

A Counterexample to the Chain Rule for Conditional HILL Entropy*

Stephan Krenn
IBM Research – Zurich
Säumerstrasse 4
8803 Rüschlikon
Switzerland
skr@ibm.zurich.com

Krzysztof Pietrzak[†]
IST Austria
Am Campus 1
3400 Klosterneuburg
Austria
pietrzak@ist.ac.at

Akshay Wadia
University of California
3771, Boelter Hall
Los Angeles, CA 90095
United States
awadia@cs.ucla.edu

Daniel Wichs
Northeastern University
360 Huntington Av., #202 WVH
Boston, MA 02115
United States
wichs@ccs.neu.edu

August 30, 2014

Abstract

Most entropy notions $H(\cdot)$ like Shannon or min-entropy satisfy a chain rule stating that for random variables X, Z and A we have $H(X|Z, A) \geq H(X|Z) - |A|$. That is, by conditioning on A the entropy of X can decrease by at most the bitlength $|A|$ of A . Such chain rules are known to hold for some *computational* entropy notions like Yao’s and unpredictability-entropy. For HILL entropy, the computational analogue of min-entropy, the chain rule is of special interest and has found many applications, including leakage-resilient cryptography, deterministic encryption and memory delegation. These applications rely on restricted special cases of the chain rule. Whether the chain rule for conditional HILL entropy holds in general was an open problem for which we give a strong negative answer: We construct joint distributions (X, Z, A) , where A is a distribution over a *single* bit, such that the HILL entropy $H^{\text{HILL}}(X|Z)$ is large but $H^{\text{HILL}}(X|Z, A)$ is basically zero.

Our counterexample just makes the minimal assumption that $\mathbf{NP} \not\subseteq \mathbf{P/poly}$. Under the stronger assumption that injective one-way function exist, we can make all the distributions efficiently samplable.

Finally, we show that some more sophisticated cryptographic objects like lossy functions can be used to sample a distribution constituting a counterexample to the chain rule making only a single invocation to the underlying object.

*This is the full version of [KPW13] that appeared in TCC 2013.

[†]This work was partly funded by the European Research Council under ERC Starting Grant 259668-*PSPC* and ERC Advanced Grant 321310-*PERCY*.

1 Introduction

Various information theoretic entropy notions are used to quantify the amount of randomness of a probability distribution. The most common is Shannon entropy, which bounds how well a distribution can be compressed.

Definition 1. *The **Shannon entropy** of random variable X with support \mathcal{X} is $H_1(X) = -\sum_{x \in \mathcal{X}} \Pr[X = x] \cdot \log_2 \Pr[X = x]$.*

Another important notion, especially in cryptographic and other computational settings, is min-entropy:

Definition 2. *The **min-entropy** of a random variable X with support \mathcal{X} is $H_\infty(X) = -\log_2 \max_{x \in \mathcal{X}} \Pr[X = x]$.*

$H_\infty(X)$ upper bounds the probability of X taking any particular value (or equivalently, the advantage of any algorithm A in guessing X) as

$$\max_x \Pr[X = x] = \max_A \Pr[A = x] = 2^{-H_\infty(X)} \quad (1)$$

For any x , $\Pr[X = x] \leq 2^{-H_\infty(X)}$. Shannon and min-entropy correspond to Rényi entropy for the limiting values of $\alpha \rightarrow 1$ and $\alpha \rightarrow \infty$, respectively.

Definition 3. *The **Rényi entropy** of order α , where $\alpha \geq 0$ and $\alpha \neq 1$, is defined as $H_\alpha(X) = \frac{1}{1-\alpha} \log(\sum_{i=1}^n p_i^\alpha)$.*

1.1 Chain Rules

One of the most useful tools for manipulating and arguing about entropies are chain rules, which come in many different flavors for different entropy notions. For Shannon entropy, we have the following simple chain rule for possibly dependent random variables X, A :

$$H_1(X|A) = H_1(X, A) - H_1(A). \quad (2)$$

Using $H_1(X, A) \geq H_1(X)$ and $H_1(A) \leq |A|$ (where $|A|$ denotes the bitlength of A) this implies (with $H = H_1$)

$$H(X|A) \geq H(X) - |A|. \quad (3)$$

Although this is a weaker statement than (2), it already captures the fact that the entropy of X decreases by at most the bitlength $|A|$ of A if we condition on A . As we will discuss below, such a chain rule not only holds for Shannon entropy, but many other information theoretic and computational entropy notions.

More generally, for many notions one can give chain rules for *conditional* entropies by considering the case where X has some entropy conditioned on Z , and bound by how much the entropy drops when additionally given A :

$$H(X|Z, A) \geq H(X|Z) - |A|. \quad (4)$$

[DORS08] define conditional min-entropy as follows:

Definition 4. For a pair (X, Z) of random variables, the **average min-entropy** of X conditioned on Z is

$$\begin{aligned}\tilde{H}_\infty(X|Z) &= -\log \mathbb{E}_{z \leftarrow Z} \max_x \Pr[X = x | Z = z] \\ &= -\log \mathbb{E}_{z \leftarrow Z} 2^{-H_\infty(X|Z=z)},\end{aligned}$$

They show that this notion satisfies a chain rule like in (4):

$$\tilde{H}_\infty(X|Z, A) \geq \tilde{H}_\infty(X|Z) - H_0(A) \geq \tilde{H}_\infty(X|Z) - |A|, \quad (5)$$

where $H_0(A) \leq |A|$ denotes the logarithm of the support-size of A . They also show that this notion naturally extends the property (1) of min-entropy as an upper bound on the guessing probability to the conditional case:

$$\max_A \Pr_{(x,z) \leftarrow (X,Z)}[A(z) = x] \leq 2^{-\tilde{H}_\infty(X|Z)}. \quad (6)$$

1.2 Computational Entropy

The standard information theoretic entropy notions refer to computationally unbounded parties, e.g., no algorithm can compress a distribution X (given Z) below its Shannon entropy $H(X|Z)$ and no algorithm can guess X (given Z) better than with probability $2^{-\tilde{H}_\infty(X|Z)}$. Under computational assumptions, in particular in cryptographic settings, one often has to deal with distribution that appear to have high entropy only for computationally bounded parties. The classical example is a pseudorandom distribution [BM84, Yao82], where $X \in \{0, 1\}^n$ is said to be pseudorandom if it cannot be distinguished from the uniform distribution over $\{0, 1\}^n$ by polynomial size distinguishers. In this case X appears to have n bits of Shannon and n bits of min-entropy.

Pseudorandomness is an elegant and tremendously useful notion, and very convenient to work with. Sometimes we do not have the luxury of dealing with a distribution that appears uniformly random, but only seems to have some kind of high entropy. Such distributions are referred to as having pseudoentropy. Stated informally, some prominent pseudoentropy notions are defined as:

HILL: X has k bits of HILL-pseudoentropy (conditioned on Z), if it cannot be distinguished from some variable Y with k bits of min-entropy (given Z) [HILL99, HLR07]. We'll define this notion precisely in Definition 7 below.

Computational Shannon: Defined like HILL-entropy, but with Shannon instead of min-entropy [HILL99, VZ12].

Unpredictability: X has k bits of unpredictability entropy conditioned on Z if no efficient adversary can guess X better than with probability 2^{-k} given Z [HLR07].¹ We'll define this notion precisely in Definition 14.

Yao: X has k bits of Yao entropy (conditioned on Z) if it cannot be *efficiently* compressed/decompressed below k bits (given Z) [Yao82, BSW03, HLR07]. We'll define this notion precisely in Definition 13.

¹Unpredictability entropy is only interesting if the conditional part Z is not empty, otherwise it coincides with min-entropy.

Note that HILL and computational Shannon entropy require “indistinguishability” from the corresponding information theoretic objects (i.e., distributions of certain min and Shannon entropy), whereas Unpredictability and Yao entropy only capture the “functional” unpredictability and incompressibility properties of the information theoretic notions.

When defining the computational entropy of a variable X , not only the *quantity* k is of interest, but also the *quality*, which specifies which class of distinguishers X fools and with which advantage. Below we formally define HILL entropy, as well as the more general *conditional* HILL entropy, but first need one more basic definition.

Definition 5 (Indistinguishability). *For two probability distributions X and Y and a Boolean circuit D , the advantage of D in distinguishing X and Y is defined as:*

$$\text{adv}_D(X, Y) = |\Pr_{x \leftarrow X}[D(x) = 1] - \Pr_{y \leftarrow Y}[D(y) = 1]|.$$

We say X and Y are (ε, s) -indistinguishable, denoted by $X \sim_{\varepsilon, s} Y$, if for every Boolean circuit D of size at most s it holds that

$$\text{adv}_D(X, Y) \leq \varepsilon.$$

*Two ensembles of distributions $X = \{X_n\}_{n \in \mathbb{N}}$ and $Y = \{Y_n\}_{n \in \mathbb{N}}$ are **indistinguishable**, denoted by $X \sim_c Y$, if for every polynomial $p = p(n)$ it holds that $X_n \sim_{1/p, p} Y_n$ for all sufficiently large n .*

With $X \sim_\delta Y$ we denote that the statistical distance of X and Y is at most δ , i.e., $X \sim_{\delta, \infty} Y$. $X \sim Y$ denotes that X and Y have the same distribution, i.e., $X \sim_0 Y$.

Definition 6 ([HILL99]). *A random variable X has **HILL entropy** k , denoted by $H_{\varepsilon, s}^{\text{HILL}}(X) \geq k$, if there exists a distribution Y satisfying $H_\infty(Y) \geq k$ and $X \sim_{\varepsilon, s} Y$.*

Definition 7 ([HLR07]). *Let (X, Z) be a joint distribution of random variables. Then X has **conditional HILL entropy** k conditioned on Z , denoted by $H_{\varepsilon, s}^{\text{HILL}}(X|Z) \geq k$, if there exists a joint distribution (Y, Z) such that $\tilde{H}_\infty(Y|Z) \geq k$, and $(X, Z) \sim_{\varepsilon, s} (Y, Z)$.²*

1.3 Model of Computation

In this paper we stick to a non-uniform model of computation when considering adversaries. In particular, note that in Definition 5 we defined indistinguishability in terms of circuits. On the other hand, when considering efficient cryptographic objects like one-way functions, we always assume that they are computable in a uniform model, i.e., by a single Turing machine running in time polynomial in its input length. We also define “efficient samplability” in a uniform sense

Definition 8. *An ensemble of distributions $\{X_n\}_{n \in \mathbb{N}}$ is **efficiently samplable** if there exists a probabilistic polynomial time Turing machine M such that $M(1^n) \sim X_n$ for all $n \in \mathbb{N}$, i.e., the output distribution of $M(1^n)$ is X_n .*

These choices are not crucial, and all the result can be adapted considering uniform adversaries and/or non-uniform cryptographic objects. We chose this particular model to get the strongest conclusion (i.e., a counterexample to the chain rule by an efficiently and uniformly samplable distribution), at the cost of a stronger, but still widely believed assumption, i.e., existence of one-way functions secure against non-uniform adversaries.

²Let us stress that using the same letter Z for the 2nd term in (X, Z) and (Y, Z) means that we require that the marginal distribution Z of (X, Z) and (Y, Z) is the same.

1.4 Chain Rules for Computational Entropy

Chain rules for conditional entropy as in (4) (and thus also for the special non-conditional case (3)) are easily seen to hold for some computational entropy notions. We give the chain rules for (conditional) Yao and unpredictability entropy in Appendix B. For HILL entropy a chain rule has been found independently by [RTTV08] and [DP08]. It is captured by the following

Lemma 1 (Chain Rule for HILL Entropy). *For any joint distribution (X, A) where $A \in \{0, 1\}^\ell$ we have that*

$$H_{\varepsilon', s'}^{\text{HILL}}(X|A) \geq H_{\varepsilon, s}^{\text{HILL}}(X) - \ell,$$

where $\varepsilon' \approx \varepsilon$ and $s' \approx s\varepsilon^2/2^{3\ell}$.³

Note that in Lemma 1, conditioning on $A \in \{0, 1\}^\ell$ not only decreases the quantity of entropy by ℓ , but also the quality goes down polynomially in ε^{-1} and 2^ℓ , and this is unavoidable.⁴ Another important point is that the lemma is only stated for the *non-conditional* case, i.e., it is of the form (3), not the more general (4). Whether or not a chain rule for *conditional* HILL entropy holds “remains an interesting open problem” [FOR12].⁵ Concretely, we would like a statement of the form

$$H_{\varepsilon', s'}^{\text{HILL}}(X|Z, A) \geq H_{\varepsilon, s}^{\text{HILL}}(X|Z) - \ell \quad (7)$$

where $\varepsilon' = \varepsilon \cdot p(2^\ell, \varepsilon^{-1})$ and $s' = s/q(2^\ell, \varepsilon^{-1})$, for some polynomial functions $p(\cdot)$ and $q(\cdot)$.

2 Our Contribution

2.1 The Counterexamples

In this paper we give a strong negative answer to the open problem outlined above, showing that (7) does not hold in general. Condition (ii) in the Theorem 1 below follows from condition (iii) by Lemma 3. We state this (redundant) condition (ii) explicitly because from (iii) it’s not obviously clear how the chain rule is contradicted.

Theorem 1 (Main). *a. If injective one-way functions secure against non-uniform adversaries exist, then there exists an ensemble of efficiently samplable (cf. Definition 8) joint distributions $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$, where A_n is a single bit, that satisfy the following conditions*

- (i) X_n has a lot of high quality HILL entropy conditioned on Z_n : For any polynomial $p(\cdot)$, there exists an n_0 s.t. for all $n \geq n_0$

$$H_{1/p(n), p(n)}^{\text{HILL}}(X_n|Z_n) \geq n$$

³The quantitative bounds for ε', s' in the statement of this lemma are from [JP14]. The result from [VZ13] can be used to get a better $s' \approx s\varepsilon^2/2^\ell$ bound (and the same $\varepsilon' \approx \varepsilon$), whenever s is large enough as a function of $1/\varepsilon$ and 2^ℓ , concretely, this bound holds if $s = \Omega(2^\ell/\varepsilon^4)$.

⁴Concretely, we know that a loss of $\max\{2^\ell, \varepsilon^{-1}\}$ in either the distinguishing advantage or the circuit size is unavoidable [TTV09].

⁵We will discuss some other restricted settings in which the chain rule for HILL entropy holds in Section 3.

(ii) The HILL entropy drops to basically 0 if we additionally condition on the single bit A_n , even if we simultaneously allow for massive degradation in the quality: For some fixed polynomial $s(\cdot)$ it holds that

$$H_{1/2, s(n)}^{\text{HILL}}(X_n | Z_n, A_n) \leq 1.$$

(iii) There exists a polynomial time Turing machine \mathbf{M} that perfectly recognizes X_n given (Z_n, A_n) , i.e., for all $n \in \mathbb{N}$ and $(x, z, a) \in \text{supp}[(X_n, Z_n, A_n)]$ and every x'

$$(\mathbf{M}(x', z, a) = 1) \iff (x' = x).$$

b. Under the weaker assumption that $\mathbf{NP} \not\subseteq \mathbf{P/poly}$, such an ensemble exists but is not necessarily efficiently samplable and condition (i) only holds for infinitely many (not all sufficiently large) $n \in \mathbb{N}$.

In condition (ii) of Theorem 1, when conditioning on A_n , we “only” get the HILL entropy down to 1, not all the way down to 0. For the large distinguishing advantage $\varepsilon = 1/2$ we consider this is optimal, as no variable can get below one bit of HILL entropy by Lemma 2 below.

Lemma 2. For any joint distribution (V, C) over $\mathcal{V} \times \mathcal{C}$ (where $|\mathcal{V}| \geq 2$) and any $s \in \mathbb{N}$

$$H_{1/2, s}^{\text{HILL}}(V|C) \geq H_{1/2, \infty}^{\text{HILL}}(V|C) \geq 1$$

The proof of Lemma 2 is in Section 7.4. Lemma 2 can be generalized to show that for any real-valued $\tau \geq 0$ where $|\mathcal{V}| \geq 2^\tau$ it holds that $H_{1-2^{-\tau}, \infty}^{\text{HILL}}(V|C) \geq \tau$, but in order to keep the number of parameters low, we’ll only use the $\tau = 1$ case in this paper. The lemma below states that the lower bound in Lemma 2 is tight whenever V can be efficiently recognized given C .

Lemma 3. For any joint distribution (V, C) , if there exists a circuit \mathbf{D} of size s that perfectly recognizes V given C , i.e.,

$$\forall (v, c) \in \text{supp}[(V, C)] : (\mathbf{D}(v', c) = 1) \iff (v' = v)$$

then $H_{1/2, s}^{\text{HILL}}(V|C) \leq 1$.

The proof of Lemma 3 is in Section 7.5.

2.2 On the Necessity of the Assumptions Used in Theorem 1

Assuming one-way functions in Theorem 1.a and $\mathbf{NP} \not\subseteq \mathbf{P/poly}$ in Theorem b is necessary, as stated in the propositions below, whose proofs can be found in Appendix 7.1.

Proposition 1. If there exists an ensemble of distributions as in the conclusions of Theorem 1.a, then there exist (not necessarily injective) one-way functions.

This holds even when just assuming condition (i) of Theorem 1.a and $H_\infty(X_n | Z_n, A_n) = 0$, which, by Lemma 4 below, is implied by condition (ii) of Theorem 1.a.

Lemma 4. For any joint distribution (V, C) over $\mathcal{V} \times \mathcal{C}$ where $|\mathcal{V}| > 2$, and any s

$$H_{1/2, s}^{\text{HILL}}(V|C) \leq 1 \implies H_{1/2, \infty}^{\text{HILL}}(V|C) \leq 1 \implies H_\infty(V|C) = 0.$$

The proof of Lemma 4 is in Section 7.4.

Proposition 2. *If there exists an ensemble of distributions as in the conclusion of Theorem b, then $\mathbf{NP} \not\subseteq \mathbf{P}/\mathbf{poly}$.*⁶

2.3 Proof Outline

The proof of Theorem 1 is given in Section 5, below we give a short outline.

For the rest of this section we fix some n (used to index the distributions $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$) and will omit the subscripts n . We construct distributions as in Theorem 1.a from any perfectly binding (bit) commitment scheme $\mathbf{com} : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ as in Definition 9 and note that such a scheme exists iff injective one-way functions exist. Concretely, the distribution (X, Z, A) constituting a counterexample to the chain rule is sampled using \mathbf{com} as follows:

- Sample the bit $A \leftarrow \{0, 1\}$ at random.
- Compute commitments $C^i \leftarrow \mathbf{com}(B^i, R^i)$ for $i = 1, \dots, 3n$, where the first $2n$ are commitments to A (i.e., $B^i = A$ for $i = 1 \dots 2n$) and the last n are commitments to $1 - A$.
- Let $S^i = (C^i, R^i, B^i)$ and output the sorted list of the C^i 's as Z and the sorted list of the n last S^i 's as X , that is

$$X = \text{sort}(S^{2n+1}, \dots, S^{3n}) \quad Z = \text{sort}(C^1, \dots, C^{3n})$$

So, X is the “opening information” for the n commitments to $1 - A$ which are “hidden” amongst the $3n$ commitments in Z . As required by Theorem 1.a.iii one can efficiently distinguish X from any $X' \neq X$ given (Z, A) . To see this, note that X is a list of n sorted tuples $(C, R, 1 - A)$ where each such tuple satisfies $C = \mathbf{com}(1 - A, R)$ for some $C \in Z$. As \mathbf{com} is perfectly binding, X is sorted and Z contains exactly n commitments to $1 - A$, it follows that there is exactly one list (namely X) satisfying all these conditions, and these conditions can all be efficiently checked given (Z, A) .

On the other hand we claim that, as required Theorem 1.a.i, X has at least n bits of HILL entropy given only Z but not A . To prove this we consider a random variable Y which is defined by picking a random n element subset of (S^1, \dots, S^{2n}) and outputting it in lexicographic order. This Y has min-entropy $\tilde{H}_\infty(Y|Z) \geq n$ since it is uniform over exponentially $\binom{2n}{n} \geq 2^n$ many possible subsets. Moreover (Y, Z) is computationally indistinguishable from (X, Z) . We will prove this by showing how a distinguisher for these distributions can be used to break the hiding property of \mathbf{com} . By Definition 7 the existence of such a (Y, Z) means that X has n bits of HILL entropy conditioned on Z as claimed.

The construction of a distribution (X, Z, A) as claimed in Theorem b assuming only $\mathbf{NP} \not\subseteq \mathbf{P}/\mathbf{poly}$ is very similar. The only difference is that instead of using a perfectly binding commitment scheme to construct a distribution (C, R, B) as above, we only assume that there exists a distribution that has some, but not all, of the properties of

⁶Unlike in Proposition 1, here we need to assume condition (iii). We can relax condition (iii) to the non-uniform setting, asking for a family of poly size circuits (instead the Turing machine M) to recognize X_n given Z_n, A_n , and this would still imply $\mathbf{NP}/\mathbf{poly} \neq \mathbf{P}/\mathbf{poly}$.

(C, R, B) . In particular, we don't require the distribution to be efficiently samplable, and we need B to be pseudorandom given C only for infinitely many (not all sufficiently large) $n \in \mathbb{N}$. We'll call an ensemble of such distributions "committing" (cf. Definition 10). We show that such an ensemble exists assuming $\mathbf{NP} \not\subseteq \mathbf{P/poly}$ in Section 5.3.

2.4 Efficient Counterexample from Lossy Functions and Deniable Encryption

To sample a distribution that constitutes a counterexample to the chain rule using injective one-way functions as in Theorem 1.a requires a linear (in the entropy gap n) number of invocations to the OWF.

In Section 6 we show how that using more sophisticated cryptographic objects, one can sample such distributions much more efficiently. In particular, we construct an ensemble of efficiently samplable distributions $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ (where $A_n \in \{0, 1\}$) making just two invocations to an ℓ -lossy function, where for any polynomial $p = p(n)$ and some fixed polynomial $s = s(n)$, for all sufficiently large n it holds that

$$H_{1/p,p}^{\text{HILL}}(X_n|Z_n) \geq \ell \quad \text{but} \quad H_{1/2,s}^{\text{HILL}}(X_n|Z_n, A_n) \leq 1. \quad (8)$$

In Appendix A we construct an ensemble $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ from any deniable encryption scheme where

$$H_{1/p,p}^{\text{HILL}}(X_n|Z_n) - H_{1/2,s}^{\text{HILL}}(X_n|Z_n, A_n) \in \omega(\log n). \quad (9)$$

Note that the counterexample to the chain rule in (9) is weaker than the distributions constructed in Theorem 1.a or from lossy functions (8), because in (9) we do not necessarily get the HILL entropy all the way down to 1 after conditioning on A_n , we just get a super-logarithmic decrease in HILL entropy. Concretely, if $\varepsilon(n)$ denotes the (negligible) correctness error of the encryption scheme, the entropy gap is $\log(1/\varepsilon(n)) - 2 = \omega(\log n)$.

This indicates that cryptographic objects achieving some kind of deniability or lossiness must embed an *efficient* counterexample to the chain rule for HILL entropy. Thus, distributions which constitute such a counterexample seem to be a useful cryptographic resource.⁷

3 Related Work

HILL entropy was introduced by [HILL99], and the conditional variant was suggested by [HLR07]. Other notions of computational entropy include Yao entropy [Yao82, BSW03], unpredictability entropy [HLR07], and metric entropy [BSW03].

Chain rules for many of these entropy notions are known. Although in this work we show that the chain rule for conditional HILL entropy does not hold in general, it does hold in some interesting and useful restricted cases. We already discussed that a chain rule for *non-conditional* HILL entropy holds as stated in Lemma 1. This rule has applications in leakage-resilient cryptography [DP08] and deterministic encryption [FOR12].

⁷This is analogous to the fact that distributions having high HILL entropy can be constructed from one-way functions, but only via inefficient constructions. Thus, despite being "equivalent", there's a significant quantitative difference between one-wayness and HILL entropy. In the same vein, we can construct a distribution constituting a counterexample to the chain rule for HILL entropy from any distribution with high HILL entropy (and thus even from OWF), but the reduction seems to require a substantial number of samples from the underlying distribution.

[CKLR11] prove a chain rule for conditional *samplable* HILL entropy, a variant of conditional HILL entropy one gets when changing Definition 7 by additionally requiring Y to be efficiently samplable given Z . They use this rule to construct “memory delegation” schemes. [FOR12] give a chain rule for *decomposable* HILL entropy,⁸ where one requires X to have high HILL entropy conditioned on any particular conditional part $Z = z$. [Sko13] introduces modular entropy, and gives a unified treatment of the chain rules for decomposable and samplable HILL entropy in terms of this new notion. [Rey11, Theorem 2 and thereafter] gives a chain rule for conditional *relaxed* HILL entropy. Here “relaxed” does not refer to the notion of HILL entropy, but rather to the type of the chain rule: We get the notion of relaxed HILL entropy by replacing (Y, Z) in Definition 7 with (Y, Z') , i.e., one does not require the marginal distribution of the conditional part Z' to be the same as in the original distribution (X, Z) . Such a rule is already implicit in the work of [GW11], who use it to prove a black-box separation.

4 Notation and Basic Definitions

We use the standard “big O” notation: $f(n) \in \mathcal{O}(g(n))$ if there exists a $c > 0$ and $n_0 \in \mathbb{N}$ s.t. $f(n) \leq c \cdot g(n)$ for all $n \geq n_0$. $f(n) \in \omega(g(n))$ if for every $c > 0$ there exists $n_0 \in \mathbb{N}$ s.t. $f(n) > c \cdot g(n)$ for all $n \geq n_0$. All logarithms are to base 2. For a set \mathcal{S} , we denote by $|\mathcal{S}|$ its cardinality. For a bitstring x , $|x|$ denotes its length and for a circuit \mathbf{A} , $|\mathbf{A}|$ denotes its size. The *support* of a distribution X is $\text{supp}[X] = \{x \mid \Pr[X = x] > 0\}$. By $x \leftarrow X$ we denote that x is assigned a value according to the distribution X . For a set \mathcal{X} , by $x \leftarrow \mathcal{X}$ we denote that x is drawn uniformly at random from \mathcal{X} . For a probabilistic algorithm \mathbf{A} , we denote with $x \leftarrow \mathbf{A}$ that x is assigned the output of \mathbf{A} using fresh random coins. For an integer m , we define $[m] = \{1, \dots, m\}$.

5 Proof of Theorem 1

In Section 5.1, we construct distributions $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ as claimed in Theorem 1.a using $3n$ invocations of a commitment scheme as defined in Definition 9 below. Such a commitment scheme can be based on any injective one-way function as we’ll discuss in Remark 1.

In Definition 10 we define “committing distributions”, which can be seen as a relaxation of commitment schemes, where we drop efficient samplability and only require the hiding property to hold for infinitely many (not all sufficiently large) $n \in \mathbb{N}$. In Section 5.2 we construct distributions as claimed in Theorem b using committing distributions, and in Section 5.3 we show that committing distributions exist if $\mathbf{NP} \not\subseteq \mathbf{P/poly}$.

Definition 9. *An efficiently computable function $\text{com} : \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}$ is a perfectly binding bit-commitment scheme if*

perfect binding: *for all $b \in \{0, 1\}, r \in \{0, 1\}^*$*

$$\text{com}(b, r) = \text{com}(b', r') \quad \Rightarrow \quad (b, r) = (b', r')$$

⁸Actually, they state their result in terms of metric entropy, a notion weaker than HILL. It implies HILL entropy, but the reduction going from metric to HILL losses quite a bit in terms of quality as shown by [BSW03].

computational hiding: for the uniform distribution $R_n \leftarrow \{0, 1\}^n$

$$\{\text{com}(0, R_n)\}_{n \in \mathbb{N}} \sim_c \{\text{com}(1, R_n)\}_{n \in \mathbb{N}}$$

equivalently, for random $B_n \leftarrow \{0, 1\}$ and $C_n = \text{com}(B_n, R_n)$

$$\{(B_n, C_n)\}_{n \in \mathbb{N}} \sim_c \{(1 - B_n, C_n)\}_{n \in \mathbb{N}}$$

Remark 1 (Constructing com from (injective) OWF). *The binding property we require in Definition 9 is stronger than the standard definition of perfect binding, where one only requires $\text{com}(b, r) = \text{com}(b', r')$ to imply $b = b'$ (and not $(b, r) = (b', r')$). The standard construction⁹ of commitment schemes from injective one-way functions satisfies this stronger notion [Gol00, Section 4.4.1.2].*

[Nao91] shows how to construct a perfectly binding bit-commitment scheme from any (not necessarily injective) OWF,¹⁰ unfortunately his construction only can be shown to achieve the standard definition of perfect binding, not the stronger we require.

Definition 10. An ensemble of joint distributions $\{(C_n, R_n, B_n)\}_{n \in \mathbb{N}}$ is **committing** if

1. For every polynomial $p(\cdot)$, the following holds for infinitely many $n \in \mathbb{N}$:

$$(B_n, C_n) \sim_{1/p(n), p(n)} (1 - B_n, C_n)$$

2. There exists an efficiently uniformly computable predicate ϕ such that for all $n \in \mathbb{N}$ and any $(c, r, b) \in \text{supp}[(C_n, R_n, B_n)]$

$$(\phi(c, r', b') = 1) \iff ((r', b') = (r, b))$$

A committing distribution (C, R, B) as in Definition 10 can be constructed from a commitment scheme as in Definition 9 by simply defining $(C, R, B) = (\text{com}(B, R), R, B)$ and $(\phi(C, R, B) = 1) \iff (\text{com}(B, R) = C)$.

5.1 Counterexample from Commitment Scheme

We now define distributions $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ for which we'll prove they satisfy the conditions claimed in Theorem 1.a. For any $n \in \mathbb{N}$, the last element A_n of the tuple (X_n, Z_n, A_n) is a uniformly random bit. For $i = 1, \dots, 3n$ define

$$B^i = \begin{cases} A_n & \text{for } 1 \leq i \leq 2n \\ 1 - A_n & \text{for } 2n + 1 \leq i \leq 3n \end{cases} \quad (10)$$

Let R^1, \dots, R^{3n} be uniform over $\{0, 1\}^n$ and

$$C^i = \text{com}(B^i, R^i) \quad \text{and} \quad S^i = (C^i, R^i, B^i).$$

The remaining two elements X_n and Z_n are now defined as

$$X_n = \text{sort}(S^{2n+1}, \dots, S^{3n}) \quad \text{and} \quad Z_n = \text{sort}(C^1, \dots, C^{3n})$$

⁹Using an injective OWF f , the commitment $\text{com}(b, r = (s, x))$ is defined as $(s, f(x), b \oplus \langle x, s \rangle)$, where $\langle x, s \rangle$ is the inner product.

¹⁰Technically, his scheme is only perfectly binding with overwhelming probability over the choice of some string used to specify the construction, this would not be an issue for us.

where `sort` outputs the given input sequence in lexicographic order. Finally, we define a distribution (\tilde{X}_n, Z_n) (to be used only in the proof) where Z_n is as above, and \tilde{X}_n is sampled by choosing an n element subset $\{i_1, \dots, i_n\}$ of $\{1, \dots, 2n\}$ at random and setting

$$\tilde{X}_n = \text{sort}(S^{i_1}, \dots, S^{i_n}) \quad (11)$$

By the following lemma, these distributions satisfy the conditions of Theorem 1.a.

Lemma 5. *If `com` used to define $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ above is a commitment scheme as in Definition 9, then this ensemble is efficiently samplable and the following two conditions hold:*

1. For any polynomial $p = p(n)$

$$H_{1/p,p}^{\text{HILL}}(X_n | Z_n) \geq n \text{ for all sufficiently large } n \in \mathbb{N} .$$

2. There exists a polynomial time Turing machine \mathbf{M} such that for all $(x, z, a) \in \text{supp}[(X_n, Z_n, A_n)]$ we have

$$(\mathbf{M}(x', z, a) = 1) \iff (x' = x)$$

Proof. Efficient samplability of (X_n, Z_n, A_n) follows since `com` can be efficiently computed. We'll omit the subscript n for the rest of the proof.

Condition 2 holds as for any $(x, z, a) \in \text{supp}[(X, Z, A)]$, given (z, a) and some $x' = (s^1, \dots, s^n)$ where $s^i = (c^i, r^i, b^i)$, one can efficiently check if $x' = x$ by verifying that

- the s^1, \dots, s^n are in lexicographic order.
- x' contains openings to n commitments from z to the bit $1 - a$, i.e., for every $i \in [n]$ it holds that c^i is contained in z and $c^i = \text{com}(1 - a, r^i)$.

As `com`(., .) is injective and z contains exactly n commitments to $1 - a$, there's only one sequence, namely $x' = x$, that satisfies the above conditions. To prove condition 1 we'll need the following

Claim 1. $(X, Z) \sim_c (\tilde{X}, Z)$ with (\tilde{X}, Z) as in (11).

of Claim. Consider hybrid distributions

$$H_{(0)} = (X, Z_{(0)}), \dots, H_{(n)} = (X, Z_{(n)})$$

where

$$Z_{(0)} = Z = \text{sort}(C^1, \dots, C^{3n})$$

and for $0 < i \leq n$

$$Z_{(i)} = \text{sort}(\bar{C}^1, \dots, \bar{C}^i, C^{i+1}, \dots, C^{3n}) \quad (12)$$

where

$$\bar{C}^j = \text{com}(1 - A, R) \text{ for a random } R \leftarrow \{0, 1\}^n . \quad (13)$$

So, $Z_{(j)}$ is derived from $Z_{(j-1)}$ by replacing the commitment C^j (to the bit A) with a new commitment \bar{C}^j (to the opposite bit $1 - A$). To prove the claim we show that

$$(X, Z_{(0)}) \sim_c (X, Z_{(n)}) \sim (\tilde{X}, Z_{(0)}) \quad (14)$$

To see $(X, Z_{\langle n \rangle}) \sim (\tilde{X}, Z_{\langle 0 \rangle})$, we first observe it holds for the marginal distribution $Z_{\langle n \rangle} \sim Z_{\langle 0 \rangle}$: Both consist of $3n$ random commitments in lexicographic order, of which exactly $2n$ open to the same random bit. Moreover $(X, Z_{\langle n \rangle}) \sim (\tilde{X}, Z_{\langle 0 \rangle})$ since X (resp. \tilde{X}) are openings of a random n element subset of the $2n$ commitments contained in $Z_{\langle n \rangle}$ (resp. $Z_{\langle 0 \rangle}$) that open to the same bit.

By a standard hybrid argument, to show $H_{\langle 0 \rangle} \sim_c H_{\langle n \rangle}$ (equivalently $(X, Z_{\langle 0 \rangle}) \sim_c (X, Z_{\langle n \rangle})$), it is sufficient to prove that $H_{\langle i-1 \rangle} \sim_c H_{\langle i \rangle}$ for all $0 < i \leq n$, which we'll now do.

Assume for contradiction that for some i there exists an efficient distinguisher D for $H_{\langle i-1 \rangle}$ and $H_{\langle i \rangle}$ with non-negligible advantage $\delta = \delta(n)$, i.e.,

$$\Pr [D(H_{\langle i \rangle}) = 1] - \Pr [D(H_{\langle i-1 \rangle}) = 1] = \delta .$$

Below, we show how to construct an adversary D' from D which given a commitment $C = \text{com}(B, R)$ (for random B, R) predicts B with probability

$$\Pr [D'(C) = B] = 1/2 + \delta/2 \tag{15}$$

thus breaking the hiding property of com . This contradicts the presumed security of com . The adversary $D'(C)$ first samples (X, Z, A) and then derives $Z' = \text{sort}(\bar{C}^1, \dots, \bar{C}^{i-1}, C, C^{i+1}, \dots, C^{3n})$ from Z by replacing C^1, \dots, C^i in the same vain as $Z_{\langle i \rangle}$ (cf. (12) and (13)) was derived from Z , except that we use the challenge C instead of \bar{C}^i for the i th slot. Note that depending on which bit B the commitment C commits to, the tuple (X, Z') has either the distribution $H_{\langle i-1 \rangle}$ (if $B = A$) or $H_{\langle i \rangle}$ (if $B = 1 - A$). Finally, $D'(C)$ invokes D and outputs $D(X, Z') \oplus A$, now (15) follows by a straightforward and standard calculation, cf. [Gol00, Claim 3.3.7.2]. \square

For (\tilde{X}, Z) as defined in (11), for all sufficiently large n

$$\tilde{H}_\infty(\tilde{X}|Z) \geq \log \binom{2n}{n} \geq n \tag{16}$$

To see this, we note that \tilde{X} is uniform over a set of size $\binom{2n}{n}$ and thus $\tilde{H}_\infty(\tilde{X}|Z) = \log \binom{2n}{n}$.¹¹ By Definition 7, (16) together with Claim 1, which states that $(X, Z) \sim_{1/p,p} (\tilde{X}, Z)$ for any polynomial $p = p(n)$ and all sufficiently large n , implies $H_{1/p,p}^{\text{HILL}}(X|Z) \geq n$, and thus proves condition (1) of Lemma 5. \square

5.2 Counterexample from Committing Distribution

We now show how to construct an ensemble $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ as in Theorem b from an ensemble $\{(C_n, R_n, B_n)\}_{n \in \mathbb{N}}$ of committing distributions as in Definition 10. The construction is basically the same as in the previous section, except for how the S^i are sampled, details follow. Let $A_n \leftarrow \{0, 1\}$ be random and define the B^i as in (10). For each $i = 1, \dots, 3n$, let S^i be a sample of (C_n, R_n, B_n) , sampled conditioned on $B_n = B^i$.

¹¹There's a minor technicality we ignored so far in order to keep things simple: \tilde{X} is only uniform over a set of size $\binom{2n}{n}$ if for $i = 1, \dots, 2n$ the R_i 's (and thus also the S_i 's) are all distinct. We note that the randomly sampled R_i 's will be all distinct with overwhelming probability, so this omission only adds an additive negligible error. We can avoid even this negligible error by initially sampling a random permutation π over $\{1, \dots, 3n\}$, and then replacing the C^i with a tuple $(C^i, \pi(i))$ throughout. This way, we enforce the $S^i = ((C^i, \pi(i)), R^i, B^i)$ to be all distinct, while the extra $\pi(i)$ does not reveal any information about i , which is necessary for the proof.

Lemma 6. *If $\{(C_n, R_n, B_n)\}_{n \in \mathbb{N}}$ is committing (cf. Definition 10), then the ensemble $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ defined above satisfies*

1. *For any polynomial $p = p(n)$*

$$H_{1/p,p}^{\text{HILL}}(X_n|Z_n) \geq n \text{ for infinitely many } n \in \mathbb{N} .$$

2. *There exists a polynomial time Turing machine \mathbf{M} such that for all $(x, y, a) \in \text{supp}[(X_n, Z_n, A_n)]$ we have*

$$(\mathbf{M}(x', z, a) = 1) \iff (x' = x) .$$

The proof of Lemma 6 is analogous to the proof of Lemma 5, with two differences. First, the distributions (X_n, Z_n, A_n) are no longer efficiently samplable, as unlike in Lemma 5, the distributions (C_n, R_n, B_n) used to define it are not efficiently samplable. Second, Claim 1 must be relaxed to $(X_n, Z_n) \sim_{1/p(n), p(n)} (\tilde{X}_n, Z_n)$ for any polynomial $p(\cdot)$ and *infinitely many* $n \in \mathbb{N}$ (as opposed to all sufficiently large n). Consequently, also condition (1) in Lemma 6 only holds for infinitely many (not all sufficiently large) $n \in \mathbb{N}$. The reason this relaxation is necessary is due to the fact that the hiding property for commitment schemes as in Definition 9 holds for all sufficiently large n , whereas in Definition 10 the analogous condition (i) for committing distributions is only required to hold for infinitely many n .

5.3 Committing Distribution from $\text{NP} \not\subseteq \text{P/poly}$

In this section we show that committing distributions exist if $\text{NP} \not\subseteq \text{P/poly}$. Unique-SAT (uSAT) is a promise problem, which asks for distinguishing unsatisfiable Boolean formulas from those with exactly one satisfying assignment. We first state an assumption (Assumption 1) on the hardness of uSAT, and in Lemma 7 construct a committing distribution under this assumption. We then show in Lemma 8 that the assumption holds if $\text{NP} \not\subseteq \text{P/poly}$.

Assumption 1 (nonuniform hardness of uSAT).

Let $\Pi_{YES}, \Pi_{NO} \subseteq \{0, 1\}^$ denote the sets of Boolean formulas that have exactly one and zero satisfying assignments, respectively. Then, there exist ensembles*

$$\begin{aligned} \{T_n\}_{n \in \mathbb{N}} & \quad \text{with} \quad \text{supp}[T_n] \subseteq \Pi_{YES} \cap \{0, 1\}^n \\ \{F_n\}_{n \in \mathbb{N}} & \quad \text{with} \quad \text{supp}[F_n] \subseteq \Pi_{NO} \cap \{0, 1\}^n \end{aligned}$$

of distributions over true and false instances, such that for any polynomial $p(\cdot)$ and any family of circuits $\{A_n\}_{n \in \mathbb{N}}$ where A_n is of size $p(n)$, for infinitely many n we have

$$\Pr[A_n(T_n) = 1] - \Pr[A_n(F_n) = 1] \leq 1/p(n) .$$

Lemma 7 (Committing Distribution from $\text{NP} \not\subseteq \text{P/poly}$). *If Appendix 1 holds, then the ensemble $\{(C_n, R_n, B_n)\}_{n \in \mathbb{N}}$ defined below is a committing distribution as in Definition 10.*

Let $\{T_n\}_{n \in \mathbb{N}}, \{F_n\}_{n \in \mathbb{N}}$ be as in Appendix 1. Sample

$$B_n \leftarrow \{0, 1\} \quad , \quad I_{B_n} \leftarrow T_n \quad , \quad I_{1-B_n} \leftarrow F_n$$

set $C_n = (I_0, I_1)$ and let R_n be the (unique) satisfying assignment of I_{B_n} .

Proof. Condition (i) in Definition 10 directly follows from Appendix 1, as any circuit predicting B_n with non-negligible advantage can be used to distinguish T_n from F_n with non-negligible advantage.¹² A predicate as required for condition (ii) in Definition 10 can be defined as $\phi((I_0, I_1), R, B) = 1$ iff R is a satisfying assignment for I_B . Note that I_B has a unique satisfying assignment, namely R , whereas I_{1-B} has no such assignment. Thus, $\phi((I_0, I_1), R', B') = 1$ iff $(R, B) = (R', B')$ as required by Definition 10. \square

Lemma 8. *Appendix 1 holds if $\mathbf{NP} \not\subseteq \mathbf{P/poly}$.*

In the proof of Lemma 8, we will use Impagliazzo's hardcore lemma [Imp95]. The lemma below follows from a variant of this lemma due to [Hol05, Lemma 2.1].¹³

Lemma 9. *For some constant c , for every $n \in \mathbb{N}$ the following holds. Let $\mathcal{R}_n \subseteq \{0, 1\}^n$ be a set with $|\mathcal{R}_n| \geq 2^{n/2}$ and $f : \mathcal{R}_n \rightarrow \{0, 1\}$ be any predicate. For any constants $\gamma, \delta \in [0, 1]$ and any $p \leq 2^{n/2} \frac{\delta^2}{32}$, if for all circuits A of size p*

$$\Pr_{x \leftarrow \mathcal{R}_n}[A(x) = f(x)] \leq 1 - \frac{\delta}{2},$$

then there exists a set $\mathcal{S} \subseteq \mathcal{R}_n$ such that for all circuits A' of size $p' = \frac{\gamma^2}{32n}p - c$

$$\Pr_{x \leftarrow \mathcal{S}}[A'(x) = f(x)] \leq \frac{1 + \gamma}{2}.$$

of Lemma 8. It was shown by [VV86] that unique-SAT is hard (i.e., not in \mathbf{BPP}) if $\mathbf{NP} \neq \mathbf{RP}$ (see also [Gol08, Section 6.2.3] for a different, more general exposition). They consider a uniform model of computation. We work in a non-uniform model, and thus need unique-SAT to be hard against circuits. A randomized reduction implies a non-uniform one (by using nondeterminism to fix some good coins), thus their reduction also shows that unique-SAT is hard in a non-uniform setting if $\mathbf{NP} \not\subseteq \mathbf{P/poly}$.

Let Π_{YES}^* and Π_{NO}^* denote any encodings of Boolean formulas with exactly one and zero satisfying assignments, respectively. We define redundant encodings Π_{YES} and Π_{NO} , where the last 5/6th of the bits can be ignored:

$$\Pi_{YES} = \bigcup_{n \in \mathbb{N}} \{x \| r : x \in \Pi_{YES}^* \cap \{0, 1\}^n, r \in \{0, 1\}^{5n}\} \quad (17)$$

Π_{NO} is defined analogously using Π_{NO}^* . Looking forward, considering such padded instances will allow us to boost the error probability of any family of poly size circuits in deciding unique-SAT for the infinitely many n' on which it errs on instances of length n'

¹² Assume $A_n(C_n)$ predicts B_n with non-negligible advantage. Equivalently, for any polynomial $p(n)$, for infinitely many n and $I_F \leftarrow F_n, I_T \leftarrow T_n$, A_n distinguishes (I_F, I_T) from (I_T, I_F) with advantage $> 1/p(n)$. Then, using the triangle inequality, A_n distinguishes either (I_F, I_F) from (I_F, I_T) , or (I_F, I_T) from (I_T, I_T) with advantage $> 1/2p(n)$. Assume the latter is the case, then A' defined as $A'(I) = A(I, I_T)$ distinguishes I_F from I_T . A' is not efficient, as I_T is not necessarily efficiently samplable, but we can simply fix some optimal value for the second argument I_T .

¹³The main goal of Holenstein's lemma (compared to the original lemma due to Impagliazzo) was to get a tight lower bound on the size of the set \mathcal{S} . For us, the size of \mathcal{S} will be irrelevant, and so we don't even mention it in the statement of the Lemma. [Hol05, Lemma 2.1] assumes $|\mathcal{R}_n| \geq 2^{n-1}$, whereas we need a smaller $|\mathcal{R}_n| \geq 2^{n/2}$. By inspection of his proof, assuming a smaller \mathcal{R}_n (as we do) doesn't make much of a difference, except that we get a smaller (but still exponential) upper bound on p (any superpolynomial upper bound would be enough for us).

from $2^{-n'}$ as in (18) to $2^{-n/6}$ (where $n = 6n'$) on instances of length n as in (19), and this larger error (in terms of instance length) is required to apply Lemma 9. Let

$$\mathcal{R}_n^* = (\Pi_{YES}^* \cup \Pi_{NO}^*) \cap \{0, 1\}^n, \quad \mathcal{R}_n = (\Pi_{YES} \cup \Pi_{NO}) \cap \{0, 1\}^n$$

and define the predicates

$$(f^*(x) = 1) \iff (x \in \Pi_{YES}^*) \quad \text{and} \quad (f(x) = 1) \iff (x \in \Pi_{YES})$$

As by assumption unique-SAT is not in **P/poly**, for any polynomial $p^*(\cdot)$ and any circuit family $\{\mathbf{A}_n^*\}_{n \in \mathbb{N}}$ of size $|\mathbf{A}_n^*| \leq p^*(n)$, for infinitely many n the circuit \mathbf{A}_n^* must err on at least one instance of \mathcal{R}_n^* . In particular, for the family $\{\mathbf{A}_n^*\}_{n \in \mathbb{N}}$ with the best advantage in predicting $f^*(n)$

$$\Pr_{x \leftarrow \mathcal{R}_n^*}[\mathbf{A}_n^*(x) = f^*(x)] \leq 1 - |\mathcal{R}_n^*|^{-1} \leq 1 - 2^{-n'} . \quad (18)$$

holds for infinitely many n' . If we define the polynomial $p(6n) = p^*(n)$, then (18) implies that for the ensemble $\{\mathbf{A}_n\}_{n \in \mathbb{N}}$ of size $|\mathbf{A}_n| \leq p(n)$ with the best advantage in predicting $f(x)$, for infinitely many n (namely, all $n = 6n'$ where (18) holds for n') we have with $\delta(n) = 2^{-n/6+1}$

$$\Pr_{x||r \leftarrow \mathcal{R}_n}[\mathbf{A}_n(x||r) = f(x||r) = f^*(x)] \leq 1 - 2^{-n/6} = 1 - \delta(n)/2 . \quad (19)$$

To see this, we first observe that we can assume that \mathbf{A}_n ignores the last $5/6n$ bits of the input.¹⁴ Thus, if $\mathbf{A}_n(x||r)$ errs on some input $x||r$, it will err on $2^{5n/6}$ inputs, namely on $x||r'$ for all $r' \in \{0, 1\}^{5n/6}$.

Note that $2^{n/2}\delta(n)^2/32 = 2^{n/6}/8$ is exponential, and thus it upper bounds the polynomial $p(n)$ for all sufficiently large n . For any such sufficiently large n for which (19) holds, we can apply Lemma 9 to conclude there exists a set $\mathcal{S}_n \subseteq \mathcal{R}_n$ for which¹⁵

$$\Pr_{x \leftarrow \mathcal{S}_n}[\mathbf{A}'_n(x) = f(x)] \leq \frac{1 + \gamma(n)}{2} \quad (20)$$

for all circuits \mathbf{A}'_n of size $p'(n) = \gamma(n)^2 p(n)/32n - c$. Setting $\gamma(n) = \sqrt{32n}/\sqrt[3]{p(n)}$ we get

$$p'(n) = \sqrt[3]{p(n)} - c \geq 1/2\gamma(n) .$$

Using (20), we will construct distributions T_n and F_n over yes and no instances such that for all all circuits \mathbf{A}'_n as above

$$\Pr_{x \leftarrow T_n}[\mathbf{A}'_n(x) = 1] - \Pr_{x \leftarrow F_n}[\mathbf{A}'_n(x) = 1] \leq 3\gamma(n) \quad (21)$$

As we can choose $1/3\gamma(n)$ to be an arbitrary large polynomial by choosing $p^*(n)$ large enough, this proves Lemma 8 (with the arbitrary polynomial $p(n)$ in Appendix 1 being $1/3\gamma(n)$).¹⁶

¹⁴As, using (17), for any circuit \mathbf{A}_n there exists an r' such that the circuit $\mathbf{A}'_n(x||r) = \mathbf{A}_n(x||r')$ (where \mathbf{A}'_n ignores r) errs with at most the same probability as $\mathbf{A}_n(x||r)$.

¹⁵We note that the size requirement $|\mathcal{R}_n| \geq 2^{n/2}$ of Lemma 9 is satisfied as the definition (17) of our padding implies $|\mathcal{R}_n| = 2^{5n/6} \cdot |(\Pi_{YES}^* \cup \Pi_{NO}^*) \cap \{0, 1\}^{n/6}|$.

¹⁶Technically, here we have assumed that Appendix 1 holds for Π_{YES} and Π_{NO} instances as in (17), where the last $5/6$ th of the bits are just random paddings. We observe that the assumption for such “padded” instances trivially implies the assumption for any encoding of instances.

It remains to show that (20) implies (21). For this, let $\mathcal{T}_n = \mathcal{S}_n \cap \Pi_{YES}$ and $\mathcal{F}_n = \mathcal{S}_n \cap \Pi_{NO}$. We note that \mathcal{S}_n must contain roughly the same number of yes and no instances:

$$|\Pr_{x \leftarrow \mathcal{S}_n}[x \in \mathcal{T}_n] - 1/2| \leq \gamma(n)/2 .$$

To see this, assume for contradiction that this does not hold, i.e., $\Pr_{x \leftarrow \mathcal{S}_n}[x \in \mathcal{T}_n]$ is either $> \frac{1+\gamma(n)}{2}$ or $< \frac{1-\gamma(n)}{2}$. In the first case, the constant function $A'_n(x) = 1$ contradicts (20), in the second case we get a contradiction using $A'_n(x) = 0$.

With this observation we can define a distribution S which has statistical distance at most $\gamma(n)$ to the uniform distribution over $\mathcal{S}_n = \mathcal{T}_n \cup \mathcal{F}_n$, where the support of S lies in \mathcal{S}_n and which is perfectly balanced in the sense that¹⁷

$$\Pr_{x \leftarrow S}[x \in \mathcal{T}_n] = \Pr_{x \leftarrow S}[x \in \mathcal{F}_n] = 1/2 \quad (22)$$

Let T_n (F_n) denote the distribution S conditioned on $x \in \mathcal{T}_n$ ($x \in \mathcal{F}_n$), this definition is used in the third equality below. In the fourth equality we use (22). The fifth equality uses the fact that S is $\gamma(n)$ close to the uniform distribution over \mathcal{S}_n and the last inequality follows by (20).

$$\begin{aligned} & \Pr_{x \leftarrow T_n}[A'_n(x) = 1] - \Pr_{x \leftarrow F_n}[A'_n(x) = 1] \\ &= (\Pr_{x \leftarrow T_n}[A'_n(x) = 1] + \Pr_{x \leftarrow F_n}[A'_n(x) = 0]) - 1 \\ &= (\Pr_{x \leftarrow T_n}[A'_n(x) = f(x)] + \Pr_{x \leftarrow F_n}[A'_n(x) = f(x)]) - 1 \\ &= (\Pr_{x \leftarrow S}[A'_n(x) = f(x)|x \in \mathcal{T}_n] + \Pr_{x \leftarrow S}[A'_n(x) = f(x)|x \in \mathcal{F}_n]) - 1 \\ &= 2\Pr_{x \leftarrow S}[A'_n(x) = f(x)] - 1 \\ &= 2\Pr_{x \leftarrow \mathcal{S}_n}[A'_n(x) = f(x)] - 1 \pm 2\gamma(n) \\ &\leq 3\gamma(n) \end{aligned}$$

Note that the above proves (21). □

6 Counterexample from Lossy Functions

In this section we give a particularly simple counterexample to the chain rule for conditional HILL entropy, which is based on lossy functions. We stress that we only need lossy functions not lossy *trapdoor* functions. Below we define ℓ -lossy functions [PW08], where for simplicity we only define the particular setting where the key-space, input and output domain are all bitstrings of length n (but everything goes through unchanged for the general definition, where these domains can be described by strings of length polynomial in n).

Definition 11 (Lossy Function). *An $\ell = \ell(n)$ lossy function consists of two efficient algorithms*

$$\begin{aligned} \text{KG} &: 1^* \times \{\text{lossy, injective}\} \rightarrow \{0, 1\}^* \\ \text{F} &: \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \end{aligned}$$

¹⁷Concretely, we let S be the uniform distribution over \mathcal{S}_n conditioned on an event that holds with probability at least $1 - \gamma$. Assume that $\Pr_{x \leftarrow \mathcal{S}_n}[x \in \mathcal{T}_n] = 1/2 + \gamma/2$ for some $\gamma > 0$, then this event is defined as follows: The event always holds if $x \in \mathcal{F}_n$, and fails with probability $\delta = 1 - (1/2 - \gamma/2)/(1/2 + \gamma/2)$ if $x \in \mathcal{T}$ (this δ satisfies $(1/2 - \gamma/2) = (1/2 + \gamma/2)(1 - \delta)$ as required to satisfy (22)). The probability that the event fails is $\Pr_{x \leftarrow \mathcal{S}_n}[x \in \mathcal{T}_n] \Pr_{x \leftarrow \mathcal{S}_n}[\text{event fails}|x \in \mathcal{T}_n] = (1/2 + \gamma/2)\delta = \gamma$.

The probabilistic key-generation KG takes as input a security parameter 1^n in unary and mode $\in \{\text{lossy}, \text{injective}\}$ and outputs a key $K \in \{0, 1\}^n$. For every $K \in \{0, 1\}^n$ it holds that $F(K, \cdot) = F_K(\cdot)$ is a function $\{0, 1\}^n \rightarrow \{0, 1\}^n$ such that for every

$$K_{\text{los}} \in \text{supp}[\text{KG}(1^n, \text{lossy})] \quad \text{and} \quad K_{\text{inj}} \in \text{supp}[\text{KG}(1^n, \text{injective})]$$

and every $x \in \{0, 1\}^n$ we have

- $|F_{K_{\text{inj}}}^{-1}(F_{K_{\text{inj}}}(x))| = 1$; i.e., $F_{K_{\text{inj}}}(\cdot)$ is injective.
- $|F_{K_{\text{los}}}^{-1}(F_{K_{\text{los}}}(x))| = 2^\ell$; i.e., each value in the output range of $F_{K_{\text{los}}}$ has 2^ℓ pre-images.

Moreover, lossy and injective keys are indistinguishable

$$\{\text{KG}(1^n, \text{lossy})\}_{n \in \mathbb{N}} \sim_c \{\text{KG}(1^n, \text{injective})\}_{n \in \mathbb{N}}$$

Theorem 2. Given an $\ell = \ell(n)$ lossy function, we can sample an ensemble of joint distributions $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$, where sampling a distribution requires only two invocations of KG and of F , such that for some fixed polynomial $s = s(n)$, any polynomial $p = p(n)$ and all sufficiently large n

$$H_{1/p, p}^{\text{HILL}}(X_n | Z_n) \geq \ell(n) \quad \text{but} \quad H_{1/2, s}^{\text{HILL}}(X_n | Z_n, A_n) \leq 1.$$

Proof. We will omit the subscripts n in this proof. First, sample an injective and lossy key

$$K^0 \leftarrow \text{KG}(1^n, \text{injective}) \quad \text{and} \quad K^1 \leftarrow \text{KG}(1^n, \text{lossy}).$$

Then sample inputs $X^0, X^1 \leftarrow \{0, 1\}^n$ and compute

$$Z^0 = F_{K^0}(X^0) \quad \text{and} \quad Z^1 = F_{K^1}(X^1).$$

Next, choose a random bit $A \leftarrow \{0, 1\}$ and define $Z = (K^A, Z^A, K^{1-A}, Z^{1-A})$. Looking forward, the tuple (X^0, Z, A) will correspond to (X_n, Z_n, A_n) in the statement of the theorem. It is instructive to observe that the min-entropy of X^0 and X^1 conditioned on Z is

$$H_\infty(X^0 | Z) = 0 \quad \text{and} \quad H_\infty(X^1 | Z) = \ell.$$

The left equation holds as X^0 can be computed (in exponential time) given Z . Concretely, we can perfectly distinguish lossy from injective keys, and thus determine A which tells us which tuple in Z is (K^0, Z^0) . Now $X^0 = F_{K^0}^{-1}(Z^0)$ is well defined as K^0 is injective. To see the equation $H_\infty(X^1 | Z) = \ell$ on the right side note that X^1 is uniform over a the set $F_{K^1}^{-1}(Z^1)$ which is of size 2^ℓ since K^1 is lossy.

As we'll show below, for HILL entropy the picture is different, because given Z the value of A is computationally hidden and thus we don't know which of the two keys is the lossy one.

We first show that $H_{1/2, s}^{\text{HILL}}(X^0 | Z, A) \leq 1$, where s is roughly the size of a circuit that computes F . Note that given (Z, A) , we know which of the tuples in Z is (K^0, Z^0) , and there is a single X' (namely X^0) such that $F_{K^0}(X') = Z^0$. Thus given (Z, A) , we can efficiently check if some value X' is equal to X^0 using a circuit of size $s \approx |F|$. This implies $H_{1/2, s}^{\text{HILL}}(X^0 | Z, A) \leq 1$ by Lemma 3.

We'll now show that $H_{p, 1/p}^{\text{HILL}}(X^0 | Z) \geq \ell$. For this, we must show there exists a random variable Y such that given Z , the distribution Y has min-entropy ℓ and Y is indistinguishable from X^0 . We claim this holds for Y that is uniform on $F_{K^1}^{-1}(Z^1)$. As this set

has size 2^ℓ , $H_\infty(Y|Z) = \ell$. Moreover (Y, Z) is indistinguishable from (X^0, Z) as we could use a distinguisher for these distributions to distinguish lossy from injective keys.¹⁸ \square

7 Proofs from Section 2

7.1 Predictability Implies Lack of HILL

We start with the following claim, showing that if a distribution is “conditionally predictable” than it cannot have much HILL entropy. This will be useful in the proofs of Proposition 1 and Proposition 2 in this section.

Claim 2. *There is a fixed polynomial $p_{eq}(\cdot)$ such that the following holds. Let (X, Z) be any distribution such that X is of bit-length m . If there exists a circuit C of size $|C| \leq s$ such that*

$$\Pr_{(x,z) \leftarrow (X,Z)}[C(z) = x] > \varepsilon$$

then, for every $\ell > 0$,

$$H_{\varepsilon-2^{-\ell}, s+p_{eq}(m)}^{\text{HILL}}(X|Z) \leq \ell.$$

In particular, for $\ell = \log(2/\varepsilon)$ we get

$$H_{\varepsilon/2, s+p_{eq}(m)}^{\text{HILL}}(X|Z) \leq \log(2/\varepsilon).$$

Proof. Let C , s and ε be as in the hypothesis of the claim and let $p_{eq}(m)$ be the size of the circuit that takes as input two m -bit strings and outputs 1 if they are equal. Define the distinguisher D which outputs $D(z, x) = 1$ iff $C(z) = x$, which is of size $|D| = s + p_{eq}(m)$. Let (Y, Z) be any distribution such that $H_\infty(Y|Z) > \ell$. Then

$$\begin{aligned} \Pr[C(Z) = Y] &\leq \mathbb{E}_{z \leftarrow Z} \Pr[Y = C(z) | Z = z] \\ &\leq \mathbb{E}_{z \leftarrow Z} \max_x \Pr[Y = x | Z = z] \leq 2^{H_\infty(Y|Z)} < 2^{-\ell}. \end{aligned}$$

Therefore we have

$$\begin{aligned} &\Pr[D(X, Z) = 1] - \Pr[D(Y, Z) = 1] \\ &= \Pr[C(Z) = X] - \Pr[C(Z) = Y] > \varepsilon - 2^{-\ell} \end{aligned}$$

We just proved that $(Y, Z) \not\sim_{(\varepsilon+2^{-\ell}, s+p_{eq}(m))} (X, Z)$ for any $H_\infty(Y|Z) > \ell$, which means $H_{\varepsilon-2^{-\ell}, s+p_{eq}(m)}^{\text{HILL}}(X|Z) \leq \ell$. \square

7.2 Proof of Proposition 1

Let $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$, where $A_n \in \{0, 1\}$, be an efficiently samplable ensemble where for every polynomial $p(n)$ and for all sufficiently large n it holds that

$$H_{1/p(n), p(n)}^{\text{HILL}}(X_n|Z_n) \geq n \quad \text{and} \quad H_\infty(X_n|Z_n, A_n) = 0 \quad (23)$$

To prove Proposition 1, we must use this to construct a OWF. Recall that the lhs of (23) is condition (i) from Theorem 1.a and the rhs is implied by condition (ii) from Theorem 1.a (as explained in the statement of Proposition 1).

¹⁸Given a pair of keys K, K' , one of which is lossy, we can sample random X, X' , then set $Z = (K, F_K(X), K', F_{K'}(X'))$. Now, depending on whether K or K' is lossy, (X, Y) is distributed like (X^0, Z) (if K is injective) or (Y, Z) (if K is the lossy key).

of Proposition 1. Let $f'(\cdot)$ denote the efficient sampling algorithm, which on input randomness R_n outputs a sample (X_n, Z_n, A_n) . Let $f(\cdot)$ denote $f'(\cdot)$, but where we drop the X_n part from the output: i.e., $f(R_n) = (Z_n, A_n)$. We claim that f is a one-way function.

To show this, we assume for contradiction there exists a (non-uniform) polynomial-size inversion algorithm $\mathbf{D} = \{\mathbf{D}_n\}_{n \in \mathbb{N}}$ that breaks one-wayness. In more detail, there is some polynomial $q(n)$ such that, for infinitely many $n \in \mathbb{N}$,

$$\Pr_{(x,z,a) \leftarrow (X_n, Z_n, A_n)}[f(\mathbf{D}_n(z, a)) = (z, a)] \geq 1/q(n). \quad (24)$$

As $H_\infty(X_n | Z_n, A_n) = 0$, x is completely determined by (z, a) , which implies that (24) is equivalent to

$$\Pr_{(x,z,a) \leftarrow (X_n, Z_n, A_n)}[f'(\mathbf{D}_n(z, a)) = (x, z, a)] \geq 1/q(n). \quad (25)$$

Let $f''(\cdot)$ be the same as $f'(\cdot)$, but where we only output the X_n part: i.e., $f''(R_n) = X_n$. Then (note that in the last inequality below we replace a with a random bit b , as $2\Pr[a = b] = 2 \cdot \frac{1}{2} = 1$, we can “compensate” for this by multiplying with a factor 2)

$$\begin{aligned} 1/q(n) &\leq \Pr_{(x,z,a) \leftarrow (X_n, Z_n, A_n)}[f'(\mathbf{D}_n(z, a)) = (x, z, a)] \\ &\leq \Pr_{(x,z,a) \leftarrow (X_n, Z_n, A_n)}[f''(\mathbf{D}_n(z, a)) = x] \\ &\leq 2 \cdot \Pr_{b \leftarrow \{0,1\}, (x,z) \leftarrow (X_n, Z_n)}[f''(\mathbf{D}_n(z, b)) = x] \end{aligned} \quad (26)$$

Define the polynomial-size circuit family $\mathbf{C} = \{\mathbf{C}_n\}_{n \in \mathbb{N}}$ via $\mathbf{C}_n(z) = f''(\mathbf{D}_n(z, b))$ where the bit $b \leftarrow \{0, 1\}$ is sampled uniformly at random. Now (26) can be stated as

$$\Pr_{(x,z) \leftarrow (X_n, Z_n)}[\mathbf{C}_n(z) = x] \geq 1/(2q(n))$$

Then, by applying Claim 2 with $\varepsilon = 1/(2q(n))$ and $\ell = \log(4q(n))$ there is some polynomial $p(n) = |\mathbf{C}_n| + p_{eq}(|X_n|)$ such that, for infinitely many n :

$$H_{1/(4q(n)), p(n)}^{\text{HILL}}(X_n | Z_n) \leq \log(4q(n)) = O(\log n).$$

This contradicts (23). Therefore, the function f must be one-way, which proves the proposition. \square

7.3 Proof of Proposition 2

Let $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ be an ensemble of distributions as in Theorem b. That is, for every polynomial $p(n)$ and for infinitely many n it holds that

$$H_{1/p(n), p(n)}^{\text{HILL}}(X_n | Z_n) \geq n \quad (27)$$

and moreover there exist a polynomial time Turing machine \mathbf{M} such that for any $n \in \mathbb{N}$ and any $(x, z, a) \in \text{supp}[(X_n, Z_n, A_n)]$

$$(\mathbf{M}(x', z, a) = 1) \iff (x' = x). \quad (28)$$

We do not assume that (X_n, Z_n, A_n) is efficiently samplable, but require that X_n, Z_n are of polynomial length and $A_n \in \{0, 1\}$. To prove Proposition 2 we must show that this implies $\mathbf{NP} \not\subseteq \mathbf{P/poly}$.

of Proposition 2. Define the NP language

$$\mathcal{L} = \{(z, a) : \exists x \text{ such that } M(x, z, a) = 1\}$$

Assume, by contradiction, that $\mathbf{NP} \subseteq \mathbf{P/poly}$. Since we can reduce the problem of finding a witness to the problem of deciding membership, there exists a family of polynomial-size circuits $\{C'_n\}_{n \in \mathbb{N}}$ such that $C'_n(z, a)$ outputs a witness x that satisfies $M(x, z, a) = 1$ for any $(z, a) \in \mathcal{L}$. Moreover, by (28), for any $(x, z, a) \in \text{supp}[(X_n, Z_n, A_n)]$ we must have $C'_n(z, a) = x$ since this value x is the unique witness for which $M(x, z, a) = 1$. Therefore

$$\begin{aligned} 1 &= \Pr_{(x,z,a) \leftarrow (X_n, Z_n, A_n)}[C'_n(z, a) = x] \\ &\leq \sum_{b \in \{0,1\}} \Pr_{(x,z) \leftarrow (X_n, Z_n)}[C'_n(z, b) = x] \\ &= 2 \cdot \Pr_{b \leftarrow \{0,1\}, (x,z) \leftarrow (X_n, Z_n)}[C'_n(z, b) = x] \end{aligned}$$

Define the circuit $C_n(z) = C'_n(z, b)$ where the bit $b \leftarrow \{0, 1\}$ is sampled uniformly at random. Then, by applying Claim 2 with $\varepsilon = 1/2$ and $\ell = 2$ there is some polynomial $p(n) = |C_n| + p_{eq}(|X_n|)$ such that, for every n :

$$H_{1/4, p(n)}^{\text{HILL}}(X_n | Z_n) \leq 2.$$

This contradicts (27) and therefore proves the proposition. \square

7.4 Proof of Lemma 2 and Lemma 4

We begin with the following claim that will allow us to prove both lemmas.

Claim 3. *For any joint distribution (V, C) over $\mathcal{V} \times \mathcal{C}$ with $|\mathcal{V}| \geq 2$ we have:*

$$H_{1/2, \infty}^{\text{HILL}}(V|C) \geq 1.$$

Furthermore, if $|\mathcal{V}| > 2$ and $\tilde{H}_\infty(V|C) > 0$ then we have:

$$H_{1/2, \infty}^{\text{HILL}}(V|C) > 1.$$

Proof. To show the first part of the claim, we must show that there exists a distribution (Y, C) satisfying:

$$\tilde{H}_\infty(Y|C) \geq 1 \quad \text{and} \quad (V, C) \sim_{1/2} (Y, C) \quad (29)$$

For each $c \in \mathcal{C}$ define the probabilities $p_c(v) := \Pr[V = v | C = c]$. Our goal is to define a valid probability distribution $p'_c(v)$ so as to minimize $\max_v (p_c(v) + p'_c(v))$. Let $p_{\max, c} := \max_{v \in \mathcal{V}} p_c(v)$ and define the values $p'_c(v)$ via

$$\begin{aligned} p'_c(v) &:= \frac{p_{\max, c} - p_c(v)}{|\mathcal{V}| p_{\max, c} - 1} && \text{if } p_{\max, c} \geq 2/|\mathcal{V}| \\ p'_c(v) &:= (2/|\mathcal{V}| - p_c(v)) && \text{otherwise} \end{aligned}$$

Notice that the values $p'_c(v)$ form a valid probability distribution over \mathcal{V} and

$$p_c(v) + p'_c(v) \leq \max\{p_{\max, c}, 2/|\mathcal{V}|\}. \quad (30)$$

Define a random variable V'_c over \mathcal{V} with $\Pr[V'_c = v] := p'_c(v)$. Define a distribution (Y, C) by sampling $(y, c) \leftarrow (Y, C)$ via the following process: Sample $(v, c) \leftarrow (V, C)$, $v' \leftarrow V'_c$ and a random bit $b \leftarrow B$. If $b = 0$, output $(y = v, c)$ and otherwise output $(y = v', c)$.

This distribution satisfies the rhs of (29) since, conditioned on $b = 0$, which holds with probability $1/2$, the distributions (V, C) and (Y, C) are identical.

Moreover for any (v, c)

$$\begin{aligned} & \Pr[Y = v | C = c] \\ &= \Pr[B = 0] \Pr[V = v | C = c] + \Pr[B = 1] \Pr[V'_c = v] \\ &= (1/2)(p_c(v) + p'_c(v)) \\ &\leq \max\{ p_{max,c}/2, 1/|\mathcal{V}| \} \end{aligned}$$

Recalling that $p_{max,c} = \max_v \Pr[V = v | C = c]$, the above implies that for each $c \in \mathcal{C}$

$$\begin{aligned} H_\infty(Y|C = c) &\geq \min\{H_\infty(V|C = c) + 1, \log(|\mathcal{V}|)\} \\ &\geq 1. \end{aligned} \tag{31}$$

This in turn implies that $\tilde{H}_\infty(Y|C) \geq 1$. Therefore, (Y, C) satisfies the lhs of (29), which proves the first part of the claim.

For the second part of the claim, notice that when $\tilde{H}_\infty(V|C) > 0$ then there must be some $c^* \in \mathcal{C}$ such that $\tilde{H}_\infty(V|C = c^*) > 0$. Moreover, since $|\mathcal{V}| > 2$, (31) tells us that for this c^* we have $H_\infty(Y|C = c^*) > 1$. Since $H_\infty(Y|C = c) \geq 1$ for all other $c \in \mathcal{C}$, this shows that $\tilde{H}_\infty(V|C) > 1$. This proves the second part of the claim. \square

Proof of Lemma 2. Lemma 2 states that for any joint distribution (V, C) over $\mathcal{V} \times \mathcal{C}$ (where $|\mathcal{V}| \geq 2$) and any $s \in \mathbb{N}$

$$H_{1/2,s}^{\text{HILL}}(V|C) \geq H_{1/2,\infty}^{\text{HILL}}(V|C) \geq 1$$

The first inequality follows directly from Definition 7.¹⁹ The second inequality follows from the first part of Claim 3.

Proof of Lemma 4. Lemma 4 states that for any joint distribution (V, C) over $\mathcal{V} \times \mathcal{C}$ where $|\mathcal{V}| > 2$, and any s

$$H_{1/2,s}^{\text{HILL}}(V|C) \leq 1 \Rightarrow H_{1/2,\infty}^{\text{HILL}}(V|C) \leq 1 \Rightarrow H_\infty(V|C) = 0$$

The first implication follows by Definition 7. For the second implication we notice that the contrapositive

$$H_\infty(V|C) > 0 \Rightarrow H_{1/2,\infty}^{\text{HILL}}(V|C) > 1$$

follows directly from the second part of Claim 3.

¹⁹To see this, we note that for any $s \in \mathbb{N}$, $(X, Z) \sim_{\varepsilon,\infty} (Y, Z)$ implies $(X, Z) \sim_{\varepsilon,s} (Y, Z)$ as the set of circuits of size s is a subset of all circuits.

7.5 Proof of Lemma 3

Lemma 3 states that if there exists a circuit D of size s where

$$\forall (v, c) \in \text{supp}[(V, C)] : (D(v', c) = 1) \iff (v' = v)$$

then $H_{1/2, s}^{\text{HILL}}(V|C) \leq 1$.

of Lemma 3. Assume for contradiction that a circuit D as in the statement of the lemma exists, but $H_{1/2, s}^{\text{HILL}}(V|C) > 1$. By Definition 7 this means that there exists a distribution (Y, C) with

$$\tilde{H}_\infty(Y|C) > 1 \quad \text{and} \quad (V, C) \sim_{1/2, s} (Y, C). \quad (32)$$

By the hypothesis of the lemma, for every $c \in \text{supp}[C]$, the circuit $D(\cdot, c)$ outputs 1 on exactly one possible value, which we'll denote by v_c . Now

$$\begin{aligned} \Pr_{(y, c) \leftarrow (Y, C)}[D(y, c) = 1] &= \Pr_{(y, c) \leftarrow (Y, C)}[y = v_c] \\ &= \mathbb{E}_{c \leftarrow C} \Pr[Y = v_c | C = c] \\ &\leq \mathbb{E}_{c \leftarrow C} \max_v \Pr[Y = v | C = c] \end{aligned}$$

Taking the logarithm and using Definition 4

$$-\log(\Pr_{(y, c) \leftarrow (Y, C)}[D(y, c) = 1]) \geq \tilde{H}_\infty(Y|C)$$

Using the left equation of (32)

$$\log(\Pr_{(y, c) \leftarrow (Y, C)}[D(y, c) = 1]) < -1$$

or equivalently

$$\Pr_{(y, c) \leftarrow (Y, C)}[D(y, c) = 1] < 1/2$$

Now

$$\left| \underbrace{\Pr_{(v, c) \leftarrow (V, C)}[D(v, c) = 1]}_{=1} - \underbrace{\Pr_{(y, c) \leftarrow (Y, C)}[D(y, c) = 1]}_{< 1/2} \right| > 1/2$$

which contradicts the right side of (32). \square

8 Conclusion

Computational entropy, most notably pseudorandomness, is an extremely useful concept in cryptography. The general idea is to exploit the fact that to computationally bounded parties, random variables can look and behave as if they had much more entropy than they actually do.

In this paper we showed that one of the most fundamental properties of entropy notions, the chain rule, does not hold for HILL entropy, arguably the most important computational entropy notion.

We gave counterexamples to the chain rule from a variety of cryptographic primitives: injective one-way functions, lossy functions and (in the conference version of this paper) also from deniable encryption. As discussed in Section 2.4, the latter two are very efficient counterexamples, using just one or two invocations of the underlying primitive. This shows that schemes achieving sophisticated cryptographic properties like deniability or lossiness inherently embed strong counterexamples to the chain rule, and we believe it might be fruitful to investigate some cryptographic objects from this perspective.

Acknowledgements

Parts of this work were done while the first author was at IST Austria, and the third author was visiting IST Austria. This work was partly funded by the European Research Council under ERC Starting Grant *259668-PSPC* and ERC Advanced Grant *321310-PERCY*.

A preliminary version of this work has appeared in [KPW13].

References

- [BM84] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudorandom bits. *SIAM Journal on Computing*, 13(4):850–864, 1984.
- [BNN011] Rikke Bendlin, Jesper Buus Nielsen, Peter Sebastian Nordholt, and Claudio Orlandi. Lower and upper bounds for deniable public-key encryption. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 125–142, Seoul, South Korea, December 4–8, 2011. Springer, Berlin, Germany.
- [BSW03] B. Barak, R. Shaltiel, and A. Wigderson. Computational Analogues of Entropy. In S. Arora, K. Jansen, J. D. P. Rolim, and A. Sahai, editors, *RANDOM-APPROX 03*, volume 2764 of *LNCS*, pages 200–215. Springer, 2003.
- [CDNO97] R. Canetti, C. Dwork, M. Naor, and R. Ostrovsky. Deniable Encryption. In B. S. Kaliski Jr., editor, *CRYPTO 97*, volume 1294 of *LNCS*, pages 90–104. Springer, 1997.
- [CKLR11] K.-M. Chung, Y. T. Kalai, F.-H. Liu, and R. Raz. Memory Delegation. In P. Rogaway, editor, *CRYPTO 11*, volume 6841 of *LNCS*, pages 151–168. Springer, 2011.
- [DF11] M. Dürmuth and D. M. Freeman. Deniable Encryption with Negligible Detection Probability: An Interactive Construction. In K. G. Paterson, editor, *EUROCRYPT 11*, volume 6632 of *LNCS*, pages 610–626. Springer, 2011. Full version including a description of the flaw available at: <http://eprint.iacr.org/2011/066.pdf>.
- [DORS08] Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM Journal on Computing*, 38(1):97–139, 2008.
- [DP08] S. Dziembowski and K. Pietrzak. Leakage-Resilient Cryptography. In *FOCS 08*, pages 293–302. IEEE Computer Society, 2008.
- [FOR12] B. Fuller, A. O’Neill, and L. Reyzin. A Unified Approach to Deterministic Encryption: New Constructions and a Connection to Computational Entropy. In R. Cramer, editor, *TCC 12*, volume 7194 of *LNCS*, pages 582–599. Springer, 2012.

- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, pages 40–49, 2013.
- [Gol00] O. Goldreich. *Foundations of Cryptography: Basic Tools*. Cambridge University Press, New York, NY, USA, 2000.
- [Gol08] O. Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108, San Jose, California, USA, June 6–8, 2011. ACM Press.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [HLR07] Chun-Yuan Hsiao, Chi-Jen Lu, and Leonid Reyzin. Conditional computational entropy, or toward separating pseudoentropy from compressibility. In Moni Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 169–186, Barcelona, Spain, May 20–24, 2007. Springer, Berlin, Germany.
- [Hol05] Thomas Holenstein. Key agreement from weak bit agreement. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 664–673, Baltimore, Maryland, USA, May 22–24, 2005. ACM Press.
- [Imp95] Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *FOCS*, pages 538–545, 1995.
- [JP14] Dimitar Jetchev and Krzysztof Pietrzak. How to fake auxiliary input. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 566–590, San Diego, CA, USA, February 24–26, 2014. Springer, Berlin, Germany.
- [KPW13] Stephan Krenn, Krzysztof Pietrzak, and Akshay Wadia. A counterexample to the chain rule for conditional HILL entropy - and what deniable encryption has to do with it. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 23–39, Tokyo, Japan, March 3–6, 2013. Springer, Berlin, Germany.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, 1991.
- [PW08] C. Peikert and B. Waters. Lossy Trapdoor Functions and Their Applications. In C. Dwork, editor, *STOC*, pages 187–196. ACM, 2008.
- [Rey11] L. Reyzin. Some Notions of Entropy for Cryptography. In S. Fehr, editor, *ICITS 11*, volume 6673 of *LNCS*, pages 138–142. Springer, 2011.
- [RTTV08] O. Reingold, L. Trevisan, M. Tuliani, and S. P. Vadhan. Dense Subsets of Pseudorandom Sets. In *FOCS 08*, pages 76–85. IEEE Computer Society, 2008.
- [Sko13] M. Skorski. Modulus Computational Entropy. *CoRR*, abs/1302.2128, 2013.

- [SW13] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. Cryptology ePrint Archive, Report 2013/454, 2013. <http://eprint.iacr.org/>.
- [TTV09] Luca Trevisan, Madhur Tulsiani, and Salil P. Vadhan. Regularity, boosting, and efficiently simulating every high-entropy distribution. In *IEEE Conference on Computational Complexity*, pages 126–136, 2009.
- [VV86] Leslie G. Valiant and Vijay V. Vazirani. Np is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
- [VZ12] Salil P. Vadhan and Colin Jia Zheng. Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 817–836, New York, NY, USA, May 19–22, 2012. ACM Press.
- [VZ13] Salil P. Vadhan and Colin Jia Zheng. A uniform min-max theorem with applications in cryptography. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 93–110, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Berlin, Germany.
- [Yao82] A. C. Yao. Theory and Applications of Trapdoor Functions (Extended Abstract). In *FOCS 82*, pages 80–91. IEEE Computer Society, 1982.

A Counterexample from Deniable Encryption

In this section we will state and prove Theorem 3 which shows that any deniable encryption scheme gives an efficient counterexample to the chain rule for conditional HILL entropy.

The existence of a sender deniable encryption scheme was a long standing open problem, and the original motivation for this work came from the observation that any deniable encryption scheme would constitute a counterexample to the chain rule for conditional HILL entropy. Thus, proving such a chain rule would have implied that sender deniable encryption does not exist. However, it turned out that the chain rule does not hold. In fact, the first counterexample we found was based on a construction of a deniable encryption scheme due to [DF11]. Although their construction is known to have a subtle but fatal flaw, this didn’t affect its usefulness for constructing the counterexample.²⁰ Very recently [SW13] constructed a sender deniable encryption scheme from *indistinguishability obfuscation*. A candidate for such obfuscation schemes has been proposed by [GGH⁺13] based on multilinear maps. The only negative result concerning the existence of deniable encryption is due to [BNNO11], who show that *non-interactive receiver* deniable encryption does not exist.

A.1 Sender Deniable PKE

Deniable encryption, first introduced by [CDNO97], offers protection against *coercion*. Suppose Alice wants to secretly send a message to Bob over a public authentic channel.

²⁰Informally, the reason for this is that for our counterexample we don’t require the faking procedure to be efficient.

They can do this using a public-key encryption (PKE) scheme where Bob sends his public-key to Alice, Alice encrypts the message using this key and sends Bob the ciphertext. Now suppose an adversary coerces one of the parties, say the sender Alice, into revealing all the secret information, i.e., the message and the random coins used to encrypt. If we use a standard PKE scheme, the adversary can verify the consistency of this information with the transcript observed on the public channel. In particular, the adversary can check that the message provided by Alice was indeed the one sent.

A deniable encryption scheme tackles this problem by providing a *faking* algorithm. The faking algorithm allows a coerced party to come up with a fake random tape that is consistent with the public transcript for any given message. Deniable encryption schemes are classified as *sender deniable*, *receiver deniable* or *bi-deniable*, depending on whether the sender, the receiver or both simultaneously can withstand coercion.

Normally PKE is a two round protocol: Bob sends a public-key, Alice answers with a ciphertext. However, for the sake of generality, we will consider protocols that can have any polynomial number of rounds. We will only consider sender deniable encryption schemes, which is without loss of generality, since a k -round receiver deniable encryption scheme implies a $k + 1$ round sender deniable encryption scheme.²¹

We model deniable encryption as a two-party protocol between a sender S and a receiver R . Both get as input a security parameter 1^n in unary and S gets a message bit $b \in \{0, 1\}$ as input (without loss of generality, we'll only consider bit-encryption). At the end of the protocol, R outputs a message b' . As usual, S and R can be probabilistic. For a deniable encryption scheme ψ , we denote by $tr_\psi^n(b, r_S, r_R)$ the transcript (i.e., all the messages exchanged over the public channel) between the sender $S(1^n, b)$ and a receiver $R(1^n)$ which have random tapes r_S and r_R , respectively. If we replace any of b, r_S, r_R with \star , this means we consider the random variable where those values are chosen uniformly at random from their respective domain, e.g., $tr_\psi^n(0, \star, \star)$ is the distribution of transcripts for message bit 0 using security parameter 1^n . A sender deniable encryption scheme is then defined as follows [CDNO97]:

Definition 12 (Sender Deniable PKE). *A protocol ψ between a sender S and a receiver R is a (ε, δ, t) -sender-deniable (bit) encryption scheme (where $\varepsilon = \varepsilon(n), \delta = \delta(n), t = t(n)$ are functions of the security parameter n) if the following three properties are satisfied:*

Correctness: *For any $b \in \{0, 1\}$, the probability that $R(1^n)$ after interacting with $S(1^n, b)$ outputs b is exactly $1 - \varepsilon$.*²²

Security: *$tr_\psi^n(0, \star, \star) \sim_{\delta, t} tr_\psi^n(1, \star, \star)$, that is, the transcript computationally hides the encrypted bit.*

²¹ This is done as follows: The receiver first uses the k -round scheme to send a random key K to the sender. The sender then uses K as the key of a one-time pad to encrypt B as $C = B \oplus K$, and sends C to the receiver. This scheme is sender deniable, because for any message B' , the sender can use the receiver deniability of the k round scheme to come up with a transcript that is consistent with having received the key $K' = C \oplus B'$ in the first phase, and thus also consistent with having send the message B' .

²² The standard definition only requires that the correct bit is output with probability *at least* (not exactly) $1 - \varepsilon$, in particular, the decryption error can be different depending on whether the bit 0 or 1 was encrypted. Requiring the error to be independent of the encrypted bit will be convenient in the proof and is without loss of generality as one always can make the error symmetric by having the receiver sometimes flip the received bit.

Deniability: *There exists an efficient faking algorithm ϕ having the following property: For a random message bit $b \leftarrow \{0, 1\}$, random r_S, r_R , let $c = \text{tr}_\psi^n(b, r_S, r_R)$, and $r'_S \leftarrow \phi(b, r_S, c)$. Then*

$$(b, r_S, c) \sim_{\delta, t} (1 - b, r'_S, c)$$

The standard asymptotic security notion requires $(\varepsilon(n), \delta(n), t(n))$ -security where for any polynomial $t(n)$ the correctness error $\varepsilon(n)$ and soundness error $\delta(n)$ are negligible. Note that a deniable encryption scheme cannot have perfect correctness $\varepsilon(n) = 0$, because the deniability property requires that for a given transcript c there must exist sender random tapes which are consistent with encrypting the bit 0 and 1 to c , respectively. To get a small correctness error, the set of “bad” random tapes for which decryption will fail must be much smaller than the set of “good” random tapes. It is exactly this ratio of good vs bad tapes that we will use for our counterexample.

Throughout, we will non-crucially assume that a transcript $c \in \text{supp}[\text{tr}_\psi^n(\cdot, \cdot, \cdot)]$ uniquely determines the bit to which the receiver will decrypt, and we’ll denote this bit by b_c .²³ For any c as above, we denote with \mathcal{R}_c (\mathcal{R}'_c) the set of sender random tapes which encrypt b_c ($1 - b_c$) to c .

$$\begin{aligned} \mathcal{R}_c &= \{r_S \mid c \in \text{supp}[\text{tr}_\psi^n(b_c, r_S, \cdot)]\}, \\ \mathcal{R}'_c &= \{r_S \mid c \in \text{supp}[\text{tr}_\psi^n(1 - b_c, r_S, \cdot)]\}. \end{aligned} \quad (33)$$

Let $\tilde{C} \sim \text{tr}_\psi^n(\star, \star, \star)$ be the distribution over random transcripts. Given $c \leftarrow \tilde{C}$, we know that the sender random tape is uniform over $\mathcal{R}_c \cup \mathcal{R}'_c$,²⁴ and decryption will fail if it’s in \mathcal{R}'_c . We can thus express the decryption error ε as

$$\mathbb{E}_{c \leftarrow \tilde{C}} [|\mathcal{R}'_c| / (|\mathcal{R}'_c| + |\mathcal{R}_c|)] = \varepsilon \quad (34)$$

Taking the inverse of the above equation we get

$$\varepsilon^{-1} = \mathbb{E}_{c \leftarrow \tilde{C}} [(|\mathcal{R}'_c| + |\mathcal{R}_c|) / |\mathcal{R}'_c|] = 1 + \mathbb{E}_{c \leftarrow \tilde{C}} [|\mathcal{R}_c| / |\mathcal{R}'_c|]. \quad (35)$$

A.2 The Counterexample from Deniable Encryption

In Theorem 3 below we assume that the deniable encryption scheme is balanced in the sense that the number of random tapes $|\mathcal{R}'_c|$ is the same for all $c \in \text{supp}[C]$. The only step in the proof of Theorem 3 where we’ll use this fact is (43). The restriction to balanced schemes is not completely without loss of generality. The proof for the general case seem more complicated and is omitted, let us just mention that it requires to understand which distributions (Z, C, B) satisfy (42) while maximising $\Pr_{(z, c, b) \leftarrow (Z, C, B)}[\mathbf{D}(z, c, b) = 1]$, which boils down to a simple optimisation problem. Unlike in the balanced case, where in (44) we upper bound this probability by $1/2$, for the general case we can only prove an upper bound of $1/\log |\mathcal{R}|$, where \mathcal{R} is the domain of the sender random tape. As a consequence, we also get a slightly weaker conclusion. When dropping the balanced requirement in the theorem below, in (36) we must replace $H_{1/2, s}^{\text{HILL}}$ with $H_{1/\log |\mathcal{R}|, s}^{\text{HILL}}$.

²³In general, c together with the receiver random tape r_R that was used to generate $c = \text{tr}_\psi^n(b, r_S, r_R)$ does determine the bit b_{c, r_R} to which the receiver will decrypt. To adapt our proof to this general case, one would have to consider the variable b_{c, R_R} (where R_R is random, conditioned on being consistent with c) instead the constant b_c throughout. This doesn’t add any technical difficulties, just requires some more bookkeeping.

²⁴Technically, one has to take union with multiplicity here as $\mathcal{R}_c \cap \mathcal{R}'_c$ must not be empty.

Theorem 3. *From a balanced (as explained in the paragraph above) $(\varepsilon(n), \delta(n), t(n))$ -secure sender deniable (bit) encryption scheme one can construct an ensemble $\{(X_n, Z_n, A_n)\}_{n \in \mathbb{N}}$ of distributions where A_n is a single bit, sampling a distribution requires two executions of the protocol and two invocation of the faking algorithm in expectation,²⁵ and where for some fixed polynomial $t(n)$ it holds²⁶*

$$H_{\delta+\varepsilon, t}^{\text{HILL}}(X_n|Z_n) - H_{1/2, s}^{\text{HILL}}(X_n|Z_n, A_n) \geq \log \varepsilon^{-1} - 2 \quad (36)$$

In particular, if $\varepsilon(n), \delta(n)$ are negligible for any polynomial $t(n)$, then by additionally conditioning on the single message bit A_n , the HILL entropy of the variable X_n decreases by a super-logarithmic $\omega(\log n)$ bits in quantity, even if we simultaneously allow for massive degradation in quality: from cryptographic strength $(\delta(n) + \varepsilon(n), t(n))$, where $\delta(n) + \varepsilon(n)$ is negligible for any polynomial $t(n)$, to $(1/2, s(n))$ for a fixed polynomial $s(n)$.

Before we prove the theorem, let us give the high level intuition. We sample a random bit B and sender and receiver random tapes R_S, R_R . We then compute the transcript $C = \text{tr}_{\psi}^n(B, R_S, R_R)$ and fake randomness $R'_S \leftarrow \phi(B, R_S, C)$. If $b_C \neq B$ (i.e., there's a decryption error), we resample until $b_C = B$ holds.

Using the notation introduced in (33), we now have $R_S \in \mathcal{R}_C$ and $R'_S \in \mathcal{R}'_C$. We show that the HILL entropy $H^{\text{HILL}}(R'_S|C)$ decreases significantly in quantity (but also quality) when additionally given B as outlined below.

Although given C we know that $R'_S \in \mathcal{R}'_C$, the deniability property implies that computationally we can't distinguish this R'_S from the uniform distribution over the larger set \mathcal{R}_C , and thus R'_S has HILL entropy $\log |\mathcal{R}_C|$. When additionally given the bit B , this is no longer true, as we can distinguish R'_S from any $r \notin \mathcal{R}'_C$ by checking if r is consistent with the transcript C and message bit $1 - B$. Thus, the HILL entropy $H_{\varepsilon, s}^{\text{HILL}}(R'_S|C, B)$ cannot be much larger than $\log |\mathcal{R}'_C|$ (concretely, it's at most $\log |\mathcal{R}'_C| + 1$ even if we allow a large distinguishing advantage of $\varepsilon = 1/2$). Summing up, revealing B decreases the HILL entropy by $\log |\mathcal{R}_C| - \log |\mathcal{R}'_C| - 1 = \log(|\mathcal{R}_C|/|\mathcal{R}'_C|) - 1$, which by (35) is at least $\geq \log \varepsilon^{-1} - 2$.

of Theorem 3. Consider a balanced $(\varepsilon(n), \delta(n), t(n))$ -secure sender-deniable bit encryption scheme. Let $\tilde{B} \leftarrow \{0, 1\}$ be a random bit, \tilde{R}_S be a random sender tape, $\tilde{C} \leftarrow \text{tr}_{\psi}^n(\tilde{B}, \tilde{R}_S, \star)$ a transcript and $\tilde{R}'_S \leftarrow \phi(\tilde{B}, \tilde{R}_S, \tilde{C})$ a fake sender random tape.

Define the distribution (B, C, R_S, R'_S) as $(\tilde{B}, \tilde{C}, \tilde{R}_S, \tilde{R}'_S)$ conditioned on $b_{\tilde{C}} = \tilde{B}$, i.e., there's no decryption error. Note that the probability that this conditions fails to hold is exactly the decryption error ε , and thus those distributions are ε close

$$(B, C, R_S, R'_S) \sim_{\varepsilon} (\tilde{B}, \tilde{C}, \tilde{R}_S, \tilde{R}'_S) . \quad (37)$$

We'll show that

$$H_{\delta+\varepsilon, t}^{\text{HILL}}(R'_S|C) - H_{1/2, s}^{\text{HILL}}(R'_S|C, B) \geq \log \varepsilon^{-1} - 2 \quad (38)$$

This proves the theorem by identifying (X_n, Z_n, A_n) in the statement of the theorem with (R'_S, C, B) . Note that, as stated in the theorem, one can sample (B, C, R_S, R'_S) by

²⁵ The exact number of expected executions/invocations is $\sum_{i=0}^{\infty} \varepsilon^i = 1 + \varepsilon + \varepsilon^2 \dots \leq 1 + 2\varepsilon \leq 2$. Alternatively, we can make this number exactly one (also in the worst case) at the prize of slightly decreasing the entropy gap by multiplying the right hand side of (36) with $(1 - \varepsilon)$.

²⁶More generally, we can replace the distinguishing advantage $1/2$ with $1 - 2^{-\tau}$ for any real-valued $\tau > 0$, but then must replace the -1 on the right-hand side with $-\tau$.

resampling $(\tilde{B}, \tilde{C}, \tilde{R}_S, \tilde{R}'_S)$ until the $b_{\tilde{C}} = \tilde{B}$ condition is met, which requires an expected $\sum_{i=0}^{\infty} \varepsilon^i = 1 + \varepsilon + \varepsilon^2 \dots \leq 2$ number tries, each requiring one execution of the protocol and one invocation of the faking algorithm.

We now prove (38). The deniability property states $(\tilde{B}, \tilde{R}_S, \tilde{C}) \sim_{\delta, t} (1 - \tilde{B}, \tilde{R}'_S, \tilde{C})$, and with (37) we further get $(B, R_S, C) \sim_{\delta+\varepsilon, t} (1 - B, R'_S, C)$. The latter also holds if we just consider the marginal distributions we get by ignoring the first term

$$(R_S, C) \sim_{\delta+\varepsilon, t} (R'_S, C) .$$

As given C , the random tape R_S is uniform over \mathcal{R}_C

$$\begin{aligned} \tilde{H}_{\infty}(R_S|C) &= -\log \mathbb{E}_{c \leftarrow C} \max_r \Pr[R_S = r|C = c] \\ &= -\log \mathbb{E}_{c \leftarrow C} 1/|\mathcal{R}_c| \\ &= \log \mathbb{E}_{c \leftarrow C} |\mathcal{R}_c| . \end{aligned}$$

By Definition 7 the two equations above imply

$$H_{\delta+\varepsilon, t}^{\text{HILL}}(R'_S|C) \geq \log \mathbb{E}_{c \leftarrow C} |\mathcal{R}_c| \quad (39)$$

Below, we prove

$$H_{1/2, s}^{\text{HILL}}(R'_S|C, B) \leq \log \mathbb{E}_{c \leftarrow C} |\mathcal{R}'_c| + 1 \quad (40)$$

But first observe that using (39) and (40) in the first, and (35) in the third step below (we discuss this step in more detail later), (38) follows as

$$\begin{aligned} H_{\delta+\varepsilon, t}^{\text{HILL}}(R'_S|C) - H_{1/2, s}^{\text{HILL}}(R'_S|C, B) &\geq \log \mathbb{E}_{c \leftarrow C} |\mathcal{R}_c| - \log \mathbb{E}_{c \leftarrow C} |\mathcal{R}'_c| - 1 \\ &= \log \mathbb{E}_{c \leftarrow C} \frac{|\mathcal{R}_c|}{|\mathcal{R}'_c|} - 1 \\ &\geq \log(\varepsilon^{-1} - 1) - 1 \geq \log \varepsilon^{-1} - 2 , \end{aligned}$$

Eq. (35) used in the third step above requires that c is sampled according to \tilde{C} , not C as above, but this doesn't matter as $C \sim \tilde{C}$.²⁷

It remains to prove (40). As $\phi(B, \cdot, C)$ maps \mathcal{R}_C to \mathcal{R}'_C and $R_S \in \mathcal{R}_C$, the fake transcript $R'_S \leftarrow \phi(B, R_S, C)$ must be in \mathcal{R}'_C .

Using (33), when given $(C, B = b_C)$, one can check if some sender random tape r is in \mathcal{R}'_C by verifying that $C \in \text{supp}[tr_{\psi}^n(1 - B, r, \cdot)]$, and this can be done in polynomial time.²⁸ Let D be a boolean circuit of polynomial size $s(n)$ which does that test:

$$\forall c \in \text{supp}[tr_{\psi}^n(\cdot, \cdot, \cdot)] : (D(r, c, b_c) = 1) \iff (r \in \mathcal{R}'_c) \quad (41)$$

Consider any distribution (Y, C, B) where the marginal distribution (C, B) is the same as in (R_S, C, B) (in particular $B = b_C$) and where Y has min-entropy

$$\tilde{H}_{\infty}(Y|C, B) = \tilde{H}_{\infty}(Y|C) > \log \mathbb{E}_{c \leftarrow C} |\mathcal{R}'_c| + 1 \quad (42)$$

²⁷ Note that $C \sim_{\varepsilon} \tilde{C}$ follows directly by (37). The stronger $C \sim \tilde{C}$ statement we use follows as we assume that the encryption scheme is balanced (cf. Footnote 22). This implies that when sampling C (recall this is done by resampling \tilde{C} until $b_{\tilde{C}} = \tilde{B}$), the probability that we resample is independent of the actual value of \tilde{C} , and thus C and \tilde{C} have the same distribution.

²⁸For this one only has to check that the messages computed by the sender using randomness r are consistent with the transcript C .

The first equality above holds as $B = b_C$ is determined by C . We can upper bound the probability that D outputs 1 on this distribution as (in step (43) below we use our assumption that the scheme is balanced in the sense that $|\mathcal{R}'_c|$ is a constant independent of c)

$$\begin{aligned}
\Pr_{(y,c,b)\leftarrow(Y,C,B)}[D(y,c,b) = 1] &\stackrel{(41)}{=} \Pr_{(y,c,b)\leftarrow(Y,C,B)}[y \in \mathcal{R}'_c] \\
&= \mathbb{E}_{c\leftarrow C} \Pr[Y \in \mathcal{R}'_c | C = c] \\
&\leq \mathbb{E}_{c\leftarrow C} \left[|\mathcal{R}'_c| \max_v \Pr[Y = v | C = c] \right] \\
&\leq \mathbb{E}_{c\leftarrow C} |\mathcal{R}'_c| \mathbb{E}_{c\leftarrow C} \left[\max_v \Pr[Y = v | C = c] \right] \quad (43)
\end{aligned}$$

If we take the logarithm of the equation above

$$\begin{aligned}
&-\log \Pr_{(y,c,b)\leftarrow(Z,C,B)}[D(y,c,b) = 1] \\
&\geq -\log \mathbb{E}_{c\leftarrow C} |\mathcal{R}'_c| - \log \mathbb{E}_{c\leftarrow C} \left[\max_v \Pr[Y = v | C = c] \right] \\
&\geq -\log \mathbb{E}_{c\leftarrow C} |\mathcal{R}'_c| + \tilde{H}_\infty(Y|C) \\
&\stackrel{(42)}{>} 1
\end{aligned}$$

Exponentiating the above we get

$$\Pr_{(y,c,b)\leftarrow(Y,C,B)}[D(y,c,b) = 1] < 1/2 \quad (44)$$

Using this we can lower bound D 's advantage as

$$\underbrace{\Pr_{(r,c,b)\leftarrow(R'_S,C,B)}[D(r,c,b) = 1]}_{=1 \text{ by (41)}} - \underbrace{\Pr_{(y,c,b)\leftarrow(Y,C,B)}[D(y,c,b) = 1]}_{<1/2 \text{ by (44)}} > 1/2$$

Which means

$$(Y, C, B) \not\sim_{1/2,s} (R'_S, C, B) \quad (45)$$

Now (40) follows by (42) and (45). \square

B Chain Rule for Unpredictability and Yao

In this section we prove chain rules of the form (4) for conditional Yao and Unpredictability entropy. Although it seems that the chain rules for these notions are folklore and the proofs are straightforward, we could not find a written account of this in the literature and thus provide it here. Below we define Yao and unpredictability entropy, which we already informally defined in Section 1.2.

Definition 13 ([Yao82, BSW03, HLR07]). *Let (X, Z) be a joint distribution of random variables. Then X has **conditional Yao entropy** k conditioned on Z , denoted by $H_{\varepsilon,s}^{\text{YAO}}(X|Z) \geq k$, if for every m , and any pair of circuits C, D of total size s where C has output length m , it holds that*

$$\Pr_{(x,z)\leftarrow(X,Z)}[D(C(x,z), z) = x] \leq 2^{m-k} + \varepsilon$$

So, C compresses X to a string of length m , and D tries to recover X from this string. Both circuits also get the conditional part Z as input, and the probability that D recovers X correctly is exponentially small in the gap $k - m$ between the Yao entropy and the length of the compressed X . The definition considers an additional “smoothness parameter” ε , which is simply added to D ’s success probability, and will be useful when comparing Yao-entropy with other pseudentropy notions.

Definition 14 ([HLR07]). *Let (X, Z) be a joint distribution of random variables. Then X has **unpredictability entropy** k conditioned on Z , denoted by $H_{\varepsilon, s}^{\text{unp}}(X|Z) \geq k$, if there exists a joint distribution (Y, Z) such that $(X, Z) \sim_{\varepsilon, s} (Y, Z)$ and for all circuits C of size s*

$$\Pr_{(y,z) \leftarrow (Y,Z)}[C(z) = y] \leq 2^{-k}$$

[HLR07] prove the following simple relations amongst HILL, Yao and unpredictability entropy

$$H_{\varepsilon, s}^{\text{YAO}}(X|Z) \geq H_{\varepsilon, s}^{\text{unp}}(X|Z) \geq H_{\varepsilon, s}^{\text{HILL}}(X|Z)$$

We first prove the chain rule for Yao entropy.

Lemma 10 (Chain Rule for conditional Yao Entropy). *For any joint distribution (X, Z, A) where $|A| = \ell$*

$$H_{\varepsilon, s2^\ell + O(2^\ell|X|)}^{\text{YAO}}(X|Z) \geq k \Rightarrow H_{\varepsilon, s}^{\text{YAO}}(X|Z, A) \geq k - \ell$$

Proof. We’ll prove the contraposition

$$H_{\varepsilon, s}^{\text{YAO}}(X|Z, A) < k - \ell \Rightarrow H_{\varepsilon, s2^\ell + O(2^\ell|X|)}^{\text{YAO}}(X|Z) < k. \quad (46)$$

The left-hand of (46) means that for some $m \in \mathbb{N}$, there exist circuits C, D of total size s where C has output length m and

$$\Pr_{(x,z,a) \leftarrow (X,Z,A)}[D(C(x, z, a), a, z) = x] > 2^{m-(k-\ell)} + \varepsilon. \quad (47)$$

We define a circuit C' with output length $m' = m + \ell$ as $C'(x, z) = (C(x, z, a), a)$ where $a \in \{0, 1\}^\ell$ is chosen so that $D(C(x, z, a), z, a) = x$ holds, and $a = 0^\ell$ if no such a exists. Now $D(C(x, z, a), z, a) = x$ implies $D(C'(x, z), z) = x$ for any (x, z, a) , using this we see that (47) implies

$$\Pr_{(x,z) \leftarrow (X,Z)}[D(C'(x, z), z) = x] > 2^{m-(k-\ell)} + \varepsilon = 2^{m'-k} + \varepsilon.$$

As C', D can be realized with total size at most $s2^\ell + O(2^\ell|X|)$,²⁹ the above is equivalent to the right-hand side of (46). \square \square

Below we state the chain rule for unpredictability entropy. Note that unlike in the chain rule for HILL or Yao entropy, here there’s basically no loss in circuit size.

Lemma 11 (Chain Rule for Unpredictability Entropy). *For any joint distribution (X, Z, A) where $|A| = \ell$*

$$H_{\varepsilon, s}^{\text{unp}}(X|Z) \geq k \Rightarrow H_{\varepsilon, s-O(\ell)}^{\text{unp}}(X|Z, A) \geq k - \ell$$

²⁹ $C'(x, z)$ computes $x_a = D(C(x, z, a), a, z)$ for all $a \in \{0, 1\}^\ell \setminus 0^\ell$ and outputs $(C(x, z, a), a)$ if $x_a = x$ for some a , and $(C(x, z, 0^\ell), 0^\ell)$ otherwise. This can be done by a circuit of size $2^\ell|C| + (2^\ell - 1)|D|$ plus $O(2^\ell|X|)$ extra gates required to check if $x = x_a$ for all $a \in \{0, 1\}^\ell$. As $|C| + |D| \leq s$ we get $|C'| + |D| \leq 2^\ell s + O(2^\ell|X|)$.

Proof. We'll prove the contraposition

$$H_{\varepsilon,s}^{\text{unp}}(X|Z, A) < k - \ell \Rightarrow H_{\varepsilon,s+O(\ell)}^{\text{unp}}(X|Z) < k \quad (48)$$

The left-hand side of (48) means that for every distribution (Y, Z, A) where $(X, Z, A) \sim_{\varepsilon,s} (Y, Z, A)$ there exists a circuit C of size at most s such that

$$\Pr_{(y,z,a) \leftarrow (Y,Z,A)}[C(z, a) = y] > 2^{-k+\ell} \quad (49)$$

Define the circuit C' as follows: $C'(z)$ picks a random $a' \leftarrow \{0, 1\}^\ell$ and outputs $C(z, a')$. Below all probabilities are over sampling $(y, z, a) \leftarrow (Y, Z, A)$ and $a' \leftarrow \{0, 1\}^\ell$, in the last step we use (49)

$$\begin{aligned} \Pr[C'(z) = y] &\geq \Pr[C'(z) = y|a = a']\Pr[a = a'] \\ &= \Pr[C(z, a) = y]2^{-\ell} \\ &> 2^{-k} \end{aligned}$$

This, together with the fact that $(X, Z, A) \sim_{\varepsilon,s} (Y, Z, A)$ implies $(X, Z) \sim_{\varepsilon,s} (Y, Z)$ (as ignoring part of a distribution cannot make distinguishing easier) and $|C'| = |C| + O(\ell)$, implies the right-hand side of (48). \square \square