Y = y] \\
&= \log \frac{1}{\Pr[\eta]} + \tilde{H}_\infty(X|\eta, Y).
\end{aligned}$$

□

Combining Lemmas 7.8 and 7.7 with Equation 3, we get

$$\begin{aligned}
\tilde{H}_\infty(W_Z|Z, \text{SS}(W_Z)) &= \tilde{H}_\infty(\text{M}|\text{SS}(\text{M}), \text{K}, \text{Hash}_\text{K}(\text{M})) \\
&\leq \log \frac{1}{\Pr[\text{SS}(\text{M}) \in \text{GoodSketch}_\text{M}]} + \\
&\quad \tilde{H}_\infty(\text{M}|ss \text{ s.t. } ss = \text{SS}(\text{M}) \text{ and } ss \in \text{GoodSketch}_\text{M}, \text{K}, \text{Hash}_\text{K}(\text{M})) \\
&\leq \log \frac{1}{\Pr[\text{SS}(\text{M}) \in \text{GoodSketch}_\text{M}]} + \\
&\quad \tilde{H}_0(\text{M}|ss \text{ s.t. } ss = \text{SS}(\text{M}) \text{ and } ss \in \text{GoodSketch}_\text{M}, \text{K}, \text{Hash}_\text{K}(\text{M})) \\
&< \log \frac{1}{\Pr[\text{SS}(\text{M}) \in \text{GoodSketch}_\text{M}]} + \max \left(1, 1 + \frac{n - \log |B_t| + h_2(2\delta)}{1 - 2\delta} - a \right).
\end{aligned}$$

We can have shown that $\tilde{H}_\infty(W_Z|Z, \text{SS}(W_Z)) \leq 2$, because the first term of the above sum is at most 1 by Proposition 7.4 and the second term is 1 by our choice of a as $a = (n - \log |B_t| + h_2(2\delta))/(1 - 2\delta)$.

It remains to show that the desired hash family exists. Note in that (because $\delta < .25$) setting any $\alpha \geq 1 - h_2(\tau) + \frac{.5 \log n + 4\delta n + 2}{n}$ and choosing an αn -universal hash function will be sufficient, because, by Lemma 2.2, $\log |B_t| \geq n h_2(\tau) - \frac{1}{2} \log n - 1$, and so

$$\begin{aligned}
a &= \frac{n - \log |B_t| + h_2(2\delta)}{1 - 2\delta} \leq n \cdot \frac{1 - h_2(\tau) + (.5 \log n + 1 + h_2(2\delta))/n}{1 - 2\delta} \\
&< n \cdot \left(1 - h_2(\tau) + \frac{.5 \log n + 1 + h_2(2\delta)}{n} + 4\delta \right) \\
&\leq n \cdot \left(1 - h_2(\tau) + \frac{.5 \log n + 4\delta n + 2}{n} \right) \\
&\leq n \cdot \alpha
\end{aligned}$$

(the second inequality is true because for any $x < 1$ and $0 < y < .5$, $x/(1 - y) < x + 2y$, because $x < (x + 2y)(1 - y)$, because $0 < y(2 - x - 2y)$; the third inequality follows from $h_2(2\delta) < 1$). Such a hash family exists by Lemma 5.4 as long as $\mu \leq 1 - \alpha - 2/n \leq h_2(\tau) - (.5 \log n + 4\delta n + 4)/n$ and $\mu \leq 1 - h_2(\tau) - 2/n$. \square

8 Conclusion

This work introduces fuzzy min-entropy as a new metric for measuring the suitability of deriving keys from a noisy probability distribution. This condition is sufficient for security if the distribution is exactly known. This setting is comparable to the traditional setting when Eve receives no auxiliary information.

Our negative results show that providing security simultaneously for a family of sources is impossible for all distributions with fuzzy min-entropy. The core of all of these proofs is constructing a family of distributions \mathcal{W} where the description of the element $W \in \mathcal{W}$ provides the adversary with information independent of what is (necessarily) leaked by the fuzzy extractor (or secure sketch). Our three results require a careful tuning between the information leaked by the fuzzy extractor and the independent information in the description. This state of affairs seems somewhat bleak, however, there are several ways to avoid these negative results:

1. Focus on providing security for high entropy distributions only. However, many noisy distributions come from nature and system designers cannot effectively adjust their parameters,
2. Assume some additional structure from the distribution such as independence between dimensions [TG04] or that random subsets of dimensions have high entropy [CFP⁺16].
3. Restrict the adversary, for example, assuming the adversary runs in polynomial time. Recently, constructions have shown fuzzy extractors secure against bounded adversaries relying on hardness of discrete log [WLH18, WL18], decoding random codes [FMR13, HRvD⁺16], security of hash functions [CFP⁺16, ABC⁺18], and general cryptographic primitives [WLG19]. Fuller, Meng, and Reyzin showed that computationally-secure sketches imply information-theoretic ones, so computationally-secure sketches are subject to the negative results in this work [FMR13, Theorem 1]. A comparable theorem is not known for computationally-secure fuzzy extractors.

Acknowledgements

The authors are grateful to Gene Itkis and Yevgeniy Dodis for helpful discussions and to Thomas Holenstein for clarifying the results of [Hol06, HR05]. The work of Benjamin Fuller was sponsored in part by US NSF grants 1012910, 1012798, and 1849904 and the United States Air Force under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the authors and are not necessarily endorsed by the United States Government. Leonid Reyzin is supported in part by US NSF grants 0831281, 1012910, 1012798, and 1422965 and The Institute of Science and Technology, Austria, where part of this work was performed. Adam Smith's work was supported in part by NSF awards 0747294, 0941553 and 1447700 and was performed partly while at Boston University's Hariri Institute for Computing and RISCs Center, and the Harvard Center for Research on Computation & Society.

References

- [ABC⁺18] Quentin Alamérou, Paul-Edmond Berthier, Chloé Cachet, Stéphane Cauchie, Benjamin Fuller, Philippe Gaborit, and Sailesh Simhadri. Pseudoentropic isometries: a new framework for fuzzy extractor reusability. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, pages 673–684. ACM, 2018.
- [AC93] Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography - I: secret sharing. *IEEE Transactions on Information Theory*, 39(4):1121–1132, 1993.
- [Ash65] Robert Ash. *Information Theory*. Interscience Publishers, 1965.
- [BBR88] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. *SIAM Journal on Computing*, 17(2):210–229, 1988.
- [BCKP14] Nir Bitansky, Ran Canetti, Yael Tauman Kalai, and Omer Paneth. On virtual grey box obfuscation for general circuits. In *Advances in Cryptology - CRYPTO 2014*, 2014.
- [BCKP17] Nir Bitansky, Ran Canetti, Yael Tauman Kalai, and Omer Paneth. On virtual grey box obfuscation for general circuits. *Algorithmica*, 79(4):1014–1051, 2017.
- [BCL⁺11] Boaz Barak, Ran Canetti, Yehuda Lindell, Rafael Pass, and Tal Rabin. Secure computation without authentication. *J. Cryptology*, 24(4):720–760, 2011.
- [BDK⁺05] Xavier Boyen, Yevgeniy Dodis, Jonathan Katz, Rafail Ostrovsky, and Adam Smith. Secure remote authentication using biometric data. In *EUROCRYPT*, pages 147–163. Springer, 2005.
- [BH09] Marina Blanton and William MP Hudelson. Biometric-based non-transferable anonymous credentials. In *Information and Communications Security*, pages 165–180. Springer, 2009.
- [BS00] Sacha Brostoff and M. Angela Sasse. Are passfaces more usable than passwords?: A field trial investigation. *People and Computers*, pages 405–424, 2000.
- [CFP⁺16] Ran Canetti, Benjamin Fuller, Omer Paneth, Leonid Reyzin, and Adam Smith. Reusable fuzzy extractors for low-entropy distributions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 117–146. Springer, 2016.
- [CK78] Imre Csiszár and János Körner. Broadcast channels with confidential messages. *IEEE Transactions on Information Theory*, 24(3):339–348, 1978.
- [CW79] Larry Carter and Mark N. Wegman. Universal classes of hash functions. *J. Comput. Syst. Sci.*, 18(2):143–154, 1979.
- [Dau04] John Daugman. How iris recognition works. *Circuits and Systems for Video Technology, IEEE Transactions on*, 14(1):21 – 30, January 2004.
- [Dau06] J. Daugman. Probing the uniqueness and randomness of iriscodes: Results from 200 billion iris pair comparisons. *Proceedings of the IEEE*, 94(11):1927–1935, 2006.

- [DHP⁺18] Pierre-Alain Dupont, Julia Hesse, David Pointcheval, Leonid Reyzin, and Sophia Yakoubov. Fuzzy password-authenticated key exchange. In *Advances in Cryptology – EUROCRYPT*, pages 393–424. Springer, 2018.
- [DORS08] 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.
- [EHMS00] Carl Ellison, Chris Hall, Randy Milbert, and Bruce Schneier. Protecting secret keys with personal entropy. *Future Generation Computer Systems*, 16(4):311–318, 2000.
- [Fan61] R.M. Fano. *Transmission of Information: A Statistical Theory of Communications*. MIT Press Classics. M.I.T. Press, 1961.
- [FF81] Peter Frankl and Zoltán Füredi. A short proof for a theorem of Harper about Hamming-spheres. *Discrete Mathematics*, 34(3):311–313, 1981.
- [FMR13] Benjamin Fuller, Xianrui Meng, and Leonid Reyzin. Computational fuzzy extractors. In *Advances in Cryptology-ASIACRYPT 2013*, pages 174–193. Springer, 2013.
- [FP19] Benjamin Fuller and Lowen Peng. When are continuous-source fuzzy extractors possible? In *IEEE International Symposium on Information Theory*, 2019.
- [GCVDD02] Blaise Gassend, Dwaine Clarke, Marten Van Dijk, and Srinivas Devadas. Silicon physical random functions. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 148–160. ACM, 2002.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *Proc. of FOCS*, 2013.
- [GGH⁺16] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM Journal on Computing*, 45(3):882–929, 2016.
- [HAD06] Feng Hao, Ross Anderson, and John Daugman. Combining crypto with biometrics effectively. *Computers, IEEE Transactions on*, 55(9):1081–1088, 2006.
- [Har66] Lawrence H Harper. Optimal numberings and isoperimetric problems on graphs. *Journal of Combinatorial Theory*, 1(3):385–393, 1966.
- [Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- [Hol06] Thomas Holenstein. *Strengthening Key Agreement using Hard-Core Sets*. PhD thesis, ETH Zurich, May 2006. Reprint as vol. 7 of *ETH Series in Information Security and Cryptography*, ISBN 3-86626-088-2, Hartung-Gorre Verlag, Konstanz, 2006.
- [HR05] Thomas Holenstein and Renato Renner. One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In Victor Shoup, editor, *Advances in Cryptology - CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 478–493. Springer, 2005.

- [HRvD⁺16] Charles Herder, Ling Ren, Marten van Dijk, Meng-Day Yu, and Srinivas Devadas. Trapdoor computational fuzzy extractors and stateless cryptographically-secure physical unclonable functions. *IEEE Transactions on Dependable and Secure Computing*, 14(1):65–82, 2016.
- [HTW14] Masahito Hayashi, Himanshu Tyagi, and Shun Watanabe. Secret key agreement: General capacity and second-order asymptotics. In *2014 IEEE International Symposium on Information Theory*, pages 1136–1140. IEEE, 2014.
- [IW12] Tanya Ignatenko and Frans M.J. Willems. Biometric security from an information-theoretical perspective. *Foundations and Trends in Communications and Information Theory*, 7(2–3):135–316, 2012.
- [JW99] Ari Juels and Martin Wattenberg. A fuzzy commitment scheme. In *Sixth ACM Conference on Computer and Communication Security*, pages 28–36. ACM, November 1999.
- [LA18] Cheuk Ting Li and Venkat Anantharam. One-shot variable-length secret key agreement approaching mutual information. *CoRR*, abs/1809.01793, 2018.
- [LT03] J.-P. M. G. Linnartz and P. Tuyls. New shielding functions to enhance privacy and prevent misuse of biometric templates. In *International Conference on Audio-and Video-Based Biometric Person Authentication*, pages 393–402, 2003.
- [Mau93] Ueli M. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39(3):733–742, 1993.
- [MG09] Rene Mayrhofer and Hans Gellersen. Shake well before use: Intuitive and secure pairing of mobile devices. *IEEE Transactions on Mobile Computing*, 8(6):792–806, 2009.
- [MRW02] Fabian Monroe, Michael K Reiter, and Susanne Wetzels. Password hardening based on keystroke dynamics. *International Journal of Information Security*, 1(2):69–83, 2002.
- [NZ93] Noam Nisan and David Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, pages 43–52, 1993.
- [PRTG02] Ravikanth Pappu, Ben Recht, Jason Taylor, and Neil Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026–2030, 2002.
- [RW04] Renato Renner and Stefan Wolf. The exact price for unconditionally secure asymmetric cryptography. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 109–125. Springer, 2004.
- [RW05] Renato Renner and Stefan Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In Bimal K. Roy, editor, *Advances in Cryptology - ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 199–216. Springer, 2005.
- [SD07] G. Edward Suh and Srinivas Devadas. Physical unclonable functions for device authentication and secret key generation. In *Proceedings of the 44th annual Design Automation Conference*, pages 9–14. ACM, 2007.

- [ST09] B. Skoric and P. Tuyls. An efficient fuzzy extractor for limited noise. Cryptology ePrint Archive, Report 2009/030, 2009.
- [TG04] Pim Tuyls and Jasper Goseling. Capacity and examples of template-protecting biometric authentication systems. In *Biometric Authentication*, volume 3087 of *Lecture Notes in Computer Science*, pages 158–170. Springer, 2004.
- [TSS⁺06] Pim Tuyls, Geert Jan Schrijen, Boris Skoric, Jan van Geloven, Nynke Verhaegh, and Rob Wolters. Read-proof hardware from protective coatings. In Louis Goubin and Mitsuru Matsui, editors, *Cryptographic Hardware and Embedded Systems - CHES 2006*, volume 4249 of *Lecture Notes in Computer Science*, pages 369–383. Springer, 2006.
- [TVW18] Himanshu Tyagi, Pramod Viswanath, and Shun Watanabe. Interactive communication for data exchange. *IEEE Trans. Information Theory*, 64(1):26–37, 2018.
- [TW15] Himanshu Tyagi and Shun Watanabe. Converses for secret key agreement and secure computing. *IEEE Transactions on Information Theory*, 61(9):4809–4827, 2015.
- [TW17] Himanshu Tyagi and Shun Watanabe. Universal multiparty data exchange and secret key agreement. *IEEE Transactions on Information Theory*, 63(7):4057–4074, 2017.
- [WCD⁺17] Joanne Woodage, Rahul Chatterjee, Yevgeniy Dodis, Ari Juels, and Thomas Ristenpart. A new distribution-sensitive secure sketch and popularity-proportional hashing. In *Annual International Cryptology Conference*, pages 682–710. Springer, 2017.
- [WL18] Yunhua Wen and Shengli Liu. Robustly reusable fuzzy extractor from standard assumptions. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 459–489. Springer, 2018.
- [WLG19] Yunhua Wen, Shengli Liu, and Dawu Gu. Generic constructions of robustly reusable fuzzy extractor. In *IACR International Workshop on Public Key Cryptography*, pages 349–378. Springer, 2019.
- [WLH18] Yunhua Wen, Shengli Liu, and Shuai Han. Reusable fuzzy extractor from the decisional diffie–hellman assumption. *Designs, Codes and Cryptography*, 86(11):2495–2512, 2018.
- [WRDI11] Ye Wang, Shantanu Rane, Stark C. Draper, and Prakash Ishwar. A theoretical analysis of authentication, privacy and reusability across secure biometric systems. *ArXiv CoRR report*, 1112.5630, 2011.
- [Wyn75] Aaron D Wyner. The wire-tap channel. *Bell System Technical Journal*, The, 54(8):1355–1387, 1975.
- [YY18] Kenji Yasunaga and Kosuke Yuzawa. On the limitations of computational fuzzy extractors. *work*, 3(1):7–9, 2018.
- [ZH93] Moshe Zviran and William J. Haga. A comparison of password techniques for multilevel authentication mechanisms. *The Computer Journal*, 36(3):227–237, 1993.