On Pseudorandom Encodings

Thomas Agrikola, Geoffroy Couteau, Yuval Ishai, Stanislaw Jarecki, Amit Sahai

► To cite this version:

Thomas Agrikola, Geoffroy Couteau, Yuval Ishai, Stanislaw Jarecki, Amit Sahai. On Pseudorandom Encodings. TCC 2020 - Theory of Cryptography Conference, Dec 2020, Durham, United States. pp.639-669, 10.1007/978-3-030-64381-2_23 . hal-03374163

HAL Id: hal-03374163

<https://hal.science/hal-03374163>

Submitted on 11 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Pseudorandom Encodings

Thomas Agrikola*
Karlsruhe Institute of Technology
thomas.agrikola@kit.edu

Geoffroy Couteau†
IRIF, Paris-Diderot University,
CNRS
couteau@irif.fr

Yuval Ishai‡
Technion
yuvali@cs.technion.ac.il

Stanisław Jarecki§
UC Irvine
stasio@ics.uci.edu

Amit Sahai¶
UCLA
sahai@cs.ucla.edu

October 6, 2020

Abstract.

We initiate a study of *pseudorandom encodings*: efficiently computable and decodable encoding functions that map messages from a given distribution to a random-looking distribution. For instance, every distribution that can be perfectly and efficiently compressed admits such a pseudorandom encoding. Pseudorandom encodings are motivated by a variety of cryptographic applications, including password-authenticated key exchange, “honey encryption” and steganography.

The main question we ask is whether *every* efficiently samplable distribution admits a pseudorandom encoding. Under different cryptographic assumptions, we obtain positive and negative answers for different flavors of pseudorandom encodings, and relate this question to problems in other areas of cryptography. In particular, by establishing a two-way relation between pseudorandom encoding schemes and efficient invertible sampling algorithms, we reveal a connection between adaptively secure multiparty computation for randomized functionalities and questions in the domain of steganography.

*Supported by ERC Project PREP-CRYPTO 724307 and by the German Federal Ministry of Education and Research within the framework of the project KASTEL_SKI in the Competence Center for Applied Security Technology (KASTEL).

†Supported by ERC Projects PREP-CRYPTO 724307 and CryptoCloud 339563. Work done in part while visiting UCLA and the Technion.

‡Supported by ERC Project NTSC (742754), NSF-BSF grant 2015782, BSF grant 2018393, ISF grant 2774/20, and a grant from the Ministry of Science and Technology, Israel and Department of Science and Technology, Government of India. Work done in part while visiting UCLA.

§Supported by the NSF SaTC award 1817143.

¶Supported in part by DARPA SAFEWARE and SIEVE awards, NTT Research, NSF Frontier Award 1413955, and NSF grant 1619348, BSF grant 2012378, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the ARL under Contract W911NF-15-C-0205. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, NTT Research, or the U.S. Government.

Table of Contents

1	Introduction	1
1.1	Flavors of pseudorandom encoding	2
1.2	Implications and applications of our results	6
1.3	Negative results for stronger notions of pseudorandom encodings	8
1.4	Open questions and subsequent work	10
2	Overview of techniques	11
3	Preliminaries	15
4	The pseudorandom encoding hypothesis	15
4.1	The pseudorandom encoding hypothesis with setup	18
4.2	Static-to-adaptive transformation	20
5	Pseudorandom encodings and invertible sampling	23
5.1	The invertible sampling hypothesis	23
5.2	The invertible sampling hypothesis with setup	24
5.3	Equivalence of pseudorandom encodings and invertible sampling	25
5.3.1	Every inverse samplable distribution can be pseudorandomly encoded	26
5.3.2	Every pseudorandomly encodable distribution can be inversely sampled	27
6	Classification of the different flavors of pseudorandom encodings	29
6.1	Deterministic encoding	29
6.1.1	Information theoretic guarantees and compression	31
6.1.2	Computational guarantees and pseudoentropy	34
6.2	Randomized encoding	38
6.2.1	(Generalized) extractable one-way functions	38
6.2.2	Information theoretic guarantees and privately verifiable GEOWFs	40
6.2.3	Computational guarantees and EOWFs with common auxiliary information	41
6.3	Static pseudorandom encodings with universal setup from IO	46
6.4	Bootstrapping pseudorandom encodings with a common random string	48
7	Relations and applications of pseudorandom encodings	51
7.1	Fully adaptively secure multi-party computation in the CRS model	51
7.1.1	Pseudorandom encodings imply UC-secure AMPC in the plain model	53
7.1.2	AMPC with global setup implies adaptive pseudorandom encodings with setup	55
7.2	Honey encryption	57
7.3	Keyless steganography	58
7.4	Covert multi-party computation	59
7.4.1	A compiler for covert oblivious transfer	60
7.4.2	A general compiler for covert protocols	61
7.5	Deniable encryption	62
7.6	Non-committing encryption	64
7.7	Super-polynomial encoding	65

1 Introduction

The problem of *compression* has been extensively studied in the field of information theory and, more recently, in computational complexity and cryptography [GS85; Wee04; TVZ05; HLR07]. Informally, given a distribution X , compression aims to efficiently encode samples from X as short strings while at the same time being able to efficiently recover these samples. While the typical information-theoretic study of compression considers the case of compressing multiple independent samples from the same source X , its study in computer science, and in particular in this work, considers the “single-shot” case. Compression in this setting is closely related to *randomness condensers* [RR99; TV00; TUZ01; DRV12] and *resource-bounded Kolmogorov complexity* [LV90; LV19] – two well-studied problems in computational complexity. Randomness condensers, which relax randomness extractors, are functions that efficiently map an input distribution into an output distribution with a higher entropy rate. A randomness condenser can be viewed as an efficient compression algorithm, without a corresponding efficient decompression algorithm. The resource-bounded Kolmogorov complexity of a string is the smallest description length of an efficient program that outputs this string. This program description can be viewed as a compressed string, such that decoding is efficiently possible, while finding the compressed string may be inefficient.

An important property of efficient compression algorithms, which combines the efficiency features of randomness condensers and resource-bounded Kolmogorov complexity, is their ability to efficiently produce “random-looking” outputs while allowing the original input to be efficiently recovered. Despite the large body of work on compression and its computational variants, this fundamental property has, to our knowledge, never been the subject of a dedicated study. In this work, we fill this gap by initiating such a study. Before formalizing the problem, we give a simple motivating example.

Consider the goal of encrypting a sample x from a distribution X (say, a random 5-letter English word from the Merriam-Webster Dictionary) using a low-entropy secret key k . Applying a standard symmetric-key encryption scheme with a key derived from k gives rise to the following brute-force attack: Try to decrypt with different keys until obtaining x' in the support of X . In the typical case that wrong keys always lead to x' outside the support of X , this attack successfully recovers x . Variants of this attack arise in different scenarios, including password-authenticated key exchange [BM92], honey encryption [JR14], subliminal communication and steganography [HPRV19], and more. A natural solution is to use perfect compression: if x can be compressed to a *uniformly random* string $\hat{x} \in \{0, 1\}^n$ before being encrypted, it cannot be distinguished from another random string $\hat{x}' \in \{0, 1\}^n$ obtained by trying the wrong key. Note, however, that compression may be an overkill for this application. Instead, it suffices to efficiently encode x into a (possibly longer) *pseudorandom* string from which x can be efficiently decoded. This more general solution motivates the question we consider in this work.

Encoding into the uniform distribution. We initiate the study of encoding distributions into a random-looking distribution. Informally, we say that a distribution ensemble X_λ admits a *pseudorandom encoding* if there exist efficient encoding and decoding algorithms (E_X, D_X) , where D_X is deterministic, such that

$$\Pr [y \leftarrow X_\lambda : D_X(E_X(y)) = y] \text{ is overwhelming and} \quad (1)$$

$$\{y \leftarrow X_\lambda : E_X(y)\} \approx U_{n(\lambda)}. \quad (2)$$

Here, “ $\approx$ ” denotes some notion of indistinguishability (we will consider both computational and statistical indistinguishability), and the probability is over the randomness of both E_X and X_λ . The polynomial $n(\lambda)$ denotes the output length of the encoding algorithm E_X . We refer to Equation (1) as *correctness* and to Equation (2) as *pseudorandomness*. It will also be useful to consider distribution ensembles parameterized by an input m from a language L . We say that

such a distribution ensemble $(X_m)_{m \in L}$ admits a pseudorandom encoding if there exist efficient algorithms (E_X, D_X) as above satisfying correctness and pseudorandomness for all $m \in L$, where E_X and D_X both additionally receive m as input. Note that we insist on the decoding algorithm being efficient. This is required for our motivating applications.¹ Note also that encoding and decoding above are *keyless*; that is, we want encoded samples to be close to uniform *even though anyone can decode them*. This is a crucial distinction from, for instance, encryption schemes with pseudorandom ciphertexts, which look uniformly distributed to everyone except the owner of the decryption key, and cannot be efficiently decrypted except by the owner of the decryption key. Here, we seek to simultaneously achieve pseudorandomness and correctness for all parties.

Our motivation for studying pseudorandom encodings stems from the fact that this very natural problem appears in a wide variety of – sometimes seemingly unrelated – problems in cryptography. We already mentioned steganography, honey encryption, and password-authenticated key exchange; we will cover more such connections in this work. Yet, this notion of encoding has to our knowledge never been studied systematically. In this work we study several natural flavors of this notion, obtain positive and negative results about realizing them, and map their connections with other problems in cryptography.

The main focus of this work is on the hypothesis that *all* efficiently samplable distributions admit a pseudorandom encoding. Henceforth, we refer to this hypothesis the *pseudorandom encoding hypothesis* (PREH).

For describing our results, it will be convenient to use the following general notion of efficiently samplable distributions. A distribution family ensemble $(X_m)_{m \in L}$ (for some language $L \subseteq \{0, 1\}^*$) is efficiently samplable if there exists a probabilistic polynomial time (PPT) algorithm S such that $S(m)$ is distributed according to X_m for every $m \in L$. In case the distribution does not depend on additional inputs, L can be considered equal to $\mathbb{N}$.

Overview of contributions. Following is a brief summary of our main contributions along with pointers to the relevant technical sections. We will give an expanded overview of the contributions and the underlying techniques in the rest of this section.

- We provide a unified study of different flavors of pseudorandom encodings (PRE) and identify computational, randomized PRE in the CRS model as a useful and achievable notion (Section 6 and Theorem 15).
- We establish a two-way relation between PRE and the previously studied notion of invertible sampling (Theorem 2). This reveals unexpected connections between seemingly unrelated problems in cryptography (e.g., between adaptively secure computation for general functionalities and “honey encryption”).
- We bootstrap “adaptive PRE” from “static PRE” (Theorem 1). As a consequence, one can base succinct adaptively secure computation on standard iO as opposed to subexponential iO [CsW19] (Corollary 3).
- We use PRE to obtain a compiler from standard secure multiparty computation (MPC) protocols to *covert* MPC protocols (Section 7.4).

1.1 Flavors of pseudorandom encoding

The notion of pseudorandom encoding has several natural flavors, depending on whether the encoding algorithm is allowed to use randomness or not, and whether the pseudorandomness property satisfies a computational or information-theoretic notion of indistinguishability. We denote the corresponding hypotheses that every efficiently samplable distribution can be pseu-

¹ Without this requirement, the problem can be solved using non-interactive commitment schemes with the additional property that commitments are pseudorandom (which exist under standard cryptographic assumptions).

dorandomly encoded according to the above variants as $\text{PREH}_{\approx_c}^{\text{rand}}$, $\text{PREH}_{\equiv_s}^{\text{rand}}$, $\text{PREH}_{\approx_c}^{\text{det}}$ and $\text{PREH}_{\equiv_s}^{\text{det}}$.²

Further, we explore relaxations which rely on a trusted setup assumption: we consider the pseudorandom encoding hypothesis in the *common reference string model*, in which a common string sampled in a trusted way from some distribution is made available to the parties. This is the most common setup assumption in cryptography and it is standard to consider the feasibility of cryptographic primitives in this model to overcome limitations in the plain model. That is, we ask whether for every efficiently samplable distribution X , there exists an efficiently samplable CRS distribution and efficient encoding and decoding algorithms (E_X, D_X) as above, such that correctness and pseudorandomness hold, where the encoding and decoding algorithm as well as the distinguisher receive the CRS as input, and the distributions in Equations (1) and (2) are additionally over the choice of the CRS.

Considering distributions which may depend on an input $m \in L$ further entails two different flavors. On the one hand, we consider the notion where inputs m are chosen adversarially but *statically* (that is, independent of the CRS) and, on the other hand, we consider the stronger notion where inputs m are chosen adversarially and *adaptively* depending on the CRS. We henceforth denote these variants by the prefix “c” and “ac”, respectively.

Static-to-adaptive transformation. The adaptive notion, where inputs may be chosen depending on the CRS, is clearly stronger than the static notion. However, surprisingly, the very nature of pseudorandom encodings allows one to apply an indirection argument similar to the one used in [HJKSWZ16; CPR17; CPV17], which yields a static-to-adaptive transformation.

Theorem (informal). *If all efficiently samplable distributions can be pseudorandomly encoded in the CRS model with a static choice of inputs, then all efficiently samplable distributions can be pseudorandomly encoded in the CRS model with an adaptive choice of inputs.*

Static-to-adaptive transformations in cryptography are generally non-trivial, and often come at a big cost in security when they rely on a “complexity leveraging” technique. This connection and its application we will discuss below are a good demonstration of the usefulness of the notion of pseudorandom encodings.

Relaxing compression. The notion of statistical deterministic pseudorandom encodings recovers the notion of optimal compression. Hence, this conflicts with the existence of one-way functions.³ In our systematic study of pseudorandom encodings, we gradually relax perfect compression in several dimensions, while maintaining one crucial property – the indistinguishability of the encoded distribution from true randomness.

Example. To illustrate the importance of this property, we elaborate on the example we outline at the beginning of the introduction, focusing more specifically on password-authenticated key exchange (PAKE). A PAKE protocol allows two parties holding a (low entropy) common password to jointly and confidentially generate a (high entropy) secret key, such that the protocol is resilient against offline dictionary attacks, and no adversary can establish a shared key with a party if he does not know the matching password. A widely used PAKE protocol due to Bellare and Merritt [BM92] has a very simple structure: the parties use their low-entropy password to encrypt the flows of a key-exchange protocol using a block cipher. When the block cipher is modeled as a random cipher, it has the property that decrypting a ciphertext (of an arbitrary

² We note that not all efficiently samplable distributions can be pseudorandomly encoded with a deterministic encoding algorithm. For instance, a distribution which has one very likely event and many less likely ones requires one specific encoding to appear with high probability. Thus, we formally restrict the deterministic variants of the pseudorandom encoding hypothesis to only hold for “compatible” samplers, which still results in interesting connections. In this overview, however, we ignore this restriction.

³ If perfect compression exists, pseudorandom generators cannot exist (observation attributed to Levin in [GS85]).

plaintext) under an incorrect secret key yields a fresh random plaintext. Thus, Bellare and Merritt point out that the security of their PAKE protocol requires that “the message to be encrypted by the password must be indistinguishable from a random number.” This is easy to achieve for Diffie-Hellman key exchange over the multiplicative group of integers modulo a prime p . However, for elliptic curve groups this is no longer the case, and one needs to resort to alternative techniques including nontrivial *point compression* algorithms that compress the representation of a random group element into a nearly uniform bitstring [BMN01].

Clearly, our relaxation of compression does not preserve the useful property of obtaining outputs that are *shorter* than the inputs. However, the remaining pseudorandomness property is good enough for many applications.

In the following, we elaborate on our weakest notion of pseudorandom encodings, that is, pseudorandom encodings allowing the encoding algorithm to be randomized and providing a computational pseudorandomness guarantee. We defer the discussion on the stronger statistical or deterministic variants to Section 1.3, where we derive negative results for most of these stronger notions, which leaves computational randomized pseudorandom encodings as the “best possible” notion that can be realized for general distributions.

1.1.1 Randomized, computational pseudorandom encodings. Computational randomized pseudorandom encodings allow the encoding algorithm to be randomized and require only computational pseudorandomness.

Relation to invertible sampling. We show a simple but unexpected connection with the notion of *invertible sampling* [CFGN96; DN00; GKMRV00]. Informally, invertible sampling refers to the task of finding, given samples from a distribution, random coins that *explain* the sample. Invertible sampling allows to *obviously* sample from distributions, that is, sampling from distributions *without knowing the corresponding secrets*. This can be useful for, e.g., sampling common reference strings without knowing the random coins or public keys without knowing the corresponding secret keys. A natural relaxation of this notion was systematically studied by Ishai, Kumarasubramanian, Orlandi and Sahai [IKOS10]. Concretely, a PPT sampler S is *inverse samplable* if there exists an alternative PPT sampler $\bar{S}$ and a PPT inverse sampler $\bar{S}^{-1}$ such that

$$\begin{aligned} \{y \leftarrow S(1^\lambda) : y\} &\approx_c \{y \leftarrow \bar{S}(1^\lambda) : y\}, \\ \{y \leftarrow \bar{S}(1^\lambda; r) : (r, y)\} &\approx_c \{y \leftarrow \bar{S}(1^\lambda) : (\bar{S}^{-1}(1^\lambda, y), y)\}. \end{aligned}$$

Note that the inverse sampling algorithm is only required to efficiently inverse-sample from another distribution $\bar{S}$, but this distribution must be computationally close to the distribution induced by S . The main question studied in [IKOS10] is whether *every* efficient sampler admits such an invertible sampler. They refer to this hypothesis as the *invertible sampling hypothesis* (ISH), and show that ISH is equivalent to adaptive MPC for general randomized functionalities that may hide their internal randomness. In this work, we show the following two-way relation with pseudorandom encoding.

Theorem (informal). *A distribution admits a pseudorandom encoding if and only if it admits invertible sampling.*

Intuitively, the efficient encoding algorithm corresponds to the inverse sampling algorithm, and decoding an encoded string corresponds to sampling with the de-randomized alternative sampler $\bar{S}$. This equivalence immediately extends to all variants of pseudorandom encodings and corresponding variants of invertible sampling we introduce in this work. Invertible sampling is itself connected to other useful cryptographic notions, such as oblivious sampling, trusted common reference string generations, and adaptively secure computation (which we will elaborate upon below).

Building on this connection, the impossibility result of [IKOS10] translates to our setting. On a high level, extractable one-way functions (EOWFs) conflict with invertible sampling because they allow to extract a “secret” (in this case a pre-image) from an image, independently of how it was computed. This conflicts with invertible sampling because invertible sampling is about sampling without knowing the secrets.

Theorem (informal, [IKOS10]). *Assuming the existence of extractable one-way functions (EOWF) and a non-interactive zero-knowledge proof system, $\text{PREH}_{\approx_c}^{\text{rand}}$ does not hold.*

This suggests that towards a realizable notion of pseudorandom encodings, a further relaxation is due. Thus, we ask whether the above impossibility result extends to the CRS model. In the CRS model, the above intuition why ISH conflicts with EOWFs fails, because the CRS can contain an obfuscated program that samples an image using some secret, but does not output this secret.

Dachman-Soled, Katz, and Rao [DKR15] (building on the universal deniable encryption construction of Sahai and Waters [SW14]) construct a so-called “explainability compiler” that implies $\text{cISH}_{\approx_c}^{\text{rand}}$ based on indistinguishability obfuscation⁴ (iO). By our equivalence theorem above, this implies pseudorandom encodings for all efficiently samplable distributions in the CRS model, with static choice of inputs, from iO. Invoking the static-to-adaptive transformation detailed above, this also applies to the adaptive variant.

Theorem (informal). *Assuming the existence of (polynomially secure) indistinguishability obfuscation and one-way functions, $\text{acPREH}_{\approx_c}^{\text{rand}}$ holds.*

Note that [IKOS10] claim that their impossibility result extends to the CRS model, whereas the above theorem seems to suggest the opposite. We show that technically the result of [IKOS10] does extend to the CRS model at the cost of assuming *unbounded auxiliary-input* extractable one-way functions, a strong flavor of EOWFs that seems very unlikely to exist but cannot be unconditionally ruled out.

Theorem (informal). *Assuming the existence of extractable one-way functions with unbounded common auxiliary input and a non-interactive zero-knowledge proof system, $\text{cPREH}_{\approx_c}^{\text{rand}}$ does not hold.*

In fact, this apparent contradiction has been the source of some confusion in previous works: the work of [IKOS10] makes an informal claim that their impossibility result for ISH extends to the CRS model. However, due to the connection between ISH and adaptively secure MPC (which we will discuss in more details later on), this claim was challenged in [DKR15]: the authors achieve a construction of adaptively secure MPC for all functionalities assuming iO, which seemingly contradicts the claim of [IKOS10]. The authors of [DKR15] therefore stated that the “impossibility result of Ishai et al. [...] does not hold in the CRS model.” Our extension clarifies that the distinction is in fact more subtle: the result of [IKOS10] *does* extend to the CRS model, but at the cost of assuming EOWF with *unbounded auxiliary inputs*. This does not contradict the constructions based on iO, because iO and EOWF with unbounded auxiliary inputs are known to be contradictory [BCPR14].

Overview. In Figure 1, we provide a general summary of the many flavors of the pseudorandom encoding hypothesis, and how they relate to a wide variety of other primitives.

⁴ Informally, an iO scheme is a PPT algorithm that takes as input a circuit C and produces another circuit $\text{iO}(C)$ such that C and $\text{iO}(C)$ compute the same function, but $\text{iO}(C)$ is unintelligible in the following sense. If two circuits C_1 and C_2 compute the same function, then $\text{iO}(C_1)$ and $\text{iO}(C_2)$ are computationally indistinguishable. The notion of iO was introduced in [BGIRSVY01] and first instantiated in [GGHRSW13].

Fig. 1. An overview of the relations between the pseudorandom encoding hypothesis and other fields of cryptography and computational complexity theory. For simplicity, our static-to-adaptive transformation only appears in the computational, randomized setting in this overview, but also applies to the other settings. (Since the deterministic variants of the pseudorandom encoding hypothesis are impossible for some pathologic samplers, the arrows between deterministic and randomized variants of the pseudorandom encoding hypothesis are to be read as if the deterministic variant is true *for some sampler*, then the corresponding randomized variant is true *for that sampler*.)

Further relaxation. We further study an additional relaxation of pseudorandom encodings, where we allow the encoding algorithm to run in super-polynomial time. We show that this relaxed variant can be achieved from cryptographic primitives similar to *extremely lossy functions* [Zha16], which can be based on the exponential hardness of the decisional Diffie-Hellman problem – a strong assumption, but (still) more standard than indistinguishability obfuscation. However, the applicability of the resulting notion turns out to be rather restricted.

1.2 Implications and applications of our results

In this section, we elaborate on the implications of the techniques we develop and the results we obtain for a variety of other cryptographic primitives.

1.2.1 New results for adaptively secure computation. As mentioned above, a sampler admits invertible sampling if and only if it can be pseudorandomly encoded. A two-way connection between invertible sampling and *adaptively* secure MPC for general randomized functionalities was established in [IKOS10]. An MPC protocol allows two or more parties to jointly evaluate a (possibly randomized) functionality $\mathcal{F}$ on their inputs without revealing anything to each other except what follows from their inputs and outputs. This should hold even in the presence of an

adversary who can corrupt *any* number of parties in an adaptive (sequential) fashion. When we write “adaptive MPC”, we mean adaptive MPC for *all* randomized functionalities. This should be contrasted with weaker notions of adaptive MPC for strict subsets of corrupted parties [BH92; CFGN96; GS12] or for adaptively well-formed functionalities⁵ [CLOS02] which can both be done from mild assumptions. The connection from [IKOS10] shows that adaptive MPC for all randomized functions is possible if and only if every PPT sampler admits invertible sampling, i.e., the invertible sampling hypothesis is true.

We show that this result generalizes to the global CRS model. More precisely, we prove the adaptive variant of the pseudorandom encoding hypothesis in the CRS model $\text{acPREH}_{\approx_c}^{\text{rand}}$ is equivalent to adaptive MPC in the global CRS model.⁶

As detailed above, the static pseudorandom encoding hypothesis $\text{cPREH}_{\approx_c}^{\text{rand}}$ in the CRS model follows from iO (and one-way functions). Applying our static-to-adaptive transformation, the same holds for the adaptive variant. Thus, we obtain the first instantiation of an adaptive explainability compiler [DKR15] without complexity leveraging and, hence, based only on polynomial hardness assumptions. The recent work of Cohen, Shelat, and Wichs [CsW19] uses such an adaptive explainability compiler to obtain succinct adaptive MPC, where “succinct” means that the communication complexity is sublinear in the complexity of the evaluated function. Due to our instantiation of $\text{acPREH}_{\approx_c}^{\text{rand}}$ from polynomial iO , we improve the results of [CsW19] by relaxing the requirement for subexponentially secure iO to polynomially secure iO in a black-box way.

Corollary (informal). *Assuming the existence of polynomially secure indistinguishability obfuscation and the adaptive hardness of the learning with errors problem, then malicious, two-round, UC-secure adaptive MPC and sublinear communication complexity is possible (in the local CRS model, for all deterministic functionalities).*

1.2.2 Steganography and covert multi-party computation. We explore the connection of the pseudorandom encoding hypothesis to various flavors of steganography. The goal of steganography, informally, is to embed secret messages in distributions of natural-looking messages, in order to hide them from external observers. While the standard setting for steganography relies on shared secret keys to encode the messages, we show that pseudorandom encodings naturally give rise to a strong form of *keyless steganography*. Namely, one can rely on pseudorandom encodings to encode any message into an innocent-looking distribution, without truly hiding the message (since anyone can decode the stream), but providing *plausible deniability*, in the sense that, even with the decoded message, it is impossible to tell apart whether this message was indeed encoded by the sender, or whether it is simply the result of decoding the innocent distribution.

Corollary (informal). *Assuming pseudorandom encodings, then there exists a keyless steganographic protocol which provides plausible deniability.*

Plausible deniability is an important security notion; in particular, an important cryptographic primitive in this area is the notion of (sender-)deniable encryption [CDNO97], which is known to exist assuming indistinguishability obfuscation [SW14]. Deniable encryption enables to “explain” ciphertexts produced for some message to any arbitrary other message by providing corresponding random coins for a faked encryption process. We view it as an interesting open problem to build deniable encryption under the pseudorandom encoding hypothesis together with more standard cryptographic primitives; we make a first step in this direction and show the following:

⁵ Adaptively well-formed functionalities do not hide internal randomness.

⁶ Together with the conflict between $\text{cPREH}_{\approx_c}^{\text{rand}}$ and EOWFs with unbounded auxiliary input, this corrects a claim made in [DKR15] that the impossibility result of adaptive MPC from [IKOS10] would not extend to the CRS model.

the *statistical* variant of pseudorandom encodings, together with the existence of public-key encryption, implies deniable encryption. Interestingly, we also show that the computational randomized pseudorandom encoding hypothesis suffices to imply non-committing encryption, a weaker form of deniable encryption allowing to explain only *simulated* ciphertexts to arbitrary messages [CFGN96].

Covert secure computation. Covert MPC [vHL05; CGOS07] is an intriguing flavor of MPC that aims at achieving the following strong security guarantee: if the output of the protocol is not “favorable,” the transcript of the interaction should not leak any information to the parties, including whether any given party was *actually taking part in the protocol*. This strong form of MPC aims at providing security guarantees when the very act of starting a computation with other parties should remain hidden. As an example [vHL05], suppose that a CIA agent who infiltrated a terrorist group wants to make a handshake with another individual to find out whether she is also a CIA agent. Here, we show that pseudorandom encodings give rise to a general compiler transforming a standard MPC protocol into a covert one, in a round-preserving way. The idea is to encode each round of the protocol such that encoded messages look random. Together with the equivalence between adaptively secure MPC and pseudorandom encodings, this gives a connection between two seemingly unrelated notions of secure computation.

Corollary (informal). *Assuming adaptively secure MPC for all functionalities, there exists a round-preserving compiler that transforms a large class of “natural” MPC protocols into covert MPC protocols (in the static, semi-honest setting).*

1.2.3 Other results. Due to our infeasibility results of $\text{PREH}_{\equiv_s}^{\text{rand}}$, distribution transforming encoders (DTEs) for all efficiently samplable distributions are infeasible. Even the computational relaxation of DTEs is infeasible assuming extractable one-way functions. Since all currently known constructions of honey encryption rely on DTEs, we conditionally refute the existence of honey encryption based on the DTE-then-encrypt framework from [JR14]. On the positive side, due to our feasibility result of $\text{acPREH}_{\approx_c}^{\text{rand}}$, computational honey encryption is feasible in the CRS model.

Theorem (informal). *Assuming $\text{acPREH}_{\approx_c}^{\text{rand}}$ and a suitable symmetric-key encryption scheme (modeled as a random cipher), computational honey encryption for all efficiently samplable distributions exists in the CRS model.*

1.3 Negative results for stronger notions of pseudorandom encodings

Below we describe how we gradually relax optimal compression via different notions of pseudorandom encodings and derive infeasibility results for all variants of pseudorandom encodings which restrict the encoding algorithm to be deterministic or require an information-theoretic pseudorandomness guarantee. This leaves computational randomized pseudorandom encodings as the best possible achievable notion.

1.3.1 Deterministic, statistical pseudorandom encodings. The notion of pseudorandom encodings with a deterministic encoding algorithm and information-theoretic indistinguishability is perhaps the simplest notion one can consider. As we will prove in this paper, this notion recovers the notion of optimal compression: since the encoding algorithm for some source X is deterministic, it can be seen with an entropy argument that the output size of E_X must be at most $H_\infty(X)$, the min-entropy of X ; otherwise, the distribution $\{E_X(X)\}$ can necessarily be distinguished from random with some statistically non-negligible advantage. Therefore, E_X is an optimal and efficient compression algorithm for X , with decompression algorithm D_X ; this is true even for the relaxation in the CRS model. The existence of efficient compression algorithms

for various categories of samplers was thoroughly studied [TVZ05]. In particular, the existence of compression algorithms for all efficiently samplable sources implies the inexistence of one-way functions (this is an observation attributed to Levin in [GS85]) since compressing the output of a pseudorandom generator to its entropy would distinguish it from a random string, and the existence of one-way functions implies the existence of pseudorandom generators [HILL99]).

Theorem (informal). *Assuming the existence of one-way functions, neither $\text{PREH}_{\equiv_s}^{\text{det}}$ nor $\text{cPREH}_{\equiv_s}^{\text{det}}$ hold.*

This is a strong impossibility result, as one-way functions dwell among the weakest assumptions in cryptography, [Imp95]. One can circumvent this impossibility by studying whether compression can be achieved for more restricted classes of distributions, as was done e.g. in [TVZ05]. Our work can be seen as pursuing an orthogonal direction. We seek to determine whether a relaxed notion of compression can be achieved for *all* efficiently samplable distributions. The relaxations we consider comprise the possibility to use randomness in the encoding algorithm, and weakening the requirement on the encoded distribution to being only computationally indistinguishable from random. Clearly, these relaxations remove one of the most important features of compression algorithms, which is that their outputs are smaller than their inputs (i.e., they compress). Nevertheless, the indistinguishability of the encoded distribution from the uniform distribution is another crucial feature of optimal compression algorithms, which has independent applications.

1.3.2 Deterministic, computational pseudorandom encodings. We now turn towards a relaxation where the encoded distribution is only required to be computationally indistinguishable from random, but the encoding algorithm is still required to be deterministic. This flavor is strongly connected to an important problem in cryptography: the problem of separating HILL entropy [HILL99] from Yao entropy [Yao82]. HILL and Yao entropy are different approaches of formalizing computational entropy, i.e., the amount of entropy a distribution appears to have from the viewpoint of a computationally bounded entity. Informally, a distribution has high HILL entropy if it is computationally close to a distribution with high min-entropy; a distribution has high Yao entropy if it cannot be compressed efficiently. Finding a distribution which, under standard cryptographic assumptions, has high Yao entropy, but low HILL entropy constitutes a long standing open problem in cryptography. Currently, only an oracle separation [Wee04] and a separation for conditional distributions [HLR07] are known. To establish the connection between $\text{PREH}_{\approx_c}^{\text{det}}$ and this problem, we proceed as follows: informally, a deterministic pseudorandom encoding must necessarily *compress its input to the HILL entropy of the distribution*. That is, the output size of the encoding cannot be much larger than the HILL entropy of the distribution. This, in turn, implies that a distribution which admits such a pseudorandom encoding cannot have high Yao entropy.

In this work, we formalize the above argument, and show that the *conditional* separation of HILL and Yao entropy from [HLR07] suffices to refute $\text{PREH}_{\approx_c}^{\text{det}}$. This separation holds under the assumption that non-interactive zero-knowledge proofs with some appropriate structural properties exist (which in turn can be based on standard assumptions such as the quadratic residuosity assumption). Thus, we obtain the following infeasibility result:

Theorem (informal). *If the quadratic residuosity assumption holds, then $\text{PREH}_{\approx_c}^{\text{det}}$ does not hold.*

Hence, we may conclude that towards a feasible variant of pseudorandom encodings for all efficiently samplable distributions, requiring the encoding algorithm to be deterministic poses a strong restriction.

1.3.3 Randomized, statistical pseudorandom encodings. We now consider the relaxation of perfect compression by allowing the encoding algorithm to be randomized while still requiring information-theoretic indistinguishability from randomness. This flavor of pseudorandom encoding was used in the context of *honey encryption* [JR14]. Honey encryption is a cryptographic primitive which has been introduced to mitigate attacks on encryption schemes resulting from the use of low-entropy passwords. Honey encryption has the property that decrypting a ciphertext with an incorrect key always yields a valid-looking plaintext which seems to come from the expected distribution, thereby mitigating brute-force attacks. This is the same property that was useful in the previous PAKE example.

The study of honey encryption was initiated in [JR14], where it was shown that honey encryption can naturally be constructed by composing a block cipher (modeled as a random cipher) with a *distribution transforming encoder* (DTE), a notion which is equivalent to our notion of pseudorandom encoding with randomized encoding and statistical pseudorandomness. The focus of [JR14] was on obtaining such DTEs for simple and useful distributions. In contrast, we seek to understand the feasibility of this notion for *arbitrary* distributions. Intuitively, it is not straightforward to encode any efficient distribution into the uniform distribution; consider for example the distribution over RSA moduli, i.e., products of two random n -bit primes. Since no efficient algorithm is known to test membership in the support of this distribution, natural approaches seem to break down. In fact, we show in this work that this difficulty is inherent: building on techniques from [BCPR14; IKOS10], we demonstrate the impossibility of (randomized, statistical) pseudorandom encodings for all efficiently samplable distributions, under a relatively standard cryptographic assumption.

Theorem (informal). *Assuming the sub-exponential hardness of the learning with errors (LWE) problem, $\text{PREH}_{\equiv_s}^{\text{rand}}$ does not hold.*

This result directly implies that under the same assumption, there exist efficiently samplable distributions (with input) for which no distribution transforming encoder exists. We view it as an interesting open problem whether this result can be extended to rule out the existence of honey encryption for arbitrary distributions under the same assumption.

1.4 Open questions and subsequent work

The most intriguing question left open by our work is whether the weakest variant of the pseudorandom encoding hypothesis $\text{cPREH}_{\approx_c}^{\text{rand}}$, which is implied by iO, also *implies* iO. Very recently, this question was settled in the affirmative by Wee and Wichs [WW20] under the LWE assumption. More concretely, by modifying a heuristic iO construction of Brakerski et al. [BDGM20], they show that iO is implied by LWE if one is additionally given an *oblivious LWE-sampler* in the CRS model. Such a sampler, given a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, generates outputs that are indistinguishable from LWE samples $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ without knowing the secrets $\mathbf{s}$ or the noise $\mathbf{e}$. The existence of an oblivious LWE sampler is nontrivial even under the LWE assumption, because $\mathbf{A}$ can be such that $\mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ is not pseudorandom; however, such a sampler still follows from the invertible sampling hypothesis [IKOS10], which we show to be equivalent to the pseudorandom encoding hypothesis. By proposing an explicit heuristic construction of (a relaxed flavor of) an oblivious LWE sampler, the end result of [WW20] is a construction of iO from a new “falsifiable” assumption.

Whether $\text{cPREH}_{\approx_c}^{\text{rand}}$ implies iO under weaker or different assumptions than LWE remains open. A potentially easier goal is using $\text{cPREH}_{\approx_c}^{\text{rand}}$ to construct public-key encryption from one-way functions. This is related to the possibility of constructing oblivious transfer from any public-key encryption in which public keys and ciphertexts are obliviously samplable [EGL85; GKMRV00], which is implied by public-key encryption and $\text{cPREH}_{\approx_c}^{\text{rand}}$. Here $\text{cPREH}_{\approx_c}^{\text{rand}}$ is used to bypass the black-box separation between public-key encryption and oblivious transfer [GKMRV00].

Finally, there is a lot of room for relaxing the intractability assumptions we use to rule out the statistical ($\text{cPREH}_{\equiv_s}^{\text{rand}}$) and deterministic ($\text{cPREH}_{\approx_c}^{\text{det}}$) flavors of pseudorandom encodings.

2 Overview of techniques

In this section, we elaborate on some of our technical results in more detail. In the following, we identify a PPT sampler S with the distribution (family) ensemble it induces.

The relation to invertible sampling. A PPT sampler S is *inverse samplable* [DN00; IKOS10], if there exists an alternative sampler $\bar{S}$ inducing a distribution which is computationally indistinguishable to the distribution induced by S such that the computations of $\bar{S}$ can be efficiently inverted. Efficiently inverting the computation of $\bar{S}$ means that there exists an efficient inverse sampler $\bar{S}^{-1}$ which, given an output of $\bar{S}$, recovers a well-distributed random tape for $\bar{S}$ to compute the given output in the following sense. The inverse sampled random tape is required to be computationally indistinguishable from the actually used random tape. More formally, a PPT sampler S is inverse samplable if there exists an efficient alternative sampler $\bar{S}$ and an efficient inverse sampler $\bar{S}^{-1}$ such that

$$\{y \leftarrow S(1^\lambda) : y\} \approx_c \{y \leftarrow \bar{S}(1^\lambda) : y\}, \quad (3)$$

$$\{y \leftarrow \bar{S}(1^\lambda; r) : (r, y)\} \approx_c \{y \leftarrow \bar{S}(1^\lambda) : (\bar{S}^{-1}(1^\lambda, y), y)\}. \quad (4)$$

We refer to Equation (3) as *closeness* and to Equation (4) as *invertibility*. If the sampler S admits an input m , the above is required to hold for all inputs m in the input space L , where $\bar{S}$ and $\bar{S}^{-1}$ both additionally receive m as input. In accordance with [IKOS10], we refer to the hypothesis that all PPT algorithms with input are inverse samplable as the *invertible sampling hypothesis*. Restricting the invertible sampling hypothesis to algorithms which do not admit inputs is denoted the *weak* invertible sampling hypothesis.

The concepts of inverse samplability and pseudorandom encodings are tightly connected. Suppose a PPT algorithm S is inverse samplable. Then, there exists an alternative and an inverse sampler $(\bar{S}, \bar{S}^{-1})$ satisfying closeness and invertibility. Invertibility guarantees that the inverse sampler $\bar{S}^{-1}$ on input of a sample y from $\bar{S}(1^\lambda)$, outputs a computationally well-distributed random tape r . Hence, with overwhelming probability over the choice of $y \leftarrow \bar{S}(1^\lambda)$ and $r \leftarrow \bar{S}^{-1}(y)$, the alternative sampler on input of r , recovers y . In other words, the inverse sampler $\bar{S}^{-1}$ can be seen as encoding a given sample y , whereas the *de-randomized* alternative sampler $\bar{S}$ given this encoding *as random tape*, is able to recover y . Looking through the lens of pseudorandom encoding, this almost proves correctness except that y is sampled according to $\bar{S}(1^\lambda)$ instead of $S(1^\lambda)$. This difference can be bridged due to closeness. We now turn towards showing pseudorandomness of the encoded distribution. Due to closeness, the distributions $\{y \leftarrow \bar{S}(1^\lambda) : (\bar{S}^{-1}(1^\lambda, y), y)\}$ and $\{y \leftarrow S(1^\lambda) : (\bar{S}^{-1}(1^\lambda, y), y)\}$ are computationally indistinguishable. Invertibility guarantees that, given a sample y from $\bar{S}(1^\lambda)$, an encoding of y is indistinguishable to uniformly chosen randomness conditioned on the fact that decoding yields y . Removing y from this distribution, almost corresponds to pseudorandomness, except that y is sampled according to $\bar{S}(1^\lambda)$ instead of $S(1^\lambda)$. Again, we are able to bridge this gap due to closeness. Note that we crucially use the fact that the initial randomness used by $\bar{S}$ resides outside of the view of an adversary. Summing up, if a PPT sampler S is inverse samplable, then it can be pseudorandomly encoded.

Interestingly, this connection turns out to be bidirectional. Suppose a PPT algorithm S can be pseudorandomly encoded. Then, there exists an efficient encoding algorithm E_S and an efficient deterministic decoding algorithm D_S satisfying correctness and pseudorandomness. Looking through the lens of invertible sampling, we identify the decoding algorithm to correspond to the alternative sampler (viewing the random tape of the alternative sampler as explicit input

to D_S) and the encoding algorithm to correspond to the inverse sampler. Pseudorandomness guarantees that $E_S(S(1^\lambda))$ is indistinguishable from uniform randomness. Hence, applying the decode algorithm D_S on uniform randomness is indistinguishable from applying D_S to outputs of $E_S(S(1^\lambda))$. Correctness guarantees that $D_S(E_S(y))$ for y sampled according to $S(1^\lambda)$ recovers y with overwhelming probability. Thus, the distribution induced by applying D_S on uniform randomness is computationally close to the distribution induced by $S(1^\lambda)$. This shows closeness. For the purpose of arguing about invertibility, consider the distribution $A := \{y \leftarrow D_S(r) : (r, y)\}$. Due to pseudorandomness r can be considered an encoded sample from $S(1^\lambda)$. Hence, A is indistinguishable to the distribution, where r is produced by $E_S(y')$ for some *independent* $y' \leftarrow S(1^\lambda)$, i.e.

$$\{y \leftarrow D_S(r) : (r, y)\} \approx_c \{y' \leftarrow S(1^\lambda), r \leftarrow E_S(y'), y \leftarrow D_S(r) : (r, y)\}.$$

Note that by correctness, y and y' are identical with overwhelming probability. Therefore, A is indistinguishable to $\{y' \leftarrow S(1^\lambda), r \leftarrow E_S(y') : (r, y')\}$. Since sampling y' via D_S applied on uniform randomness is computationally close to the above distribution due to closeness, invertibility follows. Summing up, a sampler S can be pseudorandomly encoded *if and only if* it is inverse samplable.

Likewise to the variations and relaxations described for pseudorandom encodings, we vary and relax the notion of invertible sampling. The inverse sampler can be required to be deterministic or allowed to be randomized. Further, closeness and invertibility can be required to hold information theoretically or computationally. We denote these variants as $\text{ISH}_{\approx_c}^{\text{rand}}$, $\text{ISH}_{\equiv_s}^{\text{rand}}$, $\text{ISH}_{\approx_c}^{\text{det}}$ and $\text{ISH}_{\equiv_s}^{\text{det}}$. To circumvent impossibilities in the plain model, we also define the relaxations in the common reference string model in static and adaptive flavors, denoted the prefix “c” and “ac”, respectively. The above equivalence extends to all introduced variations of the pseudorandom encoding and invertible sampling hypotheses.

The static-to-adaptive transformation. The static variant of pseudorandom encodings in the CRS model only guarantees correctness and pseudorandomness as long as the input m for the sampler S is chosen independently of the CRS. The adaptive variant, on the other hand, provides correctness and pseudorandomness even for adaptive choices of inputs. Adaptive notions always imply their static analogues. Interestingly, for pseudorandom encodings, the opposite direction is true as well. The core idea is to use an *indirection* argument (similar to [HJKSWZ16; CPR17; CPV17]) to delay CRS generation until during the actual encoding process. Thus, the advantage stemming from adaptively choosing the input is eliminated.

Suppose that the static variant of the pseudorandom encoding hypothesis in the CRS model is true and let S be some PPT sampler. Since S can be pseudorandomly encoded in the CRS model with static choice of inputs, there exist algorithms (Setup', E', D') such that static correctness and pseudorandomness hold. Further, the algorithm Setup' can also be pseudorandomly encoded as above. Let $(\text{Setup}'', E'', D'')$ be the corresponding algorithms such that static correctness and pseudorandomness hold. Note that since the sampler Setup' does not expect an input, static and adaptive guarantees are equivalent.

Then, the sampler S can be pseudorandomly encoded in the CRS model with *adaptive* choice of inputs as follows. Initially, we sample a common reference string crs'' via $\text{Setup}''(1^\lambda)$ and make it available to the parties. Given crs'' and a sample y from $S(m)$, adaptive encoding works in two phases. First, a fresh CRS crs' is sampled via $\text{Setup}'(1^\lambda)$ and pseudorandomly encoded via $r_1 \leftarrow E''(\text{crs}'', \text{crs}')$. Second, the given sample y is pseudorandomly encoded via $r_2 \leftarrow E'(\text{crs}', m, y)$. The encoding of y then consists of (r_1, r_2) . To decode, the CRS crs' is restored via $D''(\text{crs}'', r_1)$. Then, using crs' , the original sample y is recovered via $D'(\text{crs}', m, r_2)$.

Since crs' is chosen freshly during the encoding process, the input m which may depend on crs'' , cannot depend on crs' . Further, the distribution Setup'' does not expect an input. Hence, static guarantees suffice.

To realize that adaptive pseudorandomness holds, consider the encoding of $S(m)$ for some adaptively chosen message m . Since the view of $\mathcal{A}$ when choosing the message m is independent of crs' , static pseudorandomness can be applied to replace the distribution $\mathbf{E}'(\text{crs}', m, S(m))$ with uniform randomness. Further, since the sampler Setup' does not expect any input, static pseudorandomness suffices to replace the distribution $\mathbf{E}''(\text{crs}'', \text{Setup}'(1^\lambda))$ with uniform randomness. This proves adaptive pseudorandomness.

The adaptive variant of correctness follows similarly from the static variant of correctness. Consider the distribution of decoding an encoded sample of $S(m)$, where m is adaptively chosen. Since the sampler Setup' does not expect an input, static correctness can be applied to replace decoding $\mathbf{D}''(\text{crs}'', r_1)$ with the crs' sampled during encoding. Again, since crs' does not lie in the view of the adversary when choosing the message m , static correctness guarantees that decoding succeeds with overwhelming probability. This proves adaptive correctness.

On deterministic pseudorandom encoding and compression. The notion of pseudorandom encoding is inspired by the notion of compression. A tuple of deterministic functions $(\mathbf{E}_X, \mathbf{D}_X)$ is said to compress a source X_λ to length $m(\lambda)$ with decoding error $\epsilon(\lambda)$, if (i) $\Pr[\mathbf{D}_X(\mathbf{E}_X(X_\lambda)) \neq X_\lambda] \leq \epsilon(\lambda)$ and (ii) $\mathbb{E}[|\mathbf{E}_X(X_\lambda)|] \leq m(\lambda)$, see [Wee04; TVZ05]. Pseudorandom encoding partially recovers the notion of compression if we require the encoding algorithm to be deterministic. If a source X_λ can be pseudorandomly encoded with a deterministic encoding algorithm having output length $n(\lambda)$, then X_λ is compressible to length $n(\lambda)$. Note, however, that the converse direction is not true. Compression and decompression algorithms for a compressible source do not necessarily encode that source pseudorandomly. The output of a compression algorithm is not required to look pseudorandom and, in some cases, admits a specific structure which makes it easily distinguishable from uniform randomness, e.g. instances using Levin search, [TVZ05].

Clearly, the requirement for correctness, poses a lower bound on the encoding length $n(\lambda)$, [Sha48]. Conversely, requiring the encoding algorithm $\mathbf{E}_X$ to be deterministic means that the only source of entropy in the distribution $\mathbf{E}_X(X_\lambda)$ originates from the source X_λ itself. Hence, for the distributions $\mathbf{E}_X(X_\lambda)$ and the uniform distribution over $\{0, 1\}^{n(\lambda)}$ to be indistinguishable, the encoding length $n(\lambda)$ must be “sufficiently small”. We observe that correctness together with the fact that $\mathbf{E}_X$ is deterministic implies that the event $\mathbf{E}_X(\mathbf{D}_X(\mathbf{E}_X(X_\lambda))) = \mathbf{E}_X(X_\lambda)$ occurs with overwhelming probability. Applying pseudorandomness yields that $\mathbf{E}_X(\mathbf{D}_X(U_{n(\lambda)})) = U_{n(\lambda)}$ holds with overwhelming probability, wherefore we can conclude that $\mathbf{D}_X$ operates almost injectively on the set $\{0, 1\}^{n(\lambda)}$. Hence, the (smooth) min-entropy of $\mathbf{D}_X(U_{n(\lambda)})$ is at least $n(\lambda)$.

Considering information theoretical pseudorandomness, the distributions $\mathbf{D}_X(U_{n(\lambda)})$ and X_λ are statistically close. Hence, by the reasoning above, the encoding length $n(\lambda)$ is upper bounded by the (smooth) min-entropy of the source X_λ . In conclusion, if a distribution can be pseudorandomly encoded such that the encoding algorithm is deterministic satisfying statistical pseudorandomness, then this distribution is compressible to its (smooth) min-entropy. Using a technical “Splitting Lemma”, this extends to the relaxed variant of the pseudorandom encoding hypothesis in the CRS model.

Considering computational pseudorandomness, by a similar argument as above, we obtain that X_λ is computationally close to a distribution with min-entropy $n(\lambda)$. This does not yield a relation between the encoding length and the min-entropy of the source. However, we do obtain relations to computational analogues of entropy. Computational entropy is the amount of entropy a distribution appears to have from the perspective of a computationally bounded entity. The notion of HILL entropy [HILL99] is defined via the computational indistinguishability from a truly random distribution. More formally, a distribution X_λ has HILL entropy at least k , if there exists a distribution with min-entropy k which is computationally indistinguishable from X_λ . Hence, the encoding length $n(\lambda)$ is upper bounded by the HILL entropy of the source X_λ . Another important notion of computational entropy is the notion of Yao entropy [Yao82]. Yao entropy is defined via the incompressibility of a distribution. More precisely, a distribution X_λ has Yao entropy at least k if X_λ cannot be efficiently compressed to length less than k (and

successfully decompressed). If a distribution can be pseudorandomly encoded with deterministic encoding, then it can be compressed to the encoding length $n(\lambda)$. This poses an upper bound on the Yao entropy of the source. In summary, this yields

$$n(\lambda) \leq H^{\text{HILL}}(X_\lambda) \quad \text{and} \quad H^{\text{Yao}}(X_\lambda) \leq n(\lambda). \quad (5)$$

However, due to [HLR07; LMs05], if the Quadratic Residuosity Assumption (QRA) is true, then there exist distributions which have low *conditional* HILL entropy while being *conditionally* incompressible, i.e. have high conditional Yao entropy.⁷ The above observations, particularly Equation (5), can be extended to conditional HILL and conditional Yao entropy, by considering $\text{PREH}_{\approx_c}^{\text{det}}$ for PPT algorithms with input. Therefore, if the Quadratic Residuosity Assumption is true, $\text{PREH}_{\approx_c}^{\text{det}}$ cannot be true for those distributions.

Unfortunately, we do not know whether this extends to the relaxed variants of the pseudo-random encoding hypothesis admitting access to a CRS. On a high level, the problem is that the HILL entropy, in contrast to the min-entropy, does not remain untouched when additionally conditioning on some common reference string distribution, even though the initial distribution is independent of the CRS. Hence, the splitting technique can not be applied here.

⁷ Let (X, Z) be a joint distribution. The conditional computational entropy is the entropy X appears to have to a bounded adversary when additionally given Z .

3 Preliminaries

We denote by $[n]$ the set $\{1, \dots, n\}$. Throughout this paper, λ denotes a security parameter which is given as input to all algorithms. A probabilistic polynomial time (PPT) algorithm (also referred to as an *efficient* algorithm) runs in time polynomial in the (implicit) security parameter λ . In this paper, we consider non-uniform polynomial time adversaries, i.e. polynomial time adversaries receiving a polynomially bounded auxiliary input (or advice) depending only on the security parameter.⁸ A function $f(\lambda)$ is *negligible* if for any polynomial p there exists a bound $B \in \mathbb{N}$ such that, for any integer $k \geq B$, $|f(k)| \leq \frac{1}{p(k)}$. A function g is *overwhelming* if $1 - g(\lambda)$ is a negligible function.

Given a finite set A , the notation $x \leftarrow A$ means a uniformly random assignment of an element of A to the variable x . Given a probability distribution D , the notation $x \leftarrow D$ means sampling an element according to the distribution D and assigning that element to x . We denote the uniform distribution over bitstrings of length n by U_n . Let X, Y be two distributions over a set A . Then, the statistical distance between X and Y is defined as $\Delta(X, Y) := \sum_{a \in A} |\Pr[X = a] - \Pr[Y = a]|$. We say that two distributions are statistically close if their statistical distance is negligible.

A source X is a probability distribution on strings. A family of sources is a probability ensemble $(X_\lambda)_{\lambda \in \mathbb{N}}$, where X_λ is distributed on $\{0, 1\}^{p(\lambda)}$ for some polynomial p . A family of sources $(X_m)_{m \in L}$ can also be indexed by strings from some language $L \subseteq \{0, 1\}^+$, where X_m is distributed on $\{0, 1\}^{p(|m|)}$ for some polynomial p . We say that a source X_λ is efficiently samplable if there is a PPT algorithm S such that $S(1^\lambda)$ is distributed according to X_λ for all $\lambda \in \mathbb{N}$. We say that a source X_m indexed by strings is efficiently samplable if there exists a PPT algorithm S such that $S(m)$ is distributed according to X_m for all $m \in L$.

In game based proofs, it will be useful to let out_i denote the output of game $\mathbf{G}_i$. If we want to make the adversary $\mathcal{A}$ playing $\mathbf{G}_i$ explicit, we write $out_{i, \mathcal{A}}$. Unless stated otherwise, we consider stateful adversaries.

The hypotheses stated in the following are formulated for the class of all PPT algorithms S , possibly excluding pathological cases. In some cases it is sufficient to consider a weaker variant of these hypotheses which is only required to be true for a specific class of PPT algorithms S . In this case we say that the respective hypothesis holds for the class $\mathcal{S}$.

4 The pseudorandom encoding hypothesis

We study the ability to encode efficiently samplable distributions into the uniform distribution. In the following, an efficiently samplable distribution will be defined by the corresponding sampler S with input space L . A distribution defined via S can be pseudorandomly encoded if there exists an efficient potentially randomized encoding algorithm E_S and an efficient deterministic decoding algorithm D_S such that for all $m \in L$, the probability for the event $D_S(E_S(S(m))) = S(m)$ is overwhelming and the distribution $E_S(S(m))$ is indistinguishable from the uniform distribution $U_{n(\lambda)}$. We work with the hypothesis that every efficiently samplable distribution can be pseudorandomly encoded. In this section, we formally define the pseudorandom encoding hypothesis and its variations.

Definition 1 (Pseudorandom encoding hypothesis, $\text{PREH}_{\approx_c}^{\text{rand}}$). *For every PPT algorithm S , there exist efficient algorithms E_S (the encoding algorithm) with output length $n(\lambda)$ and D_S (the decoding algorithm), where D_S is deterministic and E_S is randomized satisfying the following two properties.*

Correctness. *For all inputs $m \in L$, $\epsilon_{\text{dec-error}}(\lambda) := \Pr[y \leftarrow S(m) : D_S(m, E_S(m, y)) \neq y]$ is negligible.*

⁸ Note that by a coin-fixing argument, it is sufficient to consider non-uniform deterministic adversaries. Most of our results apply for uniform PPT adversaries as well. In case we make explicit use of the non-uniformity of the adversary, we remark this explicitly.

Pseudorandomness. For all PPT adversaries $\mathcal{A}$ and all inputs $m \in L$,

$$\text{Adv}_{\mathcal{A},m}^{\text{pre}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A},m,0}^{\text{pre}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},m,1}^{\text{pre}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $\text{Exp}_{\mathcal{A},m,0}^{\text{pre}}$ and $\text{Exp}_{\mathcal{A},m,1}^{\text{pre}}$ are defined in Figure 2.

$\text{Exp}_{\mathcal{A},m,0}^{\text{pre}}(\lambda)$	$\text{Exp}_{\mathcal{A},m,1}^{\text{pre}}(\lambda)$
$r \leftarrow \{0,1\}^{p(\lambda)}$	$u \leftarrow \{0,1\}^{n(\lambda)}$
$y := S(m; r)$	return $\mathcal{A}(m, u)$
return $\mathcal{A}(m, E_S(m, y))$	

Fig. 2. The pseudorandomness experiments.

Remark 1. Definition 1 formulated for PPT algorithms S which do not admit an input m is called the *weak* $\text{PREH}_{\approx_c}^{\text{rand}}$.

Remark 2 ($\text{PREH}_{\approx_c}^{\text{det}}, \text{PREH}_{\equiv_s}^{\text{rand}}, \text{PREH}_{\equiv_s}^{\text{det}}$). Definition 1 can be tuned in two dimensions: the encode algorithm can be required to be deterministic or allowed to be randomized, and the pseudorandomness property can be required to hold statistically or computationally. We denote these variants as $\text{PREH}_{\alpha}^{\beta}$, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

Remark 3. Definition 1 demands indistinguishability between encoded samples and the uniform distribution over all bitstrings of some length $n(\lambda)$. This requirement can be relaxed to indistinguishability from the uniform distribution over some efficiently samplable and efficiently recognizable set $\mathcal{R}$ of size N , where elements in $\mathcal{R}$ can be represented with $O(\log N)$ bits.

Deterministic encoding and compatible samplers

Requiring the encoding algorithm E_S to be deterministic entails the existence of what we call “incompatible samplers” for which $\text{PREH}_{\equiv_s}^{\text{det}}$ and even $\text{PREH}_{\approx_c}^{\text{det}}$ are unconditionally false. For instance, consider the sampler S^* which on input of 1^λ uniformly chooses an element from the set $\{00, 01, 10\} \subset \{0,1\}^2$. Assume $\text{PREH}_{\approx_c}^{\text{det}}$ is true for S^* . Then, correctness requires that with overwhelming probability over the sampling process $y \leftarrow S^*(1^\lambda)$, $D_{S^*}(E_{S^*}(y)) = y$. Hence, E_{S^*} must map into the set $\{0,1\}^k$ for $k \geq 2$ (otherwise there is a correctness error of at least $\frac{1}{3}$). Pseudorandomness, on the other hand, requires that $E_{S^*}(S^*(1^\lambda))$ is computationally indistinguishable from uniform distribution over $\{0,1\}^k$. However, since E_{S^*} is a deterministic algorithm, $|\text{supp}(E_{S^*}(S^*(1^\lambda)))| = 3$. Therefore, there exists at least an element in $\{0,1\}^k \setminus \text{supp}(E_{S^*}(S^*(1^\lambda)))$. An adversary can easily determine this element by evaluating E_{S^*} on each element in the support of S^* (if the support of the sampler was super-polynomial, this would not be possible for a PPT adversary, but very well possible for an unbounded one).

Another example of such an incompatible sampler is a sampler with large support but low min-entropy (i.e. a sampler that has (at least one) very likely output and many much less likely outputs). For instance, consider the sampler S' with probability distribution

$$\begin{aligned} \Pr[S'(1^\lambda) = 0^\lambda] &= \frac{1}{2} \\ \Pr[S'(1^\lambda) = 1 \parallel x] &= \frac{1}{2^\lambda} \text{ for each } x \in \{0,1\}^{\lambda-1} \end{aligned}$$

Assume $\text{PREH}_{\approx_c}^{\text{det}}$ is true for S' . Then, correctness requires that an overwhelming fraction of the elements of the form $1 \parallel x$ for $x \in \{0,1\}^{\lambda-1}$ are correctly decodable. Furthermore, since the element 0^λ has a non-negligible probability to occur, it needs to be correctly decodable. Let the support of $E_{S'}(S'(1^\lambda))$ be (a subset of) $\{0,1\}^{\lambda-1}$. Pseudorandomness requires that $E_{S'}(S'(1^\lambda))$ is indistinguishable from uniform distribution over $\{0,1\}^{\lambda-1}$. However, since $E_{S'}$ is deterministic,

the value $E_{S'}(0^\lambda)$ occurs with probability at least $\frac{1}{2}$ and, due to correctness, all other values $E_{S'}(1 \parallel x)$ occur with far lower probability.

In order to obtain a meaningful definition of $\text{PREH}_{\approx_c}^{\text{det}}$ and $\text{PREH}_{\equiv_s}^{\text{det}}$, we restrict these hypotheses to only hold for specific classes of what we refer to as “compatible samplers” $\mathcal{S}^{\text{comp}}$ and $\mathcal{S}^{\text{stat}}$, respectively.

Definition 2 (Compatibility with deterministic encodings). *A sampler S is statistically compatible with deterministic encodings if there exists a set A whose cardinality is a power of 2, such that $\Pr[S(1^\lambda) \in A]$ is overwhelming, and S is ϵ -flat on A , i.e. for all $a \in A$ we have $|\Pr[S(1^\lambda) = a] - \frac{1}{|A|}| \leq \epsilon(\lambda)$, for some negligible function ϵ . The class $\mathcal{S}^{\text{stat}}$ contains all samplers which are statistically compatible with deterministic encodings.*

A sampler S is computationally compatible with deterministic encodings if $S \in \mathcal{S}^{\text{stat}}$ or if $|\text{supp}(S(1^\lambda))|$ is super-polynomial and the min-entropy $H_\infty(S(1^\lambda)) \in \omega(\log(\lambda))$ (i.e. the most likely event occurs with negligible probability). The class $\mathcal{S}^{\text{comp}}$ contains all samplers which are computationally compatible with deterministic encodings.

If S admits an input $z \in L$, S is statistically or computationally compatible with deterministic encodings if the corresponding criterion is met for all $z \in L$.

The above criterion for $\mathcal{S}^{\text{stat}}$ may seem unnatural. We note that by relaxing Definition 1 as noted in Remark 3, requiring high min-entropy suffices for a sampler to be statistically compatible with deterministic encodings.

We note that ϵ -flatness is a weaker criterion than statistical closeness to the uniform distribution over A . ϵ -flatness on A corresponds to closeness to the uniform distribution over A with respect to the infinity norm. Statistical closeness, however, is formalized with respect to the Manhattan norm. The deterministic pseudorandom encodings restricted to compatible samplers still has interesting connections.

Let iPRG be an injective PRG with stretch ℓ . Let S be the sampler which on input of 1^λ draws a uniform seed $s \in \{0, 1\}^\lambda$ and outputs $\text{iPRG}(s) \in \{0, 1\}^{\ell(\lambda)}$. Clearly, S is statistically compatible with deterministic encodings.

Lemma 1. *Let PRG be a PRG with stretch ℓ and let S be the sampler which on input of 1^λ produces the distribution $\text{PRG}(U_\lambda)$. Then, $S \in \mathcal{S}^{\text{stat}}$.*

Proof. Let $A' := \text{PRG}(\{0, 1\}^\lambda)$ and let $A'' \subseteq \{0, 1\}^{\ell(\lambda)} \setminus A'$ such that for $A := A' \cup A''$ we have $|A| = 2^\lambda$. We have $\Pr[S(1^\lambda) \in A] = 1$.

For flatness on A , we upper bound the probability of the most likely event in A . Due to pseudorandomness of PRG , all (non-uniform) PPT adversary distinguishing the distributions $S(1^\lambda)$ and $U_{\ell(\lambda)}$ have a negligible advantage. Assume there exists an $a \in A'$, such that $\Pr[S(1^\lambda) = a] \geq \delta(\lambda)$, for a non-negligible function δ . Then, the most likely value in A' could be used as polynomial advice string allowing a non-uniform adversary to recognize the distribution $S(1^\lambda)$ with non-negligible probability $\delta(\lambda)$. (A uniform adversary can sample a random seed and evaluate the PRG. The thereby obtained output equals the most likely event a with probability $\delta(\lambda)$. Hence, a uniform adversary has advantage at least $\delta(\lambda)^2$ in distinguishing.) Therefore, for all $a \in A'$, $\Pr[S(1^\lambda) = a] \leq \epsilon(\lambda)$ for some negligible function ϵ . Hence, for all $a \in A$, we have $|\Pr[S(1^\lambda) = a] - \frac{1}{|A|}| \leq \epsilon(\lambda) + 2^{-\lambda}$ which is negligible. $\square$

Lemma 2. *Let PRG be a PRG with stretch ℓ . Let S be the sampler which on input of $x \in L$ produces the distribution $\{y_1 \leftarrow \text{PRG}(U_{|x|}), y_2 \leftarrow D_{x, y_1} : (y_1, y_2)\}$ for any distribution D . Then, $S \in \mathcal{S}^{\text{comp}}$.*

Proof. Let $\lambda := |x|$. Using the argument of Lemma 1, all images in $\text{PRG}(\{0, 1\}^\lambda)$ have only a negligible probability to occur. Therefore, for all $x \in L$ and all $(y_1, y_2) \in \text{supp}(S(x))$, we have $\Pr[S(1^\lambda, x) = (y_1, y_2)]$ is negligible. Therefore, $S \in \mathcal{S}^{\text{comp}}$. $\square$

Impossibility of universal encoding

It is essential that the decoding algorithm D_S depends on the sampler S , since due to pseudorandomness, D_S on input of a random string needs to produce a sample that is in some sense close to the distribution produced by S . This argument does not hold necessarily for the encode algorithm. In [TVZ05], universal compression was studied. This translates to the following definition of PREH with universal encoding.

Definition 3 (PREH $^\beta_\alpha$ with universal encoding). *Let $\mathcal{S}$ be a class of sampling algorithms. We say PREH $^\beta_\alpha$ with universal sampling is true for the class $\mathcal{S}$, if there exists a universal encoding algorithm E , such that for every PPT algorithm $S \in \mathcal{S}$, there exists an efficient deterministic decoding algorithm D_S , such that correctness and pseudorandomness are satisfied.*

In contrast to universal compression, pseudorandom encoding with universal encoding is impossible.

Lemma 3. *PREH with universal encoding is not possible for arbitrary classes of samplers.*

Proof. Consider the class of sampling algorithms $\mathcal{S} = \{S_1, S_2, S_3\}$ over $\{0, 1\}^{p(\lambda)}$ with support $Y_{i,\lambda} := \text{supp}(S_i(1^\lambda))$. We require that for $i \in \{1, 2\}$, $|Y_{i,\lambda}| = 2^{k+1}$, $|Y_{1,\lambda} \cap Y_{2,\lambda}| = 2^k$, $Y_{1,\lambda} \cap Y_{2,\lambda} = Y_{3,\lambda}$ for $k \in O(\log \lambda)$, and that S_1, S_2 and S_3 produce uniform samples over their support (i.e. correspond to flat distributions). We note that $S_1, S_2, S_3 \in \mathcal{S}^{\text{stat}}$. For notational convenience we omit the dependency on λ in the following.

Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. Assume toward a contradiction, that PREH $^\beta_\alpha$ with universal encoding is true for the class $\mathcal{S}$ of sampling algorithms above. Let $\{0, 1\}^{n(\lambda)}$ be the range of E . Due to correctness, for $i \in \{1, 2\}$, there is a negligible function ϵ , such that

$$\begin{aligned} \Pr_{y \leftarrow Y_i}[D_{S_i}(E(y)) \neq y] &= \sum_{y' \in Y_i} \Pr_{y \leftarrow Y_i}[D_{S_i}(E(y)) \neq y \wedge y = y'] \\ &= \frac{1}{|Y_i|} \sum_{y' \in Y_i} \Pr_{y \leftarrow Y_i}[D_{S_i}(E(y)) \neq y \mid y = y'] \leq \epsilon(\lambda) \end{aligned}$$

Since $|Y_i|$ is polynomial, for each $y \in Y_i$, $\Pr[D_{S_i}(E(y)) = y] \geq 1 - \epsilon'(\lambda)$ for some negligible function ϵ' , where the probability is solely over the randomness of E . Hence, for all $y \in Y_3 = Y_1 \cap Y_2$,

$$\Pr[u \leftarrow E(y) : y = D_{S_1}(u) = D_{S_2}(u) = D_{S_3}(u)] \geq 1 - \epsilon''(\lambda)$$

for some negligible function ϵ'' .

Due to Theorem 2 (see Section 5.3), for $i \in \{1, 2, 3\}$, the distribution $\{u \leftarrow U_{n(\lambda)} : D_{S_i}(u)\}$ is computationally (statistically) close to the distribution produced by $S_i(1^\lambda)$. Hence, (up to some negligible fraction ϵ''') the algorithms D_{S_1} and D_{S_2} map at most half of the strings in $\{0, 1\}^{n(\lambda)}$ into the same set $Y_1 \cap Y_2 = Y_3$.

We build an adversary $\mathcal{A}$ on pseudorandomness with respect to sampler S_3 as follows. On input of u , $\mathcal{A}$ outputs 1 if and only if $D_{S_1}(u) = D_{S_2}(u)$.

$$\begin{aligned} \Pr[\text{Exp}_{\mathcal{A},0}^{\text{pre}}(\lambda) = 1] &\geq 1 - \epsilon''(\lambda) \\ \Pr[\text{Exp}_{\mathcal{A},1}^{\text{pre}}(\lambda) = 1] &\leq \frac{1}{2} - \epsilon'''(\lambda) \end{aligned}$$

Therefore, $\mathcal{A}$ has a non-negligible advantage $\text{Adv}_{\mathcal{A}}^{\text{pre}}(\lambda)$. □

4.1 The pseudorandom encoding hypothesis with setup

We obtain a natural relaxation of the pseudorandom encoding hypothesis by introducing public parameters. We refer to these public parameters as *global* or *non-programmable* common reference

string. That is, a distribution defined via S can be pseudorandomly encoded in this relaxed sense, if there exists a probabilistic setup algorithm Setup_S and encode and decode algorithms as before such that for all $m \in L$, the event $D_S(\text{crs}, E_S(\text{crs}, S(m))) = S(m)$ is overwhelming, where the probability is also over the choice of crs , and the distribution $(\text{Setup}_S(1^\lambda), E_S(\text{Setup}_S(1^\lambda), S(m)))$ is indistinguishable from the distribution $(\text{Setup}_S(1^\lambda), U_{n(\lambda)})$.

Definition 4 (Pseudorandom encoding hypothesis with setup, $\text{cPREH}_{\approx_c}^{\text{rand}}$). $\text{cPREH}_{\approx_c}^{\text{rand}}$ holds if for every PPT algorithm S , there exists a PPT algorithm Setup_S , efficient algorithms (E_S, D_S) , where D_S is deterministic and E_S is randomized (with output length $n(\lambda)$) satisfying the following two requirements.

Correctness. For all PPT adversaries $\mathcal{A}$,

$$\epsilon_{\mathcal{A}}^{\text{c-dec-error}}(\lambda) := \Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ \text{crs} \leftarrow \text{Setup}_S(1^\lambda) : D_S(\text{crs}, m, E_S(\text{crs}, m, y)) \neq y \\ y \leftarrow S(m) \end{array} \right]$$

is negligible.

Pseudorandomness. For all PPT adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{crs-pre}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A},0}^{\text{crs-pre}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},1}^{\text{crs-pre}}(\lambda) = 1] \right|$$

is negligible, where $\text{Exp}_{\mathcal{A},0}^{\text{crs-pre}}$ and $\text{Exp}_{\mathcal{A},1}^{\text{crs-pre}}$ are defined in Figure 3.

$\text{Exp}_{\mathcal{A},0}^{\text{crs-pre}}(\lambda)$	$\text{Exp}_{\mathcal{A},1}^{\text{crs-pre}}(\lambda)$
$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$
$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$
$r \leftarrow \{0, 1\}^{p(\lambda)}$	$u \leftarrow \{0, 1\}^{n(\lambda)}$
$y := S(m; r)$	return $\mathcal{A}(\text{crs}, m, u)$
return $\mathcal{A}(\text{crs}, m, E_S(\text{crs}, m, y))$	

Fig. 3. The static pseudorandomness experiments with setup.

We note that assuming non-uniform adversaries, correctness and pseudorandomness defined in Definition 4 can be equivalently defined by quantifying over all messages $m \in L$. Definition 4 is static in the sense that the inputs $m \in L$ are required to be chosen statically, i.e. independently of crs . In the following, we define the corresponding adaptive variant.

Definition 5 (Adaptive pseudorandom encoding hypothesis with setup, $\text{acPREH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S , there exists a PPT algorithm Setup_S , efficient algorithms (E_S, D_S) , where D_S is deterministic and E_S is randomized (with output length $n(\lambda)$) satisfying the following two requirements.

Correctness. For all PPT adversaries $\mathcal{A}$,

$$\epsilon_{\mathcal{A}}^{\text{ac-dec-error}}(\lambda) := \Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}_S(1^\lambda) \\ m \leftarrow \mathcal{A}(\text{crs}) : D_S(\text{crs}, m, E_S(\text{crs}, m, y)) \neq y \\ y \leftarrow S(m) \end{array} \right]$$

is negligible.

Pseudorandomness. For all PPT adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{a-crs-pre}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A},0}^{\text{a-crs-pre}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},1}^{\text{a-crs-pre}}(\lambda) = 1] \right|$$

is negligible, where $\text{Exp}_{\mathcal{A},0}^{\text{a-crs-pre}}$ and $\text{Exp}_{\mathcal{A},1}^{\text{a-crs-pre}}$ are defined in Figure 4.

$Exp_{\mathcal{A},0}^{\text{a-crs-pre}}(\lambda)$	$Exp_{\mathcal{A},1}^{\text{a-crs-pre}}(\lambda)$
$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$m \leftarrow \mathcal{A}(crs)$	$m \leftarrow \mathcal{A}(crs)$
$r \leftarrow \{0,1\}^{p(\lambda)}$	$u \leftarrow \{0,1\}^{n(\lambda)}$
$y := S(m; r)$	return $\mathcal{A}(crs, m, u)$
return $\mathcal{A}(crs, m, E_S(crs, m, y))$	

Fig. 4. The adaptive pseudorandomness experiments with setup.

Definition 6 ($\text{cPREH}_{\approx_c}^{\text{det}}, \text{cPREH}_{\equiv_s}^{\text{rand}}, \text{cPREH}_{\equiv_s}^{\text{det}}, \text{acPREH}_{\approx_c}^{\text{det}}, \text{acPREH}_{\equiv_s}^{\text{rand}}, \text{acPREH}_{\equiv_s}^{\text{det}}$). *Definitions 4 and 5 can be tuned in two dimensions: the encoding algorithm can be required to be deterministic or allowed to be randomized, and the closeness and invertibility properties can be required to hold statistically or computationally. We denote these variants as $\text{cPREH}_\alpha^\beta$ and $\text{acPREH}_\alpha^\beta$, respectively, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.*

Remark 4 (Universal setup). In Definitions 4 and 5, we allow the algorithm Setup_S to depend on the sampler S . A somewhat incomparable variant of this definition is to require the existence of a *universal setup algorithm* $\text{Setup}(1^\lambda, B)$ for $B \in \mathbb{N}$ which provides the above guarantees for all samplers which can be represented with B bits. We refer to this variant as universal $\text{cPREH}_\alpha^\beta$ and universal $\text{acPREH}_\alpha^\beta$, respectively.

Remark 5 (Common random string). We will denote the strengthening of the pseudorandom encoding hypothesis as in Definitions 4 and 5, where the setup algorithm Setup_S is required to sample uniform random strings, as the pseudorandom encoding hypothesis with a *common random string*.

Note that in Definitions 4 and 5, we implicitly only consider *legitimate* adversaries which guarantee that $m \in L$. For the sake of avoiding notational overhead, we do not make this explicit.

4.2 Static-to-adaptive transformation

Clearly, Definition 5 implies Definition 4. Interestingly, the opposite direction is true as well by an “indirection” argument similar as in [HJKSWZ16]. Similar techniques were also used in the context of non-committing encryption [CPR17] and adaptive multi-party computation [CPV17].

Theorem 1. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If $\text{cPREH}_\alpha^\beta$ is true, then $\text{acPREH}_\alpha^\beta$ is true.*

Proof. We prove the statement for the computational randomized case. The proof directly extends to the remaining cases.

Let S be a PPT sampler with input space L . Since $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true, for the PPT sampler S , there exist $(\text{Setup}'_S, E'_S, D'_S)$ with output length $n'(\lambda)$ such that correctness and pseudorandomness hold (statically) as in Definition 4. Again, since $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true, for the PPT sampler Setup'_S , there exist $(\text{Setup}'', E'', D'')$ with output length $n''(\lambda)$ such that correctness and pseudorandomness hold (statically).⁹ Note that Setup'_S does not expect an input.

In Figure 5, we define algorithms $(\text{Setup}_S, E_S, D_S)$ satisfying *adaptive* correctness and pseudorandomness, as in Definition 5.

On a high level, since crs' is chosen freshly and independently after the adversary fixes the message m , selective security suffices. Furthermore, since the distribution of crs' has no input, selective security is sufficient.

Adaptive correctness

We define a series of hybrid games to prove correctness, see Figure 6. The game hop from $\mathbf{G}_0$ to $\mathbf{G}_1$ only conceptual and $\Pr[out_0 = 1] = \Pr[out_1 = 1]$.

⁹ For notational convenience, we do not write the sampler Setup'_S as index.

$\text{Setup}_S(1^\lambda)$	$E_S(\text{crs}, m, y)$	$D_S(\text{crs}, m, r)$
$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$	$\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$	parse $r =: r_1 \parallel r_2$
$\text{crs} := \text{crs}''$	$r_1 \leftarrow E''(\text{crs}'', \text{crs}')$	$\text{crs}' := D''(\text{crs}'', r_1)$
return crs	$r_2 \leftarrow E'_S(\text{crs}', m, y)$	$y := D'_S(\text{crs}', m, r_2)$
	return $r_1 \parallel r_2$	return y

Fig. 5. Adaptive pseudorandom encodings.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$	$\mathbf{G}_3$
$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$	$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$	$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$	$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$
$m \leftarrow \mathcal{A}(\text{crs}'')$	$\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$	$\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$	$m \leftarrow \mathcal{A}(\text{crs}'')$
$y \leftarrow S(m)$	$r_1 \leftarrow E''(\text{crs}'', \text{crs}')$	$r_1 \leftarrow E''(\text{crs}'', \text{crs}')$	$\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$
// encode	$\text{crs}'_D := D''(\text{crs}'', r_1)$	$\text{crs}'_D := D''(\text{crs}'', r_1)$	$y \leftarrow S(m)$
$\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$	$m \leftarrow \mathcal{A}(\text{crs}'')$	$m \leftarrow \mathcal{A}(\text{crs}'')$	$r_2 \leftarrow E'_S(\text{crs}', m, y)$
$r_1 \leftarrow E''(\text{crs}'', \text{crs}')$	$y \leftarrow S(m)$	$y \leftarrow S(m)$	$y_D := D'_S(\text{crs}', m, r_2)$
$r_2 \leftarrow E'_S(\text{crs}', m, y)$	$r_2 \leftarrow E'_S(\text{crs}', m, y)$	$r_2 \leftarrow E'_S(\text{crs}', m, y)$	return $y_D = y$
// decode	$y_D := D'_S(\text{crs}'_D, m, r_2)$	$y_D := D'_S(\text{crs}', m, r_2)$	
$\text{crs}'_D := D''(\text{crs}'', r_1)$	return $y_D = y$	return $y_D = y$	
$y_D := D'_S(\text{crs}'_D, m, r_2)$			
return $y_D = y$			

Fig. 6. Hybrid games for the proof of adaptive correctness.

Claim. For all PPT adversaries $\mathcal{A}$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]| \leq \epsilon_{(\text{Setup}'', E'', D''), \bar{\mathcal{A}}}^{\text{c-dec-error}}(\lambda)$.

Proof. The games $\mathbf{G}_1$ and $\mathbf{G}_2$ proceed exactly identically if $\text{crs}'_D = \text{crs}'$. Let E be the event that $\text{crs}' \neq \text{crs}'_D$. We have that $\text{out}_1 = 1 \wedge \neg E \Leftrightarrow \text{out}_2 \wedge \neg E$. Due to correctness of $(\text{Setup}'', E'', D'')$,

$$\Pr \left[\begin{array}{l} \text{crs}'' \leftarrow \text{Setup}(1^\lambda) \\ \text{crs}' \leftarrow \text{Setup}'_S(1^\lambda) \\ r_1 \leftarrow E''(\text{crs}'', \text{crs}') \\ \text{crs}'_D := D''(\text{crs}'', r_1) \end{array} : \text{crs}'_D \neq \text{crs}' \right]$$

is negligible. Hence, the Difference Lemma (due to Shoup, [Sho04]) upper bounds

$$|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]| \leq \Pr[E].$$

□

The game hop from $\mathbf{G}_2$ to $\mathbf{G}_3$ only conceptual and $\Pr[\text{out}_2 = 1] = \Pr[\text{out}_3 = 1]$.

Claim. For all PPT adversaries $\mathcal{A}$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $\Pr[\text{out}_3 = 1] \geq 1 - \epsilon_{(\text{Setup}'_S, E'_S, D'_S), \bar{\mathcal{A}}}^{\text{c-dec-error}}(\lambda)$.

Proof. Due to correctness of $(\text{Setup}'_S, E'_S, D'_S)$, we have that for all PPT adversaries $\bar{\mathcal{A}}$,

$$\Pr \left[\begin{array}{l} m \leftarrow \bar{\mathcal{A}}(1^\lambda) \\ \text{crs}' \leftarrow \text{Setup}'_S(1^\lambda) \\ y \leftarrow S(m) \\ r \leftarrow E'_S(\text{crs}', m, y) \\ y_D := D'_S(\text{crs}', m, r) \end{array} : y = y_D \right]$$

is overwhelming. Consider the PPT adversary $\bar{\mathcal{A}}$ which proceeds as follows.

$\bar{\mathcal{A}}(1^\lambda)$

 $\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$
 $m \leftarrow \mathcal{A}(\text{crs}'')$
return m

Therefore, for all PPT adversaries $\mathcal{A}$, $\Pr[out_3 = 1]$ is overwhelming. $\square$

Thus, we have

$$\epsilon_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{ac-dec-error}}(\lambda) \leq \epsilon_{(\text{Setup}'', E'', D''), \bar{\mathcal{A}}}^{\text{c-dec-error}}(\lambda) + \epsilon_{(\text{Setup}'_S, E'_S, D'_S), \bar{\mathcal{A}}'}^{\text{c-dec-error}}(\lambda)$$

for some PPT adversaries $\bar{\mathcal{A}}$ and $\bar{\mathcal{A}}'$.

Adaptive pseudorandomness

We define a series of hybrid games to prove pseudorandomness, see Figure 7.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$	$\mathbf{G}_3$
$crs'' \leftarrow \text{Setup}''(1^\lambda)$	$crs'' \leftarrow \text{Setup}''(1^\lambda)$	$crs'' \leftarrow \text{Setup}''(1^\lambda)$	$crs'' \leftarrow \text{Setup}''(1^\lambda)$
$m \leftarrow \mathcal{A}(crs'')$	$m \leftarrow \mathcal{A}(crs'')$	$crs' \leftarrow \text{Setup}'_S(1^\lambda)$	$r_1 \leftarrow \{0, 1\}^{n'(\lambda)}$
$y \leftarrow S(m)$	$y \leftarrow S(m)$	$r_1 \leftarrow E''(crs'', crs')$	$m \leftarrow \mathcal{A}(crs'')$
$crs' \leftarrow \text{Setup}'_S(1^\lambda)$	$crs' \leftarrow \text{Setup}'_S(1^\lambda)$	$m \leftarrow \mathcal{A}(crs'')$	$r_2 \leftarrow \{0, 1\}^{n'(\lambda)}$
$r_1 \leftarrow E''(crs'', crs')$	$r_1 \leftarrow E''(crs'', crs')$	$r_2 \leftarrow \{0, 1\}^{n'(\lambda)}$	return $\mathcal{A}(crs'', m, r_1 \parallel r_2)$
$r_2 \leftarrow E'_S(cr s', m, y)$	$r_2 \leftarrow \{0, 1\}^{n'(\lambda)}$	return $\mathcal{A}(crs'', m, r_1 \parallel r_2)$	
return $\mathcal{A}(crs'', m, r_1 \parallel r_2)$	return $\mathcal{A}(crs'', m, r_1 \parallel r_2)$		

Fig. 7. Hybrid games for the proof of adaptive pseudorandomness.

Claim. For all PPT adversaries $\mathcal{A}$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq Adv_{(\text{Setup}'_S, E'_S, D'_S), \bar{\mathcal{A}}}^{\text{crs-pre}}(\lambda)$.

Proof. Construct an adversary $\bar{\mathcal{A}}$ on *static* pseudorandomness relative to $(\text{Setup}'_S, E'_S, D'_S)$ as follows. On input of 1^λ , $\bar{\mathcal{A}}$ samples $crs'' \leftarrow \text{Setup}''(1^\lambda)$ calls $\mathcal{A}$ on input of crs'' , and outputs the message m produced by $\mathcal{A}$. In return, $\bar{\mathcal{A}}$ receives $crs' \leftarrow \text{Setup}'_S(1^\lambda)$ and either $u := E'_S(cr s', m, S(m))$ or a uniform random string $u \leftarrow \{0, 1\}^{n'(\lambda)}$ from $Exp_{(\text{Setup}'_S, E'_S, D'_S), \mathcal{A}, b}^{\text{crs-pre}}(\lambda)$. $\bar{\mathcal{A}}$ computes $r_1 \leftarrow E''(crs'', crs')$, calls $\mathcal{A}$ on input of $(crs'', m, r_1 \parallel u)$ and returns $\mathcal{A}$'s output. See Figure 8 for a description of $\bar{\mathcal{A}}$.

$\bar{\mathcal{A}}(1^\lambda)$	$\bar{\mathcal{A}}(crs', m, u)$
$crs'' \leftarrow \text{Setup}''(1^\lambda)$	$\parallel u \leftarrow E'_S(cr s', m, S(m))$ or $u \leftarrow \{0, 1\}^{n'(\lambda)}$
$crs := crs''$	$r_1 \leftarrow E''(crs'', crs')$
$m \leftarrow \mathcal{A}(crs'')$	return $\mathcal{A}(crs'', m, r_1 \parallel u)$
return m	

Fig. 8. Adversary used for the game hop from $\mathbf{G}_0$ to $\mathbf{G}_1$.

If $\bar{\mathcal{A}}$ plays $Exp_{(\text{Setup}'_S, E'_S, D'_S), \bar{\mathcal{A}}, 0}^{\text{crs-pre}}(\lambda)$, then it perfectly simulates $\mathbf{G}_0$. On the other hand, if $\bar{\mathcal{A}}$ plays $Exp_{(\text{Setup}'_S, E'_S, D'_S), \bar{\mathcal{A}}, 1}^{\text{crs-pre}}(\lambda)$, then it perfectly simulates $\mathbf{G}_1$. Hence, we have

$$|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq Adv_{(\text{Setup}'_S, E'_S, D'_S), \bar{\mathcal{A}}}^{\text{crs-pre}}(\lambda).$$

$\square$

The game hop from $\mathbf{G}_1$ to $\mathbf{G}_2$ is only conceptual and $\Pr[out_2 = 1] = \Pr[out_1 = 1]$.

Claim. For all PPT adversaries $\mathcal{A}$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq Adv_{(\text{Setup}'', E'', D''), \bar{\mathcal{A}}}^{\text{crs-pre}}(\lambda)$.

Proof. Construct an adversary $\bar{\mathcal{A}}$ on *static* pseudorandomness relative to $(\text{Setup}'', E'', D'')$ as follows. On input of 1^λ , $\bar{\mathcal{A}}$ returns $\perp$ since the input space L of the sampler $\text{Setup}'_S(1^\lambda)$ is

$\overline{\mathcal{A}}(1^\lambda)$	$\overline{\mathcal{A}}(crs'', u)$
return $\perp$	$\parallel u \leftarrow E''(crs'', \text{Setup}'(1^\lambda)) \text{ or } u \leftarrow \{0, 1\}^{n''(\lambda)}$ $m \leftarrow \mathcal{A}(crs'')$ $r_2 \leftarrow \{0, 1\}^{n''(\lambda)}$ return $\mathcal{A}(crs'', m, u \parallel r_2)$

Fig. 9. Adversary used for the game hop from $\mathbf{G}_2$ to $\mathbf{G}_3$.

empty. In return, $\overline{\mathcal{A}}$ receives crs'' sampled via $\text{Setup}''(1^\lambda)$ and u which is either produced via $E''(crs'', \text{Setup}'(1^\lambda))$ or via uniform sampling from $\{0, 1\}^{n''(\lambda)}$. $\overline{\mathcal{A}}$ calls $\mathcal{A}$ on input of crs'' and receives a message m from $\mathcal{A}$. Finally, $\overline{\mathcal{A}}$ samples $r_2 \leftarrow \{0, 1\}^{n''(\lambda)}$, calls $\mathcal{A}$ on input of $(crs'', m, u \parallel r_2)$ and outputs his output. See Figure 9 for a description of $\overline{\mathcal{A}}$.

If $\overline{\mathcal{A}}$ plays $\text{Exp}_{(\text{Setup}'', E'', D''), \overline{\mathcal{A}}, 0}^{\text{crs-pre}}(\lambda)$, then it perfectly simulates $\mathbf{G}_2$. On the other hand, if $\overline{\mathcal{A}}$ plays $\text{Exp}_{(\text{Setup}'', E'', D''), \overline{\mathcal{A}}, 1}^{\text{crs-pre}}(\lambda)$, then it perfectly simulates $\mathbf{G}_3$. Hence, we have

$$|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \text{Adv}_{(\text{Setup}'', E'', D''), \overline{\mathcal{A}}}^{\text{crs-pre}}(\lambda).$$

□

Therefore, we have

$$\begin{aligned} \text{Adv}_{(\text{Setup}_S, E_S, D_S), \mathcal{A}}^{\text{a-crs-pre}}(\lambda) &= |\Pr[out_3 = 1] - \Pr[out_0 = 1]| \\ &\leq \text{Adv}_{(\text{Setup}'', E'', D''), \overline{\mathcal{A}}}^{\text{crs-pre}}(\lambda) + \text{Adv}_{(\text{Setup}'_S, E'_S, D'_S), \overline{\mathcal{A}}'}^{\text{crs-pre}}(\lambda) \end{aligned}$$

for some PPT adversaries $\overline{\mathcal{A}}$ and $\overline{\mathcal{A}}'$.

□

5 Pseudorandom encodings and invertible sampling

In this section we explain relation between pseudorandom encodings and invertible sampling, [DN00]. In Sections 5.1 and 5.2, we restate the invertible sampling hypothesis of [IKOS10] and define several variants thereof. In Section 5.3, we prove that a distribution can be pseudorandomly encoded if and only if it is inverse samplable. This extends to all of the introduced variations of the pseudorandom encoding hypothesis and the invertible sampling hypothesis.

5.1 The invertible sampling hypothesis

A PPT algorithm S is inverse samplable according to [DN00; IKOS10] if there exists an alternative PPT algorithm $\overline{S}$ and a corresponding inverse sampler $\overline{S}^{-1}$ such that $\overline{S}$ (on every input) induces a distribution which is computationally indistinguishable from the distribution induced by S (on identical inputs) and $\overline{S}^{-1}$ inverts the computation of $\overline{S}$. That is, $\overline{S}^{-1}$ on input of an output produced by $\overline{S}$ produces computationally well-distributed random coins for $\overline{S}$ to produce the given output.

Definition 7 (Invertible sampling hypothesis, $\text{ISH}_{\approx_c}^{\text{rand}}$, [IKOS10]). For every PPT algorithm S , there exists a PPT algorithm $\overline{S}$ (the alternate sampler) with randomness space $\{0, 1\}^{n(\lambda)}$ and an efficient randomized algorithm $\overline{S}^{-1}$ (the inverse sampler), satisfying the following two properties.

Closeness. For all PPT adversaries $\mathcal{A}$ and all inputs $m \in L$,

$$\text{Adv}_{\mathcal{A}, m}^{\text{close}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A}, m, 0}^{\text{close}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A}, m, 1}^{\text{close}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $\text{Exp}_{\mathcal{A}, m, 0}^{\text{close}}$ and $\text{Exp}_{\mathcal{A}, m, 1}^{\text{close}}$ are defined in Figure 10.

Invertibility. For all PPT adversaries $\mathcal{A}$ and all inputs $m \in L$,

$$Adv_{\mathcal{A},m}^{\text{inv}}(\lambda) := \left| \Pr[Exp_{\mathcal{A},m,0}^{\text{inv}}(\lambda) = 1] - \Pr[Exp_{\mathcal{A},m,1}^{\text{inv}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $Exp_{\mathcal{A},m,0}^{\text{inv}}$ and $Exp_{\mathcal{A},m,1}^{\text{inv}}$ are defined in Figure 10.

$Exp_{\mathcal{A},m,0}^{\text{close}}(\lambda)$	$Exp_{\mathcal{A},m,1}^{\text{close}}(\lambda)$	$Exp_{\mathcal{A},m,0}^{\text{inv}}(\lambda)$	$Exp_{\mathcal{A},m,1}^{\text{inv}}(\lambda)$
$r \leftarrow \{0,1\}^{p(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$
$y := S(m; r)$	$y := \bar{S}(m; r)$	$y := \bar{S}(m; r)$	$y := \bar{S}(m; r)$
return $\mathcal{A}(m, y)$	return $\mathcal{A}(m, y)$	return $\mathcal{A}(m, r, y)$	$\bar{r} \leftarrow \bar{S}^{-1}(m, y)$ return $\mathcal{A}(m, \bar{r}, y)$

Fig. 10. The closeness and invertibility experiments.

Definition 8 ($ISH_{\approx_c}^{\text{det}}, ISH_{\equiv_s}^{\text{rand}}, ISH_{\equiv_s}^{\text{det}}$). Definition 7 can be tuned in two dimensions: the inverse sampler can be required to be deterministic or allowed to be randomized, and the closeness and invertibility properties can be required to hold statistically or computationally. We denote these variants as ISH_{α}^{β} , where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

5.2 The invertible sampling hypothesis with setup

The invertible sampling hypothesis can be naturally relaxed by introducing public parameters (or global CRS), henceforth denoted crs . This allows the alternative sampler and the inverse sampler to use crs . Closeness and invertibility are defined against adversaries knowing the crs but choosing the inputs statically.

Definition 9 (Invertible sampling hypothesis with setup, $\text{clSH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S there exists a PPT algorithm Setup_S , a PPT algorithm $\bar{S}$ (with randomness space $\{0,1\}^{n(\lambda)}$) and an efficient randomized $\bar{S}^{-1}$ satisfying the following two properties.

Closeness. For all PPT adversaries $\mathcal{A}$,

$$Adv_{\mathcal{A}}^{\text{crs-close}}(\lambda) := \left| \Pr[Exp_{\mathcal{A},0}^{\text{crs-close}}(\lambda) = 1] - \Pr[Exp_{\mathcal{A},1}^{\text{crs-close}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $Exp_{\mathcal{A},0}^{\text{crs-close}}$ and $Exp_{\mathcal{A},1}^{\text{crs-close}}$ are defined in Figure 11.

Invertibility. For all PPT adversaries $\mathcal{A}$,

$$Adv_{\mathcal{A}}^{\text{crs-inv}}(\lambda) := \left| \Pr[Exp_{\mathcal{A},0}^{\text{crs-inv}}(\lambda) = 1] - \Pr[Exp_{\mathcal{A},1}^{\text{crs-inv}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $Exp_{\mathcal{A},0}^{\text{crs-inv}}$ and $Exp_{\mathcal{A},1}^{\text{crs-inv}}$ are defined in Figure 11.

$Exp_{\mathcal{A},0}^{\text{crs-close}}(\lambda)$	$Exp_{\mathcal{A},1}^{\text{crs-close}}(\lambda)$	$Exp_{\mathcal{A},0}^{\text{crs-inv}}(\lambda)$	$Exp_{\mathcal{A},1}^{\text{crs-inv}}(\lambda)$
$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$
$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$
$r \leftarrow \{0,1\}^{p(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$
$y := S(m; r)$	$y := \bar{S}(\text{crs}, m; r)$	$y := \bar{S}(\text{crs}, m; r)$	$y := \bar{S}(\text{crs}, m; r)$
return $\mathcal{A}(\text{crs}, y)$	return $\mathcal{A}(\text{crs}, y)$	return $\mathcal{A}(\text{crs}, r, y)$	$\bar{r} \leftarrow \bar{S}^{-1}(\text{crs}, m, y)$ return $\mathcal{A}(\text{crs}, \bar{r}, y)$

Fig. 11. The static closeness and invertibility experiments with setup.

Definition 9 is static in the sense that closeness and invertibility adversaries are required to statically choose the challenge input $m \in L$. In the following, we define the corresponding adaptive version, where adversaries are allowed to choose the challenge input $m \in L$ depending on crs .

Definition 10 (Adaptive invertible sampling hypothesis with setup, $\text{acISH}_{\approx_c}^{\text{rand}}$). For every PPT algorithm S , there exists a PPT algorithm Setup_S , a PPT algorithm $\bar{S}$ (with randomness space $\{0,1\}^{n(\lambda)}$) and an efficient randomized $\bar{S}^{-1}$ satisfying the following two properties.

(Adaptive) closeness. For all PPT adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{a-crs-close}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A},0}^{\text{a-crs-close}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},1}^{\text{a-crs-close}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $\text{Exp}_{\mathcal{A},0}^{\text{a-crs-close}}$ and $\text{Exp}_{\mathcal{A},1}^{\text{a-crs-close}}$ are defined in Figure 12.

(Adaptive) invertibility. For all PPT adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{a-crs-inv}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A},0}^{\text{a-crs-inv}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},1}^{\text{a-crs-inv}}(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where $\text{Exp}_{\mathcal{A},0}^{\text{a-crs-inv}}$ and $\text{Exp}_{\mathcal{A},1}^{\text{a-crs-inv}}$ are defined in Figure 12.

$\text{Exp}_{\mathcal{A},0}^{\text{a-crs-close}}(\lambda)$	$\text{Exp}_{\mathcal{A},1}^{\text{a-crs-close}}(\lambda)$	$\text{Exp}_{\mathcal{A},0}^{\text{a-crs-inv}}(\lambda)$	$\text{Exp}_{\mathcal{A},1}^{\text{a-crs-inv}}(\lambda)$
$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$
$m \leftarrow \mathcal{A}(\text{crs})$	$m \leftarrow \mathcal{A}(\text{crs})$	$m \leftarrow \mathcal{A}(\text{crs})$	$m \leftarrow \mathcal{A}(\text{crs})$
$r \leftarrow \{0,1\}^{p(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$	$r \leftarrow \{0,1\}^{n(\lambda)}$
$y := S(m; r)$	$y := \bar{S}(\text{crs}, m; r)$	$y := \bar{S}(\text{crs}, m; r)$	$y := \bar{S}(\text{crs}, m; r)$
return $\mathcal{A}(y)$	return $\mathcal{A}(y)$	return $\mathcal{A}(r, y)$	$\bar{r} \leftarrow \bar{S}^{-1}(\text{crs}, m, y)$ return $\mathcal{A}(\bar{r}, y)$

Fig. 12. The adaptive closeness and invertibility experiments with setup.

Definition 11 ($\text{clSH}_{\approx_c}^{\text{det}}, \text{clSH}_{\equiv_s}^{\text{rand}}, \text{clSH}_{\equiv_s}^{\text{det}}, \text{acISH}_{\approx_c}^{\text{det}}, \text{acISH}_{\equiv_s}^{\text{rand}}, \text{acISH}_{\equiv_s}^{\text{det}}$). Definitions 9 and 10 can be tuned in two dimensions: the inverse sampler can be required to be deterministic or allowed to be randomized, and the closeness and invertibility properties can be required to hold statistically or computationally. We denote these variants as $\text{clSH}_{\alpha}^{\beta}$ and $\text{acISH}_{\alpha}^{\beta}$, respectively, where $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$.

Remark 6 (Universal setup). In Definitions 9 and 10, we allow the algorithm Setup_S to depend on the sampler S . A somewhat incomparable variant of this definition is to require the existence of a *universal setup algorithm* $\text{Setup}(1^\lambda, B)$ for $B \in \mathbb{N}$ which provides the above guarantees for all samplers which can be represented with B bits. We refer to this variant as *universal $\text{clSH}_{\alpha}^{\beta}$* and *universal $\text{acISH}_{\alpha}^{\beta}$* , respectively.

Remark 7 (Common random string). We will denote the strengthening of the invertible sampling hypothesis as in Definitions 9 and 10, where the setup algorithm Setup_S is required to sample uniform random strings, as the invertible sampling hypothesis with a *common random string*.

Again, we note that in Definitions 9 and 10, we implicitly only consider *legitimate* adversaries which guarantee that $m \in L$.

5.3 Equivalence of pseudorandom encodings and invertible sampling

We prove the following theorem.

Theorem 2. Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. $\text{PREH}_{\alpha}^{\beta}$ is true if and only if $\text{ISH}_{\alpha}^{\beta}$ is true.

It is straight forward to extend Theorem 2 to the non-adaptive and adaptive variants with setup.

5.3.1 Every inverse samplable distribution can be pseudorandomly encoded

Lemma 4. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If ISH_α^β holds, then PREH_α^β holds.*

Proof. We prove this for the computational randomized case. The remaining cases are similar.

Assume $\text{ISH}_{\approx_c}^{\text{rand}}$ holds. Let S be a PPT algorithm. $\text{ISH}_{\approx_c}^{\text{rand}}$ implies that there exists an alternative sampler $\bar{S}$ (with randomness space $\{0, 1\}^{n(\lambda)}$) and a corresponding inverse sampler $\bar{S}^{-1}$ satisfying closeness and invertibility.

For $m \in L, y \in \{0, 1\}^*, r \in \{0, 1\}^{n(\lambda)}$, we define the algorithms $E_S(m, y) := \bar{S}^{-1}(m, y)$ (potentially randomized) and $D_S(m, r) := \bar{S}(m; r)$ (deterministic).

Correctness

We consider a series of hybrids, see Figure 13.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$
$r \leftarrow \{0, 1\}^{n(\lambda)}$	$r \leftarrow \{0, 1\}^{n(\lambda)}$	$r \leftarrow \{0, 1\}^{p(\lambda)}$
$y := \bar{S}(m; r)$	$y := \bar{S}(m; r)$	$y := S(m; r)$
return $\mathcal{A}(m, r, y)$	$\bar{r} \leftarrow \bar{S}^{-1}(m, y)$	$\bar{r} \leftarrow \bar{S}^{-1}(m, y)$
	return $\mathcal{A}(m, \bar{r}, y)$	return $\mathcal{A}(m, \bar{r}, y)$

Fig. 13. Hybrids used in the proof of correctness of Lemma 4.

Game $\mathbf{G}_0$ is identical to $\text{Exp}_{\mathcal{A}, m, 0}^{\text{inv}}$ and game $\mathbf{G}_1$ is identical to $\text{Exp}_{\mathcal{A}, m, 1}^{\text{inv}}$. Hence, $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{\mathcal{A}, m}^{\text{inv}}(\lambda)$.

Claim. For all PPT adversaries $\mathcal{A}$, for all $m \in L$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{\bar{\mathcal{A}}, m}^{\text{close}}(\lambda)$.

Proof. Construct an adversary $\bar{\mathcal{A}}$ on closeness. On input of (m, y) , $\bar{\mathcal{A}}$ computes $\bar{r} \leftarrow \bar{S}^{-1}(m, y)$, calls $\mathcal{A}$ on input of $(m, \bar{r}, y)$ and outputs the resulting output. If y is sampled using $\bar{S}(m; r)$ (for $r \leftarrow \{0, 1\}^{n(\lambda)}$), $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_1$ for $\mathcal{A}$. If y is sampled using $S(m; r)$ (for $r \leftarrow \{0, 1\}^{p(\lambda)}$), $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_2$ for $\mathcal{A}$. Therefore, $\Pr[out_1 = 1] = \Pr[\text{Exp}_{\mathcal{A}, m, 1}^{\text{close}}(\lambda) = 1]$ and $\Pr[out_2 = 1] = \Pr[\text{Exp}_{\bar{\mathcal{A}}, m, 0}^{\text{close}}(\lambda) = 1]$. $\square$

Thus, we have that $|\Pr[out_2 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{\bar{\mathcal{A}}, m}^{\text{close}}(\lambda) + \text{Adv}_{\mathcal{A}, m}^{\text{inv}}(\lambda)$ for some PPT adversaries $\bar{\mathcal{A}}, \mathcal{A}'$.

Consider the adversary $\mathcal{A}$ distinguishing between game $\mathbf{G}_0$ and game $\mathbf{G}_2$ that on input of (m, r, y) , outputs 0 if $\bar{S}(m; r) = y$ and outputs 1 otherwise. By definition, $\mathcal{A}$ always outputs 0 in $\mathbf{G}_0$. Hence, $\epsilon_{\text{dec-error}}(\lambda) = \Pr[y \leftarrow S(m) : \bar{S}(m, \bar{S}^{-1}(m, y)) \neq y] = \Pr[out_{2, \mathcal{A}} = 1] = |\Pr[out_{2, \mathcal{A}} = 1] - \Pr[out_{0, \mathcal{A}} = 1]|$.

Pseudorandomness

We consider a sequence of hybrids starting from $\text{Exp}_{\mathcal{A}, m, 0}^{\text{pre}}$ and concluding in $\text{Exp}_{\mathcal{A}, m, 1}^{\text{pre}}$, see Figure 14.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$
$r \leftarrow \{0, 1\}^{p(\lambda)}$	$r \leftarrow \{0, 1\}^{n(\lambda)}$	$r \leftarrow \{0, 1\}^{n(\lambda)}$
$y := S(m; r)$	$y := \bar{S}(m; r)$	return $\mathcal{A}(m, r)$
$u \leftarrow \bar{S}^{-1}(m, y)$	$u \leftarrow \bar{S}^{-1}(m, y)$	
return $\mathcal{A}(m, u)$	return $\mathcal{A}(m, u)$	

Fig. 14. Hybrids used in the proof of pseudorandomness of Lemma 4.

Claim. For all PPT adversaries $\mathcal{A}$, for all $m \in L$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq Adv_{\mathcal{A},m}^{\text{close}}(\lambda)$.

Proof. Construct a PPT adversary $\bar{\mathcal{A}}$ on the closeness property as follows. On input of (m, y) , $\bar{\mathcal{A}}$ calls $\mathcal{A}$ on input of $(m, \bar{S}^{-1}(m, y))$ and outputs the resulting output. If $y \leftarrow S(m)$, $\bar{\mathcal{A}}$ simulates game $\mathbf{G}_0$ for $\mathcal{A}$, and if $y \leftarrow \bar{S}(m)$, $\bar{\mathcal{A}}$ simulates game $\mathbf{G}_1$ for $\mathcal{A}$. Hence, $\Pr[out_0 = 1] = \Pr[Exp_{\mathcal{A},m,0}^{\text{close}}(\lambda) = 1]$ and $\Pr[out_1 = 1] = \Pr[Exp_{\mathcal{A},m,1}^{\text{close}}(\lambda) = 1]$. $\square$

Claim. For all PPT adversaries $\mathcal{A}$, for all $m \in L$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq Adv_{\mathcal{A},m}^{\text{inv}}(\lambda)$.

Proof. We construct a PPT adversary $\bar{\mathcal{A}}$ on the invertibility property. On input of (m, r, y) , $\bar{\mathcal{A}}$ calls $\mathcal{A}$ on input of (m, r) and outputs its output. If $r \leftarrow \bar{S}^{-1}(m, y)$ for $y \leftarrow \bar{S}(m)$, $\bar{\mathcal{A}}$ simulates game $\mathbf{G}_1$ for $\mathcal{A}$. If $r \leftarrow \{0, 1\}^{n(\lambda)}$, $\bar{\mathcal{A}}$ simulates game $\mathbf{G}_2$ for $\mathcal{A}$. Therefore, $\Pr[out_1 = 1] = \Pr[Exp_{\mathcal{A},m,0}^{\text{inv}}(\lambda) = 1]$ and $\Pr[out_2 = 1] = \Pr[Exp_{\mathcal{A},m,1}^{\text{inv}}(\lambda) = 1]$. $\square$

Hence, $Adv_{\mathcal{A},m}^{\text{pre}}(\lambda) = |\Pr[out_2 = 1] - \Pr[out_0 = 1]| \leq Adv_{\mathcal{A},m}^{\text{close}}(\lambda) + Adv_{\mathcal{A},m}^{\text{inv}}(\lambda)$ for some PPT adversaries $\bar{\mathcal{A}}$ and $\bar{\mathcal{A}}'$. $\square$

5.3.2 Every pseudorandomly encodable distribution can be inversely sampled

Lemma 5. Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. If $\text{PREH}_{\alpha}^{\beta}$ holds, then $\text{ISH}_{\alpha}^{\beta}$ holds.

Proof. We prove the statement for the computational randomized case. The remaining cases are similar.

Assume $\text{PREH}_{\approx_c}^{\text{rand}}$ holds. Let S be a PPT algorithm. $\text{PREH}_{\approx_c}^{\text{rand}}$ implies that for S there exist efficient algorithms $\bar{E}_S$ (potentially randomized) with output length $n(\lambda)$ and D_S (deterministic) satisfying correctness and pseudorandomness.

For $m \in L, r \in \{0, 1\}^{n(\lambda)}, y \in \{0, 1\}^*$, we define the alternative sampler as $\bar{S}(m; r) := D_S(m, r)$ (randomized) and the corresponding inverse sampler $\bar{S}^{-1}(m, y) := E_S(m, y)$ (potentially randomized).

Closeness

Let $\mathcal{A}$ be an adversary on closeness. We consider a sequence of games starting from $Exp_{\mathcal{A},m,0}^{\text{close}}$ and concluding in $Exp_{\mathcal{A},m,1}^{\text{close}}$, see Figure 15.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$	$\mathbf{G}_3$	$\mathbf{G}_4$
$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_D \leftarrow \{0, 1\}^{n(\lambda)}$
$y_S := S(m; r_S)$	$y_S := S(m; r_S)$	$y_S := S(m; r_S)$	$y_S := S(m; r_S)$	$y_D := D_S(m, r_D)$
return $\mathcal{A}(m, y_S)$	$r_D \leftarrow \bar{E}_S(m, y_S)$	$r_D \leftarrow E_S(m, y_S)$	$r_D \leftarrow \{0, 1\}^{n(\lambda)}$	return $\mathcal{A}(m, y_D)$
	$y_D := D_S(m, r_D)$	$y_D := D_S(m, r_D)$	$y_D := D_S(m, r_D)$	
	return $\mathcal{A}(m, y_S)$	return $\mathcal{A}(m, y_D)$	return $\mathcal{A}(m, y_D)$	

Fig. 15. Hybrids used in the proof of closeness of Lemma 5.

The difference between game $\mathbf{G}_0$ and game $\mathbf{G}_1$ is only conceptual, hence, $\Pr[out_0 = 1] = \Pr[out_1 = 1]$.

$\mathbf{G}_1$ and $\mathbf{G}_2$ proceed exactly identical if $y_S = y_D$. More formally, let F be the event that $y_S \neq y_D$. We have that $out_1 = 1 \wedge \neg F \Leftrightarrow out_2 \wedge \neg F$. Hence, the Difference Lemma (due to Shoup, [Sho04]) bounds $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \Pr[F]$. Correctness guarantees that for all $m \in L$, $\Pr[F] = \Pr[y_S \leftarrow S(m) : D_S(m, E_S(m, y_S)) \neq y_S] = \epsilon_{\text{dec-error}}(\lambda)$ is negligible.

Claim. For all PPT adversaries $\mathcal{A}$, for all $m \in L$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq Adv_{\mathcal{A},m}^{\text{pre}}(\lambda)$.

Proof. Construct an adversary $\bar{\mathcal{A}}$ on pseudorandomness as follows. On input of $(m, u =: r_D)$, $\bar{\mathcal{A}}$ calls $\mathcal{A}$ on input $(m, D_S(m, r_D))$ and outputs the resulting output. If $u \leftarrow E_S(m, y)$ for $y \leftarrow S(m)$, $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_2$ for $\mathcal{A}$. Otherwise, if u is uniformly random over $\{0, 1\}^{n(\lambda)}$, $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_3$ for $\mathcal{A}$. Hence, $\Pr[out_3 = 1] = \Pr[Exp_{\bar{\mathcal{A}}, m, 1}^{\text{pre}}(\lambda) = 1]$ and $\Pr[out_2 = 1] = \Pr[Exp_{\bar{\mathcal{A}}, m, 0}^{\text{pre}}(\lambda) = 1]$. $\square$

Finally, the difference between $\mathbf{G}_3$ and $\mathbf{G}_4$ is only conceptional and $\Pr[out_3 = 1] = \Pr[out_4 = 1]$. Hence, $Adv_{\mathcal{A}, m}^{\text{close}}(\lambda) = |\Pr[out_4 = 1] - \Pr[out_0 = 1]| \leq Adv_{\mathcal{A}, m}^{\text{pre}}(\lambda) + \epsilon_{\text{dec-error}}(\lambda)$ for some PPT adversary $\bar{\mathcal{A}}$.

Invertibility

We consider a sequence of hybrids, see Figure 16.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$
$r \leftarrow \{0, 1\}^{n(\lambda)}$	$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_S \leftarrow \{0, 1\}^{p(\lambda)}$
$y := D_S(m, r)$	$y_S := S(m; r_S)$	$y_S := S(m; r_S)$
$\bar{r} \leftarrow E_S(m, y)$	$r_D \leftarrow E_S(m, y_S)$	$r_D \leftarrow E_S(m, y_S)$
return $\mathcal{A}(m, \bar{r}, y)$	return $\mathcal{A}(m, r_D, y_S)$	$y_D := D_S(m, r_D)$
		return $\mathcal{A}(m, r_D, y_S)$
$\mathbf{G}_3$	$\mathbf{G}_4$	$\mathbf{G}_5$
$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_S \leftarrow \{0, 1\}^{p(\lambda)}$	$r_D \leftarrow \{0, 1\}^{n(\lambda)}$
$y_S := S(m; r_S)$	$y_S := S(m; r_S)$	$y_D := D_S(m, r)$
$r_D \leftarrow E_S(m, y_S)$	$r_D \leftarrow \{0, 1\}^{n(\lambda)}$	return $\mathcal{A}(m, r_D, y_D)$
$y_D := D_S(m, r_D)$	$y_D := D_S(m, r_D)$	
return $\mathcal{A}(m, r_D, y_D)$	return $\mathcal{A}(m, r_D, y_D)$	

Fig. 16. Hybrids used in the proof of invertibility of Lemma 5.

Claim. For all PPT adversaries $\mathcal{A}$, for all $m \in L$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq Adv_{\mathcal{A}, m}^{\text{pre}}(\lambda) + \epsilon_{\text{dec-error}}(\lambda)$.

Proof. Let $\mathcal{A}$ be an adversary distinguishing $\mathbf{G}_0$ and $\mathbf{G}_1$. Construct an adversary $\bar{\mathcal{A}}$ on the closeness property. On input of (m, y) , $\bar{\mathcal{A}}$ computes $\bar{r} \leftarrow E_S(m, y)$ and calls $\mathcal{A}$ on input $(m, \bar{r}, y)$. If $y \leftarrow \bar{S}(m)$, $\bar{\mathcal{A}}$ simulates game $\mathbf{G}_0$ for $\mathcal{A}$. Else, if $y \leftarrow S(m)$, $\bar{\mathcal{A}}$ simulates game $\mathbf{G}_1$ for $\mathcal{A}$. Hence, $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| = Adv_{\mathcal{A}, m}^{\text{close}}(\lambda)$. $\square$

The difference between $\mathbf{G}_1$ and $\mathbf{G}_2$ is purely conceptional. Hence, $\Pr[out_1 = 1] = \Pr[out_2 = 1]$. $\mathbf{G}_2$ and $\mathbf{G}_3$ behave identical if $y_D = y_S$. Let F denote the failure event $y_D \neq y_S$. We have that $out_2 = 1 \wedge \neg \Rightarrow out_3 \wedge \neg F$. The Difference Lemma (due to Shoup, [Sho04]) bounds $|\Pr[out_3 = 1] - \Pr[out_2 = 1]| \leq \Pr[F]$. Due to correctness, for all $m \in L$, $\Pr[F] = \Pr[y_S \leftarrow S(m): D_S(m, E_S(m, y_S)) \neq y_S] = \epsilon_{\text{dec-error}}(\lambda)$ is negligible.

Claim. For all PPT adversaries $\mathcal{A}$, for all $m \in L$, there exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_4 = 1] - \Pr[out_3 = 1]| \leq Adv_{\mathcal{A}, m}^{\text{pre}}(\lambda)$.

Proof. Construct a PPT adversary $\bar{\mathcal{A}}$ on the pseudorandomness property. On input of (m, u) , $\bar{\mathcal{A}}$ calls $\mathcal{A}$ on input $(m, u =: r_D, D_S(m, u) =: y_D)$ and outputs the resulting output. If $u \leftarrow E_S(m, y)$ for $y \leftarrow S(m)$, $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_3$ for $\mathcal{A}$. Otherwise, if u is uniformly random over $\{0, 1\}^{n(\lambda)}$, $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_4$ for $\mathcal{A}$. Hence, $\Pr[out_3 = 1] = \Pr[Exp_{\bar{\mathcal{A}}, m, 0}^{\text{pre}}(\lambda) = 1]$ and $\Pr[out_4 = 1] = \Pr[Exp_{\bar{\mathcal{A}}, m, 1}^{\text{pre}}(\lambda) = 1]$. $\square$

The difference between $\mathbf{G}_4$ and $\mathbf{G}_5$ is again only conceptual and $\Pr[out_4 = 1] = \Pr[out_5 = 1]$. Hence, $|\Pr[out_5 = 1] - \Pr[out_0 = 1]| \leq 2 \cdot \text{Adv}_{\mathcal{A},m}^{\text{pre}}(\lambda) + 2 \cdot \epsilon_{\text{dec-error}}(\lambda)$ for some PPT adversary $\overline{\mathcal{A}}$. $\square$

The above proof directly generalizes to the non-adaptive and adaptive variants of pseudorandom encodings and invertible sampling with public parameters. As a demonstration, consider the proof of closeness of Lemma 5. More specifically, consider the game hop from $\mathbf{G}_1$ to $\mathbf{G}_2$. In these games, the (stateful) adversary $\mathcal{A}$ picks m (either statically or adaptively after seeing crs). The error event F occurs if $y_S \neq y_D$, where $y_S \leftarrow S(m)$ and $y_D \leftarrow D_S(\text{crs}, m, E_S(\text{crs}, m, y_S))$. We can hence construct an adversary $\overline{\mathcal{A}}$ against correctness which chooses m (non-)adaptively like $\mathcal{A}$. Hence, the probability that F occurs can be upper bounded by $\epsilon_{\mathcal{A}}^{\text{c-dec-error}}(\lambda)$ or $\epsilon_{\mathcal{A}}^{\text{ac-dec-error}}(\lambda)$, respectively, which are both negligible by correctness.

Theorem 2 together with Theorem 1 yields the following corollaries.

Corollary 1. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. $\text{cPREH}_{\alpha}^{\beta}$ is true, then $\text{acISH}_{\alpha}^{\beta}$ is true.*

Corollary 2. *Let $\alpha \in \{\approx_c, \equiv_s\}$ and $\beta \in \{\text{rand}, \text{det}\}$. $\text{cISH}_{\alpha}^{\beta}$ is true, then $\text{acISH}_{\alpha}^{\beta}$ is true.*

This is an excellent example of the potential pseudorandom encodings offer. Looking through the lens of invertible sampling, it is rather unclear how the static and adaptive notions relate, whereas pseudorandom encodings directly provide a static-to-adaptive transformation.

Particularly, Corollary 2 together with [DKR15] yields the first instantiation of an *adaptive* explainability compiler without complexity leveraging and, hence, based only on polynomial hardness assumptions. The recent paper [CsW19] uses such an adaptive explainability compiler to obtain adaptive MPC with communication complexity which is sub-linear in the circuit size. Their construction relies on complexity leveraging which entails a sub-exponential loss relative to IO and one-way functions. Hence, we obtain the following corollary improving on Theorem 7 in the proceedings version of [CsW19] in a black-box way.

Corollary 3. *Assuming polynomially secure IO and the adaptive hardness of LWE, then succinct adaptive two-round MPC in the malicious setting is possible (for all deterministic n -party functionalities $f: (\{0,1\}^{\ell_{\text{in}}})^n \rightarrow \{0,1\}^{\ell_{\text{out}}}$ with circuit depth d) in the (local) CRS model such that the size of the CRS, the communication complexity, and online-computational complexity of the protocol are polynomial in $\lambda, \ell_{\text{in}}, \ell_{\text{out}}, d$ and n).*

6 Classification of the different flavors of pseudorandom encodings

In this section we classify the different variants of the pseudorandom encoding hypothesis. In Section 6.1, we study the pseudorandom encoding hypothesis with deterministic encoding algorithm and identify a relation to compression. In Section 6.2, we study the pseudorandom encoding hypothesis with randomized encoding and its conflicts with extractable one-way functions. In Section 6.3, we describe an instantiation of $\text{cPREH}_{\approx_c}^{\text{rand}}$ (with universal setup) based on indistinguishability obfuscation and one-way functions due to [SW14; DKR15]. Finally, in Section 6.4, we bootstrap $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common *random* string from the construction in Section 6.3 additionally assuming *weak* $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common *random* string.

6.1 Deterministic encoding

For the purpose of classifying pseudorandom encodings with a deterministic encoding algorithm, we first introduce some notions of entropy and computational analogues thereof.

Definition 12 (Min-entropy, [Rey11]). *The min-entropy of a distribution X is defined as $H_{\infty}(X) = -\log \max_{x \in \text{supp}(X)} \Pr[X = x]$.*

(By $\log$ we always mean the logarithm to base 2.) Intuitively, this captures the ability to guess the value of a distribution in a single attempt. However, this is a very pessimistic view. For many purposes, it suffices to work with a distribution which is statistically close (i.e. has negligible statistical distance) to a distribution with high min-entropy.¹⁰

Definition 13 (ϵ -smooth min-entropy, [Rey11]). A source X has ϵ -smooth min-entropy at least k , denoted as $H_\infty^\epsilon(X) \geq k$, if there exists a distribution Y with $\Delta(X, Y) \leq \epsilon$ such that $H_\infty(Y) \geq k$.

In many cases, some information Z that is correlated to the actual source X is known. Since for our purposes, the conditional part Z is not under adversarial control, we use the notion of average conditional min-entropy as used in [HLR07; DORS08].

Definition 14 (Average min-entropy, [HLR07; DORS08; Rey11]). Let (Y, Z) be a joint distribution. The average min-entropy of X conditioned on Z is

$$\tilde{H}_\infty(X | Z) := -\log(\mathbb{E}_{z \leftarrow Z}[\max_x \Pr[X_z = x]]).$$

Definition 15 (Average ϵ -smooth min-entropy, [DORS08; Rey11]). Let (Y, Z) be a joint distribution. The distribution Y has average ϵ -smooth min-entropy at least k conditioned on Z , denoted as $\tilde{H}_\infty^\epsilon(Y | Z) \geq k$, if there exists a joint distribution (Y', Z') with $\Delta((Y, Z), (Y', Z')) \leq \epsilon$ such that $\tilde{H}_\infty(Y' | Z') \geq k$.

Let X be an efficiently samplable distribution. In the literature, there are several computational notions of entropy. HILL entropy [HILL99] constitutes a natural computational variant of min-entropy. A source has high HILL entropy, if it is computationally indistinguishable from a source that has high min-entropy.

Definition 16 (HILL entropy, [HILL99; BSW03]). A distribution X has HILL entropy at least k , denoted by $H_{\epsilon, s}^{\text{HILL}}(X) \geq k$, if there exists a distribution Y such that $H_\infty(Y) \geq k$ and $|\Pr[x \leftarrow X : \mathcal{A}(x) = 1] - \Pr[y \leftarrow Y : \mathcal{A}(y) = 1]| \leq \epsilon$ for all $\mathcal{A}$ of size at most s .

Shannon's theorem [Sha48] states that the minimum compression length (over all compression and decompression algorithms) of a distribution equals its average entropy (up to small additive terms). Yao entropy [Yao82] constitutes the corresponding computational counterpart. Intuitively, a source has high Yao entropy, if it can not be *efficiently* compressed.

Definition 17 (Yao entropy, [Yao82; BSW03]). A distribution X has Yao entropy at least k , denoted by $H_{\epsilon, s}^{\text{Yao}}(X) \geq k$, if for every pair of circuits (E, D) of total size s with outputs of E having length ℓ , $\Pr_{x \leftarrow X}[D(E(x)) = x] \leq 2^{\ell-k} + \epsilon$.

When we omit the subscripts for H^{HILL} and H^{Yao} , we mean $H_{\epsilon, s}^{\text{HILL}}$ and $H_{\epsilon, s}^{\text{Yao}}$ for any negligible ϵ and polynomial s , respectively. Since compressibility implies distinguishability, HILL entropy implies Yao entropy. The converse, however, is believed to be false, [Wee04; HLR07].

Definition 18 (Conditional HILL entropy, [HLR07]). For a distribution (X, Z) , we say that X has HILL entropy at least k conditioned on Z , denoted by $H_{\epsilon, s}^{\text{HILL}}(X | Z) \geq k$, if there exists a collection of distributions Y_z giving rise to a joint distribution (Y, Z) , such that $\tilde{H}_\infty(Y | Z) \geq k$ and $|\Pr[(x, z) \leftarrow (X, Z) : \mathcal{A}(x, z) = 1] - \Pr[(y, z) \leftarrow (Y, Z) : \mathcal{A}(y, z) = 1]| \leq \epsilon$ for all circuits $\mathcal{A}$ of size at most s .

Conditional Yao entropy is defined by simply giving the compressor and decompressor algorithm the value z as input.

¹⁰ We note that if some distribution Y is ϵ -close to a distribution Z , $H_\infty(Y) \geq -\log(2^{-H_\infty(Z)} + \epsilon(\lambda)) = H_\infty(Z) - \log(1 + \epsilon(\lambda)2^{H_\infty(Z)})$.

Definition 19 (Conditional Yao entropy, [HLR07]). For a distribution (X, Z) , we say that X has Yao entropy at least k conditioned on Z , denoted by $H_{\epsilon, s}^{\text{Yao}}(X | Z) \geq k$, if for every pair of circuits (E, D) of total size at most s with outputs of E having length ℓ , $\Pr_{(x, z) \leftarrow (X, Z)}[D(z, E(z, x)) = x] \leq 2^{\ell-k} + \epsilon$.

6.1.1 Information theoretic guarantees and compression

Traditionally, the theory of compression mostly considers family of sources that are not indexed by strings. Let Σ be some alphabet. In this work, we always consider $\Sigma := \{0, 1\}$.

Definition 20 ([Wee04; TVZ05]). For functions $E_X: \Sigma^* \rightarrow \Sigma^*$ and $D_X: \Sigma^* \rightarrow \Sigma^*$, we say (E_X, D_X) compresses source X to length m with decoding error ϵ if

1. $\Pr[x \leftarrow X: D_X(E_X(x)) \neq x] \leq \epsilon$, and
2. $\mathbb{E}[|E_X(X)|] \leq m$.

Definition 21 ([Wee04; TVZ05]). We say source X is compressible to length (exactly) m if there exist functions E_X and D_X such that (E_X, D_X) compresses X to length (exactly) m .

Lemma 6 ([TVZ05]). Let X_λ be a source on $\{0, 1\}^\lambda$ which is compressible to length m with decoding error ϵ by algorithms (E_X, D_X) . Further, let $m_0 \in \mathbb{N}$ be a lower bound on the output length of E_X , i.e. such that for all $x \in \text{supp}(X_\lambda)$, $|E_X(x)| \geq m_0$. Then, X_λ is compressible to length $m + \epsilon(\lambda - m_0) + 1$ with decoding error 0.

Proof (sketch). To show this, [TVZ05] construct an encoding algorithm E'_X which on input of x tests if the $D_X(E_X(x)) = x$. If this is the case, E'_X outputs $0 \parallel E_X(x)$. Else, E'_X outputs $1 \parallel x$. $\square$

The statistical deterministic variant of the pseudorandom encoding hypothesis is strongly related to compression.

Theorem 3. If (weak) $\text{PREH}_{\equiv_s}^{\text{det}}$ is true for X , i.e. there exist deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, the ϵ -smooth min-entropy $H_\infty^\epsilon(X) \geq n$ for some negligible function ϵ .

Proof. Consider the distribution $Y' := D_X(U_n)$. As already seen in the proof of Theorem 2, Y' and X are statistically indistinguishable due to correctness and pseudorandomness. Hence, the statistical distance $\Delta(X, Y') \leq \delta(\lambda)$ for some negligible δ .

Due to correctness and since E_X is deterministic, $\Pr[x \leftarrow X: D_X(E_X(x)) = x] = \Pr[x \leftarrow X: E_X(D_X(E_X(x))) = E_X(x)] \geq 1 - \nu(\lambda)$ for some negligible function ν . Applying pseudorandomness, we get that the probability

$$\Pr[u \leftarrow U_n: E_X(D_X(u)) = u] \geq 1 - \nu'(\lambda)$$

for some negligible function ν' , where the probability is only over the choice of u . Therefore, D_X operates almost injectively on the set $\{0, 1\}^n$. More formally, let $V_0 \subset \{0, 1\}^n$ denote the set of all u such that $E_X(D_X(u)) = u$ and let $V_1 := \{0, 1\}^n \setminus V_0$. We have that, $|V_1|/|\{0, 1\}^n| \leq \nu'(\lambda)$. Let $\overline{V}_1$ be some arbitrary subset of $\{0, 1\}^{p(\lambda)} \setminus D_X(V_0)$ such that $|\overline{V}_1| = |V_1|$ (for some polynomial p). Let Y'' be the uniform distribution over $D_X(V_0) \cup \overline{V}_1$. Note that $|D_X(V_0) \cup \overline{V}_1| = 2^n$ and $\Pr[Y' \in \overline{V}_1] \leq \Pr[Y' \notin D_X(V_0)] \leq \nu'(\lambda)$. The statistical distance between Y' and Y'' is negligible.

$$\begin{aligned} \Delta(Y', Y'') &= \sum_{a \in D_X(V_0) \cup \overline{V}_1} |\Pr[Y' = a] - \Pr[Y'' = a]| \\ &= \underbrace{\sum_{a \in D_X(V_0)} |\Pr[Y' = a] - \Pr[Y'' = a]|}_{\leq 1 - |V_0| \cdot 2^{-n} \leq 1 - (1 - \nu'(\lambda))} + \underbrace{\sum_{a \in \overline{V}_1} |\Pr[Y' = a] - \Pr[Y'' = a]|}_{\leq \sum_{a \in \overline{V}_1} (\Pr[Y' = a] + \Pr[Y'' = a]) \leq 2 \cdot \nu'(\lambda)} \leq 3 \cdot \nu'(\lambda) \end{aligned}$$

Furthermore, since Y'' is the uniform distribution over a set of size 2^n , its min-entropy equals n . Since $\Delta(X, Y') \leq \delta(\lambda)$ and $\Delta(Y', Y'') \leq 3\nu'(\lambda)$ for negligible functions δ and ν' , we have $H_\infty^{\delta+3\nu'}(X) \geq H_\infty(Y'') = n$. $\square$

Corollary 4. *If $\text{PREH}_{\equiv_s}^{\text{det}}$ is true for some source X , there exists an $n \leq H_\infty^\epsilon(X)$ (for some negligible function ϵ), such that X is compressible to length exactly n . The decoding error can then be eliminated applying Lemma 6.*

Hence, for $\text{PREH}_{\equiv_s}^{\text{det}}$ to be true for a source X , it is necessary that X is efficiently compressible. However, this is not a sufficient criterion since a compression algorithm can have some structure which makes it easily distinguishable from uniform randomness (e.g. the compression algorithm used in the proof of Lemma 6).

The distribution induced by pseudorandom generators can provably not be compressed. However, those distributions have low ϵ -smooth min-entropy.

Lemma 7. *Let iPRG be an injective pseudorandom generator with polynomial stretch $\text{poly}(\cdot)$ and let ϵ be a negligible function. Then, $H_\infty^\epsilon(\text{iPRG}(U_\lambda)) \leq \lambda + \delta(\lambda)$ for some negligible function δ .*

Proof. Since $\text{iPRG}(U_\lambda)$ is a uniform distribution over $\text{iPRG}(\{0, 1\}^\lambda)$, it suffices to consider all ϵ -close flat distributions over $\{0, 1\}^{\text{poly}(\lambda)}$ to obtain an upper bound on $H_\infty^\epsilon(\text{iPRG}(U_\lambda))$. Let D the uniform distribution on $\text{supp}(\text{iPRG}(U_\lambda)) \cup A$, where $\text{supp}(\text{iPRG}(U_\lambda)) \cap A = \emptyset$. Let $|A| := k$. Then,

$$\begin{aligned} \Delta(\text{iPRG}(U_\lambda), D) &= \sum_{a \in \text{supp}(\text{iPRG}(U_\lambda))} \left| \frac{1}{2^\lambda} - \frac{1}{2^\lambda + k} \right| + \sum_{a \in A} \left| \frac{1}{2^\lambda + k} \right| \\ &= 1 - \frac{2^\lambda}{2^\lambda + k} + \frac{k}{2^\lambda + k} = \frac{2k}{2^\lambda + k}. \end{aligned}$$

Hence, for $\Delta(\text{iPRG}(U_\lambda), D) \leq \epsilon(\lambda)$, k is upper bounded by

$$k \leq \underbrace{k(2 - \epsilon(\lambda))}_{\geq 1} \leq 2^\lambda \cdot \epsilon(\lambda).$$

Therefore, the min-entropy of D is at most $H_\infty(D) = \log(2^\lambda + k) \leq \log(2^\lambda + 2^\lambda \epsilon) = \lambda + \log(1 + \epsilon)$, which is negligibly close to λ . Hence, $H_\infty^\epsilon(\text{iPRG}(U_\lambda)) \leq \lambda + \delta(\lambda)$ for some negligible function δ . $\square$

Lemma 8. *Let PRG be a pseudorandom generator with polynomial stretch $\text{poly}(\cdot)$ and let ϵ be a negligible function. Then, $H_\infty^\epsilon(\text{PRG}(U_\lambda)) \leq \lambda + \delta(\lambda)$ for some negligible function δ .*

Proof. Due to Lemma 1, there exists a uniform distribution D over a set of size 2^λ such that $\text{PRG}(U_\lambda)$ and D are statistically close. By a similar argument as in Lemma 7, in order to upper bound the ϵ -smooth min-entropy of $\text{PRG}(U_\lambda)$, it suffices to consider all ϵ -close flat distributions. Then, due to the computations already made in the proof of Lemma 7, Lemma 8 follows. $\square$

Therefore, we obtain the following corollary.

Corollary 5. *If one-way functions exist, $\text{PREH}_{\equiv_s}^{\text{det}}$ is false.*

Intuitively, a common reference string can not add additional entropy to $E_X(\text{crs}, X)$. I.e. if $\text{cPREH}_{\equiv_s}^{\text{det}}$ is true for some source X , the algorithms (E_X, D_X) implied by $\text{cPREH}_{\equiv_s}^{\text{det}}$ compress that source to its ϵ -smooth min-entropy. To prove this, we introduce the following technical lemma.

Lemma 9 (Splitting lemma). *Let X and Y_x be distributions giving rise to a joint distribution (Y, X) . Let $A \subset \text{supp}(Y, X)$ such that $\Pr_{(y,x) \leftarrow (Y,X)}[(y, x) \in A] \geq 1 - \mu(\lambda)$ for some negligible function μ . Further, for all x , let $p_x := \Pr_{y \leftarrow Y_x}[(y, x) \in A]$. Then, there exists negligible functions ν, ν' and a set $G \subseteq \text{supp}(X)$, such that for all $x \in G$, $p_x \geq 1 - \nu(\lambda)$ and $\Pr_{x \leftarrow X}[x \in G] \geq 1 - \nu'(\lambda)$.*

Proof. Define $G_n := \{x \in \text{supp}(X) : p_x \geq 1 - n \cdot \mu(\lambda)\}$.

$$\begin{aligned}
1 - \mu(\lambda) &\leq \Pr_{(y,x) \leftarrow (Y,X)}[(y, x) \in A] \\
&= \sum_{x \in G_n} \Pr[X = x] \cdot \underbrace{p_x}_{\leq 1} + \sum_{x \notin G_n} \Pr[X = x] \cdot \underbrace{p_x}_{< 1 - n \cdot \mu(\lambda)} \\
&< \Pr_{x \leftarrow X}[x \in G_n] + (1 - n \cdot \mu(\lambda)) \cdot \underbrace{\sum_{x \notin G_n} \Pr[X = x]}_{= 1 - \Pr_{x \leftarrow X}[x \in G_n]} \\
&= \Pr_{x \leftarrow X}[x \in G_n] + 1 - n \cdot \mu(\lambda) - \Pr_{x \leftarrow X}[x \in G_n] + n \cdot \mu(\lambda) \cdot \Pr_{x \leftarrow X}[x \in G_n]
\end{aligned}$$

Hence,

$$\Pr_{x \leftarrow X}[x \in G_n] \geq 1 - \frac{1}{n}.$$

Without loss of generality, we assume $\mu(\lambda) > 0$ for all $\lambda \in \mathbb{N}$. (If $\mu(\lambda) = 0$ for some λ , then for $G := \text{supp}(X)$, $p_x = 1$ for all $x \in G$, and $\Pr_{x \leftarrow X}[x \in G] = 1$.) Then, $\sqrt{\mu(\lambda)}$ is well defined and negligible. Let $n := \sqrt{\mu(\lambda)}^{-1}$. Then, by definition of G_n , for all $x \in G_n$, $p_x \geq 1 - n \cdot \mu(\lambda) = 1 - \sqrt{\mu(\lambda)}$. Furthermore, $\Pr_{x \leftarrow X}[x \in G_n] \geq 1 - \frac{1}{n} = 1 - \sqrt{\mu(\lambda)}$. $\square$

Theorem 4. *If (weak) $\text{cPREH}_{\equiv}^{\text{det}}$ is true for X , i.e. there exist a setup algorithm Setup_X and deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, the ϵ -smooth min-entropy $H_{\infty}^{\epsilon}(X) \geq n$ for some negligible function ϵ .*

Proof. Note that $H_{\infty}^{\epsilon}(X) = \tilde{H}_{\infty}^{\epsilon}(X \mid \text{Setup}_X(1^{\lambda}))$ as the distribution X is independent of the CRS. Hence, it suffices to upper bound the average ϵ -smooth min-entropy $\tilde{H}_{\infty}^{\epsilon}(X \mid \text{Setup}_X(1^{\lambda}))$.

The proof is similar to the proof of Theorem 3. We consider the distribution $Y'_{\text{crs}} := D_X(\text{crs}, U_n)$. Due to Theorem 2, the distributions $\{\text{crs} \leftarrow \text{Setup}_X(1^{\lambda}), u \leftarrow U_n : (\text{crs}, D_X(\text{crs}, u))\}$ and $\{\text{crs} \leftarrow \text{Setup}_X(1^{\lambda}), y \leftarrow X : (\text{crs}, y)\}$ are statistically indistinguishable. Hence,

$$\Delta((X \mid \text{Setup}_X(1^{\lambda})), (Y' \mid \text{Setup}_X(1^{\lambda}))) \leq \delta(\lambda)$$

for some negligible function δ .

By a similar argument as in Theorem 3, due to correctness and pseudorandomness we have

$$\Pr[\text{crs} \leftarrow \text{Setup}_X(1^{\lambda}), u \leftarrow U_n : E_X(\text{crs}, D_X(\text{crs}, u)) = u] \geq 1 - \nu'(\lambda)$$

for some negligible function ν' , where the probability is over the choice of crs and u .

The Splitting lemma (Lemma 9) implies that there exists $G \subseteq \text{supp}(\text{Setup}_X(1^{\lambda}))$ such that $\Pr[\text{crs} \leftarrow \text{Setup}_X(1^{\lambda}) : \text{crs} \in G] \geq 1 - \nu''(\lambda)$ and for all $\text{crs} \in G$, $\Pr[u \leftarrow U_n : E_X(\text{crs}, D_X(\text{crs}, u)) = u] \geq 1 - \nu'''(\lambda)$ for some negligible functions ν'', ν''' . Hence, conditioned on $\text{crs} \in G$, $D_X(\text{crs}, \cdot)$ operates almost injectively on the set $\{0, 1\}^n$. More formally, let $V_0^{\text{crs}} \subset \{0, 1\}^n$ denote the set of all u such that $E_X(\text{crs}, D_X(\text{crs}, u)) = u$ and let $V_1^{\text{crs}} := \{0, 1\}^n \setminus V_0^{\text{crs}}$. For $\text{crs} \in G$, we have $|V_1^{\text{crs}}|/|\{0, 1\}^n| \leq \nu'''(\lambda)$. Let $\bar{V}_1^{\text{crs}}$ be some arbitrary subset of $\{0, 1\}^{p(\lambda)} \setminus D_X(\text{crs}, V_0^{\text{crs}})$ such that $|\bar{V}_1^{\text{crs}}| = |V_1^{\text{crs}}|$ (for some polynomial p). Let Y'' be the uniform distribution over $D_X(V_0^{\text{crs}}) \cup \bar{V}_1^{\text{crs}}$. Note that $|D_X(\text{crs}, V_0^{\text{crs}}) \cup \bar{V}_1^{\text{crs}}| = 2^n$ and for $\text{crs} \in G$, we have $\Pr[Y'_{\text{crs}} \in \bar{V}_1^{\text{crs}}] \leq \Pr[Y'_{\text{crs}} \notin$

$D_X(crs, V_0^{crs})] \leq \nu'(\lambda)$. Hence,

$$\begin{aligned}
& \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \left(\sum_{a \in D_X(crs, V_0^{crs})} |\Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a]| \right. \\
& \quad \left. + \sum_{a \in \overline{V_1^{crs}}} |\Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a]| \right) \\
& \leq \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \left(1 - \frac{|V_0^{crs}|}{2^n} + \Pr[Y'_{crs} \in \overline{V_1^{crs}}] + \Pr[Y''_{crs} \in \overline{V_1^{crs}}] \right) \\
& \leq \sum_{crs \in G} \Pr[\text{Setup}_X = crs] \cdot \left(\frac{|V_1^{crs}|}{2^n} + \Pr[Y'_{crs} \notin D_X(crs, V_0^{crs})] + \frac{|\overline{V_1^{crs}}|}{2^n} \right) \\
& \leq 3 \cdot \nu'''(\lambda)
\end{aligned} \tag{6}$$

$$\begin{aligned}
& \sum_{crs \notin G} \Pr[\text{Setup}_X = crs] \cdot \left(\sum_{a \in \text{supp}(Y''_{crs})} |\Pr[Y'_{crs} = a] - \Pr[Y''_{crs} = a]| \right) \\
& \leq \sum_{crs \notin G} \Pr[\text{Setup}_X = crs] \cdot 2 \leq 2 \cdot \nu''(\lambda)
\end{aligned} \tag{7}$$

Due to Equations (6) and (7) we have that $\Delta((Y', \text{Setup}_X), (Y'', \text{Setup}_X)) \leq 2 \cdot \nu''(\lambda) + 3 \cdot \nu'''(\lambda)$.

Furthermore, since for all $crs \in \text{supp}(\text{Setup}_X(1^\lambda))$, Y''_{crs} is the uniform distribution over a set of 2^n elements,

$$\tilde{H}_\infty(Y'' \mid \text{Setup}_X(1^\lambda)) = -\log \left(\mathbb{E}_{crs \leftarrow \text{Setup}_X(1^\lambda)} \max_a \Pr[Y''_{crs} = a] \right) = n.$$

Therefore, we have $H_\infty^{\delta+2\nu''+3\nu'''}(X) = \tilde{H}_\infty^{\delta+2\nu''+3\nu'''}(X \mid \text{Setup}_X(1^\lambda)) \geq n$, where δ, ν'', ν''' are negligible functions. $\square$

Corollary 6. *If one-way functions exist, $\text{cPREH}_{\equiv_s}^{\text{det}}$ is false.*

6.1.2 Computational guarantees and pseudoentropy

We study the relation of the pseudorandom encoding hypothesis with deterministic encoding algorithm to HILL and Yao entropy. Interestingly, $\text{PREH}_{\approx_c}^{\text{det}}$ poses a lower bound on the HILL entropy of a source and an upper bound on its Yao entropy.

HILL entropy. Intuitively, the algorithms (E_X, D_X) implied by $\text{PREH}_{\approx_c}^{\text{det}}$ compress a source X to its HILL entropy.

Theorem 5. *If (weak) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i.e. there exist deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{HILL}}(X) \geq n$.*

Proof. We employ a similar strategy as in the proof of Theorem 3. Let $Y' := D_X(U_n)$. As already seen in the proof of Theorem 2, Y' and X are computationally indistinguishable due to correctness and pseudorandomness.

Due to correctness, $\Pr[x \leftarrow X : D_X(E_X(x)) = x] \geq 1 - \nu(\lambda)$ for some negligible function ν . Since E_X is required to be deterministic, we get $\Pr[x \leftarrow X : E_X(D_X(E_X(x))) = E_X(x)] \geq 1 - \nu(\lambda)$. Applying pseudorandomness, we get that the probability

$$\Pr[u \leftarrow U_n : E_X(D_X(u)) = u] \geq 1 - \nu'(\lambda)$$

for some negligible function ν' , where the probability is only over the choice of u . Let $V_0 \subset \{0, 1\}^n$ denote the set of all u such that $E_X(D_X(u)) = u$ and let $V_1 := \{0, 1\}^n \setminus V_0$. We have that, $|V_1|/|\{0, 1\}^n| \leq \nu'(\lambda)$.

Let $\overline{V}_1$ be some arbitrary subset of $\{0, 1\}^{p(\lambda)} \setminus D_X(V_0)$ such that $|\overline{V}_1| = |V_1|$ (for some polynomial p). Let Y'' be the uniform distribution over $D_X(V_0) \cup \overline{V}_1$. The min-entropy of Y'' equals n . By an argument already made in Theorem 3, the statistical distance between Y' and Y'' is negligible.

Therefore, the distributions X and Y'' are computationally indistinguishable and, hence, $H^{\text{HILL}}(X) \geq n$. $\square$

This result can be generalized to conditional HILL entropy using strong $\text{PREH}_{\approx_c}^{\text{det}}$.

Theorem 6. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (strong) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i.e. there exist two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{HILL}}(X | Z) \geq n$.*

Proof. For each $z \in \text{supp}(Z)$, consider the distribution $Y'_z := D_X(z, U_n)$. Due to Theorem 2, for all adversarially chosen z , the distributions $\{(Y'_z, z)\}$ and $\{(X_z, z)\}$ are computationally indistinguishable due to correctness and pseudorandomness.

Due to correctness, for all adversarially chosen z , $\Pr[x \leftarrow X_z : D_X(z, E_X(z, x)) = x] \geq 1 - \nu(\lambda)$ for some negligible function ν . Since E_X is required to be deterministic, we have that for all adversarially chosen z , $\Pr[x \leftarrow X_z : E_X(z, D_X(z, E_X(z, x))) = E_X(z, x)] \geq 1 - \nu(\lambda)$. Applying pseudorandomness, we get that there exists a negligible function ν' such that for all adversarially chosen z ,

$$\Pr[u \leftarrow U_n : E_X(z, D_X(z, u)) = u] \geq 1 - \nu'(\lambda),$$

where the probability is only over the choice of u . Let $V_{0,z} \subseteq \{0, 1\}^n$ denote the set of all $u \in \{0, 1\}^n$ such that $E_X(z, D_X(z, u)) = u$ holds and let $V_{1,z} := \{0, 1\}^n \setminus V_{0,z}$. We have that, $|V_{1,z}|/|\{0, 1\}^n| \leq \nu'(\lambda)$.

Let $\overline{V}_{1,z}$ be some arbitrary subset of $\{0, 1\}^{p(\lambda)} \setminus D_X(z, V_{0,z})$ such that $|\overline{V}_{1,z}| = |V_{1,z}|$ (for some polynomial p). Let Y''_z be the uniform distribution over $D_X(z, V_{0,z}) \cup \overline{V}_{1,z}$.

$$\tilde{H}_\infty(Y'' | Z) = -\log \left(\mathbb{E}_{z \leftarrow Z} \left[\max_{y \in \text{supp}(Y''_z)} \Pr[Y''_z = y] \right] \right) = n.$$

By a similar argument as in Theorem 3, for all adversarially chosen z , the statistical distance between Y'_z and Y''_z is negligible.

$$\begin{aligned} \Delta(Y'_z, Y''_z) &= \underbrace{\sum_{a \in D_X(z, V_{z,0})} |\Pr[Y'_z = a] - \Pr[Y''_z = a]|}_{\leq 1 - |V_{0,z}| \cdot 2^{-n} \leq 1 - (1 - \nu'(\lambda))} + \underbrace{\sum_{a \in \overline{V}_{1,z}} |\Pr[Y'_z = a] - \Pr[Y''_z = a]|}_{\leq \sum_{a \in \overline{V}_{1,z}} (\Pr[Y'_z = a] + \Pr[Y''_z = a]) \leq 2 \cdot \nu'(\lambda)} \\ &\leq 3 \cdot \nu'(\lambda) \end{aligned}$$

Hence, the (joint) distributions (X, Z) and (Y'', Z) are computationally indistinguishable. Therefore, $H^{\text{HILL}}(X | Z) \geq n$. $\square$

We recall that in the proof of Theorem 4, we used the observation that independent random variables A and B satisfy $H_\infty(A) = H_\infty(A | B)$. However, this does not necessarily extend to the computational case, meaning that in our case $H^{\text{HILL}}(X | \text{Setup}_X(1^\lambda), Z)$ and $H^{\text{HILL}}(X | Z)$ are not equal even though the random variables X and Setup_X are independent. Hence, in order to extend Theorem 6 to $\text{cPREH}_{\approx_c}^{\text{det}}$, we need that there exists a “good” CRS crs such that $Y'_{\text{crs},z}$ is almost injective on $\{0, 1\}^n$ and such that $\{(\text{crs}, y)\}$ and $\{(\text{crs}, D_X(\text{crs}, z, u))\}$ are computationally indistinguishable for this fixed crs . However, since the adversary is quantified after the crs , a non-uniform adversary could know the randomness which was used to generate crs compromising all security guarantees. Therefore, we can only hope to obtain an upper bound on n depending on the conditional HILL entropy $H^{\text{HILL}}(X | \text{Setup}_X(1^\lambda), Z)$.

Theorem 7. Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (strong) $\text{cPREH}_{\approx_c}^{\text{det}}$ is true for X , i.e. there exists a PPT algorithm Setup_X and two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{HILL}}(X \mid \text{Setup}_X(1^\lambda), Z) \geq n$.

Proof. The proof is similar to the proof of Theorem 6. We consider the distribution $Y'_{\text{crs},z} := D_X(\text{crs}, z, U_n)$.

Due to Theorem 2, we have that for all adversarially chosen z , the distributions $\{\text{crs} \leftarrow \text{Setup}_X(1^\lambda), y \leftarrow X_z : (\text{crs}, z, y)\}$ and $\{\text{crs} \leftarrow \text{Setup}_X(1^\lambda), u \leftarrow U_n : (\text{crs}, z, D_X(\text{crs}, z, u))\}$ are computationally indistinguishable.

Due to correctness and pseudorandomness we have that there exists a negligible function ν' , such that for all adversarially chosen z ,

$$\Pr[\text{crs} \leftarrow \text{Setup}_X(1^\lambda), u \leftarrow U_n : E_X(\text{crs}, z, D_X(\text{crs}, z, u)) = u] \geq 1 - \nu'(\lambda),$$

where the probability is over the choice of crs and u . The Splitting lemma (Lemma 9) implies that for all z ¹¹, there exists $G_z \subseteq \text{supp}(\text{Setup}_X(1^\lambda))$ such that $\Pr[\text{crs} \leftarrow \text{Setup}_X(1^\lambda) : \text{crs} \in G_z] \geq 1 - \nu'_z(\lambda)$ and for all $\text{crs} \in G_z$, $\Pr[u \leftarrow U_n : E_X(\text{crs}, z, D_X(\text{crs}, z, u)) = u] \geq 1 - \nu''_z(\lambda)$ for some negligible functions ν''_z, ν'''_z . Hence, conditioned on $\text{crs} \in G_z$, $D_X(\text{crs}, z, \cdot)$ operates almost injectively on the set $\{0, 1\}^n$.

Let $V_{0,z}^{\text{crs}} \subseteq \{0, 1\}^n$ denote the set of all $u \in \{0, 1\}^n$ such that $E_X(\text{crs}, z, D_X(\text{crs}, z, u)) = u$ holds and let $V_{1,z}^{\text{crs}} := \{0, 1\}^n \setminus V_{0,z}^{\text{crs}}$. For $\text{crs} \in G_z$, we have $|V_{1,z}^{\text{crs}}|/|\{0, 1\}^n| \leq \nu'''_z(\lambda)$.

Let $\bar{V}_{1,z}^{\text{crs}}$ be some arbitrary subset of $\{0, 1\}^{p(\lambda)} \setminus D_X(\text{crs}, z, V_{0,z}^{\text{crs}})$ such that $|\bar{V}_{1,z}^{\text{crs}}| = |V_{1,z}^{\text{crs}}|$ (for some polynomial p). Let $Y''_{\text{crs},z}$ be the uniform distribution over $D_X(\text{crs}, z, V_{0,z}^{\text{crs}}) \cup \bar{V}_{1,z}^{\text{crs}}$.

$$\tilde{H}_\infty(Y'' \mid \text{Setup}_X(1^\lambda), Z) = -\log \left(\mathbb{E}_{\text{crs} \leftarrow \text{Setup}_X(1^\lambda), z \leftarrow Z} \left[\max_{y \in \text{supp}(Y''_{\text{crs},z})} \Pr[Y''_{\text{crs},z} = y] \right] \right) = n.$$

Furthermore, a similar computation as in Theorem 4, for all adversarially chosen z , the statistical distance between $(Y'_z, \text{Setup}_X(1^\lambda))$ and $(Y''_z, \text{Setup}_X(1^\lambda))$ is negligible. Hence, the (joint) distributions $(X, \text{Setup}_X(1^\lambda), Z)$ and $(Y'', \text{Setup}_X(1^\lambda), Z)$ are computationally indistinguishable. Therefore, $H^{\text{HILL}}(X \mid \text{Setup}_X(1^\lambda), Z) \geq n$. $\square$

Yao entropy. Furthermore, the existence of algorithms (E_X, D_X) as implied by $\text{PREH}_{\approx_c}^{\text{det}}$ give an upper bound for the Yao entropy of a source.

Lemma 10. If (weak) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i.e. there exist deterministic algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{Yao}}(X) < n + \delta(\lambda)$ for some negligible δ .

Proof. Due to correctness, there exists a pair of efficient algorithms (E_X, D_X) (with E_X having output length n) such that $\Pr_{x \leftarrow X}[D_X(E_X(x)) = x] \geq 1 - \nu(\lambda)$ for some negligible function ν . Hence, by Definition 17, $H^{\text{Yao}}(X) < k$, for all k satisfying

$$\begin{aligned} 1 - \nu(\lambda) &\geq 2^{n-k} + \epsilon(\lambda) \\ \iff 2^k \cdot (1 - \nu(\lambda) - \epsilon(\lambda)) &\geq 2^n \\ \iff k &\geq n - \log(1 - \nu(\lambda) - \epsilon(\lambda)). \end{aligned}$$

$\square$

¹¹ Note that here we use that the adversary is non-uniform.

This result can be generalized to conditional Yao entropy as follows.

Lemma 11. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (strong) $\text{PREH}_{\approx_c}^{\text{det}}$ is true for X , i.e. there exist two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{Yao}}(X | Z) < n + \delta(\lambda)$ for some negligible δ .*

Proof. Due to correctness, for all adversarially chosen z , $\Pr_{x \leftarrow X_z}[D_X(z, E_X(z, x)) = x] \geq 1 - \nu(\lambda)$ for some negligible function ν . Hence, by Definition 19, $H^{\text{Yao}}(X | Z) < k$, for all k satisfying $k \geq n - \log(1 - \nu(\lambda) - \epsilon(\lambda))$. $\square$

Lemma 12. *Let (X, Z) be a joint distribution. More precisely, let Z be a distribution over words of length λ . For $z \in \text{supp}(Z)$, let X_z denote the conditional distribution when $Z = z$. If (strong) $\text{cPREH}_{\approx_c}^{\text{det}}$ is true for X , i.e. there exists a setup algorithm **Setup** and two deterministic polynomial time algorithms (E_X, D_X) with E_X having output length n satisfying correctness and pseudorandomness. Then, $H^{\text{Yao}}(X | Z) < n + \delta(\lambda)$ for some negligible δ .*

Proof. The Splitting lemma (Lemma 9) implies that for all z ,¹² there exists $G_z \subseteq \text{supp}(\text{Setup}_X(1^\lambda))$ such that $\Pr[\text{crs} \leftarrow \text{Setup}_X(1^\lambda) : \text{crs} \in G_z] \geq 1 - \nu_z''(\lambda)$ and for all $\text{crs} \in G_z$, $\Pr[u \leftarrow U_n : E_X(\text{crs}, z, D_X(\text{crs}, z, u)) = u] \geq 1 - \nu_z'''(\lambda)$ for some negligible functions ν_z'', ν_z''' .

We exploit the non-uniformity of the definition. In particular, we define the compression and decompression algorithms $E'_X(y, z; \text{aux}_\lambda) := E_X(\text{aux}_\lambda, z, y)$ and $D'_X(u, z; \text{aux}_\lambda) := D_X(\text{aux}_\lambda, z, u)$, where aux_λ denotes the non-uniform auxiliary input. $\square$

Due to [HLR07], assuming suitable non-interactive zero-knowledge proof systems and pseudorandom generators, there exists a joint distribution (X, Z) with high conditional Yao entropy but low conditional HILL entropy. The sampler S from [HLR07] takes as input a NIZK common random string σ , samples a seed from U_λ and outputs $y_1 := \text{PRG}(s)$ together with a proof y_2 that y_1 is in the image of PRG . Due to Lemma 2, $S \in \mathcal{S}^{\text{comp}}$. This yields the following corollary.

Corollary 7. *If there exist a pseudorandom generator and a (single-theorem) NIZK proof system such that (i) for an overwhelming fraction of common random strings, the number of accepting proofs for each statement are limited, and (ii) the simulated random string is independent of the statement as in [HLR07]. Then $\text{PREH}_{\approx_c}^{\text{det}}$ is false.*

Together with the results of [HLR07; LMs05], we obtain the following corollary.

Corollary 8. *A Blum integer is a natural number $N = p \cdot q$ such that p, q are primes with $p \equiv q \equiv 3 \pmod{4}$. Assume that for a randomly chosen Blum integer $N = p \cdot q$, the distributions $\{y \leftarrow \mathbb{Z}_N^\times \text{ s.t. } (\frac{y}{p}) = (\frac{y}{q}) = 1 : (N, y)\}$ and $\{y \leftarrow \mathbb{Z}_N^\times \text{ s.t. } (\frac{y}{N}) = 1 : (N, y)\}$ are computationally indistinguishable. Then, $\text{PREH}_{\approx_c}^{\text{det}}$ is false.*

On refuting $\text{cPREH}_{\approx_c}^{\text{det}}$. Theorem 7 only yields an upper bound on the encoding length n depending on $H^{\text{HILL}}(X | \text{Setup}_X(1^\lambda), Z)$. Since the Yao entropy of any source $H^{\text{Yao}}(X | \text{Setup}_X(1^\lambda), Z)$ is upper bounded by $n + \text{negl}(\lambda)$, we can not apply the result from [HLR07] to refute $\text{cPREH}_{\approx_c}^{\text{det}}$.

On refuting weak $\text{PREH}_{\approx_c}^{\text{det}}$. It is currently not known if plain HILL and plain Yao entropy can be separated as in [HLR07]. Such a separation would refute weak $\text{PREH}_{\approx_c}^{\text{det}}$.

¹² Note that here we use that the adversary is non-uniform.

6.2 Randomized encoding

In the following we prove positive and negative results on the validity of the pseudorandom encoding hypothesis with a randomized encoding algorithm. In particular, in Section 6.2.2 we refute $\text{PREH}_{\approx_s}^{\text{rand}}$ based on sub-exponential LWE, in Section 6.2.3 we refute $\text{cPREH}_{\approx_c}^{\text{rand}}$ based on extractable one-way functions with *unbounded* auxiliary input. On the positive side, in Section 6.3, give a construction of perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup based on indistinguishability obfuscation and one-way functions following [SW14; DKR15] together with Theorem 2. In Section 6.4, we bootstrap $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common random string from the construction in Section 6.3 additionally assuming *weak* $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common random string. Since $\text{cPREH}_{\approx_c}^{\text{rand}}$ with a common random string in conjunction with NIZK proof systems contradicts EOWFs with common but benign auxiliary information, this refutes even *weak* $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string.

6.2.1 (Generalized) extractable one-way functions

In this section, we define (generalized) extractable one-way functions (with respect to common auxiliary input) as in [BCPR14].

Definition 22 (Function family ensemble). A function family, indexed by a key space K_λ , is a set of functions $f^\lambda = \{f_k\}_{k \in K_\lambda}$ where each function has the same domain and range. A function family ensemble $\mathcal{F} = \{f^\lambda\}_{\lambda \in \mathbb{N}}$ is an ensemble of function families with key spaces $\{K_\lambda\}_{\lambda \in \mathbb{N}}$.

Definition 23 (Extractable one-way function family ensembles (EOWFs) without auxiliary information, [BCPR14]). A function family ensemble is called a extractable one-way function family ensemble without auxiliary information if the following two properties are satisfied.

One-wayness. For every PPT adversary $\mathcal{A}$, $\Pr[\text{Exp}_{\mathcal{A}}^{\text{ow}}(\lambda) = 1]$ is negligible, where $\text{Exp}_{\mathcal{A}}^{\text{ow}}$ is defined in Figure 17.

Extractability. For every PPT adversary X (using a random tape of length $m(\lambda)$), there exists a PPT algorithm K_X such that $\Pr[\text{Exp}_{X, K_X}^{\text{ext}}(\lambda) = 1]$ is overwhelming, where $\text{Exp}_{X, K_X}^{\text{ext}}$ is defined in Figure 17.

$\text{Exp}_{\mathcal{A}}^{\text{ow}}(\lambda)$	$\text{Exp}_{X, K_X}^{\text{ext}}(\lambda)$
$k \leftarrow K_\lambda$	$k \leftarrow K_\lambda$
$x \leftarrow \text{domain}(f_k)$	$r_X \leftarrow \{0, 1\}^{m(\lambda)}$
$y := f_k(x)$	$y = X(k; r_X)$
$x' \leftarrow \mathcal{A}(k, y)$	$x \leftarrow K_X(k, r_X)$
return $f_k(x') = y$	return $(f_k(x) = y) \vee (\forall x': f_k(x') \neq y)$

Fig. 17. One-way and extraction game.

Definition 24 (Extractable one-way function family ensembles (EOWFs) with common auxiliary information, [BCPR14]). A function family ensemble is called a one-way extractable function family ensemble with common auxiliary information if the following properties are satisfied.

One-wayness. As in Definition 23.

One-wayness (stronger). For every PPT adversary $\mathcal{A}$, for every polynomial b and for every $z \in \{0, 1\}^{b(\lambda)}$, $\Pr[\text{Exp}_{\mathcal{A}, z}^{\text{ow-aux}}(\lambda) = 1]$ is negligible, where $\text{Exp}_{\mathcal{A}, z}^{\text{ow-aux}}$ is defined in Figure 18.

Extractability. For every PPT adversary X , there exists a PPT algorithm K_X such that for every polynomial b and every $z \in \{0, 1\}^{b(\lambda)}$, $\Pr[\text{Exp}_{X, K_X, z}^{\text{ext-aux}}(\lambda) = 1]$ is overwhelming, where $\text{Exp}_{X, K_X, z}^{\text{ext-aux}}$ is defined in Figure 18.

$Exp_{A,z}^{\text{ow-aux}}(\lambda)$	$Exp_{X,K_X,z}^{\text{ext-aux}}(\lambda)$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$x \leftarrow \text{domain}(f_k)$	$r_X \leftarrow \{0,1\}^{m(\lambda)}$
$y := f_k(x)$	$y = X(k, z; r_X)$
$x' \leftarrow \mathcal{A}(k, y, z)$	$x \leftarrow K_X(k, z, r_X)$
return $f_k(x') = y$	return $(f_k(x) = y) \vee (\forall x': f_k(x') \neq y)$

Fig. 18. One-way and extraction game with common auxiliary input.

We note that the weaker notion of one-wayness without auxiliary input is sufficient for us. Assuming non-uniform adversaries, one-wayness without auxiliary input and one-wayness with auxiliary input are equivalent.

Definition 25 (Extractable one-way function family ensembles (EOWFs) with b -bounded common auxiliary information, [BCPR14]). Like Definition 24 but with a fixed polynomial b determining the length of the common auxiliary information.

Definition 26 (Generalized extractable one-way function family ensembles (GEOWFs) with common auxiliary information, [BCPR14]). A function family ensemble $\mathcal{F}$ is called a generalized one-way extractable function family ensemble with common auxiliary information, with respect to a relation $R_k^{\mathcal{F}}$ over triplets $(k, y, x) \in \mathcal{K}_\lambda \times \{0,1\}^{n(\lambda)} \times \{0,1\}^{m(\lambda)}$, if the following properties are satisfied.

$R^{\mathcal{F}}$ -Hardness. For every PPT adversary $\mathcal{A}$, for every polynomial b and every $z \in \{0,1\}^{b(\lambda)}$, $\Pr[Exp_{A,z}^{\text{g-hard-aux}}(\lambda) = 1]$ is negligible, where $Exp_{A,z}^{\text{g-hard-aux}}$ is defined in Figure 19.

$R^{\mathcal{F}}$ -Extractability. For every PPT adversary X , there exists a PPT algorithm K_X such that for every polynomial b and every $z \in \{0,1\}^{b(\lambda)}$, $\Pr[Exp_{X,K_X,z}^{\text{g-ext-aux}}(\lambda) = 1]$ is overwhelming, where $Exp_{X,K_X,z}^{\text{g-ext-aux}}$ is defined in Figure 19.

We call $\mathcal{F}$

- publicly verifiable if there exists a PPT algorithm T such that

$$T(k, f_k(x), x') = 1 \Leftrightarrow R_k^{\mathcal{F}}(f_k(x), x') = 1.$$

- privately verifiable if there exists a PPT algorithm T such that

$$T(k, x, x') = 1 \Leftrightarrow R_k^{\mathcal{F}}(f_k(x), x') = 1.$$

$Exp_{A,z}^{\text{g-hard-aux}}(\lambda)$	$Exp_{X,K_X,z}^{\text{g-ext-aux}}(\lambda)$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$x \leftarrow \text{domain}(f_k)$	$r_X \leftarrow \{0,1\}^{m(\lambda)}$
$y := f_k(x)$	$y = X(k, z; r_X)$
$x' \leftarrow \mathcal{A}(k, y, z)$	$x \leftarrow K_X(k, z, r_X)$
return $R_k^{\mathcal{F}}(y, x')$	return $(R_k^{\mathcal{F}}(y, x)) \vee (\forall x': f_k(x') \neq y)$

Fig. 19. $R^{\mathcal{F}}$ -hardness and $R^{\mathcal{F}}$ -extraction game with common auxiliary input.

Definition 27 (Generalized extractable one-way function family ensembles (GEOWFs) with b -bounded common auxiliary information, [BCPR14]). Like Definition 26 but with a fixed polynomial b determining the length of the common auxiliary information.

Privately verifiable generalized extractable one-way functions with bounded auxiliary input can be instantiated based on falsifiable assumptions due to [BCPR14].

Theorem 8 ([BCPR14]). Assuming the learning with errors problem is sub-exponentially hard, there exists a $(b(\lambda) - \omega(1))$ -bounded privately verifiable GEOWF family ensemble.

6.2.2 Information theoretic guarantees and privately verifiable GEOWFs

Assuming information theoretic indistinguishability, adversaries are unbounded and, hence, private verifiability is not a restriction.

Theorem 9. *If privately verifiable generalized extractable one-way function family ensembles without auxiliary information exist, then $\text{PREH}_{\equiv_s}^{\text{rand}}$ is false.*

Proof. The proof strategy follows the ideas of [IKOS10]. Let $\mathcal{F}$ be a privately verifiable GEOWF family ensemble (without auxiliary input) with respect to relation $R^{\mathcal{F}}$.

$\text{PREH}_{\equiv_s}^{\text{rand}}$ implies that for the algorithm S (given in Figure 20) there exists an alternative sampler $\bar{S}$ and a corresponding inverse sampler $\bar{S}^{-1}$ satisfying closeness and invertibility of Definition 7 against unbounded adversaries. Since $\mathcal{F}$ is a GEOWF family ensemble, for the algorithm $\bar{S}$, there exists an extractor $K_{\bar{S}}$ satisfying extractability from Definition 26 (without auxiliary information).

$S(k)$	$\mathcal{A}(1^\lambda, k, y)$
$x \leftarrow \text{domain}(f_k)$	$r'_S \leftarrow \bar{S}^{-1}(k, y)$
$y := f_k(x)$	$x' \leftarrow K_{\bar{S}}(k, r'_S)$
return y	return x'

Fig. 20. Description of the sampler S and of the adversary $\mathcal{A}$ on one-wayness of the privately verifiable GEOWF.

We prove that for $\mathcal{A}$ given in Figure 20, $\Pr[\text{Exp}_{\mathcal{A}}^{\text{g-hard}}(\lambda) = 1]$ is overwhelming. Let $\tilde{T}$ be an *unbounded* algorithm that given (k, y, x) , computes the relation $R_k^{\mathcal{F}}(y, x)$.¹³ We proceed over a series of hybrids, see Figure 21.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$x \leftarrow \text{domain}(f_k)$	$x \leftarrow \text{domain}(f_k)$	$y \leftarrow S(k)$
$y := f_k(x)$	$y := f_k(x)$	$r'_S \leftarrow \bar{S}^{-1}(k, y)$
$x' \leftarrow \mathcal{A}(k, y)$	$r'_S \leftarrow \bar{S}^{-1}(k, y)$	$x' \leftarrow K_{\bar{S}}(k, r'_S)$
return $\tilde{T}(k, y, x')$	$x' \leftarrow K_{\bar{S}}(k, r'_S)$	return $\tilde{T}(k, y, x')$
	return $\tilde{T}(k, y, x')$	
$\mathbf{G}_3$	$\mathbf{G}_4$	$\mathbf{G}_5$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$y \leftarrow \bar{S}(k)$	$r_{\bar{S}} \leftarrow \{0, 1\}^*$	$r_{\bar{S}} \leftarrow \{0, 1\}^*$
$r'_S \leftarrow \bar{S}^{-1}(k, y)$	$y \leftarrow \bar{S}(k; r_{\bar{S}})$	$y \leftarrow \bar{S}(k; r_{\bar{S}})$
$x' \leftarrow K_{\bar{S}}(k, r'_S)$	$r'_S \leftarrow \bar{S}^{-1}(k, y)$	$x' \leftarrow K_{\bar{S}}(k, r_{\bar{S}})$
return $\tilde{T}(k, y, x')$	$x' \leftarrow K_{\bar{S}}(k, r'_S)$	return $\tilde{T}(k, y, x')$
	return $\tilde{T}(k, y, x')$	return $\tilde{T}(k, y, x')$

Fig. 21. Hybrids used in the proof of Theorem 9.

The individual game hops are justified as follows. The difference between games $\mathbf{G}_0$ and $\mathbf{G}_1$, and games $\mathbf{G}_1$ and $\mathbf{G}_2$ are both only conceptual, hence, $\Pr[\text{out}_0 = 1] = \Pr[\text{out}_1 = 1] = \Pr[\text{out}_2 = 1]$.

Claim. There exists an unbounded adversary $\bar{\mathcal{A}}$ such that $|\Pr[\text{out}_3 = 1] - \Pr[\text{out}_2 = 1]| \leq \text{Adv}_{\bar{\mathcal{A}}, k}^{\text{close}}(\lambda)$ (for some $\bar{k} \in \mathcal{K}_\lambda$).

Proof. Construct an adversary $\bar{\mathcal{A}}$ on statistical closeness. $\bar{\mathcal{A}}$ receives as input a key $\bar{k}$ and some y that has either been sampled via $S(\bar{k})$ or via $\bar{S}(\bar{k})$. $\bar{\mathcal{A}}$ computes $r'_{\bar{S}}$ and x' as in game $\mathbf{G}_2$, calls

¹³ Since $\mathcal{F}$ is privately verifiable, an efficient algorithm computing the relation $R_k^{\mathcal{F}}$ additionally requires the preimage of y as an input. For our purpose, it suffices to consider an inefficient testing algorithm $\tilde{T}$.

the (inefficient) testing algorithm $\tilde{T}$ on input of $(\bar{k}, y, x')$ and outputs the resulting output. Hence, $\Pr[\text{Exp}_{\bar{A},k,0}^{\text{close}}(\lambda) = 1] = \Pr[\text{out}_2 = 1 \mid k = \bar{k}]$ and $\Pr[\text{Exp}_{\bar{A},k,1}^{\text{close}}(\lambda) = 1] = \Pr[\text{out}_3 = 1 \mid k = \bar{k}]$. Since $|\Pr[\text{out}_3 = 1] - \Pr[\text{out}_2 = 1]| = |\sum_{\bar{k} \in \mathcal{K}_\lambda} \Pr_{k \leftarrow \mathcal{K}_\lambda}[k = \bar{k}] \cdot (\Pr[\text{out}_3 = 1 \mid k = \bar{k}] - \Pr[\text{out}_2 = 1 \mid k = \bar{k}])| = |\sum_{\bar{k} \in \mathcal{K}_\lambda} \Pr_{k \leftarrow \mathcal{K}_\lambda}[k = \bar{k}] \cdot \text{Adv}_{\bar{A},k}^{\text{close}}(\lambda)|$, $|\Pr[\text{out}_3 = 1] - \Pr[\text{out}_2 = 1]|$ is negligible. $\square$

The difference between the games $\mathbf{G}_3$ and $\mathbf{G}_4$ is again only conceptual and $\Pr[\text{out}_3 = 1] = \Pr[\text{out}_4 = 1]$.

Claim. There exists an unbounded adversary $\bar{\mathcal{A}}$ such that $|\Pr[\text{out}_5 = 1] - \Pr[\text{out}_4 = 1]| \leq \text{Adv}_{\bar{\mathcal{A}},k}^{\text{inv}}(\lambda)$ (for some $\bar{k} \in \mathcal{K}_\lambda$).

Proof. Construct an unbounded adversary $\bar{\mathcal{A}}$ on statistical invertibility. On input of $(\bar{k}, r, y)$, $\bar{\mathcal{A}}$ calls $K_{\bar{S}}$ on input of $(\bar{k}, r)$ and obtains x' . Finally, $\bar{\mathcal{A}}$ calls the inefficient testing algorithm $\tilde{T}$ on input of $(\bar{k}, y, x')$ and outputs the resulting output. If $\bar{\mathcal{A}}$ plays the experiment $\text{Exp}_{\bar{\mathcal{A}},k,0}^{\text{inv}}$, r is the randomness used to produce y as $\bar{S}(\bar{k}; r)$. Hence, in this case, $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_5$ and $\Pr[\text{Exp}_{\bar{\mathcal{A}},k,0}^{\text{inv}}(\lambda) = 1] = \Pr[\text{out}_5 = 1 \mid k = \bar{k}]$. If, on the other hand, $\bar{\mathcal{A}}$ plays the experiment $\text{Exp}_{\bar{\mathcal{A}},k,1}^{\text{inv}}$, r is some inverse sampled randomness (i.e. the output of $\bar{S}^{-1}(\bar{k}, y)$). Hence, in this case, $\bar{\mathcal{A}}$ perfectly simulates game $\mathbf{G}_4$ and $\Pr[\text{Exp}_{\bar{\mathcal{A}},k,1}^{\text{inv}}(\lambda) = 1] = \Pr[\text{out}_4 = 1 \mid k = \bar{k}]$. Thus, by the same computations made in the reduction to closeness, the claim follows. $\square$

Summing up, we have that $|\Pr[\text{out}_5 = 1] - \Pr[\text{out}_0 = 1]|$ is negligible.

Lemma 13. $\Pr[\text{out}_5 = 1]$ is overwhelming.

Proof (of Lemma 13). Since $\mathcal{F}$ is a GEOWF family ensemble, we have that $\Pr_{\mathbf{G}_5}[R_k^{\mathcal{F}}(y, x') = 1 \vee y \notin \text{image}(f_k)]$ is overwhelming. Using a union bound, we get $\Pr_{\mathbf{G}_5}[R_k^{\mathcal{F}}(y, x') = 1 \vee y \notin \text{image}(f_k)] \leq \Pr_{\mathbf{G}_5}[R_k^{\mathcal{F}}(y, x') = 1] + \Pr_{\mathbf{G}_5}[y \notin \text{image}(f_k)]$.

In the following, we prove that $\Pr_{\mathbf{G}_5}[y \notin \text{image}(f_k)]$ is negligible. By construction, S on input k always produces values $y \in \text{image}(f_k)$. Intuitively, if $\Pr_{\mathbf{G}_5}[y \notin \text{image}(f_k)] = \Pr[k \leftarrow \mathcal{K}_\lambda, y \leftarrow \bar{S}(k) : y \notin \text{image}(f_k)]$ is non-negligible, an unbounded adversary can distinguish between outputs of $S(k)$ and outputs of $\bar{S}(k)$ simply by testing all possible preimages. Let $\bar{\mathcal{A}}$ be the adversary on closeness which outputs 1 if and only if $y \notin \text{image}(f_k)$. Hence, $\Pr_{\mathbf{G}_5}[y \notin \text{image}(f_k)] = \sum_{\bar{k} \in \mathcal{K}_\lambda} \Pr[k = \bar{k}] \cdot \text{Adv}_{\bar{\mathcal{A}},k}^{\text{close}}(\lambda)$. Since for all $k \in \mathcal{K}_\lambda$, $\text{Adv}_{\bar{\mathcal{A}},k}^{\text{close}}(\lambda)$ is negligible, $\Pr_{\mathbf{G}_5}[y \notin \text{image}(f_k)]$ is negligible. $\square$

From Theorem 9 together with Theorem 8, we obtain the following corollary.

Corollary 9. Assuming the learning with errors problem is sub-exponentially hard, $\text{PREH}_{\equiv_s}^{\text{rand}}$ is false.

Theorem 10. If privately verifiable generalized extractable one-way function family ensembles with unbounded auxiliary information exist, then $\text{cPREH}_{\equiv_s}^{\text{rand}}$ is false.

The proof is an easy modification of the proof of Theorem 9.

6.2.3 Computational guarantees and EOWFs with common auxiliary information

Assuming computational indistinguishability, the above proof strategy must be adapted since the adversary actually uses the alternative sampler of a sampler which outputs an image in the range of the EOWF. Computational indistinguishability does not suffice to force the alternative sampler to output an image in the range of the EOWF. Hence, for the purpose to force the alternative sampler to output an image of a given extractable one-way function, we follow the lines of [IKOS10] and require the original sampler to additionally provide a non-interactive zero-knowledge proof certifying that the provided image is in the range of the given EOWF.

Definition 28 (Non-interactive zero-knowledge proof system). A non-interactive zero-knowledge (NIZK) proof system for the $\mathcal{NP}$ -language L with witness relation R_L is a tuple of PPT algorithms (P, V) satisfying the following conditions for all sufficiently large λ (for some polynomial ℓ).

Correctness. For all $(x, w) \in R_L$, for all $\sigma \in \{0, 1\}^{\ell(\lambda)}$, $\Pr[V(\sigma, x, P(\sigma, x, w)) = 1] = 1$.

Statistical soundness. For all unbounded adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{soundness}}(\lambda) := \Pr[\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}, (x, \pi) \leftarrow \mathcal{A}(\sigma): V(\sigma, x, \pi) = 1 \wedge x \notin L]$$

is negligible.

Computational zero-knowledge. There exists a tuple of PPT algorithms $\text{Sim} = (\text{Sim}_1, \text{Sim}_2)$ such that for all non-uniform PPT algorithms $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{zk}}(\lambda) := |\Pr[\text{Exp}_{\mathcal{A},0}^{\text{zk}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},1}^{\text{zk}}(\lambda) = 1]|$$

is negligible, where $\text{Exp}_{\mathcal{A},0}^{\text{zk}}$ and $\text{Exp}_{\mathcal{A},1}^{\text{zk}}$ are defined in Figure 22.

$\text{Exp}_{\mathcal{A},0}^{\text{zk}}(\lambda)$	$\text{Exp}_{\mathcal{A},1}^{\text{zk}}(\lambda)$
$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$(\sigma, \tau) \leftarrow \text{Sim}_1(1^\lambda)$
$(x, w) \leftarrow \mathcal{A}(\sigma)$ s.t. $(x, w) \in R_L$	$(x, w) \leftarrow \mathcal{A}(\sigma)$ s.t. $(x, w) \in R_L$
$\pi \leftarrow P(\sigma, x, w)$	$\pi \leftarrow \text{Sim}_2(\tau, x)$
return $\mathcal{A}(\pi)$	return $\mathcal{A}(\pi)$

Fig. 22. Zero-knowledge games.

As already observed in [IKOS10], $\text{PREH}_{\approx_c}^{\text{rand}}$ in conjunction with NIZK proof systems conflicts with EOWFs.

Theorem 11 ([IKOS10]). If extractable one-way function family ensembles without auxiliary information and NIZK proof systems for $\mathcal{NP}$ exist, then $\text{PREH}_{\approx_c}^{\text{rand}}$ is false.

This result extends to $\text{cPREH}_{\approx_c}^{\text{rand}}$ at the cost of assuming unbounded common auxiliary inputs.

Theorem 12. If there exist extractable one-way function family ensembles with unbounded common auxiliary information and NIZK proof systems for $\mathcal{NP}$, then $\text{cPREH}_{\approx_c}^{\text{rand}}$ is false.

Proof. We adapt the proof from [IKOS10]. Let $\mathcal{F}$ be an extractable one-way function family ensemble with common auxiliary information. Let $L_\lambda := \{(k, y) \in \mathcal{K}_\lambda \times \{0, 1\}^{m(\lambda)} \mid \exists x \in \text{domain}(f_k): f_k(x) = y\}$. Let (P, V) be a NIZK proof system for L_λ .

$\text{cPREH}_{\approx_c}^{\text{rand}}$ implies that for the PPT algorithm S (see Figure 23), there exists a PPT algorithm Setup_S , an alternative sampler $\bar{S}$ and a corresponding inverse sampler $\bar{S}^{-1}$ satisfying closeness and invertibility as in Definition 9.

The common auxiliary input is necessary to give the adversary X and the extractor K_X access to the same common reference string. Further, it is crucial to assume *unbounded* auxiliary input since the adversary X can be considered to be a universal adversary which simply executes the code which is contained in the auxiliary input. Since the size of the adversary can not be bounded in advance, neither can the size of the auxiliary input.

Since $\mathcal{F}$ is an extractable one-way function family ensemble with common auxiliary information, we have that for the algorithm X , there exists an extractor K_X such that for every polynomial b and every $\text{crs} \in \{0, 1\}^{b(\lambda)}$ (hence, in particular, for every $z := \text{crs}$ produced by $\text{Setup}_S(1^\lambda)$), $\Pr[\text{Exp}_{X, K_X, z}^{\text{ext-aux}}(\lambda) = 1]$ is overwhelming.

We prove that adversary $\mathcal{A}$ given in Figure 23 has an overwhelming probability to break the one-wayness of $\mathcal{F}$. We proceed over a sequence of hybrids, see Figures 24 and 25.

The individual game hops are justified as follows. The difference between the games $\mathbf{G}_0$ and $\mathbf{G}_1$ is only conceptual, hence, $\Pr[\text{out}_0 = 1] = \Pr[\text{out}_1 = 1]$.

$S(k, \sigma)$	$X(k, z =: crs; \sigma \parallel r_{\bar{S}})$	$\mathcal{A}(1^\lambda, k, y)$
$x \leftarrow \text{domain}(f_k)$	$(y, \pi) \leftarrow \bar{S}(crs, (k, \sigma); r_{\bar{S}})$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$y := f_k(x)$	return y	$(\sigma, \tau) \leftarrow \text{Sim}_1(1^\lambda)$
$\pi \leftarrow P(\sigma, (k, y), x)$		$\pi \leftarrow \text{Sim}_2(\tau, (k, y))$
return (y, π)		$r'_{\bar{S}} \leftarrow \bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$
		$r'_X := \sigma \parallel r'_{\bar{S}}$
		$x' \leftarrow K_X(k, crs, r'_X)$
		return x'

Fig. 23. Description of the sampler S , of the adversary X on extractability of the EOWF and of the adversary $\mathcal{A}$ on one-wayness of the EOWF. Note that depending on definition of one-wayness, the CRS crs could also be auxiliary input to $\mathcal{A}$.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$	$\mathbf{G}_3$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$x \leftarrow \text{domain}(f_k)$	$x \leftarrow \text{domain}(f_k)$	$x \leftarrow \text{domain}(f_k)$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$
$y := f_k(x)$	$y := f_k(x)$	$y := f_k(x)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$x' \leftarrow \mathcal{A}(k, y)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$(y, \pi) \leftarrow S(k, \sigma)$
return $f_k(x') = y$	$(\sigma, \tau) \leftarrow \text{Sim}_1(1^\lambda)$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$r'_{\bar{S}} \leftarrow \bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$
	$\pi \leftarrow \text{Sim}_2(\tau, (k, y))$	$\pi \leftarrow P(crs, (k, y), x)$	$r'_X := \sigma \parallel r'_{\bar{S}}$
	$r'_{\bar{S}} \leftarrow \bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$	$r'_{\bar{S}} \leftarrow \bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$	$x' \leftarrow K_X(k, crs, r'_X)$
	$r'_X := \sigma \parallel r'_{\bar{S}}$	$r'_X := \sigma \parallel r'_{\bar{S}}$	return $f_k(x') = y$
	$x' \leftarrow K_X(k, crs, r'_X)$	$x' \leftarrow K_X(k, crs, r'_X)$	
	return $f_k(x') = y$	return $f_k(x') = y$	

Fig. 24. Hybrids used in the proof of Theorem 12.

Claim. There exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{\bar{\mathcal{A}}}^{\text{zk}}(\lambda)$.

Proof. Construct adversary $\bar{\mathcal{A}}$. Initially, $\bar{\mathcal{A}}$ receives σ as input and produces k, x, y, crs as in game $\mathbf{G}_1$ and outputs the statement (k, y) together with a witness x to the experiment. In return, $\bar{\mathcal{A}}$ receives a proof π and proceeds as in $\mathbf{G}_1$ to produce $r'_{\bar{S}}, r'_X, x'$. Finally, $\bar{\mathcal{A}}$ outputs 1 if $f_k(x') = y$ and 0 otherwise. If $\bar{\mathcal{A}}$ plays experiment $\text{Exp}_{\bar{\mathcal{A}},1}^{\text{zk}}$, σ was sampled along with some trapdoor τ and the proof π is simulated, hence, $\Pr[out_1 = 1] = \Pr[\text{Exp}_{\bar{\mathcal{A}},1}^{\text{zk}}(\lambda) = 1]$. If, on the other hand, $\bar{\mathcal{A}}$ plays experiment $\text{Exp}_{\bar{\mathcal{A}},0}^{\text{zk}}$, σ is uniformly random from $\{0, 1\}^{\ell(\lambda)}$ and the proof π is produced by P using the given witness, hence, $\Pr[out_2 = 1] = \Pr[\text{Exp}_{\bar{\mathcal{A}},0}^{\text{zk}}(\lambda) = 1]$. $\square$

The difference between $\mathbf{G}_2$ and $\mathbf{G}_3$ is again only conceptual and $\Pr[out_2 = 1] = \Pr[out_3 = 1]$.

Claim. There exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_4 = 1] - \Pr[out_3 = 1]| \leq \text{Adv}_{\bar{\mathcal{A}}}^{\text{crs-close}}(\lambda)$.

$\mathbf{G}_4$	$\mathbf{G}_5$	$\mathbf{G}_6$	$\mathbf{G}_7$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$
$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$(y, \pi) \leftarrow \bar{S}(crs, (k, \sigma))$	$r_{\bar{S}} \leftarrow \{0, 1\}^*$	$r_{\bar{S}} \leftarrow \{0, 1\}^*$	$r_{\bar{S}} \leftarrow \{0, 1\}^*$
$r'_{\bar{S}} \leftarrow \bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$	$(y, \pi) \leftarrow \bar{S}(crs, (k, \sigma); r_{\bar{S}})$	$(y, \pi) \leftarrow \bar{S}(crs, (k, \sigma); r_{\bar{S}})$	$r'_X := \sigma \parallel r_{\bar{S}}$
$r'_X := \sigma \parallel r'_{\bar{S}}$	$r'_{\bar{S}} \leftarrow \bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$	$r'_X := \sigma \parallel r_{\bar{S}}$	$y \leftarrow X(k, crs; r'_X)$
$x' \leftarrow K_X(k, crs, r'_X)$	$r'_X := \sigma \parallel r'_{\bar{S}}$	$x' \leftarrow K_X(k, crs, r'_X)$	$x' \leftarrow K_X(k, crs, r'_X)$
return $f_k(x') = y$	$x' \leftarrow K_X(k, crs, r'_X)$	return $f_k(x') = y$	return $f_k(x') = y$
	return $f_k(x') = y$		

Fig. 25. Hybrids used in the proof of Theorem 12.

Proof. Construct a PPT adversary $\bar{\mathcal{A}}$ on (static) closeness. Initially, $\bar{\mathcal{A}}$ produces (k, σ) as in game $\mathbf{G}_3$ and outputs them to the experiment. In the second phase, $\bar{\mathcal{A}}$ receives crs and (y, π) , where (y, π) has either been sampled using $S(k, \sigma)$ or using $\bar{S}(crs, (k, \sigma))$. Further, $\bar{\mathcal{A}}$ proceeds as in $\mathbf{G}_3$ producing r'_S, r'_X, x' . Finally, $\bar{\mathcal{A}}$ outputs $f_k(x') = y$. Hence, $\Pr[out_3 = 1] = \Pr[Exp_{\bar{\mathcal{A}},0}^{crs-close}(\lambda) = 1]$ and $\Pr[out_4 = 1] = \Pr[Exp_{\bar{\mathcal{A}},1}^{crs-close}(\lambda) = 1]$. $\square$

The difference between games $\mathbf{G}_4$ and $\mathbf{G}_5$ is again only conceptual, hence, $\Pr[out_4 = 1] = \Pr[out_5 = 1]$.

Claim. There exists a PPT adversary $\bar{\mathcal{A}}$, such that $|\Pr[out_6 = 1] - \Pr[out_5 = 1]| \leq Adv_{\bar{\mathcal{A}}}^{crs-inv}(\lambda)$.

Proof. Construct a PPT adversary $\bar{\mathcal{A}}$ on (static) invertibility. Initially, $\bar{\mathcal{A}}$ (statically) produces (k, σ) as in $\mathbf{G}_5$ and outputs them to the experiment. In the second phase, $\bar{\mathcal{A}}$ receives $(crs, r^*, (y, \pi))$, where (y, π) has been sampled using $\bar{S}(crs, (k, \sigma))$ and r^* either is the actual randomness used for that or the inverse sampled randomness produced via $\bar{S}^{-1}(crs, (k, \sigma), (y, \pi))$. Afterwards, $\bar{\mathcal{A}}$ proceeds as in $\mathbf{G}_5$ producing $r'_X := \sigma \parallel r^*$ and x' . Finally, $\bar{\mathcal{A}}$ outputs $f_k(x') = y$. Hence, $\Pr[out_5 = 1] = \Pr[Exp_{\bar{\mathcal{A}},1}^{crs-inv}(\lambda) = 1]$ and $\Pr[out_6 = 1] = \Pr[Exp_{\bar{\mathcal{A}},0}^{crs-inv}(\lambda) = 1]$. $\square$

The difference between games $\mathbf{G}_6$ and $\mathbf{G}_7$ is again only conceptual, hence, $\Pr[out_6 = 1] = \Pr[out_7 = 1]$.

Thus, we have that $|\Pr[out_0 = 1] - \Pr[out_7 = 1]|$ is negligible.

Lemma 14. $\Pr[out_7 = 1]$ is overwhelming.

Proof (of Lemma 14). Since $\mathcal{F}$ is an extractable one-way function, we have that $\Pr_{\mathbf{G}_7}[f_k(x') = y \vee (k, y) \notin L_\lambda] \leq \Pr_{\mathbf{G}_7}[f_k(x') = y] + \Pr_{\mathbf{G}_7}[(k, y) \notin L_\lambda]$ is overwhelming. In the following, we prove that $\Pr_{\mathbf{G}_7}[(k, y) \notin L_\lambda]$ is negligible. We proceed over a series of hybrids, see Figures 26 and 27.

$\mathbf{H}_0$	$\mathbf{H}_1$	$\mathbf{H}_2$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$\sigma \parallel r_{\bar{S}} = r_X \leftarrow \{0, 1\}^{\ell(\lambda)}$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$
$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$y = X(k, crs; r_X)$	$(y, \pi) = \bar{S}(crs, (k, \sigma))$	$(y, \pi) = \bar{S}(crs, (k, \sigma))$
return $(k, y) \notin L_\lambda$	return $(k, y) \notin L_\lambda$	return $(k, y) \notin L_\lambda \wedge V(\sigma, (k, y), \pi) = 0$

Fig. 26. Hybrids used in the proof of Lemma 14.

$\mathbf{H}_3$	$\mathbf{H}_4$
$k \leftarrow \mathcal{K}_\lambda$	$k \leftarrow \mathcal{K}_\lambda$
$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$	$\sigma \leftarrow \{0, 1\}^{\ell(\lambda)}$
$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$(y, \pi) = \bar{S}(crs, (k, \sigma))$	$(y, \pi) = S((k, \sigma))$
return $V(\sigma, (k, y), \pi) = 0$	return $V(\sigma, (k, y), \pi) = 0$

Fig. 27. Hybrids used in the proof of Lemma 14.

The individual game hops are justified as follows. The difference between $\mathbf{H}_0$ and $\mathbf{H}_1$ is only conceptual and $\Pr[out_{\mathbf{H}_0} = 1] = \Pr[out_{\mathbf{H}_1} = 1]$.

Claim. There exists a PPT adversary $\bar{\mathcal{A}}$ such that $|\Pr[out_{\mathbf{H}_2} = 1] - \Pr[out_{\mathbf{H}_1} = 1]| \leq Adv_{\bar{\mathcal{A}}}^{\text{soundness}}(\lambda)$.

Proof. The only possibility that the output of $\mathbf{H}_1$ and $\mathbf{H}_2$ differ is the case that $(k, y) \notin L_\lambda$ but $V(\sigma, (k, y), \pi) = 1$ (assuming, that the output of V is binary). Hence, $|\Pr[out_{\mathbf{H}_2} = 1] - \Pr[out_{\mathbf{H}_1} = 1]| \leq \Pr_{\mathbf{H}_1}[(k, y) \notin L_\lambda \wedge V(\sigma, (k, y), \pi) = 1]$.

Construct adversary $\bar{\mathcal{A}}$ on soundness of the NIZK proof system. On input of σ , $\bar{\mathcal{A}}$ produces $k, crs, (y, \pi)$ as in $\mathbf{H}_1$ and outputs the statement (k, y) and the proof π . Thus, $\Pr_{\mathbf{H}_1}[(k, y) \notin L_\lambda \wedge V(\sigma, (k, y), \pi) = 1] \leq Adv_{\bar{\mathcal{A}}}^{\text{soundness}}(\lambda)$. $\square$

The distribution underlying $\mathbf{H}_2$ and $\mathbf{H}_3$ is identical. The only difference between these games is that the condition that $\mathbf{H}_3$ outputs 1 is less restrictive than the one of $\mathbf{H}_2$, hence, $\Pr[out_{\mathbf{H}_2} = 1] \leq \Pr[out_{\mathbf{H}_3} = 1]$.

In contrast to the preceding games, the simulation of game $\mathbf{H}_3$ is efficient.

Claim. There exists a PPT adversary $\bar{\mathcal{A}}$ such that $|\Pr[out_{\mathbf{H}_4} = 1] - \Pr[out_{\mathbf{H}_3} = 1]| \leq Adv_{\bar{\mathcal{A}}}^{\text{crs-close}}(\lambda)$.

Proof. Construct a PPT adversary $\bar{\mathcal{A}}$ on static closeness. Initially, $\bar{\mathcal{A}}$ produces (k, σ) as in $\mathbf{H}_3$ and outputs it to the experiment. On input of $(crs, \bar{y} := (y, \pi))$, where $\bar{y}$ is either sampled via $S(k, \sigma)$ or via $\bar{S}(crs, (k, \sigma))$. Finally, $\bar{\mathcal{A}}$ outputs $V(\sigma, (k, y), \pi) = 0$. We have that $\Pr[out_{\mathbf{H}_3} = 1] = \Pr[Exp_{\bar{\mathcal{A}},1}^{\text{crs-close}}(\lambda) = 1]$ and $\Pr[out_{\mathbf{H}_4} = 1] = \Pr[Exp_{\bar{\mathcal{A}},0}^{\text{crs-close}}(\lambda) = 1]$. $\square$

Finally, due to the definition of S and the correctness of the NIZK proof system (P, V) , $\Pr[out_{\mathbf{H}_4} = 1] = 0$. This proves that $\Pr[out_{\mathbf{H}_0} = 1]$ is negligible. $\square$

This concludes the proof. $\square$

By an easy modification of the above proof, we can refute $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup.

Theorem 13. *If there exist extractable one-way function family ensembles with unbounded common auxiliary information and NIZK proof systems for $\mathcal{NP}$, then $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup is false.*

We stress that assuming *unbounded* common auxiliary input is necessary since the adversary X can be considered the universal adversary. On a technical level, the common auxiliary input which leads to a contradiction is an output of the universal setup algorithm **Setup** on input of a sufficiently large bound B on the supported circuit size. This bound B depends on the adversary size.

Discussion. We stress that we are able to prove that EOWFs with unbounded common auxiliary input (in conjunction with NIZK proof systems) implies that $\text{cPREH}_{\approx_c}^{\text{rand}}$ is false. Furthermore, due to Theorems 15 and 16 in Section 6.3, $\text{cPREH}_{\approx_c}^{\text{rand}}$ can be instantiated from indistinguishability obfuscation and one-way functions. We restate a theorem from [BCPR14].

Theorem 14 ([BCPR14]). *Assuming indistinguishability obfuscation for all circuits, neither EOWFs nor GEOWFs exist with respect to unbounded common auxiliary information.*

Thus, the above is not a contradiction because due to Theorem 14, indistinguishability obfuscation for all polynomial sized circuits does not exist assuming EOWFs with unbounded common auxiliary input.

Common but benign auxiliary information. The definition due to [BCPR14] of (G)EOWFs with common auxiliary input requires that extractability holds for all common auxiliary inputs and hence also for a *worst-case* choice of common auxiliary input. This requirement can be weakened such that the common auxiliary input is drawn from some specific distribution. This distribution is called *benign* if it is unlikely that a common auxiliary input sampled according to this distribution encodes a malicious obfuscation¹⁴. In particular, the uniform distribution over $\{0, 1\}^{b(\lambda)}$ is conjectured to be benign. This notion of (G)EOWFs with common but benign auxiliary information does not contradict IO.

¹⁴ By malicious obfuscation we mean an obfuscated circuit which renders extraction from an adversary which simply executes the obfuscated circuit on the given key infeasible, see [BCPR14].

However, by an easy modification of the proof of Theorem 12, the existence (G)EOWFs with common but benign auxiliary input (drawn uniformly at random from $\{0,1\}^{b(\lambda)}$) contradicts $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common *random* string.

Corollary 10. *If there exist extractable one-way function family ensembles with common but benign auxiliary information, particularly for auxiliary inputs drawn uniformly at random from $\{0,1\}^{b(\lambda)}$, and NIZK proof systems for $\mathcal{NP}$, then $\text{cPREH}_{\approx_c}^{\text{rand}}$, where the setup algorithm produces uniform random strings, is false.*

Key-less extractable one-way function ensembles. Definitions 23, 25 and 27 can be defined for *key-less* function ensembles, where the key space $\mathcal{K}_\lambda$ contains exactly one element (for every λ). This poses much stronger requirements on one-wayness and extractability since these properties are bound to be met for one fixed key as opposed to the keyed variants.

As observed in [IKOS10], Theorems 9 and 11 can be adapted to the *weak* variants of $\text{PREH}_{\approx_c}^{\text{rand}}$ and $\text{PREH}_{\equiv_s}^{\text{rand}}$, respectively.¹⁵ However, key-less (G)EOWFs with unbounded auxiliary input are impossible since the adversary might get a random image of the (fixed) (G)EOWF as auxiliary input. An extractor given the output of the adversary together with its random tape and the same auxiliary input must break one-wayness in order to produce a pre-image. Hence, Theorem 12 can not be adapted to refute weak $\text{cPREH}_{\approx_c}^{\text{rand}}$.

6.3 Static pseudorandom encodings with universal setup from IO

$\text{cISH}_{\approx_c}^{\text{rand}}$ (and hence $\text{cPREH}_{\approx_c}^{\text{rand}}$) is implied by the existence of an explainability compiler [DKR15] which can be built from indistinguishability obfuscation and one-way functions. Using the ideas from [SW14; DKR15], we obtain perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup. Let S be a PPT algorithm and let $U_B(C, x; r)$ be the universal circuit that accepts any circuit C which can be represented with B bits and evaluates C on input x and randomness r .

<u>Setup($1^\lambda, B$)</u>	<u>$E_S(\text{crs}, m, y; r)$</u>	<u>$D_S(\text{crs}, m, u)$</u>
$\Lambda_E \leftarrow \text{iO}(C_E[k_2, k_2])$	return $\Lambda_E(S, m, y, r)$	return $\Lambda_D(S, m, u)$
$\Lambda_D \leftarrow \text{iO}(C_D[k_1, k_2, k_2])$		
return $\text{crs} := (\Lambda_E, \Lambda_D)$		
<u>$C_E[k_2, k_3](S, m, y, r)$</u>	<u>$C_D[k_1, k_2, k_3](S, m, u)$</u>	
$e_1 := F_2(S, m, y, \text{PRG}(r))$	$(S', m', y', r') := F_2(u_1) \oplus u_2$	
$e_2 := F_3(e_1) \oplus (S, m, y, \text{PRG}(r))$	if $((S', m') = (S, m) \wedge$	
return (e_1, e_2)	$u_1 = F_2(S', m', y', r'))$ then	
	return y'	
	$x := F_1(S, m, u)$	
	return $U_B(S, m; x)$	

Fig. 28. Instantiation of perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ due to [SW14; DKR15].

Let B be an upper bound on the bitlength which is necessary to describe a sampler. Let $\ell_{in} = |m|$ be an upper bound on the bitlength of the inputs, ℓ_{out} be an upper bound on bitlength of the outputs and ℓ_r be an upper bound on the bitlength of random tape of such samplers. Let PRG be a PRG that maps $\{0,1\}^\lambda$ to $\{0,1\}^{2\lambda}$. Let $|u_1| = \ell_1 = 2B + 2\ell_{in} + 2\ell_{out} + 5\lambda$, $|u_2| = \ell_2 = B + \ell_{in} + \ell_{out} + 2\lambda$.

Theorem 15. *Let iO be a perfectly correct indistinguishability obfuscator, F_1, F_2, F_3 be puncturable PRFs satisfying the following additional properties*

¹⁵ For weak $\text{PREH}_{\approx_c}^{\text{rand}}$, the NIZK proof system must be replaced by a NIWI proof system (without CRS) and therefore the sampler S samples two images of the EOWF and a NIWI proof using one of the preimages as witness.

- F_1 is extracting when the input min-entropy is greater than $\ell_r + 2(\lambda + 1) + 2$ with error less than $2^{-(\lambda+1)}$ and has input length $\ell_1 + \ell_2 + \ell_{in} + B$ and output length ℓ_r (such a PRF exists from one-way functions since $\ell_1 + \ell_2 + \ell_{in} + B \geq \ell_r + 2(\lambda + 1) + 2$),
- F_2 is statistically injective and has input length $B + \ell_{in} + \ell_{out} + 2\lambda$ and output length ℓ_1 (such a PRF exists from one-way functions since $\ell_1 \geq 2(B + \ell_{in} + \ell_{out} + 2\lambda) + \lambda$),
- F_3 has input length ℓ_1 and output length ℓ_2 .

Then, perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ with universal setup is true.

Proof. Perfect correctness. Let $\text{crs} =: (\Lambda_E, \Lambda_D)$ be parameters in the support of $\text{Setup}(1^\lambda, B)$. Let S be a PPT algorithm represented as a polynomial sized circuit, $m \in L$ be an input for S and let $y \in \text{supp}(S(m))$. Due to perfect correctness of iO we have

$$\begin{aligned} & D_S(\text{crs}, m, E_S(\text{crs}, m, y; r)) \\ &= C_D(S, m, C_E(S, m, y, r)) \\ &= C_D(S, m, \underbrace{(F_2(S, m, y, \text{PRG}(r)))}_{=e_1}, \underbrace{(F_3(e_1) \oplus (S, m, y, \text{PRG}(r)))}_{=e_2}) = y. \end{aligned}$$

Pseudorandomness. The proof can be divided into two main steps, see Figure 29.

G_0	G_1	G_2
$m^* \leftarrow \mathcal{A}(1^\lambda)$	$m^* \leftarrow \mathcal{A}(1^\lambda)$	$m^* \leftarrow \mathcal{A}(1^\lambda)$
sample k_1, k_2, k_3	sample k_1, k_2, k_3	sample k_1, k_2, k_3
$r^* \leftarrow \{0, 1\}^*$	$r^*, u^* \leftarrow \{0, 1\}^*$	$r^*, u^* \leftarrow \{0, 1\}^*$
$x^* \leftarrow \{0, 1\}^{p(\lambda)}$	$x^* := F_1(S^*, m^*, u^*)$	$x^* := F_1(S^*, m^*, u^*)$
$y^* := S^*(m^*; x^*)$	$y^* := S^*(m^*; x^*)$	$y^* := S^*(m^*; x^*)$
$e_1^* := F_2(S^*, m^*, y^*, \text{PRG}(r^*))$	$e_1^* := F_2(S^*, m^*, y^*, \text{PRG}(r^*))$	$e_1^* := F_2(S^*, m^*, y^*, \text{PRG}(r^*))$
$e_2^* := F_3(e_1^*) \oplus (S^*, m^*, y^*, \text{PRG}(r^*))$	$e_2^* := F_3(e_1^*) \oplus (S^*, m^*, y^*, \text{PRG}(r^*))$	$e_2^* := F_3(e_1^*) \oplus (S^*, m^*, y^*, \text{PRG}(r^*))$
$\Lambda_E \leftarrow \text{iO}(C_E[k_2, k_3])$	$\Lambda_E \leftarrow \text{iO}(C_E[k_2, k_3])$	$\Lambda_E \leftarrow \text{iO}(C_E[k_2, k_3])$
$\Lambda_D \leftarrow \text{iO}(C_D[k_1, k_2, k_3])$	$\Lambda_D \leftarrow \text{iO}(C_D[k_1, k_2, k_3])$	$\Lambda_D \leftarrow \text{iO}(C_D[k_1, k_2, k_3])$
$\text{crs} := (\Lambda_E, \Lambda_D)$	$\text{crs} := (\Lambda_E, \Lambda_D)$	$\text{crs} := (\Lambda_E, \Lambda_D)$
return $\mathcal{A}(\text{crs}, e^*)$	return $\mathcal{A}(\text{crs}, e^*)$	return $\mathcal{A}(\text{crs}, u^*)$

Fig. 29. Hybrids used in the proof of pseudorandomness for Theorem 15.

Lemma 15. For all (potentially unbounded) adversaries $\mathcal{A}$, $|\Pr[\text{out}_1 = 1] - \Pr[\text{out}_0 = 1]|$ is negligible.

The proof works as the proof of IND-CPA security of the deniable encryption scheme from [SW14] or the proof of statistical functional equivalence from [DKR15].

Lemma 16. For all PPT adversaries $\mathcal{A}$, $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]|$ is negligible.

The proof is similarly as the proof of explainability of the deniable encryption scheme from [SW14] or the proof of explainability from [DKR15]. $\square$

By allowing the setup algorithm Setup_S to depend on S and replacing the universal circuit in C_D with S , we obtain $\text{cPREH}_{\approx_c}^{\text{rand}}$ without universal setup.

Let S be a PPT sampler. Let $\ell_{in} = |m|$ be an upper bound on the bitlength of the inputs, ℓ_{out} be an upper bound on bitlength of the outputs and ℓ_r be an upper bound on the bitlength of random tape of S . Let PRG be a PRG that maps $\{0, 1\}^\lambda$ to $\{0, 1\}^{2\lambda}$. Let $|u_1| = \ell_1 = 2\ell_{in} + 2\ell_{out} + 5\lambda + \ell_r$ and $|u_2| = \ell_2 = \ell_{in} + \ell_{out} + 2\lambda$.

Theorem 16. Let iO be a perfectly correct indistinguishability obfuscator, $\text{PRF}_1, \text{PRF}_2, \text{PRF}_3$ be puncturable PRFs satisfying the following additional properties

$\text{Setup}'_S(1^\lambda)$	$E'_S(\text{crs}, m, y; r)$	$D'_S(\text{crs}, m, u)$
$\Lambda_E \leftarrow \text{iO}(C'_E[S, k_2, k_2])$	return $\Lambda_E(m, y, r)$	return $\Lambda_D(m, u)$
$\Lambda_D \leftarrow \text{iO}(C'_D[S, k_1, k_2, k_2])$		
return $\text{crs} := (\Lambda_E, \Lambda_D)$		
$C'_E[S, k_2, k_3](m, y, r)$	$C'_D[S, k_1, k_2, k_3](m, u)$	
$e_1 := \text{PRF}_2(m, y, \text{PRG}(r))$	$(m', y', r') := \text{PRF}_2(u_1) \oplus u_2$	
$e_2 := \text{PRF}_3(e_1) \oplus (m, y, \text{PRG}(r))$	if $(m' = m \wedge$	
return (e_1, e_2)	$u_1 = \text{PRF}_2(m', y', r'))$ then	
	return y'	
	$x := \text{PRF}_1(m, u)$	
	return $S(m; x)$	

Fig. 30. Instantiation of perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ due to [SW14; DKR15] without universal setup.

- PRF_1 is extracting when the input min-entropy is greater than $\ell_r + 2(\lambda + 1) + 2$ with error less than $2^{-(\lambda+1)}$ and has input length $\ell_1 + \ell_2 + \ell_{\text{in}}$ and output length ℓ_r (such a PRF exists from one-way functions since $\ell_1 + \ell_2 + \ell_{\text{in}} + \ell_r \geq \ell_r + 2(\lambda + 1) + 2$),
- PRF_2 is statistically injective and has input length $\ell_{\text{in}} + \ell_{\text{out}} + 2\lambda$ and output length ℓ_1 (such a PRF exists from one-way functions since $\ell_1 \geq 2(\ell_{\text{in}} + \ell_{\text{out}} + 2\lambda) + \lambda$),
- PRF_3 has input length ℓ_1 and output length ℓ_2 .

Then, perfectly correct $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true.

The proof of Theorem 16 proceeds identically to Theorem 15. The only differences are changed input and output lengths for the pseudorandom functions.

Theorems 15 and 16 together with Theorem 1 yield the following corollary.

Corollary 11. Assuming polynomially secure IO for all circuits and one-way functions, then $\text{acPREH}_{\approx_c}^{\text{rand}}$ (with or without universal setup) is true.

6.4 Bootstrapping pseudorandom encodings with a common random string

Recall that $\text{cPREH}_{\approx_c}^{\text{rand}}$ where the setup algorithm produces uniform random strings (i.e., in the common random string model) conflicts with EOWFs with common but benign auxiliary input – an assumptions which we believe to be true. However, the very nature of pseudorandom encodings allows to bootstrap pseudorandom encodings in the common random string model from pseudorandom encodings with arbitrary setup if this setup algorithm – interpreted as a sampler – can be pseudorandomly encoded such that the corresponding setup algorithm produces common random strings. Note that this sampler does not expect an input. Via this bootstrapping, we are able to refute that even the *weak* pseudorandom encoding hypothesis with common random string.

Theorem 17. Assume (i) $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true for all PPT samplers, i.e., for all PPT algorithms S , there exists a pseudorandom encoding scheme $(\text{Setup}'_S, E'_S, D'_S)$, and (ii) weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true for the class of PPT algorithms Setup'_S such that the corresponding setup algorithm $\text{Setup}''_{\text{Setup}'_S}$ produces uniform random strings. Then, $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random strings is true for all PPT samplers.

Proof. (i) implies that for all PPT algorithms S , there exists a pseudorandom encoding scheme $(\text{Setup}'_S, E'_S, D'_S)$. Further, (ii) guarantees the existence of a pseudorandom encoding scheme $(\text{Setup}''_{\text{Setup}'_S}, E''_{\text{Setup}'_S}, D''_{\text{Setup}'_S})$ such that $\text{Setup}''_{\text{Setup}'_S}$ produces uniform random strings. Let $\{0, 1\}^{n''(\lambda)}$ be the range of $E''_{\text{Setup}'_S}$ (note that $\{0, 1\}^{n''(\lambda)}$ is allowed to depend on the sampler Setup'_S). For notational convenience, we henceforth omit the dependency on the sampler Setup'_S .

Let S be some PPT algorithm. We define $(\text{Setup}_S, E_S, D_S)$ corresponding to S as in Figure 31.

$\text{Setup}_S(1^\lambda)$	$E_S(\text{crs}, m, y)$	$D_S(\text{crs}, m, u)$
$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$	parse $\text{crs} := \text{crs}'' \parallel u''$	parse $\text{crs} := \text{crs}'' \parallel u''$
$u'' \leftarrow \{0, 1\}^{n''(\lambda)}$	$\text{crs}' := D''(\text{crs}'', u'')$	$\text{crs}' := D''(\text{crs}'', u'')$
$\text{crs} := \text{crs}'' \parallel u''$	$u \leftarrow E'_S(\text{crs}', m, y)$	$y \leftarrow D'_S(\text{crs}', m, u)$
return crs	return u	return y

Fig. 31. Algorithms for $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common *random* string.

Correctness. We use a sequence of hybrids, see Figure 32.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$
$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$
$\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$	$\text{crs}'' \leftarrow \text{Setup}''(1^\lambda)$	$y \leftarrow S(m)$
$y \leftarrow S(m)$	$u'' \leftarrow \{0, 1\}^{n''(\lambda)}$	$\text{crs}' \leftarrow \text{Setup}'_S(1^\lambda)$
$u \leftarrow E_S(\text{crs}, m, y)$	$y \leftarrow S(m)$	$u \leftarrow E'_S(\text{crs}', m, y)$
$y_D := D_S(\text{crs}, m, u)$	$\text{crs}' := D''(\text{crs}'', u'')$	$y_D \leftarrow D'_S(\text{crs}', m, u)$
return $y = y_D$	$u \leftarrow E'_S(\text{crs}', m, y)$	return $y = y_D$
	$y_D \leftarrow D'_S(\text{crs}', m, u)$	
	return $y = y_D$	

Fig. 32. Hybrids used in the proof of correctness of Theorem 17.

$\mathbf{G}_0$ encodes and decodes a sample y (using a statically chosen input m) using $(\text{Setup}_S, E_S, D_S)$. $\mathbf{G}_1$ is identical to $\mathbf{G}_0$ except that $\mathbf{G}_1$ uses the implementations of Setup_S , E_S and D_S explicitly. This difference is only conceptual.

$\mathbf{G}_2$ is identical to $\mathbf{G}_1$ except for crs' being sampled using $\text{Setup}'_S(1^\lambda)$ instead of decoding u'' using $D''(\text{crs}'', u'')$.

Due to Lemma 5, by correctness and pseudorandomness (ii), the two distributions

$$\begin{aligned} & \left\{ \text{crs}'' \leftarrow \text{Setup}''(1^\lambda), \text{crs}' \leftarrow \text{Setup}'_S(1^\lambda) : (\text{crs}'', \text{crs}') \right\} \\ & \left\{ \text{crs}'' \leftarrow \text{Setup}''(1^\lambda), \text{crs}' \leftarrow D''(\text{crs}'', U_{n''(\lambda)}) : (\text{crs}'', \text{crs}') \right\} \end{aligned}$$

are computationally indistinguishable. Hence, $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]|$ is negligible.

Due to correctness (i), for all PPT adversaries $\mathcal{A}$,

$$\Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ \text{crs}' \leftarrow \text{Setup}'_S(1^\lambda) : D'_S(\text{crs}', m, E'_S(\text{crs}', m, y)) = y \\ y \leftarrow S(m) \end{array} \right]$$

and, hence, $\Pr[\text{out}_2 = 1]$ are overwhelming.

Before we prove pseudorandomness, we introduce a technical lemma.

Lemma 17. *Let S be a PPT sampler such that $\text{cPREH}_{\approx_c}^{\text{rand}}$ is true for S and let $(\text{Setup}_S, E_S, D_S)$ be the corresponding algorithms (such that E_S has output length $n(\lambda)$). Then, for all PPT*

adversaries $\mathcal{A}$,

$$\left| \Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ crs \leftarrow \text{Setup}_S(1^\lambda) \\ y \leftarrow S(m) \\ u \leftarrow E_S(crs, m, y) \end{array} : \mathcal{A}(crs, m, y, u) = 1 \right] \right. \\ \left. - \Pr \left[\begin{array}{l} m \leftarrow \mathcal{A}(1^\lambda) \\ crs \leftarrow \text{Setup}_S(1^\lambda) \\ u \leftarrow \{0, 1\}^{n(\lambda)} \\ y \leftarrow D_S(crs, m, u) \end{array} : \mathcal{A}(crs, m, y, u) = 1 \right] \right|$$

is negligible.

Proof (Proof of Lemma 17). We proceed over a series of hybrids, see Figure 33.

$\mathbf{H}_0$ corresponds to the left-hand-side of the difference, where $\mathcal{A}$ receives a sample y from $S(m)$ and the corresponding encoded sample $u \leftarrow E_S(crs, m, y)$.

$\mathbf{H}_1$ is identical to $\mathbf{H}_0$ except that $\mathcal{A}$ receives the decoded sample y_D instead of the actual sample y_S . If $y_S = y_D$ both games behave identically, hence, by correctness, $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \epsilon_{\mathcal{A}}^{\text{c-dec-error}}(\lambda)$.

$\mathbf{H}_2$ is identical to $\mathbf{H}_1$ except that u is not produced as an encoding of y , but sampled uniformly at random from $\{0, 1\}^{n(\lambda)}$. This game hop is justified by pseudorandomness $|\Pr[out_2 = 1] - \Pr[out_1 = 1]| \leq \text{Adv}_{\mathcal{A}}^{\text{crs-pre}}(\lambda)$.

$\mathbf{H}_0$	$\mathbf{H}_1$	$\mathbf{H}_2$
$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$
$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$
$y_S \leftarrow S(m)$	$y_S \leftarrow S(m)$	$u \leftarrow \{0, 1\}^{n(\lambda)}$
$u \leftarrow E_S(crs, m, y)$	$u \leftarrow E_S(crs, m, y)$	$y_D \leftarrow D_S(crs, m, u)$
return $\mathcal{A}(crs, m, y_S, u)$	$y_D \leftarrow D_S(crs, m, u)$	return $\mathcal{A}(crs, m, y_D, u)$
	return $\mathcal{A}(crs, m, y_D, u)$	

Fig. 33. Hybrids used in the proof of Lemma 17.

This concludes the proof of Lemma 17. $\square$

Pseudorandomness. To prove pseudorandomness, we proceed over a sequence of hybrids, see Figure 34.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$
$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$	$m \leftarrow \mathcal{A}(1^\lambda)$
$crs'' \leftarrow \text{Setup}''(1^\lambda)$	$crs'' \leftarrow \text{Setup}''(1^\lambda)$	$crs'' \leftarrow \text{Setup}''(1^\lambda)$
$u'' \leftarrow \{0, 1\}^{n''(\lambda)}$	$crs' \leftarrow \text{Setup}'_S(1^\lambda)$	$crs' \leftarrow \text{Setup}'_S(1^\lambda)$
$crs' := D''(crs'', u'')$	$u'' \leftarrow E''(crs'', crs')$	$u'' \leftarrow E''(crs'', crs')$
$crs := crs'' \parallel u''$	$crs := crs'' \parallel u''$	$crs := crs'' \parallel u''$
$y \leftarrow S(m)$	$y \leftarrow S(m)$	$y \leftarrow S(m)$
$u \leftarrow E'_S(crs', m, y)$	$u \leftarrow E'_S(crs', m, y)$	$u \leftarrow \{0, 1\}^{n'(\lambda)}$
return $\mathcal{A}(crs, u)$	return $\mathcal{A}(crs, u)$	return $\mathcal{A}(crs, u)$

Fig. 34. Hybrids used in the proof of pseudorandomness of Theorem 17.

$\mathbf{G}_0$ corresponds to $\text{Exp}_{\mathcal{A},0}^{\text{crs-pre}}(\lambda)$. $\mathbf{G}_1$ is identical to $\mathbf{G}_0$ except for the following difference. $\mathbf{G}_1$ produces crs' directly via $\text{Setup}'_S(1^\lambda)$ and produces u'' via encoding crs' with $E''(crs'', crs')$. $\mathbf{G}_0$, on the other hand, samples u'' uniformly at random from $\{0, 1\}^{n''(\lambda)}$ and produces crs'' via decoding u'' as $D''(crs'', u'')$. Due to Lemma 17, $|\Pr[out_1 = 1] - \Pr[out_0 = 1]|$ is negligible.

$\mathbf{G}_2$ is identical to $\mathbf{G}_1$ except that $\mathbf{G}_2$ samples u uniformly at random from $\{0, 1\}^{n'(\lambda)}$, whereas $\mathbf{G}_1$ produces u as $\mathbf{E}'_G(\text{crs}', m, y)$. For all PPT adversaries $\mathcal{A}$, $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]|$ is negligible due to pseudorandomness (i). $\square$

Combining Theorem 17 and Corollary 10, we obtain the following corollary refuting weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string.

Corollary 12. *If there exist extractable one-way function family ensembles with common but benign auxiliary information and indistinguishability obfuscation for all circuits, then weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string is false.*

Remark 8. If weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ was true such that the setup algorithms always produce uniform random strings, then the global common reference string model could be replaced with the global common random string model (or the non-programmable random oracle model) in every setting in cryptography.

7 Relations and applications of pseudorandom encodings

In this section we describe further applications of pseudorandom encodings. This unifies several areas in cryptography. In Section 7.1, we show that adaptively secure pseudorandom encodings with setup are equivalent to fully adaptive multi-party computation in the global CRS model for all PPT functionalities. This extends the results from [IKOS10]. In Section 7.2, we show that pseudorandom encodings yields honey encryption due to [JR14] for arbitrary message distributions, even such which admit inputs. In Section 7.3, we define a keyless version of steganography. That is, in contrast to symmetric-key or public-key steganography, parties do not need any secret information in order to covertly communicate with each other. In Section 7.4, we define the notion of covert secure computation and give a general compiler which transforms secure computation protocol into a covert MPC protocol. In Section 7.5, we analyze the relation of pseudorandom encodings and certain PKE variants which are known to imply (fully) adaptive MPC.

7.1 Fully adaptively secure multi-party computation in the CRS model

Due to the equivalence of pseudorandom encodings and invertible sampling, we obtain an equivalence between pseudorandom encodings and fully adaptively secure multi-party computation. Due to [IKOS10], $\text{PREH}_{\approx_c}^{\text{rand}}$ in conjunction with some adaptively secure oblivious transfer protocol implies fully adaptively secure multi-party computation for all randomized functionalities in the plain model.

In this section we first introduce some preliminaries on adaptive multi-party computation, henceforth denoted AMPC. In Sections 7.1.1 and 7.1.2, we extend the results on the relation between ISH and AMPC from [IKOS10] to the global CRS model. This is directly possible since, due to Theorem 1 and Corollary 1, we have an instantiation of $\text{aclSH}_{\approx_c}^{\text{rand}}$.

Adaptive MPC

The definition of security of multi-party protocols follows the real/ideal model paradigm, [Can00]. A protocol Π is said to be secure if the output of the real execution of the protocol is indistinguishable from the output of an ideal computation, where a trusted third party exists.

Let $\mathcal{F}$ be a PPT functionality. Let x_i be the input to party P_i . We only consider semi-honest adversaries. Semi-honest adversaries are bound to follow the protocol specification while trying to obtain as much information as possible. There are general techniques to transform protocols which are secure with respect to semi-honest adversaries to protocols which are secure with respect to malicious adversaries at the cost of a local CRS which is inherent for malicious security, [CLOS02]. In the following, we describe the stand-alone model for adaptive MPC on a very high level. We refer the reader to [Can00; Lin09] for more details.

The ideal execution, [Can00]. The ideal execution involves an ideal PPT adversary Sim (also called the simulator), a PPT environment $\mathcal{Z}$ and a trusted third party T . The ideal execution proceeds through several stages. In the *first corruption stage*, Sim adaptively decides to corrupt a party. When Sim corrupts a party, Sim learns that party's input and $\mathcal{Z}$ learns the identity of the corrupted party. In the *computation stage*, the uncorrupted and the corrupted parties send their inputs to T . The trusted party T evaluates the (possibly randomized) function $(z_1, \dots, z_n) \leftarrow \mathcal{F}(x_1, \dots, x_n)$ and sends the output to the respective parties. In the *second corruption stage*, Sim learns the outputs of the corrupted parties and again adaptively decides to corrupt a party. Upon corruption of a party, Sim learns that party's input *and output* and $\mathcal{Z}$ learns the corrupted party's identity. In the *output stage*, the uncorrupted parties output what they received from T to $\mathcal{Z}$, the corrupted parties output $\perp$ to $\mathcal{Z}$ and Sim outputs an arbitrary string to $\mathcal{Z}$. Finally, in the *post-execution corruption stage*, as long as the environment $\mathcal{Z}$ did not halt, $\mathcal{Z}$ sends corruption requests to Sim . Sim generates an arbitrary answer based on its view so far. Furthermore, Sim may additionally corrupt parties as in the second corruption stage.

Let $\mathcal{Z}^{\text{Sim}, \mathcal{F}}(1^\lambda)$ denote the output distribution of the environment $\mathcal{Z}$ when interacting as described above with the simulator Sim and parties $P_1, \dots, P_n$ on inputs $x_1, \dots, x_n$ chosen by $\mathcal{Z}$.

The real execution, [Can00]. Initially, the environment $\mathcal{Z}$ (adaptively) chooses inputs $x_1, \dots, x_n$ and sends each party P_i its input. Before the communication rounds start, the adversary $\mathcal{A}$ receives an initial message from $\mathcal{Z}$. While there exist uncorrupted parties which did not halt, $\mathcal{A}$ may adaptively decide to corrupt new parties. Upon corruption of P_i , $\mathcal{Z}$ learns the identity of P_i , $\mathcal{A}$ learns the input x_i , P_i 's internal state (i.e. the random tape) and all messages received so far. From that time on, $\mathcal{A}$ is in control of the messages this party sends. Furthermore, if a corrupted party receives a message, $\mathcal{A}$ also learns that message. Additionally, $\mathcal{A}$ determines the order in which the uncorrupted parties are activated. If all parties are corrupted or all uncorrupted parties halted, each uncorrupted party and $\mathcal{A}$ produce outputs. The environment $\mathcal{Z}$ learns all of these outputs. Similarly to the post-execution corruption stage in the ideal execution, while $\mathcal{Z}$ did not halt, $\mathcal{Z}$ may instruct $\mathcal{A}$ to corrupt more parties or $\mathcal{A}$ may decide himself to corrupt more parties. Upon corruption, $\mathcal{A}$ learns P_i 's input, randomness and all messages received so far and $\mathcal{Z}$ learns the corrupted party's identity. Additionally, $\mathcal{A}$ sends P_i 's internal state to $\mathcal{Z}$.

Let $\mathcal{Z}^{\mathcal{A}, \Pi}(1^\lambda)$ denote the output of the environment $\mathcal{Z}$ when interacting with the adversary $\mathcal{A}$ and parties $P_1, \dots, P_n$ running the protocol Π on inputs $x_1, \dots, x_n$ chosen by $\mathcal{Z}$.

Definition 29 (Adaptive security in the stand-alone model, [Can00]). *We say a protocol Π computes functionality $\mathcal{F}$ in the adaptive semi-honest model if for every PPT adversary $\mathcal{A}$, there exists a PPT simulator Sim , such that for all PPT environments $\mathcal{Z}$,*

$$\left| \Pr[\mathcal{Z}^{\mathcal{A}, \Pi}(1^\lambda) = 1] - \Pr[\mathcal{Z}^{\text{Sim}, \mathcal{F}}(1^\lambda) = 1] \right|$$

is negligible.

Common reference string model. In the common reference string (CRS) model, all parties have access to a string which is honestly generated by a trusted third non-participating party. There are two widely used variants of the CRS model. In the *programmable* or *local* CRS model, the simulator may generate the CRS. This enables the simulator to sample the CRS along with corresponding trapdoors which results in an asymmetry between the simulator and the adversary facilitating simulation. However, as soon as two different protocols use the same programmable CRS, all security guarantees related to that CRS break down, see [CDPW07]. In the *non-programmable* or *global* CRS model, the simulator receives the CRS as input and, hence, has no additional power compared to the adversary. A global CRS can be made public and used by any number of protocols without compromising security.

Adaptive security in the global CRS model (with respect to some efficiently samplable CRS distribution Setup) is defined as Definition 29 with the difference that all parties including the

environment, the adversary and the simulator receive the CRS as input. We stress that the environment may decide on the inputs for the parties adaptively after seeing the CRS.

The universal composability (UC) framework, [Can01]. The following paragraph describes the UC framework on a very high level. We refer the reader to [Can01] for more details on the model.

In contrast to the stand-alone model for secure computation, the UC model allows the environment to be interactive. In particular, every message that is transmitted between parties is sent to the environment which can arbitrarily decide how to deal with it. Hence, the strongest possible adversary in this model is the so called dummy adversary who simply forwards all instructions he receives from the environment. The UC framework offers a set of security-preserving composition theorems which allow for a modular analysis. If a protocol Π UC-realizes a functionality $\mathcal{F}$, a protocol Π' which uses protocol Π as a subroutine can be proven to securely realize a functionality $\mathcal{F}'$ in the $\mathcal{F}$ -hybrid model, that is having access to the ideal functionality $\mathcal{F}$.

Definition 30 (Adaptive UC-security, [Can01]). *Let $\mathcal{A}$ be the dummy adversary and let θ denote the dummy protocol. We say a protocol Π UC-realizes a functionality $\mathcal{F}$ in the $\mathcal{G}$ -hybrid model in the presence of semi-honest adaptive adversaries if for all PPT environments $\mathcal{Z}$, there exists a simulator Sim , such that*

$$\left| \Pr[\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}] = 1] - \Pr[\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}] = 1] \right|$$

is negligible, where $\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}]$ denotes the output distribution of $\mathcal{Z}$ when interacting in the real world with $\mathcal{A}$ and $\pi^{\mathcal{G}}$ and $\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}]$ denotes the output distribution of $\mathcal{Z}$ when interacting in the ideal world with Sim and $\theta^{\mathcal{F}}$.

UC-security is a stronger notion than stand-alone security. More precisely, if there exists a protocol which UC-realizes a functionality according to Definition 30, then there exists a protocol which securely computes that functionality according to Definition 29.

In the UC framework, the global CRS model corresponds to the $\mathcal{F}_{\text{Setup}}^{\text{crs}}$ -hybrid model. The ideal functionality $\mathcal{F}_{\text{Setup}}^{\text{crs}}$ can be queried by any party and the adversary. Upon receiving such a query, $\mathcal{F}_{\text{Setup}}^{\text{crs}}$ gives the CRS crs to the querying party. If crs is not initialized, $\mathcal{F}_{\text{Setup}}^{\text{crs}}$ samples crs according to $\text{Setup}(1^\lambda)$.

7.1.1 Pseudorandom encodings imply UC-secure AMPC in the plain model

[IKOS10] already observed that adaptively secure oblivious transfer in conjunction with $\text{PREH}_{\approx_c}^{\text{rand}}$ yields AMPC for all randomized functionalities in the plain model.

Theorem 18 (informal, [Kil88; IPS08]). *Any deterministic functionality can be UC-realized in the OT-hybrid model in the presence of semi-honest adversaries adaptively corrupting any number of parties.*

Theorem 19 ([IKOS10]). *Assume $\text{PREH}_{\approx_c}^{\text{rand}}$ holds, then any functionality $\mathcal{F}$ can be UC-realized, in the $\mathcal{F}_{\text{OT}}$ -hybrid-model, in the presence of semi-honest adaptive adversaries corrupting any number of parties.*

Assuming $\text{acPREH}_{\approx_c}^{\text{rand}}$, the above strategy of [IKOS10] can be applied to show that every randomized functionality can be UC-realized in the (global CRS, OT)-hybrid model in the presence of semi-honest adversaries adaptively corrupting any number of parties.

Theorem 20. *Assume $\text{acPREH}_{\approx_c}^{\text{rand}}$ holds, then any functionality $\mathcal{F}$ can be UC-realized, in the (global CRS, $\mathcal{F}_{\text{OT}}$)-hybrid-model, in the presence of semi-honest adaptive adversaries corrupting any number of parties.*

Proof. The proof follows the strategy from [IKOS10]. We consider functionalities $\mathcal{F}$ giving the same output to all parties. (This is without loss of generality since evaluating the function $\mathcal{F}'((x_1, k_1), (x_2, k_2)) := \mathcal{F}_1(x_1, x_2) \oplus k_1 \parallel \mathcal{F}_2(x_1, x_2) \oplus k_2$ yields the general case, see [GL91].) I.e. the functionality $\mathcal{F}$ takes as input x_1, x_2 (the inputs of the parties P_1, P_2 , respectively) and internal randomness ρ , and outputs z to both parties. We consider the case of two parties (the proof easily extends to the case of many parties).

We view $\mathcal{F}$ as a PPT algorithm S . Due to $\text{acPREH}_{\approx_c}^{\text{rand}}$, there exist $\text{Setup}_S, \bar{S}, \bar{S}^{-1}$ satisfying adaptive closeness and invertibility. Define a deterministic functionality $\mathcal{G}$ which takes inputs (x_1, ρ_1) from party P_1 and (x_2, ρ_2) from party P_2 , and evaluates $\bar{S}(\text{crs}, (x_1, x_2); \rho_1 \oplus \rho_2)$, where crs comes from $\mathcal{F}_{\text{crs}}^{\text{Setup}_S}$. Due to Theorem 18, $\mathcal{G}$ can be UC-realized in the OT-hybrid model in the presence of semi-honest adversaries adaptively corrupting any number of parties.

Now we turn to realize $\mathcal{F}$ in the $\mathcal{G}$ -hybrid model. Party P_i chooses randomness ρ_i , feeds (x_i, ρ_i) into $\mathcal{G}$ and waits for the output. We assume that $\mathcal{Z}$ first observes the entire protocol before corrupting parties. The simulator Sim works as follows. Sim initially only receives the output z of the ideal functionality $\mathcal{F}$ as input and calls $\mathcal{F}_{\text{crs}}^{\text{Setup}_S}$ to obtain crs . If both parties are corrupted (first P_1 , then P_2), Sim outputs a uniformly random string ρ_1 to explain the internal randomness of party P_1 , computes $\rho \leftarrow \bar{S}^{-1}(\text{crs}, (x_1, x_2), z)$ and outputs $\rho_2 := \rho \oplus \rho_1$ to explain the internal randomness of party P_2 (note that Sim learns the inputs x_1, x_2 at the time of corrupting P_1, P_2 , respectively). Due to invertibility and closeness, for all environments $\mathcal{Z}$, $|\Pr[\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}] = 1] - \Pr[\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}] = 1]|$ is negligible. More precisely, consider the hybrids in Figure 35. Note that $\Pr[\text{out}_0 = 1] = \Pr[\text{real}[\mathcal{Z}, \mathcal{A}, \pi^{\mathcal{G}}] = 1]$ and $\Pr[\text{out}_3 = 1] = \Pr[\text{ideal}[\mathcal{Z}, \text{Sim}, \theta^{\mathcal{F}}] = 1]$. The difference between $\mathbf{G}_0$ and $\mathbf{G}_1$ is only conceptual. Furthermore, $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_1 = 1]| \leq \text{Adv}_{\mathcal{A}}^{\text{a-crs-inv}}(\lambda)$ and $|\Pr[\text{out}_3 = 1] - \Pr[\text{out}_2 = 1]| \leq \text{Adv}_{\mathcal{A}}^{\text{a-crs-close}}(\lambda)$.

$\mathbf{G}_0$ <hr/> $\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$ $(x_1, x_2) \leftarrow \mathcal{Z}(\text{crs})$ $\rho_1, \rho_2 \leftarrow \{0, 1\}^{n(\lambda)}$ $z := \mathcal{G}(x_1, \rho_1, x_2, \rho_2)$ upon corruption of P_1 do give (x_1, ρ_1, z) to $\mathcal{Z}$ upon corruption of P_2 do give (x_2, ρ_2) to $\mathcal{Z}$	$\mathbf{G}_1$ <hr/> $\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$ $(x_1, x_2) \leftarrow \mathcal{Z}(\text{crs})$ $\rho_1, \rho_2 \leftarrow \{0, 1\}^{n(\lambda)}$ $z := \bar{S}(\text{crs}, (x_1, x_2); \rho_1 \oplus \rho_2)$ upon corruption of P_1 do give (x_1, ρ_1, z) to $\mathcal{Z}$ upon corruption of P_2 do $\rho'_2 := \rho_2$ give (x_2, ρ'_2) to $\mathcal{Z}$
$\mathbf{G}_2$ <hr/> $\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$ $(x_1, x_2) \leftarrow \mathcal{Z}(\text{crs})$ $\rho_1, \rho_2 \leftarrow \{0, 1\}^{n(\lambda)}$ $z := \bar{S}(\text{crs}, x_1, x_2; \rho_1 \oplus \rho_2)$ upon corruption of P_1 do give (x_1, ρ_1, z) to $\mathcal{Z}$ upon corruption of P_2 do $\rho'_2 \leftarrow \rho_1 \oplus \bar{S}^{-1}(\text{crs}, (x_1, x_2), z)$ give (x_2, ρ'_2) to $\mathcal{Z}$	$\mathbf{G}_3$ <hr/> $\text{crs} \leftarrow \text{Setup}_S(1^\lambda)$ $(x_1, x_2) \leftarrow \mathcal{Z}(\text{crs})$ $\rho_1 \leftarrow \{0, 1\}^{n(\lambda)}$ $z \leftarrow S(x_1, x_2)$ upon corruption of P_1 do give (x_1, ρ_1, z) to $\mathcal{Z}$ upon corruption of P_2 do $\rho'_2 \leftarrow \rho_1 \oplus \bar{S}^{-1}(\text{crs}, (x_1, x_2), z)$ give (x_2, ρ'_2) to $\mathcal{Z}$

Fig. 35. Hybrids used in proof of Theorem 20.

□

Note that due to [DKR15], it is possible to realize UC-AMPC for every functionality only assuming the static version $\text{cPREH}_{\approx_c}^{\text{rand}}$ at the cost of additionally assuming the existence of adaptively secure two-round oblivious transfer protocol. This can be avoided due to Corollary 1.

7.1.2 AMPC with global setup implies adaptive pseudorandom encodings with setup

In the following we consider the standalone model [Can00]. Note that in the standalone model, the environment is strictly weaker than in the UC-framework.

Theorem 21. *If for all (two party) PPT functionalities $\mathcal{F}$, there exists a protocol Π that securely implements $\mathcal{F}$ in the global CRS model in the presence of adaptive semi-honest adversaries with post-execution corruption corrupting any number of parties, then $\text{acISH}_{\approx_c}^{\text{rand}}$ (without universal setup) is true.*

Proof. The proof follows the ideas of [IKOS10]. Let S be an arbitrary PPT algorithm. Let x_1 denote the input of party P_1 and x_2 denote the input of party P_2 .

Consider the randomized functionality $\mathcal{F}$ which computes $z := S(x_1, x_2; \rho)$ and outputs z to P_1 and $\perp$ to P_2 , where ρ is the internal random coins of the functionality. We denote the message space of S by L . Let Π be a protocol which securely realizes $\mathcal{F}$ in the global CRS model (where the common reference string is drawn according some distribution $\text{Setup}_{\mathcal{F}}$ possibly depending on the functionality $\mathcal{F}$) in the presence of adaptive semi-honest adversaries (with post-execution corruption). Further, let p_1 and p_2 be stateful PPT algorithms modeling the protocol messages exchanged between the parties P_1 and P_2 . More formally, $(m_{2i+1}, z_i) \leftarrow p_1(\text{crs}, x_1, m_{2i}; r_1)$ and $m_{2j} \leftarrow p_2(\text{crs}, x_2, m_{2j-1}; r_2)$ for $i \geq 0, j \geq 1$ and $r_1, r_2 \leftarrow \{0, 1\}^{\text{poly}(\lambda)}$ (for some sufficiently large polynomial poly). (Without loss of generality, we let p_1 additionally output z_i which equals $\perp$ during the protocol execution and contains the output of party P_1 after the execution.)

$\bar{S}(\text{crs}, (x_1, x_2); (r_1, r_2))$	$\bar{S}^{-1}(\text{crs}, (x_1, x_2), z; r_S)$
$z := \perp, m := \perp$	$\{m'_i\}_{i \in [\text{poly}(\lambda)]} \leftarrow \text{Sim}_1(\text{crs}; r_S)$
while $z = \perp$ do	$(z, r'_1) \leftarrow \text{Sim}_2(x_1, z)$
$(m, z) \leftarrow p_1(\text{crs}, x_1, m, r_1)$	$r'_2 \leftarrow \text{Sim}_3(x_2)$
$m \leftarrow p_2(\text{crs}, x_2, m, r_2)$	return (r'_1, r'_2)
return z	

Fig. 36. Definition of the alternative sampler $\bar{S}$ (left) and the corresponding inverse sampler $\bar{S}^{-1}$ (right).

Lemma 18. *The setup algorithm $\text{Setup}_S := \text{Setup}_{\mathcal{F}}$ together with the alternative sampler $\bar{S}$ defined in Figure 36 satisfy adaptive closeness.*

Proof. Since we only consider protocols Π with at most polynomially many rounds, $\bar{S}$ is a PPT algorithm.

Intuitively, the output distributions of S and $\bar{S}$ are computationally close because of “correctness” of the protocol Π .

Let $\mathcal{B}$ be the adversary on the protocol Π which only observes the protocol and does not interfere at all. By Definition 29, there exists a simulator Sim such that for all environments $\mathcal{Z}$, $|\Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\mathcal{A}, \Pi}(\text{crs}) = 1] - \Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1]|$ is negligible. Note that since $\mathcal{B}$ corrupts no parties, Sim also corrupts no parties.

Let $\mathcal{A}$ be an adversary on closeness. We construct an environment $\mathcal{Z}$ which distinguishes between the real and the ideal world as follows. Initially, $\mathcal{Z}$ receives crs as input, calls $\mathcal{A}$ on input of crs and obtains $(x_1, x_2) \in L$. $\mathcal{Z}$ uses x_1 as input for party P_1 and x_2 as input for P_2 . Then, $\mathcal{Z}$ executes the protocol (without any interference) and finally receives the output $(z, \perp)$. $\mathcal{Z}$ calls $\mathcal{A}$ on input z and outputs the $\mathcal{A}$ ’s output.

In the ideal world, $\mathcal{Z}$ receives the outputs z and $\perp$ for P_1 and P_2 from the trusted third party T (since no party is corrupted) which are computed by $\mathcal{F}$. Hence, if $\mathcal{Z}$ is in the ideal world, it simulates $\text{Exp}_{\mathcal{A}, 0}^{\text{a-crs-close}}$. In the real world, $\mathcal{Z}$ interacts with the actual parties running protocol Π and, hence, simulates $\text{Exp}_{\mathcal{A}, 1}^{\text{a-crs-close}}$. Therefore, $\text{Adv}_{\mathcal{A}}^{\text{a-crs-close}}(\lambda) \leq |\Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\mathcal{A}, \Pi}(\text{crs}) = 1] - \Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1]|$ for the environment $\mathcal{Z}$ and the adversary $\mathcal{B}$ which do not interfere with the protocol execution. $\square$

Let $\mathcal{B}$ be an adversary that operates as follows. During the protocol, $\mathcal{B}$ records the transcripts $\{m_i\}_{i \in [\text{poly}'(\lambda)]}$ and outputs them. At the end of the protocol, $\mathcal{B}$ corrupts P_1 obtaining the input x_1 , output z and random tape r_1 . $\mathcal{B}$ directly outputs (z, r_1) . Afterwards, $\mathcal{B}$ (post-execution) corrupts P_2 and obtains the input x_2 and the random tape r_2 and outputs r_2 .

Since we assume Π securely realizes $\mathcal{F}$ (in particular in the presence of $\mathcal{B}$), there exists a (stateful) simulator $\text{Sim} := (\text{Sim}_1, \text{Sim}_2, \text{Sim}_3)$ producing outputs that are computationally indistinguishable from the outputs $\mathcal{B}$ produces. Sim_1 simulates the execution before a corruption occurs. On input of crs and a random tape r_S , Sim_1 produces $\{m'_i\}_{i \in [\text{poly}'(\lambda)]}$. On corruption of P_1 (in the post-execution corruption stage), Sim_2 obtains x_1, z and outputs a string r'_1 corresponding to the randomness that P_1 used to produce its messages in the protocol (using input x_1) and outputs (z, r'_1) . On corruption of P_2 (in the post-execution corruption stage), Sim_3 obtains x_2 and outputs a string r'_2 corresponding to the randomness that P_2 used to produce its messages (using input x_2).

Lemma 19. *The setup algorithm $\text{Setup}_S := \text{Setup}_{\mathcal{F}}$ together with the alternative sampler $\bar{S}$ and the inverse sampler $\bar{S}^{-1}$ defined in Figure 36 satisfy adaptive invertibility.*

Proof. Since $\text{Sim}_1, \text{Sim}_2, \text{Sim}_3$ are PPT algorithms, $\bar{S}^{-1}$ is a PPT algorithm.

Recall that $\mathcal{F}(x_1, x_2; \rho) := S(x_1, x_2; \rho)$. Let $\mathcal{F}'(x_1, x_2; \rho') := \bar{S}(x_1, x_2; \rho')$. Intuitively, we consider an intermediate game, where $\mathcal{Z}$ interacts with the simulator Sim and the modified ideal functionality $\mathcal{F}'$.

Claim. For all PPT environments $\mathcal{Z}$,

$$\left| \Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1] - \Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1] \right| \leq \text{negl}(\lambda). \quad (8)$$

Proof. Let $\mathcal{Z}$ be a PPT environment. We construct an adversary $\mathcal{B}$ on adaptive closeness. $\mathcal{B}$ simulates the entire interaction between the $\mathcal{Z}$, the simulator Sim and the trusted party T , which evaluates the functionality $\mathcal{F}$ (or $\mathcal{F}'$). In particular, $\mathcal{B}$ takes control of the party T . On input of crs , $\mathcal{B}$ simulates the interaction between $\mathcal{Z}$ and Sim (given crs) until the computation stage. In the computation stage, the parties P_1 and P_2 send their inputs (which may have been adaptively chosen based on crs) to $\mathcal{B}$. $\mathcal{B}$ outputs $(x_1, x_2) \in L$ to $\text{Exp}_{\mathcal{A}, b}^{\text{a-crs-close}}$ and receives an output $z := y_b$, where $y_0 \leftarrow S(x_1, x_2)$ and $y_1 \leftarrow \bar{S}(x_1, x_2)$. $\mathcal{B}$ sends z to P_1 and P_2 and continues to simulate the interaction between $\mathcal{Z}$ and Sim . This is possible, because the randomness which is actually used by T is not ever needed to be known. Finally, $\mathcal{Z}$ halts and outputs a bit b' . $\mathcal{B}$ outputs b' . $\square$

Let $\mathcal{A}$ be an adversary on adaptive invertibility. We construct an environment $\mathcal{Z}$ that distinguishes between the real execution of Π with adversary $\mathcal{B}$ and the ideal execution with Sim and the ideal functionality $\mathcal{F}'$. Initially, $\mathcal{Z}$ receives crs and calls $\mathcal{A}$ on input of crs to obtain $(x_1, x_2) \in L$. $\mathcal{Z}$ uses x_1 as input for P_1 and x_2 as input for P_2 . After the execution of the protocol and the post-execution corruptions, $\mathcal{Z}$ receives $(r_1, r_2), z$ (either by the adversary $\mathcal{B}$ or the simulator Sim). Finally, $\mathcal{Z}$ calls $\mathcal{A}$ on input of $((r_1, r_2), z)$ and outputs $\mathcal{A}$'s output.

In the real world, z is the output of the real protocol and r_1, r_2 is the actual randomness used by the parties. Hence, by definition of $\bar{S}$, (r_1, r_2) is the randomness actually used by $\bar{S}$ to produce the output z . Therefore, $\Pr[\text{Exp}_{\mathcal{A}, 0}^{\text{a-crs-inv}}(\lambda) = 1] = \Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}(1^\lambda): \mathcal{Z}^{\mathcal{B}, \Pi}(\text{crs}) = 1]$.

In the ideal world, z is produced by the functionality $\mathcal{F}'$ (hence, by $\bar{S}$) and $(r_1, r_2) := (r'_1, r'_2)$ is produced by $\bar{S}^{-1}((x_1, x_2), z)$. Hence, $\Pr[\text{Exp}_{\mathcal{A}, 1}^{\text{a-crs-inv}}(\lambda) = 1] = \Pr[\text{crs} \leftarrow \text{Setup}_{\mathcal{F}}: \mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1]$. Hence,

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{a-crs-inv}}(\lambda) &= \left| \Pr_{\text{crs}}[\mathcal{Z}^{\mathcal{B}, \Pi}(\text{crs}) = 1] - \Pr_{\text{crs}}[\mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1] \right| \\ &\leq \left| \Pr_{\text{crs}}[\mathcal{Z}^{\mathcal{B}, \Pi}(\text{crs}) = 1] - \Pr_{\text{crs}}[\mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1] \right| \\ &\quad + \left| \Pr_{\text{crs}}[\mathcal{Z}^{\text{Sim}, \mathcal{F}}(\text{crs}) = 1] - \Pr_{\text{crs}}[\mathcal{Z}^{\text{Sim}, \mathcal{F}'}(\text{crs}) = 1] \right| \end{aligned}$$

which is negligible by assumption and Equation (8). $\square$

Therefore, for every PPT algorithm S , there exists an algorithm Setup_S , an alternative sampler $\bar{S}$ and an inverse sampler $\bar{S}^{-1}$ such that adaptive closeness and adaptive invertibility hold. This concludes the proof. $\square$

Combining Theorems 20 and 21 and since semi-honest adaptive MPC in the non-programmable common random string model implies semi-honest adaptive MPC in the plain model, we obtain the following corollary improving Corollary 12.

Corollary 13. *If there exist extractable one-way function family ensembles (without auxiliary information) and indistinguishability obfuscation for all circuits, then weak $\text{cPREH}_{\approx_c}^{\text{rand}}$ with common random string is false.*

7.2 Honey encryption

Honey encryption schemes [JR14] are symmetric key encryption schemes which offer security even against unbounded adversaries, i.e. even against adversaries which are able find the used symmetric key. This is a particularly useful notion when the key comes from a distribution with low min-entropy like passwords. On decryption with a wrong key, a honey encryption scheme behaves indistinguishably from decryption with the actually used key.

Let D_1 be an efficiently samplable key distribution over the key space $\mathcal{K}_\lambda$ and let D_2 be an efficiently samplable message distribution over the message space $\mathcal{M}_\lambda$.

Definition 31 (Honey encryption, [JR14]). *A symmetric encryption scheme (Enc, Dec) is called honey encryption scheme for key distribution D_1 and plaintext distribution D_2 if it satisfies the following property.*

Security against message recovery (with respect to D_1, D_2). *For all unbounded adversaries $\mathcal{A}$, the advantage*

$$\text{Adv}_{\mathcal{A}, D_1, D_2}^{\text{mr}}(\lambda) = \Pr[\text{Exp}_{\mathcal{A}, D_1, D_2}^{\text{mr}}(\lambda) = 1]$$

is negligibly close to $2^{-\mu_1}$, where $\mu_1 = H_\infty(D_1)$ and $\text{Exp}_{\mathcal{A}, D_1, D_2}^{\text{mr}}$ is defined as in Figure 37.

$\text{Exp}_{\mathcal{A}, D_1, D_2}^{\text{mr}}(\lambda)$	$\text{Exp}_{\mathcal{A}, 0}^{\text{dte}}(\lambda)$	$\text{Exp}_{\mathcal{A}, 1}^{\text{dte}}(\lambda)$
$k \leftarrow D_1(1^\lambda)$	$y^* \leftarrow S(1^\lambda)$	$u^* \leftarrow \{0, 1\}^{n(\lambda)}$
$m^* \leftarrow D_2(1^\lambda)$	$u^* \leftarrow \text{E}_S(m^*)$	$y^* := \text{D}_S(u^*)$
$c^* \leftarrow \text{Enc}(k, m^*)$	return $\mathcal{A}(u^*, y^*)$	return $\mathcal{A}(u^*, y^*)$
$m \leftarrow \mathcal{A}(c^*)$		
return $m = m^*$		

Fig. 37. Message recovery experiment and DTE security experiment.

Definition 32 (Distribution-transforming encoder, [JR14; JRT16]). *A distribution-transforming encoder (DTE) for a distribution sampled by a sampler S over $\mathcal{Y}_{S, \lambda}$ is a tuple of efficient algorithms (E_S, D_S) , where D_S is deterministic, such that the following properties are satisfied.*

Perfect correctness. *For all $y \in \mathcal{Y}_{S, \lambda}$, $\Pr[\text{D}_S(\text{E}_S(m)) = m] = 1$, where the probability is over the randomness of E_S .*

DTE security. *For all unbounded adversaries, $\text{Adv}_{\mathcal{A}}^{\text{dte}}(\lambda) := |\Pr[\text{Exp}_{\mathcal{A}, 0}^{\text{dte}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A}, 1}^{\text{dte}}(\lambda) = 1]|$ is negligible, where $\text{Exp}_{\mathcal{A}, 0}^{\text{dte}}$ and $\text{Exp}_{\mathcal{A}, 1}^{\text{dte}}$ are defined in Figure 37.*

The above notions are defined with respect to unbounded adversaries. This is a simplification meant to capture the fact that an adversary is able to perform work proportional to $2^{-\mu_1}$. For low-entropic key distributions, a computational variant of the above definition suffices. We refer to this as *computational* honey encryption.

Relaxing the perfect correctness requirement from Definition 32 recovers the definition of (weak) $\text{PREH}_{\equiv_s}^{\text{rand}}$ (or weak $\text{PREH}_{\approx_c}^{\text{rand}}$ in the computational case) by Lemma 17 and still suffices to imply honey encryption via the DTE-then-Encrypt framework due to [JR14] given a secret-key encryption scheme (with message space matching the range of E_S). Applying the DTE-then-Encrypt framework on $\text{acPREH}_{\equiv_s}^{\text{rand}}$ and $\text{acPREH}_{\approx_c}^{\text{rand}}$ yields honey encryption in a CRS model for high-entropic and low-entropic key distributions, respectively. Note that the adaptive versions are only necessary when considering message distributions D_2 with input.

Theorem 22. *If $\text{acPREH}_{\equiv_s}^{\text{rand}}$ (or $\text{acPREH}_{\approx_c}^{\text{rand}}$) is true an a suitable secret-key encryption scheme (modeled as random cipher) exists, then (computational) honey encryption exists for all efficiently samplable message distributions in the CRS model.*

The proof directly follows from [JR14]. Together with Theorems 15 and 16, we obtain the following corollary.

Corollary 14. *If polynomially secure IO exists and a (suitable) secret-key encryption scheme (modeled as random cipher) exists, then computational honey encryption exists for all efficiently samplable message distributions in the CRS model.*

7.3 Keyless steganography

Pseudorandom encodings yield a notion of keyless steganography. In the case of $\text{acPREH}_{\approx_c}^{\text{rand}}$ all parties need access to some public parameters, but none of the parties needs access to a secret key. We adopt the notation from [BL18]. Since decoding does not involve any secret information, any attack corresponds to an equivalent of SS-CCA-security for public-key stegosystems (PKStS), where the decoding may even be applied on the challenge. Hence, the definition of SS-CCA-security must be adapted such that the message to be hidden is not chosen by the adversary but sampled according to a predefined message distribution. Let $dl(\lambda)$ be the document length and $ol(\lambda)$ be the output length.

Definition 33 (Keyless stegosystem (KlStS) for distribution S). A keyless stegosystem for message distribution S is a triple of PPT algorithms $(\text{KlStS.Setup}, \text{KlStS.E}_S, \text{KlStS.D}_S)$, where

- $\text{KlStS.Setup}(1^\lambda)$ produces public parameters pp (without corresponding secret information),
- KlStS.E_S on input of pp , a message y sampled from S , a history $hist \in (\Sigma^{dl(\lambda)})^*$ and some state information $s \in \{0, 1\}^*$, produces a document $d \in \Sigma^{dl}$ (by being able to sample from $C_{\lambda, dl(\lambda)}$). $\text{KlStS.E}_S^C(pp, m, hist)$ denotes sampling $ol(\lambda)$ documents using KlStS.E_S one-by-one.
- KlStS.D_S on input of pp , and a sequence of documents $d_1, \dots, d_{ol(\lambda)}$, and outputs a message m' .

We require KlStS to meet the following properties.

Universality. KlStS works on any channel without prior knowledge of the distribution of the channel.

Reliability. The probability

$$\Pr [\text{KlStS.D}_S(pp, \text{KlStS.E}_S(pp, m, hist), hist) \neq m]$$

that decoding fails is negligible, where the probability is over the choice of pp , m and the random coins of KlStS.E_S .

$Exp_{\mathcal{A},C,b}^{\text{klsts-sec}}(\lambda)$	$\text{KlStS.Setup}_S(1^\lambda)$	$\text{KlStS.E}_S((crs, H), m)$
$pp \leftarrow \text{KlStS.Setup}(1^\lambda)$	$crs \leftarrow \text{Setup}_S(1^\lambda)$	$u \leftarrow \text{E}_S(crs, m)$
$m \leftarrow \mathcal{A}(pp)$	$H \leftarrow \mathcal{H}$	$u_1 \parallel \dots \parallel u_\ell =: u$
$m^* \leftarrow S(m)$	return $pp := (crs, H)$	for $i \in [\ell]$ do
$\text{hist}^* \leftarrow \mathcal{A}(1^\lambda, pp, m^*)$		do
$d_0^* \leftarrow \text{KlStS.E}_S^C(pp, m^*, \text{hist}^*)$		$d_i \leftarrow C(\lambda, \text{hist})$
$d_1^* \leftarrow C_{\lambda, dl(\lambda), \text{hist}^*}^{ol(\lambda)}$		until $\text{prefix}_m(H(d_i)) = u_i$
return $\mathcal{A}(d_b^*)$		return $(d_1, \dots, d_\ell)$
	$\text{KlStS.D}_S((crs, H), (d_1, \dots, d_\ell))$	
	for $i \in [\ell]$ do	
	$u_i := \text{prefix}_m(H(d_i))$	
	$u := u_1 \parallel \dots \parallel u_\ell$	
	return $\text{D}_S(crs, u)$	

Fig. 38. Definition of $Exp_{\mathcal{A},C,0}^{\text{klsts-sec}}(\lambda)$ and description of a keyless stegosystem from pseudorandom encodings inspired by [vH04; Hop05].

Security. KlStS is secure (on channel C), if for all PPT adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{klsts-sec}}(\lambda) := |\Pr[Exp_{\mathcal{A},C,0}^{\text{klsts-sec}}(\lambda) = 1] - \Pr[Exp_{\mathcal{A},C,1}^{\text{klsts-sec}}(\lambda) = 1]|$$

is negligible.

Applying a similar strategy as in [vH04; Hop05], we obtain the following theorem.

Theorem 23. Let S be an efficiently samplable message distribution and let $\mathcal{H}$ be a family of pairwise independent hash functions. If $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true for S , then $(\text{KlStS.Setup}_S, \text{KlStS.E}_S, \text{KlStS.D}_S)$ defined above and in Figure 38 is a keyless stegosystem for the message distribution S .

The result basically follows from the Leftover Hash Lemma [HILL99] and the ability to embed the message distribution into the uniform distribution due to $\text{acPREH}_{\approx_c}^{\text{rand}}$. Note that since in $Exp_{\mathcal{A},C,b}^{\text{klsts-sec}}(\lambda)$, the adversary does not know the challenge message m^* , it is not necessary that each encoding of a message corresponds to exactly one stegotext, see [Hop05].

7.4 Covert multi-party computation

The notion of covert secure computation was first introduced in [vHL05] as a collection of desirable properties, such as strong internal covertness, strong fairness and final covertness. Later in [CGOS07], Chandran et al. proposed a more unified, simulation-based definition. We start by recalling the model of [CGOS07].

Ideal Model. We consider n parties, each party P_i holding an input x_i . Let $\mathbf{x}$ be the vector $(x_1, \dots, x_n)$. All the participating parties send their input to the functionality, while the non-participating parties are assumed to send $\perp$. Then, if any of the parties had input a $\perp$, the functionality sets $\perp$ to be the output of the protocol. Else, let $g : (\{0, 1\}^*)^n \mapsto \{0, 1\}$ be the function which determines whether on input $\mathbf{x}$, the output is *favorable* ($g(\mathbf{x}) = 1$) or *non-favorable* ($g(\mathbf{x}) = 0$). In the latter case, the functionality sets $\perp$ to be the output of the protocol. In the former case, the functionality computes the output $f(\mathbf{x})$. The output is sent to any subset of players, chosen by the adversary. The functionality is represented on Figure 39. For an adversary $\mathcal{A}$, an execution of $\mathcal{F}_{\text{CMPC}}(f, g)$ with participation data $\mathbf{p}$ (indicating which players are taking part to the protocol) and input $\mathbf{x}$ (where the input of non-participating parties is $\perp$) is defined as the output of the parties together with the output of the adversary. It is denoted $\text{IDEAL}_{f,g,\mathcal{A}}(\mathbf{p}, \mathbf{x})$.

Fig. 39. Ideal Functionality for Covert Multi-Party Computation.

Real Model. Honest participating parties follow the specifications of the protocol. Honest non-participating parties are assumed to send uniformly random messages. We consider *static semi-honest corruption* of players, in which the corrupted players are chosen once-for-all by the adversary before the start of the protocol, and follow the specifications of the protocol. Honest participating parties compute their output as specified, non-participating and corrupted parties output $\perp$ by convention, and the adversary outputs its entire view of the execution of the protocol. For an adversary $\mathcal{A}$, an execution of a protocol $\Pi(f, g)$ in the real model with participation data $\mathbf{p}$ and input $\mathbf{x}$ is defined as the output of the parties together with the output of the adversary. It is denoted $\text{REAL}_{\Pi(f, g), \mathcal{A}}(\mathbf{p}, \mathbf{x})$.

Definition 34 (Covert security [CGOS07]). *A protocol $\Pi(f, g)$ securely implements $\mathcal{F}_{\text{CMPC}}(f, g)$ if for every probabilistic polynomial-time adversary $\mathcal{A}$ statically corrupting up to $n - 1$ players in the real model, there is an expected polynomial-time adversary S corrupting at most $n - 1$ players in the ideal model, such that for any $(\mathbf{p}, \mathbf{x}) \in \{0, 1\}^n \times (\{0, 1\}^*)^n$,*

$$\{\text{IDEAL}_{f, g, S}(\mathbf{p}, \mathbf{x})\} \approx_c \{\text{REAL}_{\Pi(f, g), \mathcal{A}}(\mathbf{p}, \mathbf{x})\}$$

We now show that pseudorandom encodings enable to convert a large class of secure computation protocols (satisfying the standard static semi-honest security notion) into protocols secure in the covert model of [CGOS07]. Covertness is a very strong security notion, protecting the participants from even leaking the information that they took part to the protocol. Interestingly, since the pseudorandom encoding hypothesis is equivalent to the existence of adaptive MPC for all functionalities, our result create a link between two extreme but seemingly unrelated notions of secure computation: adaptive multi-party computation for all functionalities implies the existence of a generic compiler for covert multi-party computation. We find this link to be intriguing.

7.4.1 A compiler for covert oblivious transfer

As a warm-up, we start from a static, semi-honest oblivious transfer protocol, and show how to compile it into a covert OT using pseudorandom encodings. Consider a two-round OT, as follows:

$\text{ROT}(1^\lambda, b; r_R)$ takes as input a bit b and randomness r_R and outputs the first protocol message OT_1 .

$\text{SOT}(\text{OT}_1, (y_0, y_1); r_S)$ takes as input the first message OT_1 , a tuple $(y_0, y_1) \in (\{0, 1\}^\lambda)^2$ and randomness r_S and outputs the second protocol message OT_2 .

$E_{OT}(OT_2, b, r_R)$ takes as input OT_2 the bit b and the randomness r_R and outputs $y \in \{0, 1\}^\lambda$.

Covert oblivious transfer was defined in [vHL05]; a covert OT protocol realizes the standard OT functionality, and satisfies two additional properties:

- (receiver indistinguishability) For any b , the distribution induced by $R_{OT}(1^\lambda, b)$ is indistinguishable from the uniform distribution (over bistrings of an appropriate size);
- (sender indistinguishability) For any OT_1 , the distribution

$$\{OT_2 \leftarrow S_{OT}(OT_1, (y_0, y_1)) : (y_0, y_1) \leftarrow (\{0, 1\}^n)^2\}$$

is indistinguishable from the uniform distribution (over bistrings of an appropriate size).

Now, let (E_1, D_1) be a pseudorandom encoding for the sampler $R_{OT}(1^\lambda, b)$ which takes as input $(1^\lambda, b)$, and let (E_2, D_2) be a pseudorandom encoding for the sampler $S_{OT}(OT_1, (y_0, y_1))$ which takes as input a string OT_1 and samples (y_0, y_1) at random. Consider the following covert OT protocol, built from an arbitrary two-round OT:

covR_{OT} $(1^\lambda, b; r_R)$ compute $a_1 \leftarrow R_{OT}(1^\lambda, b; r_R^0)$ and output $OT_1 \leftarrow E_1(a_1, 0; r_R^1)$.
covS_{OT} $(OT_1, (y_0, y_1))$ compute $a'_1 \leftarrow D_1(OT_1, 0)$, $a_2 \leftarrow S_{OT}(a'_1, (y_0, y_1))$, and output $OT_2 \leftarrow E_2(a_2, a'_1)$.
covE_{OT} (OT_2, b, r_R) reconstruct (a_1, OT_1) from $r_R = (r_R^0, r_R^1)$. compute $a'_2 \leftarrow D_2(OT_2, a_1)$ and output $E_{OT}(a'_2, b, r_R^0)$.

We now prove that the above protocol realizes the OT functionality and satisfies receiver and sender indistinguishability.

- (receiver indistinguishability) $E_1(R_{OT}(1^\lambda, b), 0) \approx_c E_1(R_{OT}(1^\lambda, 0), 0) \approx_c U$, where U denotes the uniform distribution (over bitstrings of some appropriate length). The first indistinguishability follows from the receiver security of the underlying semi-honest OT protocol, and the second follows from the pseudorandomness of E_1 .
- (sender indistinguishability) follows directly from the pseudorandomness of E_2 .
- (OT functionality) sender security and receiver security are trivially inherited from the underlying semi-honest OT: neither E_1 nor E_2 use any of the private inputs of the parties. It remains to verify that correctness is preserved. This follows from the receiver security of the underlying OT and the correctness of the encoding: it must hold for any b that $D_1(E_1(m, 0), 0) = m$ when $m \leftarrow R_{OT}(1^\lambda, b)$, since correctness of encoding guarantees that this holds when $m \leftarrow R_{OT}(1^\lambda, 0)$; but then, it must hold as well when $b = 1$, otherwise (E_1, D_1) would allow to break the receiver security of the underlying OT. The correctness of **covE_{OT}** follows directly from the correctness of the encoding, which concludes the proof.

7.4.2 A general compiler for covert protocols

Compiling a secure computation protocol into a covert MPC protocol requires some care, since both models have syntactic differences: a covert MPC protocol must, by design, have an admissibility function g defining whether the inputs of the parties are admissible, in addition to the target function f to be computed. We therefore focus on compiling secure protocols which already satisfy this syntactic requirement.

Consider an arbitrary two-party protocol $\Pi_{f,g}$, where A has input x and B has input y , which implements the following functionality $\mathcal{F}_{f,g}$: it first computes an admissibility function $g(x, y)$ on the inputs. If $g(x, y) = 1$, it outputs random shares of the target function $f(x, y)$ to the parties; else, it outputs random values to all parties. We assimilate A (resp. B) to a sampler that take as input x (resp. y) together with the transcript T of the protocol so far, and outputs the next message of A (resp. B). Eventually, we only assume from Π a weak indistinguishability-style security guarantee, namely that for any inputs x, y , no adversary passively corrupting B (resp.

A) can distinguish the transcript of an interaction with A on input x (resp. B on input y) from the transcript of an interaction with A on input 0 (resp. B on input 0). Note that the protocol outputs random shares, so the output distribution of B is independent of A 's input.

Let (E_A, D_A) be a pseudorandom encoding associated to the sampler A , and (E_B, D_B) be a pseudorandom encoding associated to the sampler B . The compiled protocol $\text{cov}\Pi_{f,g}$ is constructed as follows: at round i , given a transcript T up to round $i - 2$, the message m_B received from Bob at round $i - 1$, and an input x , Alice computes $m \leftarrow D_B(\|0)$, appends m to T , samples $m' \leftarrow A(T\|x)$, and sends $m_A \leftarrow E_A(m, T\|0)$. Bob proceeds symmetrically.

Proof Sketch. A straightforward generalization of the argument for the covert OT protocol shows that $\text{cov}\Pi_{f,g}$ satisfies an indistinguishability-style notion of covertness: at each round, it holds that

$$E_A(A(T\|x), T\|0) \approx_c E_A(A(T\|0), T\|0) \approx_c U,$$

where U denotes the uniform distribution over strings of appropriate length; the first indistinguishability follows from the security of $\Pi_{f,g}$, and the second from the pseudorandomness of E_A . Similarly, security of $\Pi_{f,g}$ and correctness of the pseudorandom encodings guarantee the correctness of $\text{cov}\Pi_{f,g}$. If the protocol $\Pi_{f,g}$ further satisfies a simulation-style security notion and securely implements the functionality $\mathcal{F}_{f,g}$, the protocol $\text{cov}\Pi_{f,g}$ can be shown to covertly implement the functionality $\mathcal{F}_{\text{CMPC}}(f, g)$.

7.5 Deniable encryption

As noted in [CDNO97], sender deniable encryption is related to adaptively secure MPC (for adaptively well-formed functionalities). In this section, we analyze if $\text{cPREH}_{\approx_c}^{\text{rand}}$ together with the existence of a PKE scheme suffices to obtain deniable encryption.

Recall, that the explainability compiler of [DKR15] which is based on the deniable encryption scheme of [SW14] corresponds to $\text{cSH}_{\approx_c}^{\text{rand}}$, where closeness actually holds information theoretically. Let cSH' denote this variant of $\text{cSH}_{\approx_c}^{\text{rand}}$, where closeness holds information theoretically.

Definition 35 (Publicly deniable encryption, [SW14]). A publicly deniable encryption scheme for message space $\mathcal{M}$ is a tuple $(\text{Gen}, \text{Enc}, \text{Dec}, \text{Explain})$ such that $(\text{Gen}, \text{Enc}, \text{Dec})$ is an IND-CPA secure encryption scheme and the following property is satisfied.

Indistinguishability of explanation. For all PPT adversaries $\mathcal{A}$,

$$\text{Adv}_{\mathcal{A}}^{\text{ind-expl}}(\lambda) := \left| \Pr[\text{Exp}_{\mathcal{A},0}^{\text{ind-expl}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A},1}^{\text{ind-expl}}(\lambda) = 1] \right|$$

is negligible, where $\text{Exp}_{\mathcal{A},0}^{\text{ind-expl}}$ and $\text{Exp}_{\mathcal{A},1}^{\text{ind-expl}}$ are defined in Figure 40.

$\text{Exp}_{\mathcal{A},0}^{\text{ind-expl}}(\lambda)$	$\text{Exp}_{\mathcal{A},1}^{\text{ind-expl}}(\lambda)$	$\text{Exp}_{\mathcal{A},0}^{\text{ind-expl}'}(\lambda)$	$\text{Exp}_{\mathcal{A},1}^{\text{ind-expl}'}(\lambda)$
$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$m^* \leftarrow \mathcal{A}(1^\lambda)$	$m^* \leftarrow \mathcal{A}(1^\lambda)$
$m^* \leftarrow \mathcal{A}(pk)$	$m^* \leftarrow \mathcal{A}(pk)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$
$u^* \leftarrow \{0, 1\}^*$	$u^* \leftarrow \{0, 1\}^*$	$u^* \leftarrow \{0, 1\}^*$	$u^* \leftarrow \{0, 1\}^*$
$c^* \leftarrow \text{Enc}(pk, m^*; u^*)$	$c^* \leftarrow \text{Enc}(pk, m^*; u^*)$	$c^* \leftarrow \text{Enc}(pk, m^*; u^*)$	$c^* \leftarrow \text{Enc}(pk, m^*; u^*)$
return $\mathcal{A}(c^*, u^*)$	$r^* \leftarrow \{0, 1\}^*$	return $\mathcal{A}(pk, c^*, u^*)$	$r^* \leftarrow \{0, 1\}^*$
	$e^* \leftarrow \text{Explain}(pk, c^*; r^*)$		$e^* \leftarrow \text{Explain}(pk, c^*; r^*)$
	return $\mathcal{A}(c^*, e^*)$		return $\mathcal{A}(pk, c^*, e^*)$

Fig. 40. Indistinguishability of explanation experiments in its adaptive variant (left) and static variant (right).

We only consider publicly deniable encryption schemes with message space $\{0, 1\}$, since a deniable encryption scheme with message space $\{0, 1\}$ implies a publicly deniable encryption scheme for message space $\{0, 1\}^n$ (for a polynomial n in λ).

For message space $\{0, 1\}$, indistinguishability of explanation is equivalent to the indistinguishability of $\text{Exp}_{\mathcal{A},0}^{\text{ind-expl}'}$ and $\text{Exp}_{\mathcal{A},1}^{\text{ind-expl}'}$ as defined in Figure 40.

clSH' in conjunction with a PKE scheme yields (publicly) sender deniable encryption.

Theorem 24. *Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme for message space $\{0, 1\}$. If clSH' holds, then there exists a publicly deniable encryption scheme for message space $\{0, 1\}$.*

Proof. clSH' implies that for the PPT algorithm Enc there is a setup algorithm $\text{Setup}_{\text{Enc}}$, an alternative sampler $\overline{\text{Enc}}$ and an inverse sampler $\overline{\text{Enc}}^{-1}$ satisfying statistical closeness and computational invertibility. We define a publicly deniable encryption scheme $(\text{Gen}', \text{Enc}', \text{Dec}', \text{Explain}')$ in Figure 41.

$\text{Gen}'(1^\lambda)$	$\text{Enc}'(pk', m)$	$\text{Dec}'(sk, c)$	$\text{Explain}'(pk', m)$
$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$	parse $(crs, pk) =: pk'$	$m \leftarrow \text{Dec}(sk, c)$	parse $(crs, pk) =: pk'$
$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$c \leftarrow \overline{\text{Enc}}(crs, (pk, m))$	return m	$u \leftarrow \overline{\text{Enc}}^{-1}(crs, (pk, m), c)$
$pk' := (crs, pk), sk' := sk$	return c		return u
return (pk', sk')			

Fig. 41. Publicly deniable encryption scheme from clSH' .

Correctness. Let $m \in \{0, 1\}^*$ be a plaintext.

$$\begin{aligned} \epsilon_1 &:= \Pr[(crs, pk), sk) = (pk', sk') \leftarrow \text{Gen}'(1^\lambda), c \leftarrow \text{Enc}(pk, m) : \text{Dec}(sk, c) \neq m] \\ \epsilon_2 &:= \Pr[(crs, pk), sk) = (pk', sk') \leftarrow \text{Gen}'(1^\lambda), c \leftarrow \overline{\text{Enc}}(pk, m) : \text{Dec}(sk, c) \neq m] \end{aligned}$$

Consider an *unbounded* adversary $\mathcal{A}$ on closeness which on input of (crs, pk, c) , computes sk (using exhaustive search) and outputs 1 if $\text{Dec}(sk, c) \neq m$ and 0 otherwise. The advantage of this adversary is $\text{Adv}_{\mathcal{A},m}^{\text{crs-close}}(\lambda) = |\epsilon_1 - \epsilon_2|$. Hence, $\epsilon_2 \leq \text{Adv}_{\mathcal{A},m}^{\text{crs-close}}(\lambda) + \epsilon_1$ and therefore negligible due to statistical closeness and correctness of $(\text{Gen}, \text{Enc}, \text{Dec})$.

IND-CPA security. For message space $\{0, 1\}$, IND-CPA security of $(\text{Gen}', \text{Enc}', \text{Dec}')$ is equivalent to the indistinguishability between the games $\mathbf{G}_0$ and $\mathbf{G}_3$ of Figure 42.

$\mathbf{G}_0$	$\mathbf{G}_1$	$\mathbf{G}_2$	$\mathbf{G}_3$
$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$	$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$	$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$	$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$
$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$
$c^* \leftarrow \overline{\text{Enc}}(pk, 0)$	$c^* \leftarrow \overline{\text{Enc}}(pk, 0)$	$c^* \leftarrow \text{Enc}(pk, \mathbb{1})$	$c^* \leftarrow \overline{\text{Enc}}(pk, 1)$
return $\mathcal{A}((crs, pk), c^*)$	return $\mathcal{A}((crs, pk), c^*)$	return $\mathcal{A}((crs, pk), c^*)$	return $\mathcal{A}((crs, pk), c^*)$

Fig. 42. Hybrids used in the proof of IND-CPA security for Theorem 24.

Claim. For all (unbounded) adversaries $\mathcal{A}$ there exists an (unbounded) adversary $\overline{\mathcal{A}}$ such that $|\Pr[out_1 = 1] - \Pr[out_0 = 1]| \leq \text{Adv}_{\mathcal{A}}^{\text{crs-close}}(\lambda)$.

Proof. Let $\mathcal{A}$ be an unbounded adversary distinguishing $\mathbf{G}_0$ and $\mathbf{G}_1$. Construct an adversary $\overline{\mathcal{A}}$ on closeness. Initially, $\overline{\mathcal{A}}$ samples $(pk, sk) \leftarrow \text{Gen}(1^\lambda)$ and outputs $m := (pk, 0)$ to the experiment. In return, $\overline{\mathcal{A}}$ receives (crs, y) , where y is either sampled using $\text{Enc}(m)$ or $\overline{\text{Enc}}(m)$. $\overline{\mathcal{A}}$ calls $\mathcal{A}$ on input of $((crs, pk), y)$. Hence, $\Pr[out_b = 1] = \Pr[\text{Exp}_{\mathcal{A},1-b}^{\text{crs-close}}(\lambda) = 1]$ for $b \in \{0, 1\}$. $\square$

The game hop from $\mathbf{G}_1$ to $\mathbf{G}_2$ is justified by the IND-CPA security of $(\text{Gen}, \text{Enc}, \text{Dec})$. The game hop from $\mathbf{G}_2$ to $\mathbf{G}_3$ is justified by statistical closeness.

Indistinguishability of explanation. We need to prove indistinguishability between the games described in Figure 43.

Claim. For all PPT adversaries $\mathcal{A}$, $\text{Adv}_{\mathcal{A}}^{\text{ind-expl}'}(\lambda)$ is negligible.

$Exp_{\mathcal{A},0}^{\text{ind-expl}'}(\lambda)$	$Exp_{\mathcal{A},1}^{\text{ind-expl}'}(\lambda)$
$m^* \leftarrow \mathcal{A}(1^\lambda)$	$m^* \leftarrow \mathcal{A}(1^\lambda)$
$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$	$crs \leftarrow \text{Setup}_{\text{Enc}}(1^\lambda)$
$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$	$(pk, sk) \leftarrow \text{Gen}(1^\lambda)$
$u^* \leftarrow \{0,1\}^*$	$u^* \leftarrow \{0,1\}^*$
$c^* \leftarrow \overline{\text{Enc}}(pk, m^*; u^*)$	$c^* \leftarrow \overline{\text{Enc}}(pk, m^*; u^*)$
return $\mathcal{A}((crs, pk), c^*, u^*)$	$r^* \leftarrow \{0,1\}^*$
	$e^* \leftarrow \overline{\text{Enc}}^{-1}(pk, c^*; r^*)$
	return $\mathcal{A}((crs, pk), c^*, e^*)$

Fig. 43. Unwrapped indistinguishability of explanation games used in the proof of Theorem 24.

Proof. Let $\mathcal{A}$ be an adversary distinguishing $Exp_{\mathcal{A},0}^{\text{ind-expl}'}$ and $Exp_{\mathcal{A},1}^{\text{ind-expl}'}$ above. Construct an adversary $\overline{\mathcal{A}}$ on invertibility. Initially, $\overline{\mathcal{A}}$ calls $\mathcal{A}$, obtains m^* , samples $(pk, sk) \leftarrow \text{Gen}(1^\lambda)$ and outputs $m := (pk, m^*)$ to the experiment. In return, $\overline{\mathcal{A}}$ receives (crs, r, y) from the experiment, where y is sampled using $\overline{\text{Enc}}(pk, m^*)$ and r either is the randomness used or itself sampled from $\overline{\text{Enc}}^{-1}(pk, m^*, y)$. $\overline{\mathcal{A}}$ calls $\mathcal{A}$ on input of $((crs, pk), y, r)$. Hence, $\Pr[Exp_{\mathcal{A},b}^{\text{ind-expl}'}(\lambda) = 1] = \Pr[Exp_{\mathcal{A},b}^{\text{crs-inv}}(\lambda) = 1]$ for $b \in \{0,1\}$. $\square$

$\square$

Whether $\text{clSH}_{\approx_c}^{\text{rand}}$ (without statistical closeness) implies deniable encryption remains open.

7.6 Non-committing encryption

Non-committing encryption is a powerful notion which is known to imply adaptive MPC for well-formed functionalities [CFGN96; CLOS02].

On a high level, if $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true, then any IND-CPA secure PKE scheme is a simulatable PKE scheme [DN00] (in the CRS model). The strategy of [CDMW09] translates to the CRS model (where inputs may be chosen depending on the CRS) directly.

Definition 36 (Non-committing bit encryption encryption in the global CRS model).

A non-committing bit encryption scheme in the global CRS model is a tuple of PPT algorithms $(\text{Setup}, \text{Gen}, \text{Enc}, \text{Dec}, \text{Sim})$, such that $(\text{Gen}, \text{Enc}, \text{Dec})$ is a PKE scheme and the following distributions are computationally indistinguishable.

$$\begin{aligned} &\{crs \leftarrow \text{Setup}(1^\lambda), (pk, sk) := \text{Gen}(1^\lambda, crs; r_{\text{Gen}}), c := \text{Enc}(crs, pk, b; r_{\text{Enc}}) : (crs, pk, c, r_{\text{Gen}}, r_{\text{Enc}})\} \\ &\{crs \leftarrow \text{Setup}(1^\lambda), (pk, c, r_{\text{Gen}}^0, r_{\text{Gen}}^1, r_{\text{Enc}}^0, r_{\text{Enc}}^1) \leftarrow \text{Sim}(1^\lambda; crs) : (crs, pk, c, r_{\text{Gen}}^b, r_{\text{Enc}}^b)\} \end{aligned}$$

Following the lines of [CDMW09], if $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true and there exists an IND-CPA secure PKE scheme, then there exists a non-committing encryption scheme.

Theorem 25. Let $(\text{Gen}', \text{Enc}', \text{Dec}')$ be an IND-CPA secure PKE scheme for message space $\{0,1\}^\lambda$. If $\text{acPREH}_{\approx_c}^{\text{rand}}$ holds, then there exists a non-committing encryption scheme.

Proof (sketch). Let OGen be the algorithm which on input of 1^λ calls $\text{Gen}'(1^\lambda)$ and outputs only $(pk, \perp)$. Further let OEnc be the algorithm which on input of $(1^\lambda, pk)$ samples $m \leftarrow \{0,1\}^\lambda$ and outputs $c \leftarrow \text{Enc}'(1^\lambda, pk, m)$. If $\text{acPREH}_{\approx_c}^{\text{rand}}$ is true, then there exists alternative and inverse sampler for OGen and OEnc , denoted by $(\overline{\text{OGen}}, \overline{\text{OGen}}^{-1})$ and $(\overline{\text{OEnc}}, \overline{\text{OEnc}}^{-1})$, respectively. Note that the alternative sampler and the inverse sampler need access to the CRS. If the setup algorithm Setup is trivial, this corresponds to the notion of simulatable encryption due to [DN00] and yields non-committing encryption directly due to [CDMW09]. Figures 44 and 45 shows the construction of non-committing encryption in the global CRS model. Note that assuming

$\text{Gen}(1^\lambda, \text{crs})$	$\text{Enc}(\text{crs}, pk, b)$	$\text{Dec}(\text{crs}, sk, c)$
$M_0, M_1 \leftarrow \{0, 1\}^\lambda$ $T \leftarrow [4\lambda]$ s.t. $ T = \lambda$ $(pk_i, sk_i) \leftarrow \begin{cases} \text{Gen}'(1^\lambda) & \text{if } i \in T \\ \text{OGen}(\text{crs}) & \text{otherwise} \end{cases}$ $pk := (M_0, M_1, pk_1, \dots, pk_{4\lambda})$ $sk := (T, (sk_i)_{i \in T})$ return (pk, sk)	$S \leftarrow [4\lambda]$ s.t. $ S = \lambda$ $c_i \leftarrow \begin{cases} \text{Enc}'(pk_i, M_b) & \text{if } i \in S \\ \text{OEnc}(\text{crs}, pk_i) & \text{otherwise} \end{cases}$ return $c := (c_i)_{i \in [4\lambda]}$	$J := \{\text{Dec}'(sk_i, c_i) \mid i \in T\}$ if $M_0 \in J$ then return 0 else return 1

Fig. 44. Non-committing encryption scheme in the global CRS model based on [CDMW09].

$\text{Sim}(\text{crs})$
$M_0, M_1 \leftarrow \{0, 1\}^\lambda$ $S_0, T_0 \leftarrow [4\lambda]$ s.t. $ S_0 = T_0 = \lambda$ $S_1, T_1 \leftarrow [4\lambda] \setminus (S_0 \cup T_0)$ s.t. $ S_0 \cap T_0 = S_1 \cap T_1 $ $(pk_i, sk_i) \leftarrow \begin{cases} \text{Gen}'(1^\lambda; r_{\text{Gen}}^{(i)}) & \text{if } i \in T_0 \cup S_0 \cup T_1 \cup S_1 \\ \text{OGen}(\text{crs}; r_{\text{OGen}}^{(i)}) & \text{otherwise} \end{cases}$ $c_i \leftarrow \begin{cases} \text{Enc}'(pk_i, M_0; r_{\text{Enc}'}^{(i)}) & \text{if } i \in S_0 \\ \text{Enc}'(pk_i, M_1; r_{\text{Enc}'}^{(i)}) & \text{if } i \in S_1 \\ \text{OEnc}(\text{crs}, pk_i; r_{\text{OEnc}}^{(i)}) & \text{otherwise} \end{cases}$ define $r_{\text{Gen}}^b := (T_b, (u_{\text{Gen}'}^{(b,i)})_{i \in [4\lambda]})$ and $r_{\text{Enc}}^b := (S_b, (u_{\text{Enc}'}^{(b,i)})_{i \in [4\lambda]})$ $u_{\text{Gen}'}^{(b,i)} \leftarrow \begin{cases} r_{\text{Gen}'}^{(i)} & \text{if } i \in T_b \\ \text{OGen}^{-1}(\text{crs}, pk_i) & \text{if } i \in (T_0 \cup T_1 \cup S_0 \cup S_1) \setminus T_b \\ r_{\text{OGen}}^{(i)} & \text{otherwise} \end{cases}$ $u_{\text{Enc}'}^{(b,i)} \leftarrow \begin{cases} r_{\text{Enc}'}^{(i)} & \text{if } i \in S_b \\ \text{OEnc}^{-1}(\text{crs}, (pk_i, M_{1-b}), c_i) & \text{if } i \in S_{1-b} \\ r_{\text{OEnc}}^{(i)} & \text{otherwise} \end{cases}$ return $(pk := (M_0, M_1, (pk_i)_{i \in [4\lambda]}), c := (c_i)_{i \in [4\lambda]}, r_{\text{Gen}}^0, r_{\text{Gen}}^1, r_{\text{Enc}}^0, r_{\text{Enc}}^1)$

Fig. 45. The simulator Sim for the non-committing encryption scheme in the global CRS model described in Figure 44. Sim is based on the simulator given in [CDMW09].

$\text{adaptive acPREH}_{\approx_c}^{\text{rand}}$ is necessary since the inputs to the sampler OEnc are sampled via $\text{OGen}(\text{crs})$ during the simulation and, hence, depend on the CRS. The indistinguishability between the real and the simulated distribution follows the same ideas as in [CDMW09]. $\square$

7.7 Super-polynomial encoding

Extremely lossy functions due to [Zha16] are functions which can be set up in two computationally indistinguishable modes – an injective mode and a extremely lossy mode, where the range of the function is merely polynomial. A slight relaxation of this notion is what we call very lossy functions (VLFs). The difference to ELFs is that we require indistinguishability between functions with exponential range and *super-polynomial* range. The existence of these functions implies a relaxation of $\text{cPREH}_{\approx_c}^{\text{rand}}$.

Definition 37 (Very lossy function). A Very lossy function consists of an algorithm $\text{VLF.Gen}'$ and a computable function $N(M)$ such that $\log N$ is polynomial in $\log M$. $\text{VLF.Gen}'$ takes as input natural numbers $M, r \in [M]$ and a flag $b \in \{\text{inj}, \text{lossy}\}$, and outputs the description of a function $f: [M] \rightarrow [N]$ such that

- f is computable in time polynomial in $\log M$,
- if $b = \text{inj}$, f is injective with overwhelming probability (in $\log M$),
- for all $r \in [M]$, if $b = \text{lossy}$, $|f([M])| \leq r$ with overwhelming probability (in $\log M$),

- there exists a super-polynomial function q such that for all PPT adversaries $\mathcal{A}$ and any $r \in [q(\log M), M]$,

$$|\Pr[\mathcal{A}(\text{VLF.Gen}'(M, r, \text{inj})) = 1] - \Pr[\mathcal{A}(\text{VLF.Gen}'(M, r, \text{lossy})) = 1]|$$

is negligible.

Definition 38 (Strong regularity, [Zha16]). A VLF VLF is strongly regular if for all $r \in [M]$, with overwhelming probability over the choice of $f \leftarrow \text{VLF.Gen}(M, r)$ we have that the distribution $\{x \leftarrow [M]: f(x)\}$ is statistically close to uniform distribution over $f([M])$.

A VLF is called strongly efficiently enumerable, if there exists a (potentially randomized) algorithm running in polynomial time in $\log M$ and r which given $f \leftarrow \text{VLF.Gen}(M, r, \text{lossy})$ and r outputs a set $S \subseteq [N]$ such that with overwhelming probability (over the choice of f and the randomness of the algorithm), $S = f([M])$. If VLF is strongly regular, then it also is strongly efficiently enumerable, [Zha16]. [Zha16] shows that strongly regular ELF's are implied by the exponential decisional Diffie-Hellman (DDH) assumption.

We note that in contrast to ELF's, an adversary learning r does not harm security.

Assuming the sub-exponential hardness of the decisional Diffie-Hellman (DDH) problem, the bounded adversary extremely lossy function instantiation from [Zha16] is a strongly regular very lossy function according to Definitions 37 and 38.

Theorem 26. If strongly regular very lossy functions exist, then $\text{acPREH}_{\approx_c}^{\text{rand}}$ with super-polynomial encoding is true.

Proof. Let S be a PPT sampler with input space L . For $m \in L$, let $\ell_r(|m|)$ denote the polynomial which upper bounds the number of random bits $S(m)$ takes, i.e. the random tape of S is uniform over $\{0, 1\}^{\ell_r} = [2^{\ell_r}]$. Further, let ξ be a super-polynomial function. The setup algorithm, the encoding and decoding algorithms are defined in Figure 46. We prove the equivalent properties

Setup($1^\lambda, 2^{\ell_r}$)	$E_S(\text{crs}, m, y)$	$D_S(\text{crs}, m, r)$
$f \leftarrow \text{ELF.Gen}(2^{\ell_r}, \xi, \text{lossy})$	$\mathcal{R} := \emptyset$	$y := S(m; f(r))$
return $\text{crs} := f$	for $r' \in \text{image}(f)$ do	return y
	if $S(m; f(r')) = y$ then	
	$\mathcal{R} += r'$	
	$r \leftarrow \mathcal{R}$	
	return r	

Fig. 46. Description of the universal setup algorithm, the super-polynomial time encoding algorithm and the polynomial time decoding algorithm.

closeness and invertibility.

Invertibility. The encoding algorithm produces perfectly distributed inverse sampled random tapes. Hence, adaptive invertibility follows.

Closeness. We start from the game $\text{Exp}_{\mathcal{A}}^{\text{a-crs-close}}$ and switch the VLF to injective mode $f \leftarrow \text{VLF.Gen}(2^{\ell_r}, \xi, \text{inj})$. This is computationally indistinguishable for the adversary. The inverse sampler will not work anymore, but since the adversary is polynomially bounded, he can not call the inverse sampler anyway. Strong regularity implies that for $f \leftarrow \text{VLF.Gen}(2^{\ell_r}, \xi, \text{inj})$ the distribution $\{x \leftarrow [2^{\ell_r}]: f(x)\}$ is statistically close to uniform distribution over $f([2^{\ell_r}])$. Hence, $S(m; f(r))$ and $S(m; r)$ for uniform r from $[2^{\ell_r}]$ are statistically close. $\square$

Unfortunately, $\text{acPREH}_{\approx_c}^{\text{rand}}$ with super-polynomial encoding algorithm (or, equivalently $\text{acISH}_{\approx_c}^{\text{rand}}$ with super-polynomial inverse sampler) does not suffice to imply adaptive MPC even if the simulator is allowed to run in super-polynomial time, [Pas03]. This is because plugging $\text{cPREH}_{\approx_c}^{\text{rand}}$ with super-polynomial encoding algorithm into the proof of Theorem 19, the game hop from $\mathbf{G}_2$ to $\mathbf{G}_3$ can not be made, since the simulation of these games requires super-

polynomial time and, hence, a reduction to closeness against PPT adversaries is not possible. We do, however, obtain a non-standard notion of adaptive MPC with super-polynomial simulation, namely for any functionality $\mathcal{F}$ we are able to adaptively realize a functionality $\overline{\mathcal{F}}$ which produces a computationally (against PPT adversaries) indistinguishable output distribution to the original functionality. We view it as an interesting problem to further study this notion of adaptive MPC.

Acknowledgments. We thank Daniel Wichs for suggesting the unconditional static-to-adaptive transformation, as well as anonymous reviewers for extremely useful feedback.

Bibliography

- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. “On the existence of extractable one-way functions”. In: *46th Annual ACM Symposium on Theory of Computing*. Ed. by David B. Shmoys. New York, NY, USA: ACM Press, 2014, pp. 505–514. DOI: [10.1145/2591796.2591859](https://doi.org/10.1145/2591796.2591859) (cit. on pp. 5, 6, 10, 38, 39, 45).
- [BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. “Candidate iO from Homomorphic Encryption Schemes”. In: *EUROCRYPT 2020, Part I*. Springer, 2020, pp. 79–109. DOI: [10.1007/978-3-030-45721-1_4](https://doi.org/10.1007/978-3-030-45721-1_4) (cit. on p. 10).
- [BGIRSVY01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. “On the (Im)possibility of Obfuscating Programs”. In: *Advances in Cryptology – CRYPTO 2001*. Ed. by Joe Kilian. Vol. 2139. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2001, pp. 1–18. DOI: [10.1007/3-540-44647-8_1](https://doi.org/10.1007/3-540-44647-8_1) (cit. on p. 5).
- [BH92] Donald Beaver and Stuart Haber. “Cryptographic Protocols Provably Secure Against Dynamic Adversaries”. In: *EUROCRYPT 1992*. 1992, pp. 307–323 (cit. on p. 7).
- [BL18] Sebastian Berndt and Maciej Liskiewicz. “On the Gold Standard for Security of Universal Steganography”. In: *Advances in Cryptology – EUROCRYPT 2018, Part I*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10820. Lecture Notes in Computer Science. Tel Aviv, Israel: Springer, Heidelberg, Germany, 2018, pp. 29–60. DOI: [10.1007/978-3-319-78381-9_2](https://doi.org/10.1007/978-3-319-78381-9_2) (cit. on p. 58).
- [BM92] Steven M. Bellovin and Michael Merritt. “Encrypted Key Exchange: Password-Based Protocols Secure against Dictionary Attacks”. In: *1992 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, May 1992, pp. 72–84. DOI: [10.1109/RISP.1992.213269](https://doi.org/10.1109/RISP.1992.213269) (cit. on pp. 1, 3).
- [BMN01] Colin Boyd, Paul Montague, and Khanh Quoc Nguyen. “Elliptic Curve Based Password Authenticated Key Exchange Protocols”. In: *Information Security and Privacy, 6th Australasian Conference, ACISP 2001, Sydney, Australia, July 11-13, 2001, Proceedings*. Ed. by Vijay Varadharajan and Yi Mu. Vol. 2119. Lecture Notes in Computer Science. Springer, 2001, pp. 487–501. DOI: [10.1007/3-540-47719-5_38](https://doi.org/10.1007/3-540-47719-5_38) (cit. on p. 4).
- [BSW03] Boaz Barak, Ronen Shaltiel, and Avi Wigderson. “Computational Analogues of Entropy”. In: *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, 6th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2003 and 7th International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM 2003, Princeton, NJ, USA, August 24-26, 2003, Proceedings*. Ed. by Sanjeev Arora, Klaus Jansen, José D. P. Rolim, and Amit Sahai. Vol. 2764. Lecture Notes in Computer Science. Springer, 2003, pp. 200–215. ISBN: 3-540-40770-7. DOI: [10.1007/978-3-540-45198-3_18](https://doi.org/10.1007/978-3-540-45198-3_18). URL: https://doi.org/10.1007/978-3-540-45198-3_18 (cit. on p. 30).
- [Can00] Ran Canetti. “Security and Composition of Multiparty Cryptographic Protocols”. In: *Journal of Cryptology* 13.1 (Jan. 2000), pp. 143–202. DOI: [10.1007/s001459910006](https://doi.org/10.1007/s001459910006) (cit. on pp. 51, 52, 55).
- [Can01] Ran Canetti. “Universally Composable Security: A New Paradigm for Cryptographic Protocols”. In: *42nd Annual Symposium on Foundations of Computer*

- Science*. Las Vegas, NV, USA: IEEE Computer Society Press, 2001, pp. 136–145. DOI: [10.1109/SFCS.2001.959888](https://doi.org/10.1109/SFCS.2001.959888) (cit. on p. 53).
- [CDMW09] Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. “Improved Non-committing Encryption with Applications to Adaptively Secure Protocols”. In: *Advances in Cryptology – ASIACRYPT 2009*. Ed. by Mitsuru Matsui. Vol. 5912. Lecture Notes in Computer Science. Tokyo, Japan: Springer, Heidelberg, Germany, 2009, pp. 287–302. DOI: [10.1007/978-3-642-10366-7_17](https://doi.org/10.1007/978-3-642-10366-7_17) (cit. on pp. 6, 64, 65).
- [CDNO97] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. “Deniable Encryption”. In: *Advances in Cryptology – CRYPTO’97*. Ed. by Burton S. Kaliski Jr. Vol. 1294. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 1997, pp. 90–104. DOI: [10.1007/BFb0052229](https://doi.org/10.1007/BFb0052229) (cit. on pp. 7, 62).
- [CDPW07] Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. “Universally Composable Security with Global Setup”. In: *TCC 2007: 4th Theory of Cryptography Conference*. Ed. by Salil P. Vadhan. Vol. 4392. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, 2007, pp. 61–85. DOI: [10.1007/978-3-540-70936-7_4](https://doi.org/10.1007/978-3-540-70936-7_4) (cit. on p. 52).
- [CFGN96] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. “Adaptively Secure Multi-Party Computation”. In: *28th Annual ACM Symposium on Theory of Computing*. Philadelphia, PA, USA: ACM Press, 1996, pp. 639–648. DOI: [10.1145/237814.238015](https://doi.org/10.1145/237814.238015) (cit. on pp. 4, 7, 8, 64).
- [CGOS07] Nishanth Chandran, Vipul Goyal, Rafail Ostrovsky, and Amit Sahai. “Covert Multi-Party Computation”. In: *48th Annual Symposium on Foundations of Computer Science*. Providence, RI, USA: IEEE Computer Society Press, 2007, pp. 238–248. DOI: [10.1109/FOCS.2007.21](https://doi.org/10.1109/FOCS.2007.21) (cit. on pp. 8, 59, 60).
- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. “Universally composable two-party and multi-party secure computation”. In: *34th Annual ACM Symposium on Theory of Computing*. Montréal, Québec, Canada: ACM Press, 2002, pp. 494–503. DOI: [10.1145/509907.509980](https://doi.org/10.1145/509907.509980) (cit. on pp. 7, 51, 64).
- [CPR17] Ran Canetti, Oxana Poburinnaya, and Mariana Raykova. “Optimal-Rate Non-Committing Encryption”. In: *Advances in Cryptology – ASIACRYPT 2017, Part III*. Ed. by Tsuyoshi Takagi and Thomas Peyrin. Vol. 10626. Lecture Notes in Computer Science. Hong Kong, China: Springer, Heidelberg, Germany, 2017, pp. 212–241. DOI: [10.1007/978-3-319-70700-6_8](https://doi.org/10.1007/978-3-319-70700-6_8) (cit. on pp. 3, 12, 20).
- [CPV17] Ran Canetti, Oxana Poburinnaya, and Muthuramakrishnan Venkitasubramanian. “Better Two-Round Adaptive Multi-party Computation”. In: *PKC 2017: 20th International Conference on Theory and Practice of Public Key Cryptography, Part II*. Ed. by Serge Fehr. Vol. 10175. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, 2017, pp. 396–427. DOI: [10.1007/978-3-662-54388-7_14](https://doi.org/10.1007/978-3-662-54388-7_14) (cit. on pp. 3, 12, 20).
- [CsW19] Ran Cohen, abhi shelat, and Daniel Wichs. “Adaptively Secure MPC with Sublinear Communication Complexity”. In: *Advances in Cryptology – CRYPTO 2019, Part II*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11693. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2019, pp. 30–60. DOI: [10.1007/978-3-030-26951-7_2](https://doi.org/10.1007/978-3-030-26951-7_2) (cit. on pp. 2, 6, 7, 29).
- [DKR15] Dana Dachman-Soled, Jonathan Katz, and Vanishree Rao. “Adaptively Secure, Universally Composable, Multiparty Computation in Constant Rounds”. In: *TCC 2015: 12th Theory of Cryptography Conference, Part II*. Ed. by Yevgeniy Dodis and Jesper Buus Nielsen. Vol. 9015. Lecture Notes in Computer Science.

- Warsaw, Poland: Springer, Heidelberg, Germany, 2015, pp. 586–613. DOI: [10.1007/978-3-662-46497-7_23](https://doi.org/10.1007/978-3-662-46497-7_23) (cit. on pp. 5–7, 29, 38, 46–48, 54, 62).
- [DN00] Ivan Damgård and Jesper Buus Nielsen. “Improved Non-committing Encryption Schemes Based on a General Complexity Assumption”. In: *Advances in Cryptology – CRYPTO 2000*. Ed. by Mihir Bellare. Vol. 1880. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2000, pp. 432–450. DOI: [10.1007/3-540-44598-6_27](https://doi.org/10.1007/3-540-44598-6_27) (cit. on pp. 4, 11, 23, 64).
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. “Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data”. In: *SIAM J. Comput.* 38.1 (2008), pp. 97–139. DOI: [10.1137/060651380](https://doi.org/10.1137/060651380). URL: <https://doi.org/10.1137/060651380> (cit. on p. 30).
- [DRV12] Yevgeniy Dodis, Thomas Ristenpart, and Salil P. Vadhan. “Randomness Condensers for Efficiently Samplable, Seed-Dependent Sources”. In: *TCC 2012: 9th Theory of Cryptography Conference*. Ed. by Ronald Cramer. Vol. 7194. Lecture Notes in Computer Science. Taormina, Sicily, Italy: Springer, Heidelberg, Germany, 2012, pp. 618–635. DOI: [10.1007/978-3-642-28914-9_35](https://doi.org/10.1007/978-3-642-28914-9_35) (cit. on p. 1).
- [EGL85] Shimon Even, Oded Goldreich, and Abraham Lempel. “A Randomized Protocol for Signing Contracts”. In: *Commun. ACM* 28.6 (1985), pp. 637–647. DOI: [10.1145/3812.3818](https://doi.org/10.1145/3812.3818) (cit. on p. 10).
- [GGHRSW13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. “Candidate Indistinguishability Obfuscation and Functional Encryption for all Circuits”. In: *54th Annual Symposium on Foundations of Computer Science*. Berkeley, CA, USA: IEEE Computer Society Press, 2013, pp. 40–49. DOI: [10.1109/FOCS.2013.13](https://doi.org/10.1109/FOCS.2013.13) (cit. on p. 5).
- [GKMRV00] Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. “The Relationship between Public Key Encryption and Oblivious Transfer”. In: *FOCS 2000*. 2000, pp. 325–335. DOI: [10.1109/SFCS.2000.892121](https://doi.org/10.1109/SFCS.2000.892121) (cit. on pp. 4, 10).
- [GL91] Shafi Goldwasser and Leonid A. Levin. “Fair Computation of General Functions in Presence of Immoral Majority”. In: *Advances in Cryptology – CRYPTO’90*. Ed. by Alfred J. Menezes and Scott A. Vanstone. Vol. 537. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 1991, pp. 77–93. DOI: [10.1007/3-540-38424-3_6](https://doi.org/10.1007/3-540-38424-3_6) (cit. on p. 54).
- [GS12] Sanjam Garg and Amit Sahai. “Adaptively Secure Multi-Party Computation with Dishonest Majority”. In: *Advances in Cryptology – CRYPTO 2012*. Ed. by Reihaneh Safavi-Naini and Ran Canetti. Vol. 7417. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2012, pp. 105–123. DOI: [10.1007/978-3-642-32009-5_8](https://doi.org/10.1007/978-3-642-32009-5_8) (cit. on p. 7).
- [GS85] Andrew V. Goldberg and Michael Sipser. “Compression and Ranking”. In: *17th Annual ACM Symposium on Theory of Computing*. Providence, RI, USA: ACM Press, 1985, pp. 440–448. DOI: [10.1145/22145.22194](https://doi.org/10.1145/22145.22194) (cit. on pp. 1, 3, 9).
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. “A Pseudorandom Generator from any One-way Function”. In: *SIAM J. Comput.* 28.4 (1999), pp. 1364–1396. DOI: [10.1137/S0097539793244708](https://doi.org/10.1137/S0097539793244708). URL: <https://doi.org/10.1137/S0097539793244708> (cit. on pp. 9, 13, 30, 59).
- [HJKSWZ16] Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. “How to Generate and Use Universal Samplers”. In: *Advances in Cryptology – ASIACRYPT 2016, Part II*. Ed. by Jung Hee Cheon and Tsuyoshi Takagi. Vol. 10032. Lecture Notes in Computer Science. Hanoi,

- Vietnam: Springer, Heidelberg, Germany, 2016, pp. 715–744. DOI: [10.1007/978-3-662-53890-6_24](https://doi.org/10.1007/978-3-662-53890-6_24) (cit. on pp. 3, 12, 20).
- [HLR07] Chun-Yuan Hsiao, Chi-Jen Lu, and Leonid Reyzin. “Conditional Computational Entropy, or Toward Separating Pseudoentropy from Compressibility”. In: *Advances in Cryptology – EUROCRYPT 2007*. Ed. by Moni Naor. Vol. 4515. Lecture Notes in Computer Science. Barcelona, Spain: Springer, Heidelberg, Germany, 2007, pp. 169–186. DOI: [10.1007/978-3-540-72540-4_10](https://doi.org/10.1007/978-3-540-72540-4_10) (cit. on pp. 1, 6, 9, 14, 30, 31, 37).
- [Hop05] Nicholas Hopper. “On Steganographic Chosen Coverttext Security”. In: *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*. 2005, pp. 311–323. DOI: [10.1007/11523468_26](https://doi.org/10.1007/11523468_26). URL: https://doi.org/10.1007/11523468_26 (cit. on p. 59).
- [HPRV19] Thibaut Horel, Sunoo Park, Silas Richelson, and Vinod Vaikuntanathan. “How to Subvert Backdoored Encryption: Security Against Adversaries that Decrypt All Ciphertexts”. In: *ITCS 2019: 10th Innovations in Theoretical Computer Science Conference*. Ed. by Avrim Blum. Vol. 124. San Diego, CA, USA: LIPIcs, 2019, 42:1–42:20. DOI: [10.4230/LIPIcs.ITCS.2019.42](https://doi.org/10.4230/LIPIcs.ITCS.2019.42) (cit. on p. 1).
- [IKOS10] Yuval Ishai, Abishek Kumarasubramanian, Claudio Orlandi, and Amit Sahai. “On Invertible Sampling and Adaptive Security”. In: *Advances in Cryptology – ASIACRYPT 2010*. Ed. by Masayuki Abe. Vol. 6477. Lecture Notes in Computer Science. Singapore: Springer, Heidelberg, Germany, 2010, pp. 466–482. DOI: [10.1007/978-3-642-17373-8_27](https://doi.org/10.1007/978-3-642-17373-8_27) (cit. on pp. 4–7, 10, 11, 23, 40–42, 46, 51, 53–55).
- [Imp95] Russell Impagliazzo. “A Personal View of Average-Case Complexity”. In: *Proceedings of the Tenth Annual Structure in Complexity Theory Conference, Minneapolis, Minnesota, USA, June 19-22, 1995*. 1995, pp. 134–147. DOI: [10.1109/SCT.1995.514853](https://doi.org/10.1109/SCT.1995.514853). URL: <https://doi.org/10.1109/SCT.1995.514853> (cit. on p. 9).
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. “Founding Cryptography on Oblivious Transfer - Efficiently”. In: *Advances in Cryptology – CRYPTO 2008*. Ed. by David Wagner. Vol. 5157. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2008, pp. 572–591. DOI: [10.1007/978-3-540-85174-5_32](https://doi.org/10.1007/978-3-540-85174-5_32) (cit. on p. 53).
- [JR14] Ari Juels and Thomas Ristenpart. “Honey Encryption: Security Beyond the Brute-Force Bound”. In: *Advances in Cryptology – EUROCRYPT 2014*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Vol. 8441. Lecture Notes in Computer Science. Copenhagen, Denmark: Springer, Heidelberg, Germany, 2014, pp. 293–310. DOI: [10.1007/978-3-642-55220-5_17](https://doi.org/10.1007/978-3-642-55220-5_17) (cit. on pp. 1, 6, 8, 10, 51, 57, 58).
- [JRT16] Joseph Jaeger, Thomas Ristenpart, and Qiang Tang. “Honey Encryption Beyond Message Recovery Security”. In: *Advances in Cryptology – EUROCRYPT 2016, Part I*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9665. Lecture Notes in Computer Science. Vienna, Austria: Springer, Heidelberg, Germany, 2016, pp. 758–788. DOI: [10.1007/978-3-662-49890-3_29](https://doi.org/10.1007/978-3-662-49890-3_29) (cit. on p. 57).
- [Kil88] Joe Kilian. “Founding Cryptography on Oblivious Transfer”. In: *20th Annual ACM Symposium on Theory of Computing*. Chicago, IL, USA: ACM Press, 1988, pp. 20–31. DOI: [10.1145/62212.62215](https://doi.org/10.1145/62212.62215) (cit. on p. 53).
- [Lin09] Andrew Y. Lindell. “Adaptively Secure Two-Party Computation with Erasures”. In: *Topics in Cryptology – CT-RSA 2009*. Ed. by Marc Fischlin. Vol. 5473. Lecture Notes in Computer Science. San Francisco, CA, USA: Springer, Heidelberg,

- Germany, 2009, pp. 117–132. DOI: [10.1007/978-3-642-00862-7_8](https://doi.org/10.1007/978-3-642-00862-7_8) (cit. on p. 51).
- [LMs05] Matt Lepinski, Silvio Micali, and abhi shelat. “Fair-Zero Knowledge”. In: *TCC 2005: 2nd Theory of Cryptography Conference*. Ed. by Joe Kilian. Vol. 3378. Lecture Notes in Computer Science. Cambridge, MA, USA: Springer, Heidelberg, Germany, 2005, pp. 245–263. DOI: [10.1007/978-3-540-30576-7_14](https://doi.org/10.1007/978-3-540-30576-7_14) (cit. on pp. 6, 14, 37).
- [LV19] Ming Li and Paul M. B. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications, 4th Edition*. Texts in Computer Science. Springer, 2019. ISBN: 978-3-030-11297-4. DOI: [10.1007/978-3-030-11298-1](https://doi.org/10.1007/978-3-030-11298-1) (cit. on p. 1).
- [LV90] Ming Li and Paul M. B. Vitányi. “Handbook of Theoretical Computer Science (Vol. A)”. In: ed. by Jan van Leeuwen. Cambridge, MA, USA: MIT Press, 1990. Chap. Kolmogorov Complexity and Its Applications, pp. 187–254. ISBN: 0-444-88071-2. URL: <http://dl.acm.org/citation.cfm?id=114872.114876> (cit. on p. 1).
- [Pas03] Rafael Pass. “Simulation in Quasi-Polynomial Time, and Its Application to Protocol Composition”. In: *Advances in Cryptology – EUROCRYPT 2003*. Ed. by Eli Biham. Vol. 2656. Lecture Notes in Computer Science. Warsaw, Poland: Springer, Heidelberg, Germany, 2003, pp. 160–176. DOI: [10.1007/3-540-39200-9_10](https://doi.org/10.1007/3-540-39200-9_10) (cit. on p. 66).
- [Rey11] Leonid Reyzin. “Some Notions of Entropy for Cryptography - (Invited Talk)”. In: *ICITS 11: 5th International Conference on Information Theoretic Security*. Ed. by Serge Fehr. Vol. 6673. Lecture Notes in Computer Science. Amsterdam, The Netherlands: Springer, Heidelberg, Germany, 2011, pp. 138–142. DOI: [10.1007/978-3-642-20728-0_13](https://doi.org/10.1007/978-3-642-20728-0_13) (cit. on pp. 29, 30).
- [RR99] Ran Raz and Omer Reingold. “On Recycling the Randomness of States in Space Bounded Computation”. In: *31st Annual ACM Symposium on Theory of Computing*. Atlanta, GA, USA: ACM Press, 1999, pp. 159–168. DOI: [10.1145/301250.301294](https://doi.org/10.1145/301250.301294) (cit. on p. 1).
- [Sha48] Claude Elwood Shannon. “A mathematical theory of communication”. In: *Bell system technical journal* 27.3 (1948), pp. 379–423 (cit. on pp. 13, 30).
- [Sho04] Victor Shoup. *Sequences of games: a tool for taming complexity in security proofs*. Cryptology ePrint Archive, Report 2004/332. <http://eprint.iacr.org/2004/332>. 2004 (cit. on pp. 21, 27, 28).
- [SW14] Amit Sahai and Brent Waters. “How to use indistinguishability obfuscation: deniable encryption, and more”. In: *46th Annual ACM Symposium on Theory of Computing*. Ed. by David B. Shmoys. New York, NY, USA: ACM Press, 2014, pp. 475–484. DOI: [10.1145/2591796.2591825](https://doi.org/10.1145/2591796.2591825) (cit. on pp. 5–7, 29, 38, 46–48, 62).
- [TUZ01] Amnon Ta-Shma, Christopher Umans, and David Zuckerman. “Loss-less condensers, unbalanced expanders, and extractors”. In: *33rd Annual ACM Symposium on Theory of Computing*. Crete, Greece: ACM Press, 2001, pp. 143–152. DOI: [10.1145/380752.380790](https://doi.org/10.1145/380752.380790) (cit. on p. 1).
- [TV00] Luca Trevisan and Salil P. Vadhan. “Extracting Randomness from Samplable Distributions”. In: *41st Annual Symposium on Foundations of Computer Science*. Redondo Beach, CA, USA: IEEE Computer Society Press, 2000, pp. 32–42. DOI: [10.1109/SFCS.2000.892063](https://doi.org/10.1109/SFCS.2000.892063) (cit. on p. 1).
- [TVZ05] Luca Trevisan, Salil P. Vadhan, and David Zuckerman. “Compression of Samplable Sources”. In: *Computational Complexity* 14.3 (2005), pp. 186–227. DOI: [10.1007/s00037-005-0198-6](https://doi.org/10.1007/s00037-005-0198-6). URL: <https://doi.org/10.1007/s00037-005-0198-6> (cit. on pp. 1, 9, 13, 18, 31).

- [vH04] Luis von Ahn and Nicholas J. Hopper. “Public-Key Steganography”. In: *Advances in Cryptology – EUROCRYPT 2004*. Ed. by Christian Cachin and Jan Camenisch. Vol. 3027. Lecture Notes in Computer Science. Interlaken, Switzerland: Springer, Heidelberg, Germany, 2004, pp. 323–341. DOI: [10.1007/978-3-540-24676-3_20](https://doi.org/10.1007/978-3-540-24676-3_20) (cit. on p. 59).
- [vHL05] Luis von Ahn, Nicholas J. Hopper, and John Langford. “Covert two-party computation”. In: *37th Annual ACM Symposium on Theory of Computing*. Ed. by Harold N. Gabow and Ronald Fagin. Baltimore, MA, USA: ACM Press, 2005, pp. 513–522. DOI: [10.1145/1060590.1060668](https://doi.org/10.1145/1060590.1060668) (cit. on pp. 8, 59, 61).
- [Wee04] Hoeteck Wee. “On Pseudoentropy versus Compressibility”. In: *19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA*. IEEE Computer Society, 2004, pp. 29–41. ISBN: 0-7695-2120-7. DOI: [10.1109/CCC.2004.1313782](https://doi.org/10.1109/CCC.2004.1313782). URL: <https://doi.org/10.1109/CCC.2004.1313782> (cit. on pp. 1, 9, 13, 30, 31).
- [WW20] Hoeteck Wee and Daniel Wichs. *Candidate Obfuscation via Oblivious LWE Sampling*. Cryptology ePrint Archive, Report 2020/1042. <https://eprint.iacr.org/2020/1042>. 2020 (cit. on p. 10).
- [Yao82] Andrew Chi-Chih Yao. “Theory and Applications of Trapdoor Functions (Extended Abstract)”. In: *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*. IEEE Computer Society, 1982, pp. 80–91. DOI: [10.1109/SFCS.1982.45](https://doi.org/10.1109/SFCS.1982.45). URL: <https://doi.org/10.1109/SFCS.1982.45> (cit. on pp. 9, 13, 30).
- [Zha16] Mark Zhandry. “The Magic of ELFs”. In: *Advances in Cryptology – CRYPTO 2016, Part I*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9814. Lecture Notes in Computer Science. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, 2016, pp. 479–508. DOI: [10.1007/978-3-662-53018-4_18](https://doi.org/10.1007/978-3-662-53018-4_18) (cit. on pp. 6, 65, 66).