

Incompressible Encryption with Everlasting Security

Shany Ben-David	Eylon Yogev
shany.ben-david@biu.ac.il	eylon.yogev@biu.ac.il
Bar-Ilan University	Bar-Ilan University

May 28, 2025

Abstract

Recently, the concept of incompressible encryption has emerged as a powerful enhancement to key-leakage resilience. In an incompressible encryption scheme, an adversary who intercepts ciphertexts is forced to dedicate a significant amount of memory to store them in full if they wish to extract any information about the plaintext later when the secret key becomes available. Given two messages, the security game involves two adversaries: the first adversary receives an encryption of one of the messages and produces a compressed state. Then, the second adversary, given both the secret key and the compressed state, attempts to determine which message was encrypted.

Several positive results exist in incompressible cryptography. On the one hand, there are constructions based on minimal assumptions but with a poor rate (i.e., rate tends to 0). On the other hand, there are rate-1 constructions that achieve optimal efficiency but rely on strong cryptographic assumptions, such as obfuscation.

A stronger security notion, known as *everlasting security*, has been proposed for incompressible encryption. In this formulation, the second adversary, who receives the compressed state and the secret key, is allowed to be computationally *unbounded*. While this notion is conceptually appealing, no constructions of everlasting incompressible encryption are currently known, regardless of the underlying assumption or even in idealized models.

In this work, we give the first construction of everlasting incompressible encryption. In fact, we show that everlasting incompressible encryption is *inherent* in any sufficiently secure public-key encryption scheme. Specifically, we prove that any public-key encryption scheme with subexponential security (when instantiated with an appropriate security parameter) already satisfies the definition of everlasting incompressible encryption with subexponential security. Furthermore, our scheme achieves rate-1, improving upon existing results even for the weaker notion of standard incompressible encryption.

Keywords: public-key encryption; incompressible encryption; everlasting security; forward security

1 Introduction

Public-key encryption is a fundamental pillar of modern cryptography, enabling secure communication over untrusted networks without the need for a prior exchange of secret keys. Over the years, increasingly strong security notions have been developed to address various adversarial threats and real-world constraints. One such notion is forward security, which ensures that even if the secret decryption key is compromised at a later stage, previously encrypted messages remain private. Forward security is particularly useful in scenarios where key exposure is a realistic threat, such as in long-term encrypted storage or communication over extended periods. However, this notion typically relies on either multi-round protocols or periodic key updates, both of which can be impractical in many scenarios.

Recently, the concept of incompressible encryption has emerged as a powerful enhancement to key-leakage resilience [Dzi06; GWZ22]. It serves as a key example of the broader trend of incompressible cryptography, a rapidly growing area that has garnered significant attention in recent years [Dzi06; DGO19; GLW20; MW20; GZ21; GWZ22; GKR25]. In an incompressible encryption scheme, an adversary who intercepts ciphertexts is forced to dedicate a significant amount of memory to store them in full if they wish to extract any information about the plaintext later when the secret key becomes available. This requirement imposes a tangible cost on the adversary, as maintaining large storage resources can be impractical or even prohibitive, especially for high-volume data or long retention periods. The idea behind incompressible encryption is to make ciphertexts resistant to succinct representation, ensuring that adversaries cannot efficiently store them in a compressed form while retaining the ability to extract private information in the future when the key is leaked.

Beyond its fundamental importance, incompressible encryption has also served as a useful tool for achieving various results. [BGKNPR24] presented a transformation from incompressible symmetric-key encryption (SKE) to leakage-resilient incompressible SKE. On a different lane, [ABY25] leveraged the existence of incompressible public-key encryption to demonstrate the non-existence of a computational notion of instance compression.

The implications of incompressible encryption are particularly relevant in settings where adversaries have intermittent access to secret keys, such as in cloud storage security, data breaches, or state-sponsored surveillance. By increasing the storage burden on attackers, incompressible encryption provides an additional layer of security, complementing traditional encryption schemes that rely solely on computational hardness assumptions.

The security game. Several notions of incompressibility have been proposed, each offering distinct security guarantees. In this work, we focus on the recent formalization by Guan, Wichs, and Zhandry [GWZ22]. Their framework retains the standard syntax of public-key encryption in the standard model but introduces a fundamentally different security game. Informally, they require that any efficient adversary $(\mathbf{A}_1, \mathbf{A}_2)$ has only a negligible probability of success in the following game:

1. The challenger first gives the adversary the public key.
2. \mathbf{A}_1 then produces two messages m_0, m_1 .
3. The challenger encrypts m_b for a random $b \in \{0, 1\}$ and sends the ciphertext c to \mathbf{A}_1 .
4. The adversary produces a state s of size somewhat smaller than c .
5. The challenger then reveals the secret key to \mathbf{A}_2 .
6. \mathbf{A}_2 , given the small state s and the secret key, guesses b (and wins if the guess is correct).

The rate of the construction is the ratio between the message size and the ciphertext size. In a

rate-1 scheme, the size of the ciphertext is $(1 + o(1)) \cdot |m|$ for large enough messages. It is, of course, desirable to construct schemes with rate approaching 1.

Constructions of incompressible encryption. Guan, Wichs, Zhandry [GWZ22] provided two constructions of incompressible encryption. The first construction assumed only any standard public-key encryption scheme, but did not have a good rate (rate approaches 0). The second construction assumed the existence of indistinguishability obfuscation, but achieved a rate-1 incompressible encryption, where the ciphertext size is $S + \text{poly}(\lambda)$, where S is the storage bound size, λ is the security parameter, and the message length can be as large as roughly S . Later, Branco, Dötting, and Dujmovi [BDD22] constructed rate-1 incompressible encryption assuming the existence of programmable hash proof systems (HPS), incompressible encodings and pseudorandom generators.

Everlasting security. A stronger security notion for incompressible encryption, known as *everlasting security*, was introduced in [DKMMMMQT25]. In this formulation, the second adversary \mathbf{A}_2 is considered to run in a long *offline* phase after the secret key has been leaked, and is thus considered to have more computational power. In particular, they allow the adversary \mathbf{A}_2 to be computationally *unbounded*. Thus, schemes that satisfy this notion are sometimes referred to as having *everlasting security*. Notably, if \mathbf{A}_1 were unbounded, it could simply decrypt the ciphertext and store only the plaintext, making the notion infeasible. However, since \mathbf{A}_2 does not have access to the ciphertext, allowing it to run in unbounded time does not invalidate the definition.

Despite the conceptual appeal of this stronger notion, no constructions of everlasting incompressible encryption were known prior to this work, regardless of the underlying assumptions or even in idealized models. This leads us to ask:

*Do incompressible encryption schemes with **everlasting security** exist?*

The work of Harnik and Naor [HN10] shows strong limitations for encryption schemes with a similar flavor of everlasting security. However, crucially their model allows the second adversary \mathbf{A}_2 to get the entire ciphertext but only a compressed version (created by \mathbf{A}_1) of the key agreement protocol performed at the beginning. Their limitations do not apply to the notion of everlasting security considered in this paper.

A positive result has been recently obtained in [DKMMMMQT25] for a quantum analog of the everlasting notion, referred to as the *Quantum Decoherence Model*. In this model, both the encryption and decryption algorithms themselves are *quantum algorithms* (not only the adversary), and the first adversary, \mathbf{A}_1 , outputs only a limited number of *qubits*, while the second adversary, \mathbf{A}_2 , remains computationally unbounded. The security of their construction relies on the quantum random oracle model (QROM). Thus, despite this progress, we still lack a construction of *standard* everlasting incompressible encryption where the encryption algorithm is executed by a classical computer.

1.1 Our results

We answer the above question affirmatively (and in the standard model). Surprisingly, our main result shows that everlasting incompressible encryption is *inherent* in any sufficiently secure public-key encryption scheme. Specifically, we prove that any public-key encryption scheme with subexponential security (when instantiated with an appropriate security parameter) **already satisfies the definition of everlasting incompressible encryption** with subexponential security. Our result requires no additional cryptographic assumptions and no additional tools. This also means that

we can achieve an everlasting public-key encryption scheme with rate-1 (the rate of an encryption scheme is the message length over the ciphertext length).

Theorem 1.1 (Informal). *Any subexponential secure (secret-key/public-key) encryption scheme, applied with security parameter $\lambda' = \lambda \cdot \text{poly}(S)$ is an everlasting incompressible (secret-key/public-key) encryption scheme, with subexponential security.*

Moreover, if the original encryption scheme is rate-1, then the incompressible encryption scheme is also rate-1.

See Corollary 3.3 and Corollary 4.4 for a precise theorem statement with additional parameters. An additional benefit of our schemes is that they are plausibly post-quantum secure, provided that the underlying public-key encryption scheme is post-quantum secure.

Key size vs. rate. One of the main complexity measures of an encryption scheme is its rate. The rate is defined as the ratio between the message length and the ciphertext length, i.e., $\frac{|m|}{\text{Enc}(\text{pk}, m)}$. This ratio is a function that depends on two parameters: the message length $|m|$ and the security parameter λ . A scheme is considered rate-1 if, for a fixed λ , the ratio approaches 1 as $|m|$ tends to infinity, and λ is fixed.

While our scheme obtains the optimal rate-1 (assuming the underlying scheme has rate-1), the size of the private key is large. In particular, if the underlying scheme had private key size λ , our incompressible scheme has private key size $\lambda' = \lambda \cdot \text{poly}(S)$, which, in particular, is much larger than S . It is very desirable to obtain schemes with short private keys, with size that is independent of S .

However, to this date, there are no incompressible schemes that obtain both rate-1 and short private keys, even without everlasting security. Furthermore, it was argued (informally) that rate-1 and short private keys are at odds, and it could be the case that no scheme can have both properties simultaneously.

Table 1 summarizes existing incompressible encryption schemes in terms of their rate, key sizes, whether they offer everlasting security, and the underlying cryptographic assumptions. The work of [GWZ22] provides two constructions: one achieves short keys but suffers from rate that approaches 0, while the other achieves rate-1 with a short public key but a large private key, and relies on indistinguishability obfuscation, a strong and non-standard assumption. Then, the work of [BDD22] made notable progress by constructing a rate-1 incompressible encryption scheme based only on standard assumptions (DDH and LWE). However, this construction requires both the public and private keys to be large. Our scheme continues this line of work by additionally providing everlasting security.

On the incompressibility of deployed schemes. Our results have shown so far that encryption schemes instantiated with a sufficiently large security parameter are incompressible with everlasting security, for any storage bound S . This general statement can also be applied to concrete schemes used in practice. For example, consider the widely used AES block cipher, which is believed to provide (approximately) λ bits of security when instantiated with a key of length λ^1 . What can we say about AES in terms of its incompressibility?

The work of [BDD22] highlights the challenges in proving the incompressibility of ideal ciphers within the standard model. They construct a scheme that is incompressible (though not everlastingly

¹The notion λ bits of security means that any adversary that runs in time T succeeds in the security game with advantage at most $T \cdot 2^{-\lambda}$.

Work	Rate-1	Short Keys		Everlasting Security	Assumptions
		Private	Public		
[GWZ22]	✗	✓	✓	✗	PKE
[GWZ22]	✓	✗	✓	✗	iO
[BDD22]	✓	✗	✗	✗	DDH + LWE
This work	✓	✗	✗	✓	rate-1 subexp-PKE

Table 1: Comparison of our incompressible encryption scheme with prior works regarding encryption rate, key size, security, and underlying assumptions. A key is considered short if its size is polynomial in the security parameter λ and does not depend on the message length or compression parameters.

secure) in the ideal cipher model (ICM), and show that any concrete instantiation of the ideal cipher renders the scheme insecure.

In contrast, we provide positive results for ciphers in the standard model. Leveraging the same techniques as in our main theorem, we show that AES, as deployed in practice, is incompressible with everlasting security against adversaries whose storage capacity is slightly below λ , assuming AES provides λ bits of security. More generally, we prove that

Corollary 1.2 (Informal). *Any secret-key encryption scheme with λ bits of security (Definition 2.2) has everlasting incompressible security for storage bound $S \leq \lambda - o(\lambda)$.*

See Corollary 4.3 for the formal theorem with the precise parameters.

1.2 Discussion and open problems

Our results can be interpreted in multiple ways. On the one hand, they provide a positive perspective, revealing that public-key encryption inherently offers additional security guarantees. On the other hand, they raise concerns about whether the current definition of incompressible encryption fully captures the properties one might expect from such a scheme, suggesting that refinements to the definition may be necessary.

Storage-rate. One possible refinement involves redefining the notion of rate in the context of incompressible encryption. The standard ciphertext rate is natural for standard encryption schemes. However, we propose a different notion of ciphertext rate that is suitable for incompressible encryption schemes, where the storage bound is considered to be larger than the message size. Incompressible encryption schemes introduce an additional parameter, S , representing the adversary’s available storage bound. As a result, the ciphertext size becomes a function of three parameters: $|m|$, S , and λ .

To analyze the asymptotic behavior of the rate, we must decide which parameters remain fixed and which tend to infinity. If we fix λ and S and take $|m|$ to infinity, we get the standard rate definition. However, in the case where the storage bound is larger than the message size (e.g., when using a bit encryption scheme), we might instead fix $|m|$ and λ , and define the rate to be the ratio $\frac{S}{\text{Enc}(\text{pk}, m)}$, where S grows to infinity, and $|m|$ and λ are fixed. We denote this limit of rate by *storage-rate*.

Under this alternative definition, a rate-1 scheme would produce ciphertexts only slightly larger than the adversary’s storage bound (which is not the case with the current definition), ensuring that an adversary must store nearly the entire ciphertext to later decrypt it.

We demonstrate the difference between the two rate definitions. First, consider a scheme with ciphertext size $|m| + S \cdot \text{poly}(\lambda)$. Under the standard rate definition, this scheme achieves rate-1, but does not achieve storage-rate-1. Alternatively, consider a scheme with ciphertext size $S + |m| \cdot \text{poly}(\lambda)$. This scheme does not achieve rate-1, but does achieve storage-rate-1. Finally, consider a scheme with ciphertext size $|m| + S + \text{poly}(\lambda)$. This scheme achieves both rate-1 and storage-rate-1.

We additionally give positive results for everlasting incompressible schemes that have both constant storage-rate and have (standard) rate-1. We apply our techniques to an exponential secure public-key encryption with cipher text length $m + O(\lambda)$ (when instantiated with an appropriate security parameter) and get an everlasting incompressible encryption scheme with a constant storage rate and rate-1. That is, an adversary must store a constant fraction of the ciphertext, in order to later decrypt, even if the message is a single bit long.

Corollary 1.3 (Informal). *Any exponentially secure secret-key encryption scheme with cipher text length $m + O(\lambda)$, applied with security parameter $\lambda' = O(S + \lambda)$ is an everlasting incompressible encryption scheme with storage-rate- $O(1)$ and (standard) rate-1.*

See Corollary 3.3 for the precise theorem statement. Our scheme does not satisfy storage-rate-1, which is left as an open problem.

Other possible refinements could consider factors such as the size of the public and private keys or the efficiency of the encryption and decryption processes. Additionally, alternative security formulations might involve a scenario where the adversary is required to compress multiple ciphertexts rather than just a single one, as suggested in [GWZ23].

1.3 Paper organization

The paper is organized as follows:

- Section 2 provides the necessary preliminary definitions.
- Section 3 presents our main result on public-key incompressible encryption. The section begins with a formal definition of everlasting security, followed by a high-level overview of our techniques in Section 3.1.
- Section 4 contains our result on secret-key incompressible encryption. It opens with a definition of symmetric-key encryption with everlasting security, and continues with an overview of our approach in Section 4.1.
- Section 5 establishes general claims about learning distributions that are used throughout the paper.

2 Preliminaries

In this section we present the basic notions and the cryptographic primitives and tools that are used in this work. For an integer $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$. For a distribution X we denote by $x \leftarrow X$ the process of sampling a value x from the distribution X . Similarly, for a set \mathcal{X} we denote by $x \leftarrow \mathcal{X}$ the process of sampling a value x from the uniform distribution over \mathcal{X} . For a distribution X , we denote by $X(x)$ the probability that x is sampled from the distribution X .

Definition 2.1 (Public-Key Encryption Scheme). *Public-key encryption (PKE) scheme for a message space \mathcal{M} is a tuple of algorithms $(\text{PKE.Gen}, \text{PKE.Enc}, \text{PKE.Dec})$ where PKE.Gen and PKE.Enc are probabilistic polynomial-time algorithms, and PKE.Dec is a polynomial-time algorithm with the following properties:*

- Correctness. For every $\lambda \in \mathbb{N}$, messages $\mathbf{m} \in \mathcal{M}$,

$$\Pr [\mathbf{m} = \text{PKE.Dec}(\mathbf{pk}, \mathbf{sk}, \text{PKE.Enc}(\mathbf{pk}, \mathbf{m})) \mid (\mathbf{pk}, \mathbf{sk}) \leftarrow \text{PKE.Gen}(1^\lambda)] \geq 1 - \text{negl}(\lambda).$$

- Security game. For every adversary \mathbf{A} , we define the experiment $\text{Dist}_{\mathbf{A}}^{\text{IE}}(\lambda)$ as follows:
 1. Generate $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{PKE.Gen}(1^\lambda)$.
 2. The adversary \mathbf{A} is then provided the public key \mathbf{pk} and submits an auxiliary input \mathbf{aux} , and messages $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$.
 3. Sample $b \leftarrow \{0, 1\}$ uniformly at random.
 4. Sample $\mathbf{ct} \leftarrow \text{PKE.Enc}(\mathbf{pk}, \mathbf{m}_b)$.
 5. The adversary \mathbf{A} then receives $(\mathbf{pk}, \mathbf{aux}, \mathbf{ct})$, and outputs a guess b' .
 6. If $b = b'$ then the adversary succeeds and the experiment outputs 1. Otherwise, the experiment outputs 0.
- T -security. We say that the scheme is T -secure if for every $\lambda \in \mathbb{N}$, and for every adversary \mathbf{A} such that $|\mathbf{A}| \leq T(\lambda)$, it holds that,

$$\Pr [\text{Dist}_{\mathbf{A}}^{\text{PKE}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda).$$

- Sub-exponential security. A PKE scheme is defined as sub-exponentially secure if there exists some $\epsilon \in (0, 1)$ such that the scheme is 2^{λ^ϵ} -secure.
- Exponential security. A PKE scheme is defined as exponentially secure if there exists some $\epsilon \in (0, 1)$ such that the scheme is $2^{\epsilon \cdot \lambda}$ -secure.
- Rate. We define the rate by $\frac{|\mathbf{m}|}{|\text{PKE.Enc}(\mathbf{pk}, \mathbf{m})|}$, which is the size of a message divided by the ciphertext encrypting the message. We say that the scheme has rate-1 if the rate approaches 1 for large enough messages.

Definition 2.2 (Symmetric-Key Encryption Scheme). *Symmetric-key encryption (SKE) scheme for a message space \mathcal{M} is a tuple of algorithms $(\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$ where SKE.Gen and SKE.Enc are probabilistic polynomial-time algorithms, and SKE.Dec is a polynomial-time algorithm with the following properties:*

- **Correctness.** For every $\lambda \in \mathbb{N}$, messages $m \in \mathcal{M}$,

$$\Pr [m = \text{SKE.Dec}(\text{sk}, \text{SKE.Enc}(\text{sk}, m)) \mid \text{sk} \leftarrow \text{SKE.Gen}(1^\lambda)] \geq 1 - \text{negl}(\lambda).$$

- **Security game.** For every adversary \mathbf{A} , we define the experiment $\text{Dist}_\mathbf{A}^{\text{IE}}(\lambda)$ as follows:
 1. Generate $\text{sk} \leftarrow \text{SKE.Gen}(1^\lambda)$.
 2. The adversary $\mathbf{A}^{\text{IE.Enc}(\text{sk}, \cdot)}$ submits an auxiliary input aux and messages $m_0, m_1 \in \mathcal{M}$.
 3. Sample $b \leftarrow \{0, 1\}$ uniformly at random.
 4. Sample $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, m_b)$.
 5. The adversary $\mathbf{A}^{\text{IE.Enc}(\text{sk}, \cdot)}$ then receives (aux, ct) , and outputs a guess b' .
 6. If $b = b'$ then the adversary succeeds and the experiment outputs 1. Otherwise, the experiment outputs 0.
- **T-security.** We say that the scheme is T -secure if for every $\lambda \in \mathbb{N}$, and for every adversary \mathbf{A} such that $|\mathbf{A}| \leq T(\lambda)$, it holds that,

$$\Pr [\text{Dist}_\mathbf{A}^{\text{SKE}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda).$$

- **Sub-exponential security.** A SKE scheme is defined as sub-exponentially secure if there exists some $\epsilon \in (0, 1)$ such that the scheme is 2^{λ^ϵ} -secure.
- **Exponential security.** A SKE scheme is defined as exponentially secure if there exists some $\epsilon \in (0, 1)$ such that the scheme is $2^{\epsilon\lambda}$ -secure.
- **λ bits of security.** We say that the scheme has λ bits of security if for every $\lambda \in \mathbb{N}$ and for every adversary \mathbf{A} that runs in time T , it holds that,

$$\Pr [\text{Dist}_\mathbf{A}^{\text{SKE}}(\lambda) = 1] \leq \frac{1}{2} + \frac{T}{2^\lambda}.$$

- **Rate.** We define the rate by $\frac{|m|}{|\text{SKE.Enc}(\text{sk}, m)|}$, which is the size of a message divided by the ciphertext encrypting the message. We say that the scheme has rate-1 if the rate approaches 1 for large enough messages.

Statistical distance. We measure the statistical distance between two random variables in terms of their total variation distance, as in the definition below.

Definition 2.3 (Statistical distance). The **statistical distance** between two random variables X and Y taking values in a finite set \mathcal{S} is defined as

$$\Delta(X, Y) := \frac{1}{2} \sum_{a \in \mathcal{S}} |\Pr[X = a] - \Pr[Y = a]|.$$

Equivalently, the statistical distance can also be defined as

$$\Delta(X, Y) := \max_{\mathcal{S}' \subseteq \mathcal{S}} |\Pr[X \in \mathcal{S}'] - \Pr[Y \in \mathcal{S}']|.$$

3 Everlasting PKE from any PKE

In this section, we present our main result. We begin by formally defining incompressible encryption with everlasting security and stating our key theorems.

- In Section 3.1, we provide a high-level overview of our proof.
- In Section 3.2, we provide a full proof of the main theorem and the proof of the corollaries in Section 3.3.
- Our main proof relies on several lemmas and claims related to learning distributions and predicting from estimated distributions, which are deferred to Section 5. This section may be of independent interest.

We begin with a formal definition of incompressible encryption with everlasting security.

Definition 3.1 (public-key encryption with everlasting security). *An everlasting incompressible public-key encryption (IE) scheme for a message space \mathcal{M} is a tuple of algorithms $(\text{IE.Gen}, \text{IE.Enc}, \text{IE.Dec})$ where IE.Gen and IE.Enc are probabilistic polynomial-time algorithms, and IE.Dec is a polynomial-time algorithm with the following properties:*

- Correctness. For every $\lambda, S \in \mathbb{N}$, messages $\mathbf{m} \in \mathcal{M}$,

$$\Pr [\mathbf{m} = \text{IE.Dec}(\text{pk}, \text{sk}, \text{IE.Enc}(\text{pk}, \mathbf{m})) \mid (\text{pk}, \text{sk}) \leftarrow \text{IE.Gen}(1^\lambda, 1^S)] \geq 1 - \text{negl}(\lambda).$$

- Everlasting security game. For every adversary $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2)$, we define the experiment $\text{Dist}_{\mathbf{A}}^{\text{IE}}(\lambda)$ as follows:

1. The adversary \mathbf{A}_0 , on input 1^λ , outputs a space bound 1^S .
2. Generate $(\text{pk}, \text{sk}) \leftarrow \text{IE.Gen}(1^\lambda, 1^S)$.
3. Sample $b \leftarrow \{0, 1\}$ uniformly at random.
4. The adversary \mathbf{A}_1 is then provided the public key pk and submits an auxiliary input aux , and messages $\mathbf{m}_0, \mathbf{m}_1 \in \mathcal{M}$.
5. The adversary \mathbf{A}_1 then receives $\text{ct} \leftarrow \text{IE.Enc}(\text{pk}, \mathbf{m}_b)$, and submits a state st of size at most S .
6. The adversary \mathbf{A}_2 receives $(\text{pk}, \text{sk}, \text{aux}, \text{st})$ and outputs a guess b' .
7. If $b = b'$ then the adversary succeeds and the experiment outputs 1. Otherwise, the experiment outputs 0.

- (T, S_{\max}) -security. We say that the scheme is (T, S_{\max}) -secure if for every $\lambda \in \mathbb{N}$, and for every adversary $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2)$ such that,

- \mathbf{A}_0 outputs $S \leq S_{\max}(\lambda)$,
- \mathbf{A}_1 is of size $T(\lambda)$,
- \mathbf{A}_2 is unbounded,

the following holds,

$$\Pr [\text{Dist}_{\mathbf{A}}^{\text{IE}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda).$$

- T -security. We say that the scheme is T -secure if it is (T, S_{\max}) -secure for any S_{\max} .

- Sub-exponential security. An IE scheme is defined as sub-exponentially secure if it satisfies T -security for some $T(\lambda) \geq 2^{\lambda^\epsilon}$, where $\epsilon \in (0, 1)$.
- Exponential security. An IE scheme is defined as exponentially secure if it satisfies T -security for some $T(\lambda) \geq 2^{\epsilon \cdot \lambda}$, where $\epsilon \in (0, 1)$.
- Standard-rate. We define the standard-rate by $\frac{|m|}{|\text{IE.Enc}(\text{pk}, m)|}$, which is the size of a message divided by the ciphertext encrypting the message. We say that the scheme has rate-1 if the rate approaches 1 for large enough messages.
- Storage-rate. We define the storage-rate by $\frac{S}{|\text{IE.Enc}(\text{pk}, m)|}$, which is the compression parameter of the scheme divided by the ciphertext encrypting the message. We say that the scheme has rate-1 if the rate approaches 1 for large enough S .

Theorem 3.2. If a PKE scheme is T -secure for some $T : \mathbb{N} \rightarrow \mathbb{N}$ (Definition 2.1), then the PKE is also (T^IE, S_{\max}) -secure everlasting IE (Definition 3.1) for any $T^\text{IE} : \mathbb{N} \rightarrow \mathbb{N}$ and $S \in \mathbb{N}$ such that $T(\lambda) \geq O(2^{S_{\max} + o(S_{\max})} \cdot 2^{O(\log^2 \lambda)} \cdot T^\text{IE}(\lambda))$ for every $\lambda \in \mathbb{N}$.

The following corollaries follow by Theorem 3.2. We give a formal proof in Section 3.3.

Corollary 3.3. Given a PKE scheme with public-key size $L_{\text{pk}}(\lambda)$, private-key size $L_{\text{sk}}(\lambda)$, and ciphertext size $L_{\text{ct}}(\ell_m, \lambda)$ for message length ℓ_m ,

1. if the PKE has sub-exponential security, then there exists a sub-exponentially secure everlasting IE scheme that has public-key size $L_{\text{pk}}(S^c \cdot \lambda)$, private-key size $L_{\text{sk}}(S^c \cdot \lambda)$, and ciphertext size $L_{\text{ct}}(\ell_m, S^c \cdot \lambda)$ for some constant c , where S is the compression parameter of the scheme.

Moreover, if the PKE scheme has rate-1, then the IE scheme has rate-1.

2. if the PKE scheme has exponential security, then there exists an exponentially secure everlasting IE scheme that has public-key size $L'_{\text{pk}} = L_{\text{pk}}(O(S + \lambda))$, private-key size $L'_{\text{sk}} = L_{\text{sk}}(O(S + \lambda))$, and ciphertext size $L'_{\text{ct}} = L_{\text{ct}}(\ell_m, O(S + \lambda))$, where S is the compression parameter of the scheme.

3.1 High-level overview

We give a high-level overview of our proof. Assume the existence of a PKE scheme with T -security. We begin by assuming that $T \gg \tilde{O}(2^{S + \text{polylog}(\lambda)})$, and we will later address how to prove security for larger values of S . The proof of the theorem includes the following key ideas. Let $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2)$ be an adversary that breaks the scheme. We use \mathbf{A} to construct an adversary that breaks the security of the PKE scheme. Recall that \mathbf{A}_2 gets $(\text{pk}, \text{sk}, \text{st})$ as input.

Removing the dependency on sk . In Section 3.4, we show that we can assume without loss of generality, that \mathbf{A}_2 does not use the private key. That is, we show that if there exists an adversary that breaks the scheme when given the private key sk , then there is another adversary \mathbf{A}'_2 that does not need access to sk and can still break the scheme with exactly the same probability. This result is not just a technical proof tool, it also has conceptual significance. It shows that the seemingly weaker definition of everlasting security, where \mathbf{A}_2 is denied access to sk , is actually equivalent to the standard definition where \mathbf{A}_2 does receive sk in the second phase (as in Definition 3.1).

To establish this equivalence, we crucially rely on the fact that \mathbf{A}_2 is computationally unbounded (and indeed our \mathbf{A}'_2 is inefficient). This sets our techniques quite apart from the techniques used

for incompressible encryption without everlasting security, since the argument does not extend to the setting of computational IE, where \mathbf{A}_2 is required to be efficient.

The construction of \mathbf{A}'_2 proceeds as follows: given $(\mathbf{pk}, \mathbf{st})$, it samples a random secret key \mathbf{sk} using the key generation algorithm PKE.Gen , conditioned on the fixed public key \mathbf{pk} . It then simulates $\mathbf{A}_2(\mathbf{pk}, \mathbf{sk}, \mathbf{st})$. We show that $\mathbf{A}'_2(\mathbf{pk}, \mathbf{st})$ succeeds with the same probability as $\mathbf{A}_2(\mathbf{pk}, \mathbf{sk}, \mathbf{st})$. The formal proof appears in Lemma 3.7.

From hereon, we assume without loss of generality that \mathbf{A}_2 does not receive \mathbf{sk} as input.

Constructing the PKE adversary. We start the proof by fixing a random \mathbf{pk} where the adversary has a success probability of $\frac{1}{2} + \epsilon$. Since \mathbf{pk} is fixed, for simplicity, we assume that \mathbf{A}_2 gets only \mathbf{st} as an input. We then proceed with the following steps.

- **Defining the output distributions of \mathbf{A}_1 :** We first consider a distribution D_b for each $b \in \{0, 1\}$, which describes the distribution of the adversary \mathbf{A}_1 given an encryption of \mathbf{m}_b . Formally,

$$D_b := \left\{ \mathbf{st} \mid \begin{array}{l} \mathbf{ct} \leftarrow \text{PKE.Enc}(\mathbf{pk}, \mathbf{m}_b) \\ \mathbf{st} \leftarrow \mathbf{A}_1(\mathbf{pk}, \mathbf{ct}) \end{array} \right\}.$$

These distributions are efficiently samplable, and in fact, our PKE adversary will rely on them to distinguish between $\mathbf{ct} \leftarrow \text{PKE.Enc}(\mathbf{pk}, \mathbf{m}_0)$ and $\mathbf{ct} \leftarrow \text{PKE.Enc}(\mathbf{pk}, \mathbf{m}_1)$. Given a challenge ciphertext $\mathbf{ct} \leftarrow \text{PKE.Enc}(\mathbf{pk}, \mathbf{m}_b)$, the adversary will compute $\mathbf{st} \leftarrow \mathbf{A}_1(\mathbf{pk}, \mathbf{ct})$, which corresponds to a sample from D_b . To succeed, the adversary must distinguish between D_0 and D_1 .

- **Statistical distance between D_0 and D_1 is large:** Given that \mathbf{A}_2 successfully distinguishes between encryptions of \mathbf{m}_0 and \mathbf{m}_1 given only their state, it can be viewed as a successful distinguisher between D_0 and D_1 . Therefore, the distributions D_0 and D_1 must differ significantly, in statistical distance. Otherwise, if the statistical distance between them were small, then even an unbounded adversary receiving a sample $\mathbf{st} \leftarrow D_b$, where b is chosen uniformly at random, would not be able to distinguish between the two distributions.

Specifically, we show that the advantage of the adversary \mathbf{A}_2 defines a lower bound on the statistical distance: if \mathbf{A}_2 succeeds with probability $\frac{1}{2} + \epsilon$, then $\Delta(D_0, D_1) \geq 2\epsilon$. This follows from the following general claim:

Claim 1. *For every pair of distributions D_0, D_1 over \mathcal{X} , and for any unbounded distinguisher A ,*

$$\Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] \leq \frac{1}{2} + \frac{\Delta(D_0, D_1)}{2}.$$

See Claim 5.2 for the formal claim and proof.

- **Learning the weights of D_0 and D_1 :** While \mathbf{A}_2 can distinguish between D_0 and D_1 , it is an inefficient algorithm which we cannot simulate. We begin by defining an algorithm $\text{Predict}_{D_0, D_1}$ that achieves roughly the same advantage. This algorithm is efficient only if it is given an explicit representation of D_0 and D_1 (represented as a list of probabilities, one for each element in the domain). In our setting, we do not have this explicit representation, and we later show how to get an approximation of this data. The algorithm is defined as follows,

$$\text{Predict}_{D_0, D_1}(\mathbf{st}) := \begin{cases} 0 & \text{if } \Pr[\mathbf{st} \in D_0] \geq \Pr[\mathbf{st} \in D_1], \\ 1 & \text{if } \Pr[\mathbf{st} \in D_0] < \Pr[\mathbf{st} \in D_1]. \end{cases}$$

We prove the following general claim,

Claim 2. For every two distributions D_0, D_1 ,

$$\Pr \left[\text{Predict}_{D_0, D_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] = \frac{1}{2} + \frac{\Delta(D_0, D_1)}{2}.$$

See Claim 5.2 for the full formal claim. Note that to predict the distribution, our algorithm requires the explicit representation of D_0 and D_1 as list of probabilities over the support. Unfortunately, we don't have this tuple for D_0, D_1 , as we only have an efficient circuit that describes those distributions. To bridge this gap, we learn an approximation tuple \tilde{D}_0, \tilde{D}_1 that describes an approximation of D_0, D_1 , and show that this approximation will suffice to get (roughly) the same advantage, up to a small loss.

We use a learning algorithm, denoted Learn^D , which, given oracle access to samples from a distribution D over a domain \mathcal{X} , draws approximately $|\mathcal{X}|$ samples and outputs an explicit representation of an approximate distribution \tilde{D} . With all but negligible probability, it holds that $\Delta(D, \tilde{D}) \leq \text{negl}(\lambda)$ (see [Can20, Theorem 1]). In our setting, the supports of D_0 and D_1 are of size 2^S , so getting close approximation for each distribution requires roughly $\tilde{O}(2^{S+\text{polylog}(\lambda)})$ samples. Consequently, with overwhelming probability, we have $\Delta(D_b, \tilde{D}_b) \leq \text{negl}(\lambda)$ for each $b \in \{0, 1\}$. Since $\tilde{O}(2^{S+\text{polylog}(\lambda)}) \ll T$, this sampling is feasible for the PKE adversary.

- **Distinguishing between D_0 and D_1 :** Once the distributions have been approximated, we can use the algorithm $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}$ that, given a sample st , outputs 0 if $\Pr[\text{st} \in \tilde{D}_0] \geq \Pr[\text{st} \in \tilde{D}_1]$ and 1 otherwise. This algorithm distinguishes between \tilde{D}_0 and \tilde{D}_1 with probability $\frac{1}{2} + \frac{\Delta(\tilde{D}_0, \tilde{D}_1)}{2}$. Importantly, $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}$ operates efficiently. Given the tuples describing \tilde{D}_0 and \tilde{D}_1 , it can directly retrieve the probability $\Pr[\text{st} \in \tilde{D}_b]$ for any state st from the stored tuple on \tilde{D}_b . This allows the algorithm to efficiently determine whether $\Pr[\text{st} \in \tilde{D}_0] \geq \Pr[\text{st} \in \tilde{D}_1]$.

Upon receiving $\text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b)$, the adversary computes $\text{st}_b \leftarrow \mathbf{A}_1(\text{pk}, \text{ct})$, and then emulates $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st})$, and outputs the same. Note that st corresponds to a sample from D_b . To prove that the PKE adversary succeeds with high probability, we show that $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st})$ can distinguish between D_0 and D_1 with probability $\frac{1}{2} + \frac{\Delta(D_0, D_1)}{2} - \text{negl}(\lambda)$.

In more detail, we show that:

- The success probability of $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st})$ when applied to samples from D_0 or D_1 is negligibly close to its success when applied to samples from \tilde{D}_0 or \tilde{D}_1 (Claim 5.7).
- The success probability of $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}$ on samples from \tilde{D}_0 and \tilde{D}_1 is $\frac{1}{2} + \frac{\mathbb{E}[\Delta(\tilde{D}_0, \tilde{D}_1)]}{2}$ (Claim 5.8).
- The expected statistical distance between the approximate distributions is negligibly close to that between the original ones: $\frac{\mathbb{E}[\Delta(\tilde{D}_0, \tilde{D}_1)]}{2} \geq \frac{\mathbb{E}[\Delta(D_0, D_1)]}{2} - \text{negl}(\lambda)$ (Claim 5.6).

Taken together, these claims imply that $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st})$ distinguishes between D_0 and D_1 with probability at least $\frac{1}{2} + \frac{\Delta(D_0, D_1)}{2} - \text{negl}(\lambda)$.

Using this approach, an adversary against the PKE scheme can break the encryption as follows:

1. For every $b \in \{0, 1\}$, $\tilde{D}_b \leftarrow \text{Learn}^{D_b}$.
2. Upon receiving a challenge ciphertext ct , the adversary computes $\text{st} \leftarrow \mathbf{A}_1(\text{ct})$.

3. Output $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st})$.

To satisfy the condition that $T \gg 2^S$, our **IE** construction is an instantiation of the PKE scheme with security parameter λ' , which depends on both λ and S . Then, we show that if for example we start with $T(\lambda) = 2^{\lambda^\epsilon}$ for some $\epsilon \in (0, 1)$, we obtain a new scheme with security $T'(\lambda) = T(\lambda') = 2^{\lambda^\epsilon + S}$, as required.

3.2 Proof of Theorem 3.2

We now give a formal proof of Theorem 3.2.

Let $(\text{PKE.Gen}, \text{PKE.Enc}, \text{PKE.Dec})$ be a PKE scheme that is secure against adversaries of size $T(\lambda)$. Let $\mathbf{A}^{\text{IE}} = (\mathbf{A}_1^{\text{IE}}, \mathbf{A}_2^{\text{IE}})$ be an adversary, where \mathbf{A}_1^{IE} is of size at most $T^{\text{IE}}(\lambda)$, that breaks the IE scheme with probability $\epsilon(\lambda)$, for compression parameter S . By Lemma 3.7, we assume without loss of generality that \mathbf{A}^{IE} breaks the scheme without receiving sk .

Fix $\lambda \in \mathbb{N}$. Let $\mathcal{PK} := \{\text{pk} \mid (\text{pk}, \cdot) \in \text{Img}(\text{IE.Gen}(1^\lambda, 1^S))\}$, and let $\mathcal{X}_{\text{pk}} := \{(\text{aux}, m_0, m_1) \mid (\text{aux}, m_0, m_1) \in \text{Img}(\mathbf{A}_1^{\text{IE}}(\text{pk}))\}$. For every $\text{pk} \in \mathcal{PK}$, $(\text{aux}, m_0, m_1) \in \mathcal{X}_{\text{pk}}$, and $b \in \{0, 1\}$, let $D_{\text{pk}, m_b, \text{aux}}$ be the following distribution,

$$D_{m_b}^{\text{pk}, \text{aux}} := \left\{ \text{st} \in \{0, 1\}^S \mid \begin{array}{l} \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}, \text{aux}, \text{ct}) \end{array} \right\}.$$

Let \mathbf{A}^{PKE} be the following adversary to the PKE scheme.

$\mathbf{A}^{\text{PKE}}(\text{pk})$:

1. Emulate $\mathbf{A}_1^{\text{IE}}(\text{pk})$ to get (aux', m_0, m_1) .
2. Set $\text{aux} := (\text{aux}', m_0, m_1)$.
3. Output (aux, m_0, m_1) .

$\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}, \text{ct})$:

1. Parse $\text{aux} := (\text{aux}', m_0, m_1)$.
2. Set $t := 2^S \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)}$.
3. For every $b \in \{0, 1\}$, set $\tilde{D}_b \leftarrow \text{Learn}^{D_{m_b}^{\text{pk}, \text{aux}}}(t)$, where $\text{Learn}^{D_{m_b}^{\text{pk}, \text{aux}}}$ is the algorithm from Definition 5.4.
4. Compute $\text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}', \text{pk}, \text{ct})$.
5. Let $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}$ be the algorithm described in Claim 5.2 for distributions \tilde{D}_0, \tilde{D}_1 .
6. Output $\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st})$.

In the following, we prove that the success probability of \mathbf{A}^{PKE} in breaking the PKE scheme is equal to the success probability of \mathbf{A}^{IE} in breaking the IE scheme, up to a negligible loss in the security parameter. More formally, we show that:

$$\Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}, \text{ct}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \\ b \leftarrow \{0, 1\} \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}^{\text{PKE}}(\text{pk}) \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \end{array} \right] \geq \epsilon(\lambda) - \text{negl}(\lambda).$$

Fix some $\text{pk} \in \mathcal{PK}$, and $(\text{aux}, m_0, m_1) \in \mathcal{X}_{\text{pk}}$, and let $\text{aux}' := (\text{aux}, m_0, m_1)$. By construction,

$$\begin{aligned} & \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}', \text{ct}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \end{array} \right] \\ &= \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}, \text{aux}, \text{ct}) \end{array} \right]. \end{aligned}$$

By the definition of $D_{m_0}^{\text{pk}, \text{aux}}$ and $D_{m_1}^{\text{pk}, \text{aux}}$,

$$\begin{aligned} & \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}, \text{aux}, \text{ct}) \end{array} \right] \\ &= \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \\ \text{st} \leftarrow D_{m_b}^{\text{pk}, \text{aux}} \end{array} \right]. \end{aligned}$$

It holds that,

$$\begin{aligned} & \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \\ \text{st} \leftarrow D_{m_b}^{\text{pk}, \text{aux}} \end{array} \right] \\ &\geq \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(\text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \\ \text{st} \leftarrow \tilde{D}_b \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_{m_0}^{\text{pk}, \text{aux}}) \\ + \Delta(\tilde{D}_1, D_{m_1}^{\text{pk}, \text{aux}}) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \end{array} \right] \\ &= \frac{1}{2} + \mathbb{E} \left[\frac{\Delta(\tilde{D}_0, \tilde{D}_1)}{2} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_{m_0}^{\text{pk}, \text{aux}}) \\ + \Delta(\tilde{D}_1, D_{m_1}^{\text{pk}, \text{aux}}) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \end{array} \right] \\ &= \frac{1}{2} + \frac{1}{2} \cdot \mathbb{E} \left[\Delta(\tilde{D}_0, \tilde{D}_1) \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_{m_0}^{\text{pk}, \text{aux}}) \\ + \Delta(\tilde{D}_1, D_{m_1}^{\text{pk}, \text{aux}}) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \end{array} \right] \\ &\geq \frac{1}{2} + \frac{\Delta(D_{m_0}^{\text{pk}, \text{aux}}, D_{m_1}^{\text{pk}, \text{aux}})}{2} - \frac{3}{2} \cdot \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_{m_0}^{\text{pk}, \text{aux}}) \\ + \Delta(\tilde{D}_1, D_{m_1}^{\text{pk}, \text{aux}}) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_{m_0}^{\text{pk}, \text{aux}}}(\text{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{m_1}^{\text{pk}, \text{aux}}}(\text{t}) \end{array} \right], \end{aligned}$$

where the first inequality is by Claim 5.7, the first equality is by linearity of expectation, the second

equality is by Claim 5.8, and the second inequality is by Claim 5.6, Overall, we get that,

$$\begin{aligned}
& \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{ct}) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \mathbf{m}_b) \end{array} \right] \\
& \geq \frac{1}{2} + \frac{\Delta(D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}, D_{\mathbf{m}_1}^{\text{pk}, \text{aux}})}{2} - \frac{3}{2} \cdot \mathbb{E} \left[\begin{array}{c} \Delta(\tilde{D}_0, D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}) \\ + \Delta(\tilde{D}_1, D_{\mathbf{m}_1}^{\text{pk}, \text{aux}}) \end{array} \mid \begin{array}{c} \tilde{D}_0 \leftarrow \text{Learn}^{D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_{\mathbf{m}_1}^{\text{pk}, \text{aux}}}(\mathbf{t}) \end{array} \right] \\
& = \frac{1}{2} + \frac{\Delta(D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}, D_{\mathbf{m}_1}^{\text{pk}, \text{aux}})}{2} - \frac{3}{2} \cdot \sum_{b \in \{0, 1\}} \mathbb{E} \left[\Delta(\tilde{D}_b, D_{\mathbf{m}_b}^{\text{pk}, \text{aux}}) \mid \tilde{D}_b \leftarrow \text{Learn}^{D_{\mathbf{m}_b}^{\text{pk}, \text{aux}}}(\mathbf{t}) \right], \quad (1)
\end{aligned}$$

where the second inequality is by linearity of expectation. We apply Claim 5.5 with $\mathcal{X} := \{0, 1\}^S$, $\xi := 2^{-\log^2 \lambda}$, and $\delta := 2^{-\log^2 \lambda}$. Note that,

$$\begin{aligned}
\mathbf{t} &= 2^S \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)} \\
&= 2^S \cdot \left(2^{\log^2 \lambda}\right)^2 \cdot 2^{2 \cdot \log \log \lambda} \\
&= 2^S \cdot \left(2^{\log^2 \lambda}\right)^2 \cdot \log \left(2^{\log^2 \lambda}\right) \\
&= |\mathcal{X}| \cdot \frac{1}{\xi^2} \cdot \log \left(\frac{1}{\delta}\right) \\
&\geq \Omega \left(\frac{|\mathcal{X}| + \log \left(\frac{1}{\delta}\right)}{\xi^2} \right)
\end{aligned}$$

Therefore, by Claim 5.5, for every $b \in \{0, 1\}$,

$$\mathbb{E} \left[\Delta(\tilde{D}_b, D_{\mathbf{m}_b}^{\text{pk}, \text{aux}}) \mid \tilde{D}_b \leftarrow \text{Learn}^{D_{\mathbf{m}_b}^{\text{pk}, \text{aux}}}(\mathbf{t}) \right] \leq \xi + \delta.$$

Therefore, by plugging in the above to Equation 1,

$$\begin{aligned}
& \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}, \text{ct}) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \mathbf{m}_b) \end{array} \right] \\
& \geq \frac{1}{2} + \frac{\Delta(D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}, D_{\mathbf{m}_1}^{\text{pk}, \text{aux}})}{2} - \frac{3}{2} \cdot \sum_{b \in \{0, 1\}} (\xi + \delta) \\
& = \frac{1}{2} + \frac{\Delta(D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}, D_{\mathbf{m}_1}^{\text{pk}, \text{aux}})}{2} - 3\xi - 3\delta. \quad (2)
\end{aligned}$$

By Claim 5.2, for every unbounded distinguisher A , and specifically for $A := \mathbf{A}_2^{\text{IE}}$,

$$\Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ \text{st} \leftarrow D^{\text{pk}, \mathbf{m}_b} \end{array} \right] \leq \frac{1}{2} + \frac{\Delta(D_{\mathbf{m}_0}^{\text{pk}, \text{aux}}, D_{\mathbf{m}_1}^{\text{pk}, \text{aux}})}{2}$$

Therefore, by plugging in the above into Equation 2,

$$\begin{aligned}
& \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}', \text{ct}) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \mathbf{m}_b) \end{array} \right] \\
& \geq \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ \text{st} \leftarrow D_{\mathbf{m}_b}^{\text{pk}, \text{aux}} \end{array} \right] - 3\xi - 3\delta.
\end{aligned}$$

By the definition of $D_{m_b}^{\text{pk}, \text{aux}}$, we get that,

$$\begin{aligned} & \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}', \text{ct}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \end{array} \right] \\ & \geq \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}, \text{aux}, \text{ct}) \end{array} \right] - 3\xi - 3\delta. \end{aligned} \quad (3)$$

Since the above is true for any fixed $\text{pk} \in \mathcal{PK}$, $(\text{aux}, m_0, m_1) \in \mathcal{X}_{\text{pk}}$, and $\text{aux}' := (\text{aux}, m_0, m_1)$, it holds that,

$$\begin{aligned} & \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}', \text{ct}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \\ b \leftarrow \{0, 1\} \\ (\text{aux}', m_0, m_1) \leftarrow \mathbf{A}^{\text{PKE}}(\text{pk}) \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \end{array} \right] \\ & = \Pr \left[\mathbf{A}^{\text{PKE}}(\text{pk}, \text{aux}', \text{ct}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \\ b \leftarrow \{0, 1\} \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ \text{aux}' := (\text{aux}, m_0, m_1) \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \end{array} \right] \\ & \geq \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{IE.Gen}(1^\lambda, 1^S) \\ b \leftarrow \{0, 1\} \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}, \text{aux}, \text{ct}) \end{array} \right] - 3\xi - 3\delta \\ & = \epsilon(\lambda) - 3\xi - 3\delta \\ & = \epsilon(\lambda) - 6 \cdot 2^{-\log^2 \lambda} \\ & = \epsilon(\lambda) - \text{negl}(\lambda), \end{aligned}$$

where the first equality is by the construction of \mathbf{A}^{PKE} , the first inequality is by Equation 3, the second equality is by the success probability of \mathbf{A}^{IE} on the IE scheme, and the third equality is by $\xi, \delta = 2^{-\log^2 \lambda}$.

To complete the proof, in what follows, we bound the size of \mathbf{A}^{PKE} .

Claim 3.4. \mathbf{A}^{PKE} can be implemented with a circuit of size $O(2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot T^{\text{IE}}(\lambda))$.

Proof. The \mathbf{A}^{PKE} algorithm consists of three main steps:

1. *Emulating \mathbf{A}_1^{IE} :* This step simulates \mathbf{A}_1^{IE} , whose size is at most $T^{\text{IE}}(\lambda)$.
2. *Learn $D_{m_b}^{\text{pk}, \text{aux}}$ - Sampling Elements:* The algorithm samples t elements from $D_{m_0}^{\text{pk}, \text{aux}}$ and $D_{m_1}^{\text{pk}, \text{aux}}$. Each sample requires emulating \mathbf{A}_1^{IE} , resulting in a total circuit size of $2 \cdot t \cdot |\mathbf{A}_1^{\text{IE}}|$.
3. *Comparing Sample Counts:* Finally, the algorithm checks whether st appears more frequently in samples from $D_{m_0}^{\text{pk}, \text{aux}}$ or $D_{m_1}^{\text{pk}, \text{aux}}$. This check can be performed in $O(t \cdot S)$ time, and can be implemented with circuit of size $\tilde{O}(t \cdot S)$.

Overall, we got that,

$$\begin{aligned} |\mathbf{A}^{\text{PKE}}| &= \tilde{O}(\mathbf{t} \cdot S) + \mathbf{t} \cdot |\mathbf{A}_1^{\text{IE}}| \\ &\leq \tilde{O}(\mathbf{t} \cdot S) + \mathbf{t} \cdot T^{\text{IE}}(\lambda). \end{aligned}$$

Since $\mathbf{t} = 2^S \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)}$.

$$\begin{aligned} |\mathbf{A}^{\text{PKE}}| &= \tilde{O}(\mathbf{t} \cdot S) + \mathbf{t} \cdot T^{\text{IE}}(\lambda) \\ &= \tilde{O}(2^S \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)} \cdot S) + 2^S \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)} \cdot T^{\text{IE}}(\lambda) \\ &= O(2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot T^{\text{IE}}(\lambda)). \end{aligned}$$

□

3.3 Proofs of the Corollaries

In this section, we provide a proof for Corollary 3.3.

Proof of Corollary 3.3. We start by proving Item 1 of the corollary.

Fix a PKE scheme $(\text{PKE.Gen}, \text{PKE.Enc}, \text{PKE.Dec})$, that is sub-exponentially secure. Therefore, there exists some $\epsilon \in (0, 1)$ such that for $T(\lambda) := 2^{\lambda^\epsilon}$, the PKE scheme is T -secure. Let $\mathsf{L}_{\text{pk}}(\lambda)$ be the public-key size, let $\mathsf{L}_{\text{sk}}(\lambda)$ be the secret-key size, and let $\mathsf{L}_{\text{ct}}(\ell_m, \lambda)$ be the ciphertext size for message length ℓ_m . Let $\epsilon' = \frac{\epsilon}{2}$. In what follows we give a construction, which we later prove as everlasting IE with T^{IE} -security for $T^{\text{IE}}(\lambda) := 2^{\lambda^{\epsilon'}}$.

Construction 3.5. The construction is as follows,

- $\text{IE.Gen}(1^\lambda, 1^S)$:
 1. Set $\tilde{\lambda} := (S + o(S) + O(\log^2 \lambda) + \lambda^{\epsilon'} + O(1))^{\frac{1}{\epsilon}}$.
 2. Output $(\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^{\tilde{\lambda}})$.
- $\text{IE.Enc}(\text{pk}, m)$: Output $\text{PKE.Enc}(\text{pk}, m)$.
- $\text{IE.Dec}(\text{sk}, \text{ct})$: Output $\text{PKE.Dec}(\text{sk}, \text{ct})$.

Fix some $S \in \mathbb{N}$. The above construction forms a PKE scheme that is T' -secure, where $T'(\lambda) := T(\tilde{\lambda}) = 2^{\tilde{\lambda}^\epsilon}$, and with efficiency parameters $\mathsf{L}'_{\text{pk}}(\lambda) = \mathsf{L}_{\text{pk}}(\tilde{\lambda})$, $\mathsf{L}'_{\text{sk}}(\lambda) = \mathsf{L}_{\text{sk}}(\tilde{\lambda})$, and $\mathsf{L}'_{\text{ct}}(\ell_m, \lambda) = \mathsf{L}_{\text{ct}}(\ell_m, \tilde{\lambda})$. In what follows, we prove that there exists some c such that $\tilde{\lambda} \leq S^c \cdot \lambda$. This will allow us to bound the efficiency parameters of the PKE scheme such that,

$$\begin{aligned} \mathsf{L}'_{\text{pk}}(\lambda) &\leq \mathsf{L}_{\text{pk}}(S^c \cdot \lambda), \\ \mathsf{L}'_{\text{sk}}(\lambda) &\leq \mathsf{L}_{\text{sk}}(S^c \cdot \lambda), \\ \mathsf{L}'_{\text{ct}}(\ell_m, \lambda) &\leq \mathsf{L}_{\text{ct}}(\ell_m, S^c \cdot \lambda). \end{aligned}$$

For large enough λ and $S > 1$,

$$\begin{aligned}
\tilde{\lambda} &= (S + o(S) + O(\log^2 \lambda) + \lambda^{\epsilon'} + O(1))^{\frac{1}{\epsilon}} \\
&\leq (2 \cdot S + \lambda^{2\epsilon'})^{\frac{1}{\epsilon}} \\
&\leq (S \cdot \lambda^{2\epsilon'})^{\frac{1}{\epsilon}} \\
&= S^{\frac{1}{\epsilon}} \cdot \lambda^{\frac{2\epsilon'}{\epsilon}} \\
&= S^{\frac{1}{\epsilon}} \cdot \lambda,
\end{aligned}$$

where the last equality is since $\epsilon' = \frac{\epsilon}{2}$. Therefore, for $c := \frac{1}{\epsilon}$, it holds that,

$$\tilde{\lambda} \leq S^c \cdot \lambda,$$

as required.

In what follows, we prove that the T' -secure PKE scheme is also (T^{IE}, S) -secure everlasting IE. By the definition of $\tilde{\lambda}$,

$$\tilde{\lambda} = (S + o(S) + O(\log^2 \lambda) + \lambda^{\epsilon'} + O(1))^{\frac{1}{\epsilon}}.$$

Raising both sides of the equation to the power of ϵ , we get that,

$$\tilde{\lambda}^\epsilon = S + o(S) + O(\log^2 \lambda) + \lambda^{\epsilon'} + O(1),$$

Therefore,

$$\begin{aligned}
2^{\tilde{\lambda}^\epsilon} &= 2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot 2^{\lambda^{\epsilon'}} \cdot 2^{O(1)} \\
&= O(2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot 2^{\lambda^{\epsilon'}}).
\end{aligned}$$

Since $T'(\lambda) = 2^{\tilde{\lambda}^\epsilon}$, and $T^{\text{IE}}(\lambda) = 2^{\lambda^{\epsilon'}}$, we get that,

$$T'(\lambda) = O(2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot T^{\text{IE}}(\lambda)).$$

Therefore, we can apply Theorem 3.2 on the T' -secure PKE scheme, and get that the scheme is also (T^{IE}, S) -secure everlasting IE scheme, as required.

Next, we prove Item 2 of the corollary.

Fix a PKE scheme, $(\text{PKE.Gen}, \text{PKE.Enc}, \text{PKE.Dec})$, that is exponentially secure. Therefore, there exists some $\epsilon \in (0, 1)$ such that for $T(\lambda) := 2^{\epsilon \cdot \lambda}$, the PKE scheme is T -secure. Let $L_{\text{pk}}(\lambda)$ be the public-key size, let $L_{\text{sk}}(\lambda)$ be the secret-key size, and let $L_{\text{ct}}(\ell_m, \lambda)$ be the ciphertext size for message length ℓ_m . Let $\epsilon' = \frac{\epsilon}{2}$. In what follows we give a construction, which we later prove as everlasting IE with T^{IE} -security for $T^{\text{IE}}(\lambda) := 2^{\epsilon' \cdot \lambda}$.

Construction 3.6. The construction is as follows,

- $\text{IE.Gen}(1^\lambda, 1^S)$:
 1. Set $\tilde{\lambda} := \frac{1}{\epsilon} \cdot (S + o(S) + O(\log^2 \lambda) + \epsilon' \cdot \lambda + O(1))$.
 2. Output $(\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^{\tilde{\lambda}})$.

- IE.Enc(pk, m) : Output PKE.Enc(pk, m) .
- IE.Dec(sk, ct) : Output PKE.Dec(sk, ct) .

Fix some $S \in \mathbb{N}$. The above construction forms a PKE scheme that is T' -secure, where $T'(\lambda) := T(\tilde{\lambda}) = 2^{\epsilon \tilde{\lambda}}$, and with efficiency parameters $L'_{\text{pk}}(\lambda) = L_{\text{pk}}(\tilde{\lambda})$, $L'_{\text{sk}}(\lambda) = L_{\text{sk}}(\tilde{\lambda})$, and $L'_{\text{ct}}(\ell_m, \lambda) = L_{\text{ct}}(\ell_m, \tilde{\lambda})$. By construction,

$$\begin{aligned}\tilde{\lambda} &= \frac{1}{\epsilon} \cdot (S + o(S) + O(\log^2 \lambda) + \epsilon' \cdot \lambda + O(1)) \\ &= O(S + \lambda),\end{aligned}$$

Therefore,

$$\begin{aligned}L'_{\text{pk}}(\lambda) &\leq L_{\text{pk}}(O(S + \lambda)), \\ L'_{\text{sk}}(\lambda) &\leq L_{\text{sk}}(O(S + \lambda)), \\ L'_{\text{ct}}(\ell_m, \lambda) &\leq L_{\text{ct}}(\ell_m, O(S + \lambda)).\end{aligned}$$

as required.

In what follows, we prove that the T' -secure PKE scheme is also (T^{IE}, S) -secure everlasting IE. By the definition of $\tilde{\lambda}$,

$$\tilde{\lambda} = \frac{1}{\epsilon} \cdot (S + o(S) + O(\log^2 \lambda) + \epsilon' \cdot \lambda + O(1)).$$

Therefore,

$$\epsilon \cdot \tilde{\lambda} = S + o(S) + O(\log^2 \lambda) + \epsilon' \cdot \lambda + O(1).$$

Raising 2 to the above exponent gives,

$$\begin{aligned}2^{\epsilon \cdot \tilde{\lambda}} &= 2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot 2^{\epsilon' \cdot \lambda} \cdot 2^{O(1)} \\ &= O(2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot 2^{\epsilon' \cdot \lambda}).\end{aligned}$$

Since $T'(\lambda) = 2^{\epsilon \cdot \tilde{\lambda}}$, and $T^{\text{IE}}(\lambda) = 2^{\epsilon' \cdot \lambda}$, we get that,

$$T'(\lambda) = O(2^{S+o(S)} \cdot 2^{O(\log^2 \lambda)} \cdot T^{\text{IE}}(\lambda)).$$

Therefore, we can apply Theorem 3.2 on the T' -secure PKE scheme, and get that the scheme is also (T^{IE}, S) -secure everlasting IE scheme, as required. \square

3.4 Removing the private key

In this subsection, we prove that if there exists an adversary to the IE scheme that breaks the IE scheme with some probability, then there exists another adversary - who never receives the private key sk - that breaks the scheme with exactly the same probability.

This result can be interpreted as showing that a weaker notion of everlasting security for IE, in which the second-phase (unbounded) adversary \mathbf{A}_2 does not receive the private key sk , is in fact equivalent to the standard notion defined in Definition 3.1. While the weaker notion appears less demanding since it restricts \mathbf{A}_2 from seeing sk , Lemma 3.7 shows that achieving security in this setting suffices to guarantee security even when \mathbf{A}_2 is given sk in the second phase.

Lemma 3.7. *If there exists an adversary $\mathbf{A}^{\text{IE}} = (\mathbf{A}_0^{\text{IE}}, \mathbf{A}_1^{\text{IE}}, \mathbf{A}_2^{\text{IE}})$ such that for every $\lambda \in \mathbb{N}$,*

$$\Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right] = \epsilon(\lambda),$$

then there exists an algorithm \mathbf{A}'_2 such that for every $\lambda \in \mathbb{N}$,

$$\Pr \left[\mathbf{A}'_2(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right] = \epsilon(\lambda).$$

Proof. Let \mathbf{A}'_2 be the following algorithm:

- Given as input $(\text{pk}, \text{aux}, \text{st})$.
- Sample $(\text{pk}', \text{sk}) \leftarrow \text{IE.Gen}(1^\lambda, 1^S)$ conditioned on $\text{pk}' = \text{pk}$.
- Emulate $\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st})$ and output the same.

We show that \mathbf{A}'_2 , without giving the secret key, breaks the scheme with the same probability as \mathbf{A}_2^{IE} that is giving the secret key as input. By assumption,

$$\epsilon(\lambda) = \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right].$$

Next, we break down the sampling of (pk, sk) in the above experiment into two steps: first, sampling pk , and then sampling sk given pk . We get that,

$$\begin{aligned} & \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right] \\ &= \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{pk}', \text{sk}) \leftarrow \text{IE.Gen}(1^\lambda, 1^S) \mid \text{pk}' = \text{pk} \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right]. \end{aligned}$$

As sk is only used by \mathbf{A}_2^{IE} , we can defer its sampling and move it to the end of the experiment.

$$\begin{aligned}
& \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{pk}', \text{sk}) \leftarrow \text{IE.Gen}(1^\lambda, 1^S) \mid \text{pk}' = \text{pk} \\ (\text{aux}, \text{m}_0, \text{m}_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right] \\
&= \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, \text{m}_0, \text{m}_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \\ (\text{pk}', \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \mid \text{pk}' = \text{pk} \end{array} \right].
\end{aligned}$$

By the definition of \mathbf{A}_2' , we get that,

$$\begin{aligned}
& \Pr \left[\mathbf{A}_2^{\text{IE}}(\text{pk}, \text{sk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, \text{m}_0, \text{m}_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \\ (\text{pk}', \text{sk}) \leftarrow \text{PKE.Gen}(1^\lambda) \mid \text{pk}' = \text{pk} \end{array} \right] \\
&= \Pr \left[\mathbf{A}_2'(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, \text{m}_0, \text{m}_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right].
\end{aligned}$$

Overall, by the above equations, we get that,

$$\Pr \left[\mathbf{A}_2'(\text{pk}, \text{aux}, \text{st}) = b \mid \begin{array}{l} (\text{pk}, \cdot) \leftarrow \text{PKE.Gen}(1^\lambda) \\ (\text{aux}, \text{m}_0, \text{m}_1) \leftarrow \mathbf{A}_1^{\text{IE}}(\text{pk}) \\ b \leftarrow \{0, 1\} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{IE}}(\text{aux}, \text{ct}) \end{array} \right] = \epsilon(\lambda).$$

□

4 Everlasting SKE from any SKE

In this section, we present our result on symmetric-key encryption with everlasting security. We begin by formally defining symmetric-key encryption with everlasting security and stating our key theorems.

- Section 4.1 provides a high-level overview of our approach.
- Section 4.2 contains the full proof of the main theorem, while the proofs of the corollaries appear in Section 4.3.
- The main proof relies on several claims concerning learning distributions and prediction from estimated distributions. These technical results, which may be of independent interest, are deferred to Section 5.

We begin with a formal definition of symmetric-key encryption with everlasting security.

Definition 4.1 (symmetric-key encryption with everlasting security). *An everlasting symmetric-key encryption (SKE) scheme for a message space \mathcal{M} is a tuple of algorithms $(\text{IE.Gen}, \text{IE.Enc}, \text{IE.Dec})$ where IE.Gen and IE.Enc are probabilistic polynomial-time algorithms, and IE.Dec is a polynomial-time algorithm with the following properties:*

- Correctness. For every $\lambda, S \in \mathbb{N}$, messages $m \in \mathcal{M}$,

$$\Pr [m = \text{IE.Dec}(\text{sk}, \text{IE.Enc}(\text{sk}, m)) \mid \text{sk} \leftarrow \text{IE.Gen}(1^\lambda, 1^S)] \geq 1 - \text{negl}(\lambda).$$

- Everlasting security game. For every adversary $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2)$, we define the experiment $\text{Dist}_{\mathbf{A}}^{\text{IE}}(\lambda)$ as follows:

1. The adversary \mathbf{A}_0 , on input 1^λ , outputs a space bound 1^S .
2. Generate $\text{sk} \leftarrow \text{IE.Gen}(1^\lambda, 1^S)$.
3. Sample $b \leftarrow \{0, 1\}$ uniformly at random.
4. The adversary $\mathbf{A}_1^{\text{IE.Enc}(\text{sk}, \cdot)}$ submits an auxiliary input aux and messages $m_0, m_1 \in \mathcal{M}$.
5. The adversary $\mathbf{A}_1^{\text{IE.Enc}(\text{sk}, \cdot)}$ then receives $\text{ct} \leftarrow \text{IE.Gen}(\text{sk}, m_b)$, and submits a state st of size at most S .
6. The adversary \mathbf{A}_2 receives $(\text{sk}, m_0, m_1, \text{aux}, \text{st})$ and outputs a guess b' .
7. If $b = b'$ then the adversary succeeds and the experiment outputs 1. Otherwise, the experiment outputs 0.

- (T, S_{\max}) -security. We say that the scheme is (T, S_{\max}) -secure if for every $\lambda \in \mathbb{N}$, and for every adversary $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2)$ such that,

- \mathbf{A}_0 outputs $S \leq S_{\max}(\lambda)$,
- \mathbf{A}_1 is of size $T(\lambda)$,
- \mathbf{A}_2 is unbounded,

the following holds,

$$\Pr [\text{Dist}_{\mathbf{A}}^{\text{IE}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda).$$

- T -security. We say that the scheme is T -secure if it is (T, S_{\max}) -secure for any S_{\max} .

- Sub-exponential security. An IE scheme is defined as sub-exponentially secure if it satisfies T -security for some $T(\lambda) \geq 2^{\lambda^\epsilon}$, where $\epsilon \in (0, 1)$.
- Exponential security. An IE scheme is defined as exponentially secure if it satisfies T -security for some $T(\lambda) \geq 2^{\epsilon \cdot \lambda}$, where $\epsilon \in (0, 1)$.
- Standard-rate. We define the standard-rate by $\frac{|m|}{|\text{IE.Enc}(\text{pk}, m)|}$, which is the size of a message divided by the ciphertext encrypting the message. We say that the scheme has rate-1 if the rate approaches 1 for large enough messages.
- Storage-rate. We define the storage-rate by $\frac{S}{|\text{IE.Enc}(\text{pk}, m)|}$, which is the compression parameter of the scheme divided by the ciphertext encrypting the message. We say that the scheme has rate-1 if the rate approaches 1 for large enough S .

Theorem 4.2. *If a SKE scheme is T -secure for some $T : \mathbb{N} \rightarrow \mathbb{N}$ (Definition 2.2), then the SKE is also $(T^{\text{IE}}, S_{\max})$ -secure everlasting SKE (Definition 4.1) for any $T^{\text{IE}} : \mathbb{N} \rightarrow \mathbb{N}$ and $S_{\max} \in \mathbb{N}$ such that $T(\lambda) \geq 2^{(S_{\max} + O(\log S_{\max}) + 3 \cdot \log^2 \lambda)} \cdot T^{\text{IE}}(\lambda)$ for every $\lambda \in \mathbb{N}$.*

The following corollaries follow by Theorem 4.2. We give a formal proof in Section 4.3.

Corollary 4.3. *If a SKE scheme is $(2^{\lambda - \log^2 \lambda})$ -secure (Definition 2.2) with secret-key size $L_{\text{sk}}(\lambda) = \lambda$, and ciphertext size $L_{\text{ct}}(\lambda) = \lambda$ for messages of length λ , then the SKE is also everlasting SKE with $(\lambda^{\log \lambda}, S_{\max})$ -security (Definition 4.1) for any $S_{\max} : \mathbb{N} \rightarrow \mathbb{N}$ such that $S_{\max}(\lambda) \leq \lambda - \log^4(\lambda)$.*

Corollary 4.4. *Given a SKE scheme with secret-key size $L_{\text{sk}}(\lambda)$, and ciphertext size $L_{\text{ct}}(\ell_m, \lambda)$ for message length ℓ_m ,*

1. *if the SKE scheme has sub-exponentially security, then there exists a sub-exponentially secure everlasting SKE scheme that has secret-key size $L_{\text{sk}}(S^c \cdot \lambda)$, and ciphertext size $L_{\text{ct}}(\ell_m, S^c \cdot \lambda)$ for some constant c , where S is the compression parameter of the scheme.*

Moreover, if the SKE scheme has rate-1, then the everlasting SKE scheme has rate-1.

2. *if the SKE scheme has exponential security, then there exists an exponentially secure everlasting SKE scheme that has secret-key size $L'_{\text{sk}} = L_{\text{sk}}(O(S + \lambda))$, and ciphertext size $L'_{\text{ct}} = L_{\text{ct}}(\ell_m, O(S + \lambda))$, where S is the compression parameter of the scheme.*

Moreover, if we apply SKE scheme with $L_{\text{ct}} = \ell_m + O(\lambda)$, we get everlasting SKE with $L'_{\text{ct}} = \ell_m + O(S + \lambda)$.

4.1 High-level overview

The proof proceeds along similar lines to the proof of Theorem 3.2, with one key difference. In Theorem 3.2, the first step is to remove sk from the security game. This is justified by Lemma 3.7, which shows that if there exists an adversary that breaks the everlasting security of the PKE scheme with some probability, then there exists another adversary, who never receives sk , that achieves the same success probability. This simplification allows us to prove security against a weaker adversary, one that does not rely on access to sk .

In the current setting, we would like to repeat this step by proving an analogous claim for SKE. However, a fundamental difference between PKE and SKE in this context prevents us from doing

so. In the PKE setting, once the public key \mathbf{pk} is known, the adversary \mathbf{A}_2 can compute as many encryptions as it wishes without access to the secret key. In contrast, in the SKE setting, if the secret key \mathbf{sk} is not provided to \mathbf{A}_2 , the adversary faces significantly stronger restrictions. In particular, it cannot compute encryptions of messages on its own: \mathbf{A}_2 does not have access to an encryption oracle, and the internal state is not large enough to contain any ciphertexts. Furthermore, the ciphertexts may be entirely independent of the plaintext messages, making it infeasible for the adversary to simulate encryptions. This inherent limitation makes it unclear how to eliminate \mathbf{sk} from the security game in the SKE setting.

Our proof still closely follows the structure of the public-key case, but we now need to augment the distributions to account for the secret key. Since the distribution depends on the secret key, the adversary we construct in the reduction (i.e., the adversary for the SKE scheme) cannot compute it directly, as it does not have access to the secret key. However, this adversary can simulate oracle access to the distribution by leveraging its access to the encryption oracle, which is provided in the secret-key security game. This simulation is sufficient for the proof to carry through.

4.2 Proof of Theorem 4.2

We now give a formal proof of Theorem 4.2.

Let $(\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$ be a SKE scheme that is secure against adversaries of size $T(\lambda)$. Let $\mathbf{A}^{\text{IE}} = (\mathbf{A}_1, \mathbf{A}_2)$ be an adversary, where \mathbf{A}_1 is of size at most $T^{\text{IE}}(\lambda)$, that breaks the IE scheme with probability $\epsilon(\lambda)$, for compression parameter S_{\max} . Formally,

$$\Pr \left[\mathbf{A}_2(\mathbf{sk}, m_0, m_1, \text{aux}, \text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \mathbf{sk} \leftarrow \text{SKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\mathbf{sk}, \cdot)} \\ \text{ct} \leftarrow \text{SKE.Enc}(\mathbf{sk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\mathbf{sk}, \cdot)}(\text{aux}, \text{ct}) \end{array} \right] = \epsilon(\lambda).$$

For every $b \in \{0, 1\}$, let D_b be the following distribution,

$$D_b := \left\{ (\mathbf{sk}, m_0, m_1, \text{aux}, \text{st}) \mid \begin{array}{l} \mathbf{sk} \leftarrow \text{SKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\mathbf{sk}, \cdot)} \\ \text{ct} \leftarrow \text{SKE.Enc}(\mathbf{sk}, m_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\mathbf{sk}, \cdot)}(\text{aux}, \text{ct}) \end{array} \right\}.$$

Therefore,

$$\Pr \left[\mathbf{A}_2(\mathbf{sk}, m_0, m_1, \text{aux}, \text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ (\mathbf{sk}, m_0, m_1, \text{aux}, \text{st}) \leftarrow D_b \end{array} \right] = \epsilon(\lambda).$$

By the above and by Claim 5.2 we get that,

$$\frac{1}{2} + \frac{\Delta(D_0, D_1)}{2} \geq \epsilon(\lambda). \quad (4)$$

Let G be the following distribution,

$$G := \left\{ (\mathbf{sk}, m_0, m_1, \text{aux}) \mid \begin{array}{l} \mathbf{sk} \leftarrow \text{SKE.Gen}(1^\lambda) \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\mathbf{sk}, \cdot)} \end{array} \right\}.$$

Fix $(\text{sk}, \text{m}_0, \text{m}_1, \text{aux}) \in G$, and for every $b \in \{0, 1\}$ let $B_{(\text{sk}, \text{m}_0, \text{m}_1, \text{aux}), b}$ be the following distribution,

$$B_{(\text{sk}, \text{m}_0, \text{m}_1, \text{aux}), b} := \left\{ \text{st} \mid \begin{array}{l} \text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, \text{m}_b) \\ \text{st} \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\text{sk}, \cdot)}(\text{aux}, \text{ct}) \end{array} \right\}.$$

Note that for every $b \in \{0, 1\}$,

$$D_b = \left\{ (x, y) \mid \begin{array}{l} x \leftarrow G \\ y \leftarrow B_{x, b} \end{array} \right\}.$$

By Claim 5.9,

$$\Delta(D_0, D_1) = \mathbb{E}_{x \leftarrow G}[\Delta(B_{x, 0}, B_{x, 1})].$$

Therefore, by Equation 4,

$$\frac{1}{2} + \frac{\mathbb{E}_{x \leftarrow G}[\Delta(B_{x, 0}, B_{x, 1})]}{2} \geq \epsilon(\lambda). \quad (5)$$

Let Predict be the algorithm from Claim 5.10 with $\mathcal{X} := \{0, 1\}^{S_{\max}}$, and $\xi := \frac{1}{2^{\log^2 \lambda}}$. Let \mathbf{A}_{SKE} be the following adversary to the SKE scheme.

$\mathbf{A}_{\text{SKE}}^{\text{SKE.Enc}(\text{sk}, \cdot)}$:

1. Emulate $\mathbf{A}_1^{\text{SKE.Enc}(\text{sk}, \cdot)}$ to get $(\text{aux}', \text{m}_0, \text{m}_1)$.
2. Set $\text{aux} := (\text{aux}', \text{m}_0, \text{m}_1)$.
3. Output $(\text{aux}, \text{m}_0, \text{m}_1)$.

$\mathbf{A}_{\text{SKE}}^{\text{SKE.Enc}(\text{sk}, \cdot)}(\text{aux}, \text{ct})$:

1. Parse $\text{aux} := (\text{aux}', \text{m}_0, \text{m}_1)$.
2. Emulate $\text{st} \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\text{sk}, \cdot)}(\text{aux}', \text{ct})$.
3. Emulate $b' \leftarrow \text{Predict}(\text{st})$, where for each oracle query $B_{x, b}$ the adversary
 - (a) Query the oracle with m_b to get $\text{ct} \leftarrow \text{SKE.Enc}(\text{sk}, \text{m}_b)$.
 - (b) Emulate $\text{st} \leftarrow \mathbf{A}_1^{\text{SKE.Enc}(\text{sk}, \cdot)}(\text{aux}', \text{ct})$, and return st .
4. Output b' .

It holds that,

$$\begin{aligned} & \Pr \left[\mathbf{A}_{\text{SKE}}^{\text{IE.Enc}(\text{sk}, \cdot)}(\text{aux}, \text{ct}) = b \mid \begin{array}{l} \text{sk} \leftarrow \text{SKE.Gen}(1^\lambda) \\ b \leftarrow \{0, 1\} \\ (\text{aux}, \text{m}_0, \text{m}_1) \leftarrow \mathbf{A}_{\text{SKE}}^{\text{SKE.Enc}(\text{sk}, \cdot)} \\ \text{ct} \leftarrow \text{PKE.Enc}(\text{pk}, \text{m}_b) \end{array} \right] \\ &= \Pr \left[\text{Predict}^{B_{x, 0}, B_{x, 1}}(\text{st}) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x = (\text{sk}, \text{m}_0, \text{m}_1, \text{aux}) \leftarrow G \\ \text{st} \leftarrow B_{b, x} \end{array} \right] \\ &\geq \frac{1}{2} + \frac{\mathbb{E}_{x \leftarrow G}[\Delta(B_{x, 0}, B_{x, 1})]}{2} - \xi \\ &\geq \epsilon(\lambda) - \xi \\ &= \epsilon(\lambda) - \text{negl}(\lambda). \end{aligned}$$

where the first equality is by construction, the first inequality is by Claim 5.10, the second inequality is by Equation 5, and the second equality is since $\xi = \frac{1}{2^{\log^2 \lambda}}$. Overall, we get that,

$$\Pr \left[\mathbf{A}_{\text{SKE}}^{\text{IE.Enc}(\text{sk}, \cdot)}(\text{aux}, \text{ct}) = b \mid \begin{array}{l} \text{sk} \leftarrow \text{SKE.Gen}(1^\lambda) \\ b \leftarrow \{0, 1\} \\ (\text{aux}, m_0, m_1) \leftarrow \mathbf{A}_{\text{SKE}}^{\text{SKE.Enc}(\text{sk}, \cdot)} \\ \text{ct} \leftarrow \text{SKE.Enc}(\text{pk}, m_b) \end{array} \right] \geq \epsilon(\lambda) - \text{negl}(\lambda).$$

To complete the proof, in what follows, we bound the size of \mathbf{A}_{SKE} .

Claim 4.5. \mathbf{A}_{SKE} can be implemented with a circuit of size $2^{S_{\max} + O(\log S_{\max})} \cdot 2^{3 \cdot \log^2 \lambda} \cdot T^{\text{IE}}(\lambda)$.

Proof. The \mathbf{A}_{SKE} algorithm consists of three main steps:

1. *Emulating \mathbf{A}_1* : This step simulates \mathbf{A}_1 , whose size is at most $T^{\text{IE}}(\lambda)$.
2. *Sampling Elements from $B_{x,b}$* : Each sample requires to emulate \mathbf{A}_1 , whose size is at most $T^{\text{IE}}(\lambda)$.
3. *Emulating Predict*: By Claim 5.10, for $t := \Omega\left(\frac{|\mathcal{X}| + \log\left(\frac{1}{\xi}\right)}{\xi^2}\right)$, the algorithm Predict runs in time $O(t \cdot \log |\mathcal{X}|)$ and makes t oracles queries. We answer each query by simulating $T^{\text{IE}}(\lambda)$. Therefore, this algorithm can be implemented with circuit size $\tilde{O}(t \cdot \log |\mathcal{X}|) + t \cdot T^{\text{IE}}(\lambda)$.

We got that,

$$\begin{aligned} |\mathbf{A}_{\text{SKE}}| &= \tilde{O}(t \cdot \log |\mathcal{X}|) + t \cdot |\mathbf{A}_1| \\ &\leq \tilde{O}(t \cdot \log |\mathcal{X}|) + t \cdot T^{\text{IE}}(\lambda). \end{aligned}$$

Note that since $\mathcal{X} = \{0, 1\}^{S_{\max}}$, and $\xi = \frac{1}{2^{\log^2 \lambda}}$, it holds that, $t = 2^{S_{\max}} \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)}$, we get that,

$$\begin{aligned} t &= \Omega\left(\frac{|\mathcal{X}| + \log\left(\frac{1}{\xi}\right)}{\xi^2}\right) \\ &= \Omega\left((2^{S_{\max}} + \log^2 \lambda) \cdot 2^{2 \cdot \log^2 \lambda}\right) \\ &\leq 2^{S_{\max}} \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)} \end{aligned}$$

By the above,

$$\begin{aligned} |\mathbf{A}_{\text{SKE}}| &= \tilde{O}(t \cdot \log |\mathcal{X}|) + t \cdot T^{\text{IE}}(\lambda) \\ &= \tilde{O}(2^{S_{\max}} \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)} \cdot S_{\max}) + 2^{S_{\max}} \cdot 2^{2 \cdot (\log^2 \lambda + \log \log \lambda)} \cdot T^{\text{IE}}(\lambda) \\ &\leq 2^{S_{\max} + O(\log S_{\max})} \cdot 2^{3 \cdot \log^2 \lambda} \cdot T^{\text{IE}}(\lambda). \end{aligned}$$

□

4.3 Proofs of the Corollaries

In this section, we provide a proof for Theorem 4.2.

Proof of Corollary 4.3. Let $(\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$ be a SKE scheme that is secure against adversaries of size $T(\lambda) = 2^{\lambda - \log^2 \lambda}$. Let $S(\lambda) \leq \lambda - \log^4 \lambda$, and let $T^{\text{IE}}(\lambda) = \lambda^{\log \lambda}$.

To prove that the SKE scheme is everlasting SKE scheme with (T^{IE}, S) -security, we need to prove that,

$$T(\lambda) \geq 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda)} \cdot T^{\text{IE}}(\lambda). \quad (6)$$

Then, by Theorem 4.2, we can conclude that the scheme is everlasting SKE with $(T^{\text{IE}}, S_{\max})$ -security. By definition of T^{IE} and S_{\max} ,

$$\begin{aligned} & 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda)} \cdot T^{\text{IE}}(\lambda) \\ & \leq 2^{(\lambda - \log^4 \lambda + O(\log(\lambda - \log^4 \lambda)) + 3 \cdot \log^2 \lambda)} \cdot \lambda^{\log \lambda} \\ & \leq 2^{(\lambda - \log^4 \lambda + O(\log \lambda) + 4 \cdot \log^2 \lambda)} \\ & \leq 2^{(\lambda - \log^2 \lambda)} = T(\lambda). \end{aligned}$$

Overall, we get that Equation 6 holds, as required. \square

Proof of Corollary 4.4. We start by proving Item 1 of the corollary.

Fix a SKE scheme, $(\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$, that is sub-exponentially secure. Therefore, there exists some $\epsilon \in (0, 1)$ such that for $T(\lambda) := 2^{\lambda^\epsilon}$, the SKE scheme is T -secure. Let $L_{\text{pk}}(\lambda)$ be the public-key size, let $L_{\text{sk}}(\lambda)$ be the secret-key size, and let $L_{\text{ct}}(\ell_m, \lambda)$ be the ciphertext size for message length ℓ_m . Let $\epsilon' = \frac{\epsilon}{2}$. In what follows we give a construction, which we later prove as everlasting IE with T^{IE} -security for $T^{\text{IE}}(\lambda) := 2^{\lambda^{\epsilon'}}$.

Construction 4.6. The construction is as follows,

- $\text{IE.Gen}(1^\lambda, 1^S)$:
 1. Set $\tilde{\lambda} := (S + O(\log S) + 3 \cdot \log^2 \lambda + \lambda^{\epsilon'})^{\frac{1}{\epsilon}}$.
 2. Output $(\text{pk}, \text{sk}) \leftarrow \text{SKE.Gen}(1^{\tilde{\lambda}})$.
- $\text{IE.Enc}(\text{pk}, m)$: Output $\text{SKE.Enc}(\text{pk}, m)$.
- $\text{IE.Dec}(\text{sk}, \text{ct})$: Output $\text{SKE.Dec}(\text{sk}, \text{ct})$.

Fix some $S \in \mathbb{N}$. The above construction forms a SKE scheme that is T' -secure, where $T'(\lambda) := T(\tilde{\lambda}) = 2^{\tilde{\lambda}^\epsilon}$, and with efficiency parameters $L'_{\text{pk}}(\lambda) = L_{\text{pk}}(\tilde{\lambda})$, $L'_{\text{sk}}(\lambda) = L_{\text{sk}}(\tilde{\lambda})$, and $L'_{\text{ct}}(\ell_m, \lambda) = L_{\text{ct}}(\ell_m, \tilde{\lambda})$. In what follows, we prove that there exists some c such that $\tilde{\lambda} \leq S^c \cdot \lambda$. This will allow us to bound the efficiency parameters of the SKE scheme such that,

$$\begin{aligned} L'_{\text{pk}}(\lambda) & \leq L_{\text{pk}}(S^c \cdot \lambda), \\ L'_{\text{sk}}(\lambda) & \leq L_{\text{sk}}(S^c \cdot \lambda), \\ L'_{\text{ct}}(\ell_m, \lambda) & \leq L_{\text{ct}}(\ell_m, S^c \cdot \lambda). \end{aligned}$$

For large enough λ and $S > 1$,

$$\begin{aligned}
\tilde{\lambda} &= (S + O(\log S) + 3 \cdot \log^2 \lambda + \lambda^{\epsilon'})^{\frac{1}{\epsilon}} \\
&\leq (2 \cdot S + \lambda^{2\epsilon'})^{\frac{1}{\epsilon}} \\
&\leq (S \cdot \lambda^{2\epsilon'})^{\frac{1}{\epsilon}} \\
&= S^{\frac{1}{\epsilon}} \cdot \lambda^{\frac{2\epsilon'}{\epsilon}} \\
&= S^{\frac{1}{\epsilon}} \cdot \lambda,
\end{aligned}$$

where the last equality is since $\epsilon' = \frac{\epsilon}{2}$. Therefore, for $c := \frac{1}{\epsilon}$, it holds that,

$$\tilde{\lambda} \leq S^c \cdot \lambda,$$

as required.

In what follows, we prove that the T' -secure SKE scheme is also (T^{IE}, S) -secure everlasting IE. By the definition of $\tilde{\lambda}$,

$$\tilde{\lambda} = (S + O(\log S) + 3 \cdot \log^2 \lambda + \lambda^{\epsilon'})^{\frac{1}{\epsilon}}.$$

Raising both sides of the equation to the power of ϵ , we get that,

$$\tilde{\lambda}^\epsilon = S + O(\log S) + 3 \cdot \log^2 \lambda + \lambda^{\epsilon'}.$$

Therefore,

$$\begin{aligned}
2^{\tilde{\lambda}^\epsilon} &= 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda + \lambda^{\epsilon'})} \\
&= 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda)} \cdot 2^{\lambda^{\epsilon'}}.
\end{aligned}$$

Since $T'(\lambda) = 2^{\tilde{\lambda}^\epsilon}$, and $T^{\text{IE}}(\lambda) = 2^{\lambda^{\epsilon'}}$, we get that,

$$T'(\lambda) = 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda)} \cdot T^{\text{IE}}(\lambda).$$

Therefore, we can apply Theorem 3.2 on the T' -secure SKE scheme, and get that the scheme is also (T^{IE}, S) -secure everlasting IE scheme, as required.

Next, we prove Item 2 of the corollary. Fix a SKE scheme, $(\text{SKE.Gen}, \text{SKE.Enc}, \text{SKE.Dec})$, that is exponentially secure. Therefore, there exists some $\epsilon \in (0, 1)$ such that for $T(\lambda) := 2^{\epsilon \cdot \lambda}$, the SKE scheme is T -secure. Let $L_{\text{sk}}(\lambda)$ be the secret-key size, and let $L_{\text{ct}}(\ell_m, \lambda)$ be the ciphertext size for message length ℓ_m . Let $\epsilon' = \frac{\epsilon}{2}$. In what follows we give a construction, which we later prove as everlasting IE with T^{IE} -security for $T^{\text{IE}}(\lambda) := 2^{\epsilon' \cdot \lambda}$.

Construction 4.7. The construction is as follows,

• $\text{IE.Gen}(1^\lambda, 1^S)$:

1. Set $\tilde{\lambda} := \frac{1}{\epsilon} \cdot (S + O(\log S) + 3 \cdot \log^2 \lambda + \epsilon' \cdot \lambda)$.
2. Output $(\text{pk}, \text{sk}) \leftarrow \text{SKE.Gen}(1^{\tilde{\lambda}})$.

- IE.Enc(pk, m) : Output SKE.Enc(pk, m) .
- IE.Dec(sk, ct) : Output SKE.Dec(sk, ct) .

Fix some $S \in \mathbb{N}$. The above construction forms a SKE scheme that is T' -secure, where $T'(\lambda) := T(\tilde{\lambda}) = 2^{\epsilon \cdot \tilde{\lambda}}$, and with efficiency parameters $L'_{\text{pk}}(\lambda) = L_{\text{pk}}(\tilde{\lambda})$, $L'_{\text{sk}}(\lambda) = L_{\text{sk}}(\tilde{\lambda})$, and $L'_{\text{ct}}(\ell_m, \lambda) = L_{\text{ct}}(\ell_m, \tilde{\lambda})$. By construction,

$$\begin{aligned}\tilde{\lambda} &= \frac{1}{\epsilon} \cdot (S + O(\log S) + 3 \cdot \log^2 \lambda + \epsilon' \cdot \lambda) \\ &= O(S + \lambda),\end{aligned}$$

Therefore,

$$\begin{aligned}L'_{\text{pk}}(\lambda) &\leq L_{\text{pk}}(O(S + \lambda)), \\ L'_{\text{sk}}(\lambda) &\leq L_{\text{sk}}(O(S + \lambda)), \\ L'_{\text{ct}}(\ell_m, \lambda) &\leq L_{\text{ct}}(\ell_m, O(S + \lambda)).\end{aligned}$$

as required.

In what follows, we prove that the T' -secure SKE scheme is also (T^{IE}, S) -secure everlasting SKE. By the definition of $\tilde{\lambda}$,

$$\tilde{\lambda} = \frac{1}{\epsilon} \cdot (S + O(\log S) + 3 \cdot \log^2 \lambda + \epsilon' \cdot \lambda).$$

Therefore,

$$\epsilon \cdot \tilde{\lambda} = S + O(\log S) + 3 \cdot \log^2 \lambda + \epsilon' \cdot \lambda.$$

Raising 2 to the above exponent gives,

$$\begin{aligned}2^{\epsilon \cdot \tilde{\lambda}} &= 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda + \epsilon' \cdot \lambda)} \\ &= 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda)} \cdot 2^{\epsilon' \cdot \lambda}.\end{aligned}$$

Since $T'(\lambda) = 2^{\epsilon \cdot \tilde{\lambda}}$, and $T^{\text{IE}}(\lambda) = 2^{\epsilon' \cdot \lambda}$, we get that,

$$T'(\lambda) = 2^{(S + O(\log S) + 3 \cdot \log^2 \lambda)} \cdot T^{\text{IE}}(\lambda).$$

Therefore, we can apply Theorem 4.2 on the T' -secure SKE scheme, and get that the scheme is also (T^{IE}, S) -secure everlasting IE scheme, as required. \square

5 Estimating and Predicting Distributions

In this section, we present general claims and proofs concerning statistical distance, distinguishing distributions, and estimating distributions.

Claim 5.1 (Triangle inequality). *For every three distributions D_1, D_2, D_3 ,*

$$\Delta(D_1, D_3) \leq \Delta(D_1, D_2) + \Delta(D_2, D_3).$$

Proof. Let D_1, D_2, D_3 be distributions over the domain \mathcal{X} . By the definition of statistical distance,

$$\begin{aligned} \Delta(D_1, D_3) &= \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} |D_1(x) - D_3(x)| \\ &= \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} |D_1(x) - D_2(x) + D_2(x) - D_3(x)| \\ &\leq \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} |D_1(x) - D_2(x)| + |D_2(x) - D_3(x)| \\ &= \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} |D_1(x) - D_2(x)| + \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} |D_2(x) - D_3(x)| \\ &= \Delta(D_1, D_2) + \Delta(D_2, D_3). \end{aligned}$$

□

Claim 5.2. *For every pair of distributions D_0, D_1 over \mathcal{X} , and for any unbounded distinguisher A ,*

$$\Pr \left[A(x) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] \leq \frac{1}{2} + \frac{\Delta(D_0, D_1)}{2}.$$

Moreover, for a distinguisher Predict that is defined as follows,

$$\forall x \in \mathcal{X}, \text{ Predict}(x) := \begin{cases} 0 & \text{if } D_0(x) \geq D_1(x), \\ 1 & \text{if } D_0(x) < D_1(x). \end{cases}$$

It holds that,

$$\Pr \left[\text{Predict}(x) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] = \frac{1}{2} + \frac{\Delta(D_0, D_1)}{2}.$$

Proof. Fix D_0, D_1 over \mathcal{X} , and distinguisher A . By the law of total probability,

$$\begin{aligned} &\Pr \left[A(x) = b \mid \begin{array}{c} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] \\ &= \frac{1}{2} \cdot (\Pr[A(x) = 0 \mid x \leftarrow D_0] + \Pr[A(x) = 1 \mid x \leftarrow D_1]) \\ &= \frac{1}{2} \cdot \left(\sum_{x \in \mathcal{X}} D_0(x) \cdot \Pr[A(x) = 0] + D_1(x) \cdot \Pr[A(x) = 1] \right) \\ &= \frac{1}{2} \cdot \left(\sum_{x \in \mathcal{X}} D_0(x) \cdot \Pr[A(x) = 0] + D_1(x) \cdot (1 - \Pr[A(x) = 0]) \right) \\ &\leq \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} \max(D_0(x), D_1(x)), \end{aligned}$$

where the inequality is since the expression is maximized when the following holds for every $x \in \mathcal{X}$,

- If $D_0(x) \geq D_1(x)$, then $\Pr[A(x) = 0] = 1$.
- If $D_0(x) < D_1(x)$, then $(1 - \Pr[A(x) = 0]) = 1$, which means that $\Pr[A(x) = 1] = 1$.

We get that,

$$\Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] \geq \frac{1}{2} \cdot \sum_{x \in \mathcal{X}} \max(D_0(x), D_1(x)) ,$$

and that the following distinguisher **Predict** maximizes the above probability,

$$\forall x \in \mathcal{X}, \text{ Predict}(x) := \begin{cases} 0 & \text{if } D_0(x) \geq D_1(x), \\ 1 & \text{if } D_0(x) < D_1(x). \end{cases}$$

We left to prove that,

$$\frac{1}{2} \cdot \sum_{x \in \mathcal{X}} \max(D_0(x), D_1(x)) = \frac{1}{2} + \frac{\Delta(D_0, D_1)}{2},$$

which implies that every algorithm (even unbounded) has a success probability of at most $\frac{1}{2} + \frac{\Delta(D_0, D_1)}{2}$, and that there exists a $\text{poly}(|\mathcal{X}|)$ size circuit that equals this probability, as required. To prove the above, we prove that $\sum_{x \in \mathcal{X}} \max(D_0(x), D_1(x)) = 1 + \Delta(D_0, D_1)$. Let $A \subseteq \mathcal{X}$ be the group of all elements $x \in \mathcal{X}$ such that $D_0(x) \geq D_1(x)$. Therefore,

$$\begin{aligned} \sum_{x \in \mathcal{X}} \max\{D_0(x), D_1(x)\} &= \sum_{x \in A} D_0(x) + \sum_{x \in \mathcal{X} \setminus A} D_1(x) \\ &= \sum_{x \in A} D_0(x) + \sum_{x \in \mathcal{X} \setminus A} D_1(x) + \sum_{x \in A} D_1(x) - \sum_{x \in A} D_1(x) \\ &= \sum_{x \in \mathcal{X}} D_1(x) + \sum_{x \in A} D_0(x) - \sum_{x \in A} D_1(x) \\ &= 1 + \sum_{x \in A} D_0(x) - \sum_{x \in A} D_1(x). \end{aligned}$$

In what follows, we prove that $\sum_{x \in A} D_0(x) - \sum_{x \in A} D_1(x) = \Delta(D_0, D_1)$.

$$\begin{aligned} 2 \cdot \Delta(D_0, D_1) &= \sum_{x \in \mathcal{X}} |D_0(x) - D_1(x)| \\ &= \sum_{x \in A} (D_0(x) - D_1(x)) + \sum_{x \in \mathcal{X} \setminus A} (D_1(x) - D_0(x)) \\ &= \sum_{x \in A} D_0(x) - \sum_{x \in A} D_1(x) + \sum_{x \in \mathcal{X} \setminus A} D_1(x) - \sum_{x \in \mathcal{X} \setminus A} D_0(x) \\ &= \sum_{x \in A} D_0(x) - \sum_{x \in A} D_1(x) + \left(1 - \sum_{x \in A} D_1(x)\right) - \left(1 - \sum_{x \in A} D_0(x)\right) \\ &= 2 \cdot \left(\sum_{x \in A} D_0(x) - \sum_{x \in A} D_1(x) \right) \end{aligned} \tag{7}$$

where Equation 7 is since for every $b \in \{0, 1\}$, it holds that

$$1 = \sum_{x \in \mathcal{X}} D_b(x) = \sum_{x \in A} D_b(x) + \sum_{x \in \mathcal{X} \setminus A} D_b(x),$$

and therefore, $\sum_{x \in \mathcal{X} \setminus A} D_b(x) = 1 - \sum_{x \in A} D_b(x)$. \square

Claim 5.3. *Let $D_0, D_1, \tilde{D}_0, \tilde{D}_1$ be distributions over the domain \mathcal{X} . For any unbounded distinguisher A ,*

$$\Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] \geq \Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow \tilde{D}_b \end{array} \right] - \Delta(D_0, \tilde{D}_0) - \Delta(D_1, \tilde{D}_1).$$

Proof. For every $b \in \{0, 1\}$,

$$\begin{aligned} & \Pr [A(x) = b \mid x \leftarrow D_b] \\ &= \sum_{x \in \mathcal{X}} D_b(x) \cdot \Pr[A(x) = b] \\ &= \sum_{x \in \mathcal{X}} \left(\tilde{D}_b(x) \cdot \Pr[A(x) = b] - (\tilde{D}_b(x) - D_b(x)) \cdot \Pr[A(x) = b] \right) \\ &\geq \sum_{x \in \mathcal{X}} \left(\tilde{D}_b(x) \cdot \Pr[A(x) = b] - |\tilde{D}_b(x) - D_b(x)| \cdot \Pr[A(x) = b] \right) \\ &\geq \sum_{x \in \mathcal{X}} \left(\tilde{D}_b(x) \cdot \Pr[A(x) = b] - |D_b(x) - \tilde{D}_b(x)| \right) \\ &= \sum_{x \in \mathcal{X}} \tilde{D}_b(x) \cdot \Pr[A(x) = b] - \sum_{x \in \mathcal{X}} |D_b(x) - \tilde{D}_b(x)| \\ &= \Pr [A(x) = b \mid x \leftarrow \tilde{D}_b] - 2 \cdot \Delta(D_b, \tilde{D}_b). \end{aligned}$$

By the law of total probability,

$$\Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] = \sum_{b \in \{0, 1\}} \frac{1}{2} \cdot \Pr [A(x) = b \mid x \leftarrow D_b].$$

Overall, we get that,

$$\begin{aligned} \Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] &\geq \sum_{b \in \{0, 1\}} \frac{1}{2} \cdot \left(\Pr [A(x) = b \mid x \leftarrow \tilde{D}_b] - 2 \cdot \Delta(D_b, \tilde{D}_b) \right) \\ &= \Pr \left[A(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow \tilde{D}_b \end{array} \right] - \Delta(D_0, \tilde{D}_0) - \Delta(D_1, \tilde{D}_1). \end{aligned}$$

\square

Definition 5.4. *Let D be a distribution over the domain \mathcal{X} , and let Learn^D be the following oracle algorithm that given oracle access to the distribution D and input $\mathbf{t} \in \mathbb{N}$, approximates D as follows:*

$\text{Learn}^D(\mathbf{t})$:

1. For $i \in [t]$, sample $x_i \leftarrow D$.
2. For every $x \in \mathcal{X}$, set $\tilde{D}(x) := \Pr_{i \leftarrow [t]}[x_i = x]$.

Claim 5.5. For every distribution D over the domain \mathcal{X} , and for every $\xi, \delta \in (0, 1)$, if $t := \Omega(\frac{|\mathcal{X}| + \log(1/\delta)}{\xi^2})$, then

$$\mathbb{E} [\Delta(\tilde{D}, D) \mid \tilde{D} \leftarrow \text{Learn}^D(t)] \leq \xi + \delta.$$

Proof. It holds that,

$$\begin{aligned} & \mathbb{E} [\Delta(\tilde{D}, D) \mid \tilde{D} \leftarrow \text{Learn}^D(t)] \\ & \leq \Pr[\Delta(\tilde{D}, D) \leq \xi] \cdot \xi + \Pr[\Delta(\tilde{D}, D) > \xi] \cdot 1 \\ & \leq \xi + \Pr[\Delta(\tilde{D}, D) > \xi]. \end{aligned}$$

By Claim A.1,

$$\Pr[\Delta(\tilde{D}, D) > \xi] \leq \delta.$$

Therefore,

$$\mathbb{E} [\Delta(\tilde{D}, D) \mid \tilde{D} \leftarrow \text{Learn}^D(t)] \leq \xi + \delta.$$

□

Claim 5.6. For every two distributions D_0, D_1 over the domain \mathcal{X} , and for every $t \in \mathbb{N}$,

$$\mathbb{E} \left[\Delta(\tilde{D}_0, \tilde{D}_1) \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(t) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(t) \end{array} \right] \geq \Delta(D_0, D_1) - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_0) \\ + \Delta(\tilde{D}_1, D_1) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(t) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(t) \end{array} \right].$$

Proof. Fix some distributions $D_0, D_1, \tilde{D}_0, \tilde{D}_1$. By the triangle inequality (Claim 5.1),

$$\begin{aligned} \Delta(D_0, D_1) & \leq \Delta(D_0, \tilde{D}_0) + \Delta(\tilde{D}_0, D_1) \\ & \leq \Delta(D_0, \tilde{D}_0) + \Delta(\tilde{D}_0, \tilde{D}_1) + \Delta(\tilde{D}_1, D_1). \end{aligned}$$

Therefore,

$$\Delta(\tilde{D}_0, \tilde{D}_1) \geq \Delta(D_0, D_1) - \Delta(D_0, \tilde{D}_0) - \Delta(\tilde{D}_1, D_1).$$

We get that,

$$\begin{aligned} & \mathbb{E} \left[\Delta(\tilde{D}_0, \tilde{D}_1) \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(t) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(t) \end{array} \right] \\ & \geq \mathbb{E} \left[\Delta(D_0, D_1) - \Delta(D_0, \tilde{D}_0) - \Delta(\tilde{D}_1, D_1) \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(t) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(t) \end{array} \right] \end{aligned}$$

By linearity of expectation, we get that,

$$\begin{aligned} & \mathbb{E} \left[\Delta(\tilde{D}_0, \tilde{D}_1) \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(t) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(t) \end{array} \right] \\ & \geq \Delta(D_0, D_1) - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_0) \\ + \Delta(\tilde{D}_1, D_1) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(t) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(t) \end{array} \right]. \end{aligned}$$

□

Claim 5.7. For every two distributions D_0, D_1 over the domain \mathcal{X} , and for every $\mathbf{t} \in \mathbb{N}$,

$$\Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow D_b \end{array} \right] \geq \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow \tilde{D}_b \end{array} \right] \\ - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_0) \\ + \Delta(\tilde{D}_1, D_1) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right].$$

Proof of Claim 5.7. For every $x_1 \dots, x_t \in \mathcal{X}$, let D_{x_1, \dots, x_t} be the distribution such that, for every $x \in \mathcal{X}$, $D_{x_1, \dots, x_t}(x) := \Pr_{i \leftarrow [t]}[x_i = x]$, and let $\mathcal{D} := \{D_{x_1, \dots, x_t} \mid x_1 \dots, x_t \in \mathcal{X}\}$. By the law of total probability,

$$\Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow D_b \end{array} \right] \\ = \sum_{\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}} \Pr \left[\begin{array}{l} \tilde{D}'_0 = \tilde{D}_0 \\ \tilde{D}'_1 = \tilde{D}_1 \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \cdot \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right]$$

Fix two distributions $\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}$. By Claim 5.3,

$$\Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow D_b \end{array} \right] \\ \geq \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow \tilde{D}_b \end{array} \right] - \Delta(\tilde{D}_0, D_0) - \Delta(\tilde{D}_1, D_1).$$

Therefore, by the law of total probability,

$$\Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow D_b \end{array} \right] \\ \geq \sum_{\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}} \Pr \left[\begin{array}{l} \tilde{D}'_0 = \tilde{D}_0 \\ \tilde{D}'_1 = \tilde{D}_1 \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \cdot \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow \tilde{D}_b \end{array} \right] \\ - \sum_{\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}} \Pr \left[\begin{array}{l} \tilde{D}'_0 = \tilde{D}_0 \\ \tilde{D}'_1 = \tilde{D}_1 \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \cdot \Delta(\tilde{D}_0, D_0) \\ - \sum_{\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}} \Pr \left[\begin{array}{l} \tilde{D}'_0 = \tilde{D}_0 \\ \tilde{D}'_1 = \tilde{D}_1 \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \cdot \Delta(\tilde{D}_1, D_1) \\ = \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow \tilde{D}_b \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{D}_0, D_0) \\ + \Delta(\tilde{D}_1, D_1) \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right].$$

□

Claim 5.8. For every two distributions D_0, D_1 over the domain \mathcal{X} , and for every $\mathbf{t} \in \mathbb{N}$,

$$\Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow \tilde{D}_b \end{array} \right] = \frac{1}{2} + \mathbb{E} \left[\frac{\Delta(\tilde{D}_0, \tilde{D}_1)}{2} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right].$$

Proof. For every $x_1, \dots, x_t \in \mathcal{X}$, let D_{x_1, \dots, x_t} be the distribution such that, for every $x \in \mathcal{X}$, $D_{x_1, \dots, x_t}(x) := \Pr_{i \leftarrow [t]}[x_i = x]$, and let $\mathcal{D} := \{D_{x_1, \dots, x_t} \mid x_1, \dots, x_t \in \mathcal{X}\}$. By the law of total probability,

$$\begin{aligned} & \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow \tilde{D}_b \end{array} \right] \\ &= \sum_{\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}} \Pr \left[\begin{array}{l} \tilde{D}'_0 = \tilde{D}_0 \\ \tilde{D}'_1 = \tilde{D}_1 \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \cdot \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow \tilde{D}_b \end{array} \right] \end{aligned}$$

Fix two distributions $\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}$. By Claim 5.2,

$$\Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow \tilde{D}_b \end{array} \right] = \frac{1}{2} + \frac{\Delta(\tilde{D}_0, \tilde{D}_1)}{2}.$$

Therefore,

$$\begin{aligned} & \Pr \left[\text{Predict}_{\tilde{D}_0, \tilde{D}_1}(x) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \\ x \leftarrow \tilde{D}_b \end{array} \right] \\ &= \sum_{\tilde{D}_0, \tilde{D}_1 \in \mathcal{D}} \Pr \left[\begin{array}{l} \tilde{D}'_0 = \tilde{D}_0 \\ \tilde{D}'_1 = \tilde{D}_1 \end{array} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \cdot \left(\frac{1}{2} + \frac{\Delta(\tilde{D}_0, \tilde{D}_1)}{2} \right) \\ &= \frac{1}{2} + \mathbb{E} \left[\frac{\Delta(\tilde{D}_0, \tilde{D}_1)}{2} \mid \begin{array}{l} \tilde{D}_0 \leftarrow \text{Learn}^{D_0}(\mathbf{t}) \\ \tilde{D}_1 \leftarrow \text{Learn}^{D_1}(\mathbf{t}) \end{array} \right] \end{aligned}$$

□

Claim 5.9. For every two distributions D_0, D_1 , if there exists a distribution G , and set of distributions $\{B_{0,x}\}_x$ and $\{B_{1,x}\}_x$ such that, for every $b \in \{0, 1\}$,

$$D_b := \left\{ (x, y) \mid \begin{array}{l} x \leftarrow G \\ y \leftarrow B_{x,b} \end{array} \right\}.$$

Then,

$$\Delta(D_0, D_1) = \mathbb{E}_{x \leftarrow G}[\Delta(B_{x,0}, B_{x,1})].$$

Proof. By definition of statistical distance,

$$\Delta(D_0, D_1) = \frac{1}{2} \sum_{(x', y')} |\Pr[(x, y) = (x', y') \mid (x, y) \leftarrow D_0] - \Pr[(x, y) = (x', y') \mid (x, y) \leftarrow D_1]|.$$

By the definition of D_0 and D_1 ,

$$\begin{aligned} & \frac{1}{2} \sum_{(x', y')} |\Pr[(x, y) = (x', y') \mid (x, y) \leftarrow D_0] - \Pr[(x, y) = (x', y') \mid (x, y) \leftarrow D_1]| \\ &= \frac{1}{2} \sum_{(x', y')} \left| \Pr \left[(x, y) = (x', y') \mid \begin{array}{l} x \leftarrow G \\ y \leftarrow B_{0,x} \end{array} \right] - \Pr \left[(x, y) = (x', y') \mid \begin{array}{l} x \leftarrow G \\ y \leftarrow B_{1,x} \end{array} \right] \right| \\ &= \frac{1}{2} \sum_{(x', y')} |\Pr[x = x' \mid x \leftarrow G] \cdot \Pr[y = y' \mid y \leftarrow B_{0,x}] - \Pr[x = x' \mid x \leftarrow G] \cdot \Pr[y = y' \mid y \leftarrow B_{1,x}]| \\ &= \frac{1}{2} \sum_{(x', y')} \Pr[x = x' \mid x \leftarrow G] \cdot |\Pr[y \leftarrow B_{0,x}] - \Pr[y = y' \mid y \leftarrow B_{1,x}]| \\ &= \sum_{x'} \Pr[x = x' \mid x \leftarrow G] \cdot \frac{1}{2} \sum_{y'} |\Pr[y = y' \mid y \leftarrow B_{0,x}] - \Pr[y = y' \mid y \leftarrow B_{1,x}]|. \end{aligned}$$

By the definition of statistical distance,

$$\begin{aligned} & \sum_x \Pr[x = x' \mid x \leftarrow G] \cdot \frac{1}{2} \sum_y |\Pr[y = y' \mid y \leftarrow B_{0,x}] - \Pr[y = y' \mid y \leftarrow B_{1,x}]| \\ &= \sum_x \Pr[x = x' \mid x \leftarrow G] \cdot \Delta(B_{0,x}, B_{1,x}) \\ &= \mathbb{E}_{x \leftarrow G}[\Delta(B_{0,x}, B_{1,x})], \end{aligned}$$

where the last equality is by the definition of expected value. Overall, we get that,

$$\Delta(D_0, D_1) = \mathbb{E}_{x \leftarrow G}[\Delta(B_{0,x}, B_{1,x})],$$

as required. \square

Claim 5.10. *For every domain \mathcal{X} and $\xi \in (0, 1)$, there exist an algorithm **Predict** such that for every distribution G over the domain \mathcal{X} and set of distributions $\{B_{0,x}\}_{x \in \mathcal{X}}$ and $\{B_{1,x}\}_{x \in \mathcal{X}}$, given oracle access to samples from $B_{x,0}, B_{x,1}$ it holds that for ,*

$$\Pr \left[\text{Predict}^{B_{x,0}, B_{x,1}}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow G \\ y \leftarrow B_{b,x} \end{array} \right] = \frac{1}{2} + \frac{\mathbb{E}_{x \leftarrow G}[\Delta(B_{x,0}, B_{x,1})]}{2} - \xi.$$

Moreover, **Predict** makes $t := \Omega\left(\frac{|\mathcal{X}| + \log\left(\frac{1}{\xi}\right)}{\xi^2}\right)$ oracle queries, and runs in time $O(t \cdot \log(|\mathcal{X}|))$.

Proof. Let **Predict** be the following algorithm,

Predict ^{$B_{x,0}, B_{x,1}$} (y):

1. For every $b \in \{0, 1\}$, set $\tilde{B}_b \leftarrow \text{Learn}^{B_{x,b}}(\mathbf{t})$, where $\text{Learn}^{B_{x,b}}(\mathbf{t})$ is the algorithm from Definition 5.4.
2. Output $b' \leftarrow \text{Predict}_{\tilde{B}_0, \tilde{B}_1}(y)$, where $\text{Predict}_{\tilde{B}_0, \tilde{B}_1}$ is the algorithm described in Claim 5.2 for distributions \tilde{B}_0, \tilde{B}_1 .

By the above,

$$\begin{aligned} & \Pr \left[\text{Predict}^{B_{x,0}, B_{x,1}}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow G \\ y \leftarrow B_{x,b} \end{array} \right] \\ &= \mathbb{E}_{x \leftarrow G} \left[\Pr \left[\text{Predict}_{\tilde{B}_0, \tilde{B}_1}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \\ y \leftarrow B_{x,b} \end{array} \right] \right]. \end{aligned} \quad (8)$$

Fix x . It holds that,

$$\begin{aligned} & \Pr \left[\text{Predict}_{\tilde{B}_0, \tilde{B}_1}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \\ y \leftarrow B_{x,b} \end{array} \right] \\ & \geq \Pr \left[\text{Predict}_{\tilde{B}_0, \tilde{B}_1}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \\ y \leftarrow \tilde{B}_b \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{B}_0, B_{x,0}) \\ + \Delta(\tilde{B}_1, B_{x,1}) \end{array} \mid \begin{array}{l} \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \end{array} \right] \\ &= \frac{1}{2} + \mathbb{E} \left[\frac{\Delta(\tilde{B}_0, \tilde{B}_1)}{2} \mid \begin{array}{l} \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{B}_0, B_{x,0}) \\ + \Delta(\tilde{B}_1, B_{x,1}) \end{array} \mid \begin{array}{l} \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \end{array} \right] \\ &= \frac{1}{2} + \frac{1}{2} \cdot \mathbb{E} \left[\Delta(\tilde{B}_0, \tilde{B}_1) \mid \begin{array}{l} \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \end{array} \right] - \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{B}_0, B_{x,0}) \\ + \Delta(\tilde{B}_1, B_{x,1}) \end{array} \mid \begin{array}{l} \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \end{array} \right] \\ & \geq \frac{1}{2} + \frac{\Delta(B_{x,0}, B_{x,1})}{2} - \frac{3}{2} \cdot \mathbb{E} \left[\begin{array}{l} \Delta(\tilde{B}_0, B_{x,0}) \\ + \Delta(\tilde{B}_1, B_{x,1}) \end{array} \mid \begin{array}{l} \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \end{array} \right] \\ &= \frac{1}{2} + \frac{\Delta(B_{x,0}, B_{x,1})}{2} - \frac{3}{2} \cdot \sum_{b \in \{0,1\}} \mathbb{E} \left[\Delta(\tilde{B}_b, B_{x,b}) \mid \tilde{B}_b \leftarrow \text{Learn}^{B_{x,b}}(\mathbf{t}) \right], \end{aligned}$$

where the first inequality is by Claim 5.7, the first equality is by linearity of expectation, the second equality is by Claim 5.8, the second inequality is by Claim 5.6, and the third equality is by linearity of expectation. Next, we apply Claim 5.5 with $\xi' = \xi/3$ and $\delta' := \xi/3$. Note that,

$$\mathbf{t} = \Omega \left(\frac{|\mathcal{X}| + \log \left(\frac{1}{\xi} \right)}{\xi^2} \right) = \Omega \left(\frac{|\mathcal{X}| + \log \left(\frac{1}{\delta'} \right)}{\xi'^2} \right).$$

Therefore, by Claim 5.5, for every $b \in \{0, 1\}$,

$$\mathbb{E} \left[\Delta(\tilde{B}_b, B_{x,b}) \mid \tilde{B}_b \leftarrow \text{Learn}^{B_{x,b}}(\mathbf{t}) \right] \leq \xi' + \delta' = \frac{2}{3}\xi.$$

Therefore,

$$\begin{aligned}
& \Pr \left[\text{Predict}_{\tilde{B}_0, \tilde{B}_1}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \\ y \leftarrow B_{x,b} \end{array} \right] \\
&= \frac{1}{2} + \frac{\Delta(B_{x,0}, B_{x,1})}{2} - \frac{3}{2} \cdot \sum_{b \in \{0,1\}} \mathbb{E} \left[\Delta(\tilde{B}_b, B_{x,b}) \mid \tilde{B}_b \leftarrow \text{Learn}^{B_{x,b}}(\mathbf{t}) \right] \\
&\geq \frac{1}{2} + \frac{\Delta(B_{x,0}, B_{x,1})}{2} - \frac{3}{2} \cdot \frac{2}{3} \xi \\
&= \frac{1}{2} + \frac{\Delta(B_{x,0}, B_{x,1})}{2} - \xi.
\end{aligned}$$

Overall, by plugging in the above into Equation 8,

$$\begin{aligned}
& \Pr \left[\text{Predict}^{B_{x,0}, B_{x,1}}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ x \leftarrow G \\ y \leftarrow B_{b,x} \end{array} \right] \\
&= \mathbb{E}_{x \leftarrow G} \left[\Pr \left[\text{Predict}_{\tilde{B}_0, \tilde{B}_1}(y) = b \mid \begin{array}{l} b \leftarrow \{0, 1\} \\ \tilde{B}_0 \leftarrow \text{Learn}^{B_{x,0}}(\mathbf{t}) \\ \tilde{B}_1 \leftarrow \text{Learn}^{B_{x,1}}(\mathbf{t}) \\ y \leftarrow B_{x,b} \end{array} \right] \right] \\
&\geq \mathbb{E}_{x \leftarrow G} \left[\frac{1}{2} + \frac{\Delta(B_{x,0}, B_{x,1})}{2} - \xi \right] \\
&\geq \frac{1}{2} + \frac{\mathbb{E}_{x \leftarrow G}[\Delta(B_{x,0}, B_{x,1})]}{2} - \xi,
\end{aligned}$$

where the last inequality is by linearity of expectation. □

Acknowledgments

Shany Ben-David is supported by the European Research Union (ERC, CRYPTOPROOF, 101164375), and by the Clore Israel Foundation. Eylon Yogev is supported by the Israel Science Foundation (Grant No. 2302/22) and the European Research Union (ERC, CRYPTOPROOF, 101164375). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- [ABY25] Gal Arnon, Shany Ben-David, and Eylon Yogev. “Instance Compression, Revisited”. In: *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part IV*. Lecture Notes in Computer Science. 2025, pp. 244–273.
- [BDD22] Pedro Branco, Nico Döttling, and Jesko Dujmovic. “Rate-1 Incompressible Encryption from Standard Assumptions”. In: *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part II*. Ed. by Eike Kiltz and Vinod Vaikuntanathan. Vol. 13748. Lecture Notes in Computer Science. Springer, 2022, pp. 33–69.
- [BGKNPR24] Kaartik Bhushan, Rishab Goyal, Venkata Koppula, Varun Narayanan, Manoj Prabhakaran, and Mahesh Sreekumar Rajasree. “Leakage-Resilient Incompressible Cryptography: Constructions and Barriers”. In: *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part VII*. Lecture Notes in Computer Science. 2024, pp. 201–234.
- [Can20] Clément L. Canonne. *A Toolkit for Learning and Testing Probability Distributions*. <https://github.com/ccanonne/probabilitydistributiontoolbox/blob/master/learning.pdf>. Accessed: 2025-02-09. 2020.
- [DGO19] Ivan Damgård, Chaya Ganesh, and Claudio Orlandi. “Proofs of Replicated Storage Without Timing Assumptions”. In: *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11692. Lecture Notes in Computer Science. Springer, 2019, pp. 355–380.
- [DKMMMMQT25] Nico Döttling, Alexander Koch, Sven Maier, Jeremias Mechler, Anne Müller, Jörn Müller-Quade, and Marcel Tiepelt. *The Quantum Decoherence Model: Everlasting Composable Secure Computation and More*. Cryptology ePrint Archive, Paper 2025/220. 2025. URL: <https://eprint.iacr.org/2025/220>.
- [Dzi06] Stefan Dziembowski. “On Forward-Secure Storage”. In: *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*. Ed. by Cynthia Dwork. Vol. 4117. Lecture Notes in Computer Science. Springer, 2006, pp. 251–270.
- [GKRV25] Rishab Goyal, Venkata Koppula, Mahesh Sreekumar Rajasree, and Aman Verma. “Incompressible Functional Encryption”. In: *16th Innovations in Theoretical Computer Science Conference, ITCS 2025, January 7-10, 2025, Columbia University, New York, NY, USA*. LIPIcs. 2025, 56:1–56:22.

- [GLW20] Rachit Garg, George Lu, and Brent Waters. “New Techniques in Replica Encodings with Client Setup”. In: *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part III*. Ed. by Rafael Pass and Krzysztof Pietrzak. Vol. 12552. Lecture Notes in Computer Science. Springer, 2020, pp. 550–583.
- [GWZ22] Jiaxin Guan, Daniel Wichs, and Mark Zhandry. “Incompressible Cryptography”. In: *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part I*. EUROCRYPT ’22. 2022, pp. 700–730.
- [GWZ23] Jiaxin Guan, Daniel Wichs, and Mark Zhandry. “Multi-instance Randomness Extraction and Security Against Bounded-Storage Mass Surveillance”. In: *Theory of Cryptography - 21st International Conference, TCC 2023, Taipei, Taiwan, November 29 - December 2, 2023, Proceedings, Part III*. Ed. by Guy N. Rothblum and Hoeteck Wee. Vol. 14371. Lecture Notes in Computer Science. Springer, 2023, pp. 93–122.
- [GZ21] Jiaxin Guan and Mark Zhandry. “Disappearing Cryptography in the Bounded Storage Model”. In: *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part II*. Ed. by Kobbi Nissim and Brent Waters. Vol. 13043. Lecture Notes in Computer Science. Springer, 2021, pp. 365–396.
- [HN10] Danny Harnik and Moni Naor. “On the Compressibility of NP Instances and Cryptographic Applications”. In: *SIAM J. Comput.* 39.5 (2010), pp. 1667–1713.
- [MW20] Tal Moran and Daniel Wichs. “Incompressible Encodings”. In: *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part I*. Ed. by Daniele Micciancio and Thomas Ristenpart. Vol. 12170. Lecture Notes in Computer Science. Springer, 2020, pp. 494–523.

A Learning distributions – a Folklore Lemma

We use the following “folklore” lemma from Learning theory, which quantifies the number of queries needed to learn a distribution with high probability up to a small error in statistical distance. The proof are reproduced verbatim from [Can20, Theorem 1].

Claim A.1 ([Can20, Theorem 1]). *Let D be a distribution over the domain \mathcal{X} . For every $\delta, \xi \in (0, 1)$, if $\mathbf{t} := \Omega(\frac{|\mathcal{X}| + \log(1/\delta)}{\xi^2})$, then it holds that*

$$\Pr \left[\Delta(D, \tilde{D}) \leq \xi \mid \tilde{D} \leftarrow \text{Learn}^D(\mathbf{t}) \right] \geq 1 - \delta,$$

where Learn^D is the algorithm described in Definition 5.4.

Proof. Consider the empirical distribution \tilde{D} obtained by drawing n independent samples s_1, \dots, s_n from the underlying distribution D :

$$\tilde{D}(x) = \frac{1}{n} \sum_{j=1}^n \mathbb{1}_{\{s_j=x\}}, \quad x \in \mathcal{X}. \quad (9)$$

We will analyze the behavior of the empirical distribution \tilde{D} . The event where $\Delta(D, \tilde{D}) > \xi$ means that there exists a subset $S \subseteq \mathcal{X}$ such that $\tilde{D}(S) > D(S) + \xi$. There are $2^{|\mathcal{X}|}$ such subsets. Let us union bound over the sets.

Fix any $S \subseteq \mathcal{X}$. We have,

$$\tilde{D}(S) = \sum_{x \in S} \tilde{D}(x) \stackrel{(9)}{=} \frac{1}{n} \sum_{x \in S} \sum_{j=1}^n \mathbb{1}_{\{s_j=x\}},$$

and so, letting $X_j := \sum_{i \in S} \mathbb{1}_{\{s_j=i\}}$ for $j \in [n]$, we have $\tilde{D}(S) = \frac{1}{n} \sum_{j=1}^n X_j$ where the X_j 's are i.i.d. Bernoulli random variable with parameter $D(S)$. In what follows, we use the Hoeffding bound,

$$\Pr \left[\tilde{D}(S) > D(S) + \xi \right] = \Pr \left[\frac{1}{n} \sum_{j=1}^n X_j > \mathbb{E} \left[\frac{1}{n} \sum_{j=1}^n X_j \right] + \xi \right] \leq e^{-2\xi^2 n},$$

and therefore $\Pr[\tilde{D}(S) > D(S) + \xi] \leq \frac{\delta}{2^{|\mathcal{X}|}}$ for any $n \geq \frac{|\mathcal{X}| \ln 2 + \log(1/\delta)}{2\xi^2}$. A union bound over these 2^ξ possible sets S concludes the proof:

$$\Pr \left[\exists S \subseteq \mathcal{X} \text{ s.t. } \tilde{D}(S) > D(S) + \xi \right] \leq 2^{|\mathcal{X}|} \cdot \frac{\delta}{2^{|\mathcal{X}|}} = \delta$$

and we are done. \square