The Malice of ELFs: Practical Anamorphic-Resistant Encryption without Random Oracles

Gennaro Avitabile[1], Vincenzo Botta[2], Emanuele Giunta[3⋆], Marcin Mielniczuk[4], and Francesco Migliaro[5⋆⋆]

[1] IMDEA Software Institute, Madrid, Spain. gennaro.avitabile@imdea.org
[2] Sapienza University of Rome, Rome, Italy. vincenzo.botta@uniroma1.it
[3] ETH Zurich, Zurich, Switzerland. emanuele.giunta@inf.ethz.ch
[4] University of Warsaw, Warsaw, Poland. m.mielniczuk@uw.edu.pl
[5] Università di Catania, Catania, Italy. francesco.migliaro@phd.unict.it

Abstract. The concept of Anamorphic Encryption (Persiano, Phan and Yung, EUROCRYPT'22), aims to enable private communication in settings where the usage of encryption is heavily controlled by a central authority (henceforth called the dictator) who can obtain users' secret keys. Since then, various works have improved our understanding of AE in several aspects, including its limitations. In this regard, two recent works at CRYPTO'25 constructed various Anamorphic-Resistant Encryption (ARE) schemes, i.e., schemes admitting at most $O(\log(\lambda))$ bits of covert communication.
However, those results are still unsatisfactory, each coming with at least one of the following issues: (1) use of cryptographic heavy hammers such as indistinguishability obfuscation (iO); (2) abuse of the original definition to define overly powerful dictators; (3) reliance on the Random Oracle Model (ROM). In particular, proofs in the ROM are controversial as they fail to account for anamorphic schemes making non-black-box usage of the hash function used to instantiate the Random Oracle.
In this work, we overcome all of these limitations. First, we describe an anamorphic-resistant encryption scheme approaching practicality by relying only on public-key encryption and Extremely Lossy Functions (ELFs), both known from the (exponential) DDH assumption. Moreover, assuming Fully Unique NIZKs (known from iO), we provide another construction, which we later use to realize the first *definitive* ARE; that is, a *single* scheme that *simultaneously* achieves the strongest level of anamorphic resistance against each of the possible levels of anamorphic security.

1 Introduction

Anamorphic Encryption (AE) [32] allows private communication in settings where the usage of encryption is heavily controlled by a central authority. It enables a sender to transmit a confidential message to a receiver without being detected by the authority, even though the receiver's secret decryption key is known by the authority itself. The quest for such a powerful primitive is motivated by scenarios such as dictatorship, where encryption is not banned altogether in order to preserve security against outsiders, but citizens might be forced to surrender their secret keys.

AE achieves its unique functionality through two modes of operation: regular and anamorphic. In the regular mode, it behaves like a standard public-key encryption scheme. In the anamorphic mode, however, a public key apk is generated alongside two secret keys: a standard-looking key ask and a covert "double key" dk. Bob needs to privately share dk with Alice *once*, and to use apk as his public key. Should Bob ever be compelled to reveal his secret key, he only discloses ask. A key feature of AE is that the pair $(\mathsf{apk}, \mathsf{ask})$ remains compatible with conventional encryption, ensuring plausible deniability. Meanwhile, Alice can use dk as a symmetric key to embed a hidden message within the ciphertext, which remains undetectable even if ask is exposed. Thus, an anamorphic ciphertext contains two messages: a regular message m (which can be decrypted by the dictator using ask) and an anamorphic one $\widehat{m}$ (recoverable only with dk). The primary security requirement is that anamorphic ciphertexts must be indistinguishable from regular ones.

⋆ Part of the work done while at IMDEA Software Institute, Madrid, Spain.
⋆⋆ Work done while visiting IRIF, Université Paris Cité, Paris, France.

Designing new anamorphic-friendly encryption schemes from scratch, as pointed out by [32], is of limited practical value. Indeed, a dictator can simply standardize any encryption scheme of their liking, possibly with the explicit goal of restricting anamorphic communication as much as possible. As a result, much of the literature has focused on demonstrating existing schemes to be anamorphic [32,1,12,27,33], often crucially exploiting specific properties of the underlying PKE. Many of such constructions are both efficient and support an exponentially large space of anamorphic messages. However, since the dictator can choose the encryption scheme adversarially, all such PKEs could end up being prohibited. This makes black-box[6] constructions of anamorphic encryption especially appealing, as they would remain usable regardless of the PKE imposed by the dictator.

On the limitations of black-box AE. Unfortunately, Catalano et al. [13] proved that any black-box AE scheme can transmit at most $O(\log \lambda)$ anamorphic bits per ciphertext. This result was further extended in [14] which showed that black-box AE is actually *impossible*; that is there exists a (contrived yet secure) PKE for which there is no secure black-box AE.

Moreover, they prove the $O(\log \lambda)$ bound of [13] to hold even in a setting where the dictator is given less power. In particular, the bound still holds for black-box AE schemes achieving a weaker notion they call *semi-adaptive* AE (explained later).

The above results, however, do not rule out the existence of *non-black-box* AE schemes that achieve high transmission rates generically for any PKE. Rather, they imply that any high-rate AE must inherently exploit the structure of the underlying PKE. The following questions were therefore left open to the dictator:

> Is it possible to construct an Anamorphic-Resistant Encryption (ARE) scheme? Specifically, do there exist PKEs with the following properties?
>
> (1) The PKE is semantically secure.
>
> (2) Any AE scheme making *non-black-box* use of the underlying PKE cannot transmit more than $O(\log \lambda)$ anamorphic bits per ciphertext.

Anamorphic-Resistant Encryption. Two recent works [18,10] provided positive answers to the above question. In [18], an elegant compiler is provided, augmenting any PKE in a way that makes it impossible to send more than $O(\log \lambda)$ anamorphic bits per ciphertext. While very efficient, their result heavily relies on the Random Oracle Model (ROM). Furthermore, they assume the dictator to initially generate a set of *public parameters* used by the regular encryption scheme (we will refer to this as the *public parameters model*). This is easily justified, as the authority is already assumed to have the power to enforce the use of a PKE of its choice[7].

In [10], the authors propose two compilers achieving a stronger notion of anamorphic resistance, namely that *any* AE for the resulting PKE is insecure. The first is in the public parameter model, but requires indistinguishability obfuscation (iO)[2]. The second one instead does not require public parameters nor obfuscation, but relies on the ROM. While these results may appear to have settled the question of anamorphic-resistance,the state of affairs is actually more nuanced. In the following, we explain why it remains relevant to consider anamorphic-resistance (i) against *semi-adaptive* schemes, and (ii) without relying on random oracles.

Adaptive and semi-adaptive AE. The original definition of anamorphic security (henceforth *adaptive* AE to distinguish it from the weaker, semi-adaptive variant) subtly grants enormous power to the adversary (i.e., the dictator). At the definitional level, it provides the adversary with the secret key *from the very beginning* of the game. The adversary can request encryptions of arbitrary pairs of regular and anamorphic messages, and is supposed to guess whether the AE scheme is operating in regular or anamorphic mode. In regular

[6] An AE scheme is *black-box* if it accesses the underlying PKE solely through oracle calls.

[7] Many traditional PKE have public parameters in practice. E.g., while formally in group-based schemes each public key should include a uniformly sampled group description, in practice common predefined groups with a fixed generators are used.

mode, anamorphic messages are simply ignored during encryption. When extending this notion to the public parameters model, the adversary is additionally provided with the corresponding trapdoor[8] at the beginning of the experiment.

The results of [14,10] crucially exploit this power. The knowledge of the user's secret key allows the adversary to adaptively obtain encryptions of messages that depend on the secret key and the trapdoor. These are ciphertexts that real users would never want to (or even be able to) produce. By artificially implanting *weak messages* in a given PKE, [14,10] rule out the existence of adaptive AE. These weak messages are only computable given the secret key (thus preserving IND-CPA), are indistinguishable from regular messages, and admit only polynomially many ciphertexts (thus preventing anamorphic communication). However, again, real users would never produce ciphertexts of weak messages. This highlights why it is important to consider anamorphic resistance against weaker notions of AE.

As mentioned before, [14] introduces the weaker yet meaningful notion of *semi-adaptive* AE. The key difference is that the secret key sk is given to the adversary distinguishing between the real and anamorphic modes only *after* encryption queries have been performed. Such notion generalizes naturally to the public parameter model, where the adversary is also provided with the trapdoor td along with sk *only after* the query phase. In this way, $(\mathsf{sk}, \mathsf{td})$ can still be used to distinguish between the two modes, while ensuring all the queried pairs $(m, \widehat{m})$ are independent of secret information.

This notion more accurately models a (perhaps more realistic) dictator who passively collects ciphertexts from the network. In this setting, messages should never depend on the associated sk or the trapdoor. Later, with access to all secret keys and the trapdoor, the dictator can attempt to determine whether users exchanged anamorphic messages. All of the above makes the semi-adaptive setting a more meaningful target for anamorphic-resistance.

Anamorphic resistance in the ROM. Finally, we argue that constructing schemes that provide anamorphic-resistance in the ROM is unsatisfactory. In general, relying on the ROM suffers from the well-known issue that its heuristic instantiations might be insecure. This concern is supported by a long line of work ([8,9,3,23,29], and more recently [7,25]) showing that, for some constructions, no secure instantiation is possible.

However, in the case of ARE, the use of random oracles raises an additional concern. A PKE is anamorphic resistant when *every* possible strategy to establish (high-rate) anamorphic communication fails in some way. Using the ROM oversimplifies the task, as it fails to model anamorphic schemes that, in the real world, make non-black-box use of the hash function used to instantiate the random oracle. Thus, a result in the ROM does not rule out anamorphic schemes that could leverage tools which take as input the circuit of the hash function, such as garbling [41], proof systems for NP [6], (fully)-homomorphic encryption [22], obfuscation [2] and more.

Open questions addressed by this work. Given the current state of affairs and the recent results of [18,10], the following refinement of the question about anamorphic-resistance naturally arises:

> ***Q1:*** *Is it possible to construct a PKE, without relying on the Random Oracle Model, that satisfies the following properties?*
>
> 1. *The PKE is semantically secure;*
> 2. *Any semi-adaptively secure AE for it cannot transmit more than $O(\log \lambda)$ anamorphic bits per ciphertext.*

We point out that the above level of anamorphic resistance would be optimal due to [14, Theorem 2], which builds a black-box compiler achieving a logarithmic rate for any underlying PKE in the semi-adaptive setting.

Note, however, that a positive answer to Q1 would actually be incomparable with the results of [10]. Indeed, an ARE as in Q1 would limit the bits any semi-adaptively secure AE can communicate to $O(\log \lambda)$ per ciphertext. Conversely, the constructions in [10] prevent any adaptively secure AE (a stronger notion) from sending even 1 anamorphic bit (a stronger bound). A more compelling question is, thus, whether one

[8] E.g., a chosen elliptic curve could be maliciously generated. While cryptanalytic scrutiny can check for hidden trapdoors, such efforts can never fully rule them out.

could *simultaneously* achieve the guarantees of Q1 and those of [10], that is, construct a single *definitive* ARE achieving the *worst of both worlds* (for the citizens). Precisely:

Q2: *Is it possible to construct a PKE as in* ***Q1*** *that further satisfies:*

3. *There exists no* adaptively *secure AE for it.*

1.1 Our Contributions

We answer both Q1 and Q2 affirmatively. Namely, we provide the first Anamorphic-Resistant Encryption scheme in the public parameters model without random oracles, in the semi-adaptive AE setting. Concretely, we provide two compilers turning any PKE[9] into an anamorphic-resistant one. More in details:

1. Our first compiler (Section 3) is based on Extremely Lossy Functions (ELFs) [42] and Fully Unique NIZKs[10] (FUNIZK), known from iO [38,40,39]. We prove it to turn any *perfectly correct* PKE into an ARE as per Q1.
2. Our second compiler (Section 4) removes the need of iO. It is based on Trapdoor Lossy Functions [31] and (an adaptation of) the Trapdoor ELF in [43], both implied by exponential DDH. We prove it applies to *any* PKE and returns an ARE as per Q1. Notably, we give the first practical ARE that does not rely on random oracles.
3. In Section 5, we show that our first FUNIZK-based compiler preserves anamorphic resistance in the sense of [10]. Combining this with the random-oracle-free compiler of [10] we obtain the first "worst-of-both-worlds compiler". That is, one that takes any (perfectly correct) PKE and turns it into an ARE featuring the best achievable level of anamorphic resistance in both the adaptive and the semi-adaptive AE setting.

In Table 1, we provide a summary of our results comparing them with the other works on AREs.

	PKE model	**AE security**	**ROM-free**	**iO-free**	$\lvert\widehat{M}\rvert$
[18]	PPM	Semi-Adaptive	✗	✓	$\mathsf{poly}(\lambda)$
[10], Construction 1	PPM	Adaptive	✓	✗	0
[10], Construction 2	Plain	Adaptive	✗	✓	0
This work, Section 3	PPM	Semi-Adaptive	✓	✗	$\mathsf{poly}(\lambda)$
This work, Section 4	PPM	Semi-Adaptive	✓	✓	$\mathsf{poly}(\lambda)$
This work, Section 5	PPM	Semi-Adaptive Adaptive	✓	✗	$\mathsf{poly}(\lambda)$ 0

Table 1. PKE model refers to whether the PKE is in the public parameters model (PPM) or the plain one. ROM-free and iO-free respectively indicate whether a random oracle or indistinguishability obfuscation is used in the constructions. $\lvert\widehat{M}\rvert$ is the size of the anamorphic message space.

[9] Precisely, for the first compiler the PKE must be perfectly correct, whereas the second one admits PKEs with negligible correctness error.

[10] I.e., non-interactive zero-knowledge proofs (in this work, we use the word proof also for computationally-sound proof systems) in which any statement admits *one* accepting proof at most.

1.2 Technical Overview

Intuitively, the only way to realize anamorphic encryption is by manipulating random coins used to encrypt a given regular message. Building on this observation, our principle to design AREs will be to:

- *scramble* the random coins before encryption, in a way that makes it hard to manipulate the scrambled output;
- *prove* that scrambled random coins were used to produce the ciphertext.

This approach closely mirrors that of [18], where the random coins are derived using a random oracle. To prove ciphertexts are well-formed, the preimage is "sent to the dictator", i.e., encrypted under a separate public key included in the public parameters, for which the dictator holds the corresponding secret key. Our main challenge is to instantiate this template without random oracles.

First construction based on ELFs and Fully Unique NIZKs. To address the first challenge – scrambling the random coins – we rely on Extremely Lossy Functions (ELFs) [42]. In a nutshell, ELFs are families of functions with large domains that can be generated either in *injective* or *lossy* mode. In the latter case, the image size is *polynomially small*. Indistinguishability for the two modes is only required to hold against time-bounded adversaries and large enough image size (though still $\mathsf{poly}(\lambda)$) in lossy mode.

To illustrate how ELFs can limit anamorphic communication[11], consider a simplified scenario: Suppose the sender, given a function f and an anamorphic message $\widehat{m}$, selects some r, and the receiver has to recover $\widehat{m}$ using only $f(r)$. If an efficient sender/receiver pair exists for a large (super-polynomial) message space and a public injective f, then ELF security can be broken. A distinguisher, given f, can attempt to "encode" a random $\widehat{m}$ and later "decode" it from $f(r)$. In injective mode, this succeeds by assumption. However, in lossy mode, decoding will likely fail since $|\mathrm{Im}\, f| \leq \mathsf{poly}(\lambda)$, while the message space is much larger.

Starting with a PKE scheme $(\mathsf{E}^*.\mathsf{Gen}, \mathsf{E}^*.\mathsf{Enc}, \mathsf{E}^*.\mathsf{Dec})$, the same reasoning applies when sending a ciphertext of the form $\mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}, m; f(r))$ for a given regular message m, and asking the receiver to extract $\widehat{m}$ from it. However, an actual AE could deviate from the prescribed scheme, attempting to avoid the ELF altogether. To prevent such attempts, we rely on non-interactive zero-knowledge proofs (NIZKs). By requiring the sender to prove that its ciphertexts are correctly computed, we ensure adherence to the prescribed use of ELFs. However, now the proof itself must also be transmitted, and as NIZKs are typically randomized, the random coins used by the prover to generate the proof could be manipulated to carry anamorphic bits, merely shifting the location of the anamorphic channel. To foreclose this possibility, we instead require the NIZK to have a *unique* and *deterministic proofs* for each statement (we call this NIZK fully unique), a property achieved for instance in [38,40,39]. Since for each valid statement there exists only one accepting proof, attaching the proof does not give any room to embed extra anamorphic bits.

Finally, we need to ensure that IND-CPA-security is preserved. Since the NIZK is zero-knowledge, it does not leak any information about the encrypted message or the randomness. However, even when the ELF is instantiated in injective mode, its output is not necessarily uniformly distributed [12]. Nevertheless, since f is injective, it preserves the min-entropy of r. To obtain an (almost) uniform string then, we compose f with a randomness extractor h (e.g., a universal hash function (UHF)). The final PKE scheme produces ciphertexts of the form

$$(e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}, m; h \circ f(r)), \pi)$$

where π is a proof that e is well-formed.

[11] In [10], ELFs are used in a fundamentally different way. They are used to embed weak messages with only polynomially many ciphertexts, enabling an active attack through repeated queries. In contrast, we use ELFs merely to scramble the randomness: the dictator only needs to verify ciphertext correctness to guarantee that any anamorphic channel has low rate.

[12] In general, $\mathrm{Im}\, f$ could be *sparse* in the set of strings with a given length.

Second Construction based on Trapdoor ELF. The role of the FUNIZK in the previous approach is to give the dictator a way to verify that the random coins during the encryption are in $\operatorname{Im} f$. A simple way to remove the FUNIZK would be to assume that f also admits a trapdoor that allows to efficiently invert it. If we could provide the scrambled random coins ρ to the dictator, they would then be able to test whether $\rho \in \operatorname{Im} f$ by simply attempting to invert f.

Sending such ρ without breaking IND-CPA is straightforward with a trapdoor lossy function F. Indeed, if the ciphertexts are of the form

$$(\mathsf{E.Enc}(\mathsf{pk}, m; h \circ f(r)),\ F \circ f(r))$$

where h is again a UHF, the dictator can always extract $f(r)$ from the second component using the trapdoor for F. For IND-CPA, on the other hand, we can switch to lossy mode in a hybrid. Then, $F(f(r))$ only leaks a fraction of the min-entropy of $f(r)$, so we can still extract good randomness through universal hashing as long as r is long enough.

However, plain Trapdoor ELFs [43] are insufficient to replicate the previous proof for anamorphic resistance:

1. In our proof technique the ELF distinguisher needs to check well-formedness of the ciphertext, however the security of Trapdoor ELFs does not allow the distinguisher to have the inversion trapdoor;
2. For lossy-mode f, the membership in $\operatorname{Im} f$ cannot be tested, even with a trapdoor.

Our main technical contribution is to adapt the Trapdoor ELF from [43] to obtain a "partial trapdoor" that preserves ELF security when leaked, but allows testing membership in an approximation of $\operatorname{Im} f$ in both injective and lossy mode.

To illustrate the main idea, we focus on the simplified task of adapting the Trapdoor Lossy Function of [31], which [43] builds on. Given a group $\mathbb{G}$ and an $m \times n$ matrix in the exponent $[A]$ [13], the (Trapdoor) Lossy Function of [31] is defined as

$$f\colon \{0,1\}^n \to \mathbb{G}^m\colon \qquad f(\mathbf{x}) = [A\mathbf{x}]\,,$$

where, in injective mode, A is sampled uniformly, while in lossy mode it is a random rank-1 matrix. Our strategy is to modify this by taking A as the product of two $m \times k$ and $k \times n$ matrices B, C, i.e., $A = BC$, where $m > k > n$. The function is defined as before, given $[BC]$. Here, B serves as the partial trapdoor and is always uniformly sampled, while C is either full rank or rank-1, depending on the mode. Importantly, knowing B does not help in learning the rank of C.

The trick now is to observe that testing membership in $[\operatorname{Im} B]$ of a value computed *only* as a function of $[BC]$ suffices to almost always imply its membership also in the (much smaller!) set $[\operatorname{Im} BC]$, regardless of the rank of C. This, informally, holds because the product BC loses all information on $\operatorname{Im} B \setminus \operatorname{Im} BC$. Since there are too many possible ways to place $\operatorname{Im} B$ in a way that contains $\operatorname{Im} BC$[14], guessing a point (whose discrete logarithm lies) in $\operatorname{Im} B \setminus \operatorname{Im} BC$ is hard. Thus, membership in $\operatorname{Im} B$ almost implies membership in $\operatorname{Im} BC$.

We finally note that adapting this trick to the full TELF proposed by Zhandry in [43] presents additional challenges, we defer to Section 4 for a detailed discussion.

One ARE to rule them all. Finally, we address the question of building the *definitive* ARE. We first prove that our first (iO-based) PKE-to-ARE compiler, when given a PKE that admits zero anamorphic bits in the adaptive AE setting, produces an ARE with the same property. By combining this result with the random-oracle-free construction of [10], we obtain an ARE achieving the best possible level of anamorphic resistance in both the adaptive and semi-adaptive AE settings. Unfortunately, for technical reasons (see Remark 4), we are unable to prove the same theorem using our second (DDH-based) construction. Thus, although the first construction relies on much heavier tools, it has an additional interesting feature compared to the second one.

[13] For a matrix $A = [a_{ij}]$, $[A]$ denotes the matrix $g^{a_{ij}}$.

[14] This holds because $k = \dim \operatorname{Im} B$ is much larger than $\dim \operatorname{Im} BC = \mathsf{rk}(C)$ for B of full rank.

1.3 Related Work

In this section, we discuss the related work on anamorphic encryption, apart the recent works on ARE discussed in Section 1. Anamorphic encryption is similar yet distinct from other notions studied in the literature; see [32] for a detailed discussion and comparison. In [27,12] the notion has been extended to achieve stronger privacy requirements regarding parties having access also to dk. Several works have introduced various flavors of robustness, a property designed to prevent the decryption of unintended anamorphic messages. Robust AE was first proposed in [1] and later extended to other settings in [37]. A stronger variant, termed anamorphic unforgeability, was subsequently introduced in [18]. In [37,10], a strong connection between AE and Algorithm Substitution Attacks (ASA) [4] is shown, allowing to reinterpret the results of the latter for the former and vice versa.

In [17], anamorphic encryption was extended to a one-to-many setting, enabling different receivers to extract distinct anamorphic messages from the same regular-looking ciphertext.

In [26], the authors introduce anamorphic signatures, where the anamorphic message is now hidden in a regular-looking digital signature. Furthermore, in [24], the security properties for anamorphic signatures are strengthened.

In [36], the notion of Anamorphic Key Exchange (AKE) is introduced. Thanks to AKE, it is possible to covertly share a double key within a regular-looking key exchange algorithm.

2 Preliminaries

We denote the security parameter by λ. For a deterministic algorithm $\mathcal{A}$, we use $y := \mathcal{A}(x)$ to denote the output of deterministic algorithms. For a probabilistic algorithm $\mathcal{A}$, we write $y := \mathcal{A}(x;r)$ to denote its output with random coins r. We write $y \leftarrow \mathcal{A}(x)$ when the random coins r are implicitly assumed to be uniformly random, and we denote sampling uniformly from a set by $x \leftarrow^{\$} X$. When we wish to ignore the output of an algorithm, we write $_ \leftarrow \mathcal{A}(x)$ or $(y, _) \leftarrow \mathcal{A}(x)$. For a probabilistic $\mathcal{A}$, its support $\mathsf{Supp}(\mathcal{A}(x))$ on input x is the set of all y s.t. $\Pr[y \leftarrow \mathcal{A}(x)] > 0$.

Given two events A, B we will write $\Pr[A, B]$ to denote their conjunction, i.e., $\Pr[A, B] = \Pr[A \cap B]$. Given two random variables X, Y we denote their statistical distance by $\Delta(X, Y)$. We write $[n] = \{1, \ldots, n\}$.

Given a matrix A, we denote its rank by $\mathsf{rk}(A)$. We denote the space of $n \times m$ matrices with entries in a field $\mathbb{F}$ by $\mathbb{F}^{n,m}$, and we write $\mathbb{F}^{n,m;k}$ to denote the subset of matrices of rank equal to k.

2.1 Public-key Encryption

We revise definitions and notation for public key encryption in the public parameters model as proposed by [18]. In our context, it is sensible to assume a global setup procedure $\mathsf{E.Init}$ returning (possibly trapdoored) public parameters.

Formally, a PKE scheme in the public parameters model is a quadruple of algorithms $\mathsf{E} = (\mathsf{E.Init}, \mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ with the following syntax, assuming without loss of generality that pp is embedded in pk and sk by the key generation algorithm:

- $\mathsf{E.Init}(1^\lambda) \to (\mathsf{pp}, \mathsf{td})$ samples parameters pp along with a trapdoor td.
- $\mathsf{E.Gen}(\mathsf{pp}) \to (\mathsf{pk}, \mathsf{sk})$ creates public and secret encryption keys.
- $\mathsf{E.Enc}(\mathsf{pk}, m) \to \mathsf{ct}$ encrypts a message m into a ciphertext ct.
- $\mathsf{E.Dec}(\mathsf{sk}, \mathsf{ct}) \to m$ decrypts a ciphertext ct.

We require the following variant of IND-CPA security (Figure 1), in which the adversary is given the public parameters pp.

Definition 1. *A PKE scheme* E *is* IND-CPA *secure if for any* PPT *adversary* $\mathcal{A}$ *it holds that:*

$$\mathsf{Adv}_{\mathcal{A},\mathsf{E}}(\lambda) := \left|\Pr\left[\text{IND-CPA}^{\mathcal{A}}_{\mathsf{E}}(\lambda) = 1\right] - 1/2\right| \leq \mathsf{negl}(\lambda).$$

where $\text{IND-CPA}^{\mathcal{A}}_{\mathsf{E}}$ *is specified in Figure 1.*

$\text{IND-CPA}^{\mathcal{A}}_{\mathsf{E}}(\lambda)$

$b \xleftarrow{\$} \{0,1\}$
$(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{E.Init}(1^\lambda)$
$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{E.Gen}(\mathsf{pp})$
$(m_0, m_1) \leftarrow \mathcal{A}(1^\lambda, \mathsf{pp}, \mathsf{pk}, c)$
$c \leftarrow \mathsf{E.Enc}(\mathsf{pk}, m_b)$
$b' \leftarrow \mathcal{A}(1^\lambda, \mathsf{pp}, \mathsf{pk}, c)$
return $b = b'$.

Fig. 1. IND-CPA security game with public parameters.

We also require *correctness*, i.e.,

$$\Pr\left[\mathsf{E.Dec}(\mathsf{sk}, \mathsf{E.Enc}(\mathsf{pk}, m)) \neq m \;\middle|\; \begin{array}{l}(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{E.Init}(1^\lambda) \\ (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{E.Gen}(\mathsf{pp})\end{array}\right] \leq \mathsf{negl}(\lambda).$$

In some cases we will require *perfect correctness*, i.e., that for any $(\mathsf{pp}, \mathsf{td}) \in \mathsf{Supp}(\mathsf{E.Init}(1^\lambda))$, any $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{Supp}(\mathsf{E.Gen}(\mathsf{pp}))$ and any message m and random coins r it holds that $m = \mathsf{E.Dec}(\mathsf{sk}, \mathsf{E.Enc}(\mathsf{pk}, m; r))$. In our second construction we will require that testing $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{Supp}(\mathsf{E.Gen}(\mathsf{pp}))$ for any valid pp is efficient. Note that we can assume it without the loss of generality, up to appending to sk the random coins used by $\mathsf{E.Gen}(\mathsf{pp})$.

2.2 Anamorphic Encryption

In the following we recall the definition of (receiver) semi-adaptive anamorphic encryption [14], mildly augmented to take into account the presence of public parameters as done in [18,10] for the case of the adaptive definition. We extended the notion of [14] by further providing td at the end of the execution. Furthermore, [14, Theorem 2] stating that rejection sampling provides semi-adaptive security continues to hold according to our extended definition.

Definition 2 (Anamorphic Triplet). *An anamorphic triplet* $\mathsf{AT} = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *consists of efficient algorithms such that:*

- $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}) \leftarrow \mathsf{AT.Gen}(\mathsf{pp})$*: where* $\mathsf{apk}, \mathsf{ask}$ *are the anamorphic public and secret keys, respectively,* dk *is the double key and* pp *are (possibly empty) public parameters.*
- $c \leftarrow \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$*: where* $m \in M$ *and* $\widehat{m} \in \widehat{M}$ *are, respectively, the standard and anamorphic messages encrypted in* c*. The set* $\widehat{M}$ *is called the anamorphic message space, while* M *is the regular message space. Both sets are decidable (i.e., it is possible to efficiently check if an element belongs to the set) and efficiently sampleable (i.e., it is possible to efficiently sample an element from the set).*
- $\widehat{m}/\perp := \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{dk}, c)$*: with* $\widehat{m}$ *the anamorphic message encrypted in* c*.*

Definition 3 (Semi-Adaptive AE). *A PKE scheme* E *equipped with an anamorphic triplet* AT *is said to be semi-adaptive anamorphic if for every* PPT *adversary* $\mathcal{A}$ *has negligible advantage, defined as*

$$\mathsf{Adv}^{\mathsf{SA\text{-}Anam}}_{\mathcal{A},\mathsf{E},\mathsf{AT}}(1^\lambda) := \left|\Pr\left[\mathsf{SA\text{-}RealG}_{\mathsf{E}}(1^\lambda, \mathcal{A}) = 1\right] - \Pr\left[\mathsf{SA\text{-}AnamorphicG}_{\mathsf{AT}}(1^\lambda, \mathcal{A}) = 1\right]\right|.$$

where $\mathsf{SA\text{-}RealG}_{\mathsf{E}}$ *and* $\mathsf{SA\text{-}AnamorphicG}_{\mathsf{AT}}$ *are presented in Figure 2.*

Adaptive AE is obtained modifying the definition of semi-adaptive AE giving $(\mathsf{pk}, \mathsf{sk}, \mathsf{pp}, \mathsf{td})$ to $\mathcal{A}$ from the very beginning of the game and allowing $\mathcal{A}$ to freely perform queries to the oracle.

$\mathsf{SA\text{-}RealG}_{\mathsf{E}}(1^\lambda, \mathcal{A})$

1 : $(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{E.Init}(1^\lambda)$
2 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{E.Gen}(\mathsf{pp})$
3 : $\mathsf{call} = \mathsf{false}$
4 : **return** $\mathcal{A}^{\mathcal{O}_{\mathsf{real}}, \mathcal{O}_{\mathsf{key}}}(\mathsf{pp}, \mathsf{pk})$

$\mathcal{O}_{\mathsf{real}}(m, \widehat{m})$

1 : **if** call is false:
2 : **return** $\mathsf{E.Enc}(\mathsf{pk}, m)$

$\mathcal{O}_{\mathsf{key}}$

1 : $\mathsf{call} = \mathsf{true}$
2 : **return** $(\mathsf{sk}, \mathsf{td})$

$\mathsf{SA\text{-}AnamorphicG}_{\mathsf{AT}}(1^\lambda, \mathcal{A})$

1 : $(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{E.Init}(1^\lambda)$
2 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}) \leftarrow \mathsf{AT.Gen}(\mathsf{pp})$
3 : $\mathsf{call} = \mathsf{false}$
4 : **return** $\mathcal{A}^{\mathcal{O}_{\mathsf{anam}}, \mathcal{O}_{\mathsf{key}}}(\mathsf{pp}, \mathsf{apk})$

$\mathcal{O}_{\mathsf{anam}}(m, \widehat{m})$

1 : **if** call is false:
2 : **return** $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$

$\mathcal{O}_{\mathsf{key}}$

1 : $\mathsf{call} = \mathsf{true}$
2 : **return** $(\mathsf{ask}, \mathsf{td})$

Fig. 2. Semi-adaptive security game with public parameters.

Definition 4 (Adaptive AE). *An* IND-CPA*-secure PKE scheme* $\mathsf{E} = (\mathsf{E.Init}, \mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *is an adaptive AE scheme if there exists an anamorphic triplet* $\mathsf{AT} = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$*, such that any* PPT *adversary* $\mathcal{A}$ *has negligible advantage, defined as:*

$$\mathsf{Adv}^{\mathsf{Anam}}_{\mathcal{A},\mathsf{E},\mathsf{AT}}(1^\lambda) \coloneqq \left|\Pr\left[\mathsf{RealG}_{\mathsf{E}}(1^\lambda, \mathcal{A}) = 1\right] - \Pr\left[\mathsf{AnamorphicG}_{\mathsf{AT}}(1^\lambda, \mathcal{A}) = 1\right]\right|$$

where $\mathsf{RealG}_{\mathsf{E}}$ *and* $\mathsf{AnamorphicG}_{\mathsf{AT}}$ *are presented in Figure 3.*

$\mathsf{RealG}_{\mathsf{E}}(1^\lambda, \mathcal{A})$

1 : $(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{E.Init}(1^\lambda)$
2 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{E.Gen}(\mathsf{pp})$
3 : **return** $\mathcal{A}^{\mathcal{O}_{\mathsf{real}}}(\mathsf{pp}, \mathsf{td}, \mathsf{pk}, \mathsf{sk})$

$\mathcal{O}_{\mathsf{real}}(m, \widehat{m})$

1 : **return** $\mathsf{E.Enc}(\mathsf{pk}, m)$

$\mathsf{AnamorphicG}_{\mathsf{AT}}(1^\lambda, \mathcal{A})$

1 : $(\mathsf{pp}, \mathsf{td}) \leftarrow \mathsf{E.Init}(1^\lambda)$
2 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}) \leftarrow \mathsf{AT.Gen}(\mathsf{pp})$
3 : **return** $\mathcal{A}^{\mathcal{O}_{\mathsf{anam}}}(\mathsf{pp}, \mathsf{td}, \mathsf{apk}, \mathsf{ask})$

$\mathcal{O}_{\mathsf{anam}}(m, \widehat{m})$

1 : **return** $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$

Fig. 3. Adaptive AE security game with public parameters.

Finally, we recall ε-correctness on average [13], a weaker notion of correctness for anamorphic triplets.

Definition 5. *An anamorphic triplet is* ε*-correct on average w.r.t. a PKE* $(\mathsf{E.Init}, \mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *if, for a* negligible ε*, sampling parameters* $\mathsf{pp} \leftarrow \mathsf{E.Init}(1^\lambda)$*, keys* $(\mathsf{ask}, \mathsf{apk}, \mathsf{dk}) \leftarrow \mathsf{AT.Gen}(\mathsf{pp})$*, and uniformly sampled* $m \leftarrow M, \widehat{m} \leftarrow \widehat{M}$*, then*

$$\Pr\left[\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{dk}, \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})) \neq \widehat{m}\right] \leq \varepsilon(\lambda).$$

2.3 Universal Hash Functions

Universal Hash Functions families (UHF) [11] are information-theoretical objects ensuring that any given pair of distinct points collides with low probability. The formal definition follows.

Definition 6. *Let* $\mathcal{H}$ *be a finite family of functions of type* $h\colon X \to Y$*. The family* $\mathcal{H}$ *is a Universal Hash Family if*

$$\forall x, y \in X\colon \Pr_{h \leftarrow\$ \mathcal{H}}[h(x) = h(y)] \leq \frac{1}{|Y|}.$$

A common usage of UHF is to deploy them as randomness extractors. This is formally justified by the Leftover Hash Lemma, presented below.

Lemma 1 (Generalized Leftover Hash Lemma [19]). *Assume $\mathcal{H}$ is a UHF family taking values in $\{0,1\}^m$ and let $h \leftarrow^{\$} \mathcal{H}$. Then, for any random variables X and Y, it holds that*

$$\Delta((h, h(X), Y), (h, U, Y)) \leq \frac{1}{2}\sqrt{2^{-\mathrm{H}_\infty(X \mid Y)+m}},$$

with U uniformly distributed in $\{0,1\}^m$.

We will also use the following standard inequality for statistical distance:

Lemma 2. *Let X and Y be random variables and F any randomized function, it holds that $\Delta(F(X), F(Y)) \leq \Delta(X, Y)$.*

2.4 Lossy Trapdoor Functions

Introduced by [31], lossy trapdoor functions (LTFs) are functions that can be instantiated in one of two indistinguishable modes: injective or lossy. Moreover, in injective mode the existence of a secret trapdoor allows to efficiently invert the function. In what follows we use the same notation as in [40] which further specifies the function's input length at setup time.

Definition 7. *A Lossy Trapdoor Function is a tuple of algorithms $\mathsf{LTF} = (\mathsf{GenInj}, \mathsf{GenLos}, \mathsf{Eval}, \mathsf{Inv})$ such that*

- $(k, \mathsf{td}) \leftarrow \mathsf{LTF.GenInj}(1^\lambda, 1^\ell)$ *with k a function index and td a trapdoor.*
- $k \leftarrow \mathsf{LTF.GenLos}(1^\lambda, 1^\ell)$ *with k a function index.*
- $y := \mathsf{LTF.Eval}(k, x)$ *evaluates the function on input x,*
- $x := \mathsf{LTF.Inv}(\mathsf{td}, y)$ *inverts the function on image y.*

Moreover, sampling $(k_0, \mathsf{td}) \leftarrow \mathsf{LTF.GenInj}(1^\lambda, 1^\ell)$ and $k_1 \leftarrow \mathsf{LTF.GenLos}(1^\lambda, 1^\ell)$, the following properties hold:

- $\mathsf{LTF.Eval}(k_b, \cdot) : \{0,1\}^\ell \to \{0,1\}^{\mathsf{poly}(\lambda)}$ *for any $b \in \{0,1\}$.*
- ***Injectivity**: $\mathsf{LTF.Eval}(k_0, \cdot)$ is an injective function with overwhelming probability.*
- ***Lossiness**: There exists a polynomial μ such that $\mathsf{LTF.Eval}(k_1, \cdot)$ has image size smaller than $2^{\mu(\lambda)}$.*
- ***Indistinguishability**: For any PPT adversary $\mathcal{A}$*

$$\mathsf{Adv}_{\mathcal{A}}(\lambda) \;:=\; \left|\Pr\left[\mathcal{A}(1^\lambda, 1^\ell, k_0) = 0\right] - \Pr\left[\mathcal{A}(1^\lambda, 1^\ell, k_1) = 0\right]\right| \;\leq\; \mathsf{negl}(\lambda).$$

Note that, as opposed to the original and more general definition of [31], the one above requires the image size in lossy mode to be independent of the input size[15] (albeit still exponentially large). Constructions achieving this high lossiness level exists from DDH, see [31, §5.3].

2.5 Extremely Lossy Functions

Extremely lossy function (ELFs), first introduced in [42], are families of functions which can be sampled to either be injective or have polynomially small image size. Distinguishing the two modes with sufficient (polynomial) time then cannot be hard. ELFs however guarantees that time-bounded adversaries cannot distinguish injective function from extremely lossy ones with significant advantage as long as the image in lossy mode is large enough (but still polynomial).

Definition 8. *An ELF consists of an algorithm $\mathsf{ELF.Gen}$ such that, for integers M, R, $\mathsf{ELF.Gen}(M, R)$ returns the description of a function $f : [M] \to [N]$ for $M < N < \mathsf{poly}(M)$ such that*

[15] Formally, this actually only needs to hold for an upper bound of the image size.

1. $f : [M] \to [N]$ *can be computed in time* $\mathsf{poly}(\log M)$.
2. $f \leftarrow \mathsf{ELF.Gen}(M, M)$ *is injective.*
3. $f \leftarrow \mathsf{ELF.Gen}(M, R)$*, then* $|\mathrm{Im}\, f| < R$.
4. *For any polynomials* t, δ *there exists a polynomial* Q *such that for any* t*-time machine* $\mathcal{A}$ *and* R *with* $Q(\log M) \leq R \leq M$ *we have that, sampling* $f_0 \leftarrow \mathsf{ELF.Gen}(M, M)$ *and* $f_1 \leftarrow \mathsf{ELF.Gen}(M, R)$

$$\mathsf{Adv}_{\mathcal{A}} \coloneqq \left|\Pr\left[\mathcal{A}(f_0) = 0\right] - \Pr\left[\mathcal{A}(f_1) = 0\right]\right| \leq \frac{1}{\delta(\log M)}.$$

Our definition mildly deviates from the one of [42] as we require $\mathsf{ELF.Gen}(M, M)$ to always return an injective function[16]. ELFs can be defined to support a trapdoor. Syntax and security properties are introduced below.

Definition 9. *A trapdoor ELF is a tuple of algorithms* $\mathsf{TELF} = (\mathsf{GenInj}, \mathsf{GenLos}, \mathsf{Inv})$*. For any integer* M*,* $\mathsf{TELF.GenInj}(M)$ *returns* (f, td) *such that*

1. $f : [M] \to [N]$ *for some* $N \leq \mathsf{poly}(M)$ *is a function computable in time polynomial in* $\log M$.
2. *For any* $x \in [M]$ *then* $x \coloneqq \mathsf{TELF.Inv}(\mathsf{td}, f(x))$.

For any integer M, R*,* $\mathsf{TELF.GenLos}(M, R)$ *returns* f *such that*

3. $f : [M] \to [N]$ *for some* $N \leq \mathsf{poly}(M)$ *is a function computable in time polynomial in* $\log M$.
4. $|\mathrm{Im}\, f| \leq R$.

Finally, for any t, δ *polynomials in* $\log M$*, there exists a polynomial* q *such that* $R \geq q(\log M)$ *implies that, sampling* $(f_0, \mathsf{td}) \leftarrow \mathsf{TELF.GenInj}(M)$ *and* $f_1 \leftarrow \mathsf{TELF.GenLos}(M, R)$*, any* t*-time adversary has advantage*

$$\mathsf{Adv}_{\mathcal{A}} = \left|\Pr\left[\mathcal{A}(f_0) = 1\right] - \Pr\left[\mathcal{A}(f_1) = 1\right]\right| \leq \frac{1}{\delta(\log M)}.$$

With abuse of notation we will identify the function f with its description. When such identification would be ambiguous, we refer to f as the function description, and $\mathsf{ELF.Eval}(f, x) \to y$ or $\mathsf{TELF.Eval}(f, x) \to y$ as the efficient procedures evaluating f on x.

2.6 Cryptographic Groups

Following [42], we now recall the definition of a cryptographic group.

Definition 10. *A cryptographic group is a procedure* $\mathsf{GRP.Gen}$ *such that for any integer* $\lambda \in \mathbb{N}$*,* $\mathsf{GRP.Gen}(\lambda)$ *returns* $(\mathbb{G}, g, p)$ *where*

- $(\mathbb{G}, \cdot)$ *is a cyclic group of order* $p = |\mathbb{G}|$ *and generator* g *with* $2^{\lambda} \leq p < 2^{\lambda+1}$.
- *Membership and group operations are computable in time polynomial in* λ.
- *Elements in* $\mathbb{G}$ *are represented by string whose length is polynomial in* λ.

Whenever the group $\mathbb{G}$ and a generator g being used are clear from context we adopt the notation $[a] = g^a$. The notation is naturally extended to vectors and matrices by applying the group exponentiation entry-wise. In order to instantiate TELF, we will need to assume exponential hardness of DDH (or equivalently Matrix-DDH) as defined below. Note they can be proven to hold generically, and are reasonable to assume for elliptic curve groups, where known attacks are currently only the generic ones.

Definition 11. *The exponential DDH assumption holds for* $\mathsf{GRP.Gen}$ *if there exists a polynomial* $q(\cdot, \cdot)$ *such that for any* t, ε *setting* $\lambda \geq \log q(t, 1/\varepsilon)$ *then for any* t*-time adversary, sampling* $(\mathbb{G}, g, p) \leftarrow \mathsf{GRP.Gen}(\lambda)$ *and scalars* $a, b, c \leftarrow^{\$} \mathbb{F}_p$

$$\left|\Pr\left[\mathcal{A}(\mathbb{G}, g, p, [a], [b], [c]) = 1\right] - \Pr\left[\mathcal{A}(\mathbb{G}, g, p, [a], [b], [ab]) = 1\right]\right| \leq \varepsilon.$$

[16] For the construction of [42] we can assume it by Lemma 5.

Definition 12. *The exponential Matrix-DDH assumption holds for* $\mathsf{GRP.Gen}$ *if there exists a polynomial* q *such that for any* t, ε, n, m *the following holds. Setting* $\lambda \geq \log q(t, n, m, 1/\varepsilon)$*, for any* t*-time adversary* $\mathcal{A}$*, sampling* $(\mathbb{G}, g, p) \leftarrow \mathsf{GRP.Gen}(\lambda)$ *and matrices* $A \leftarrow^{\$} \mathbb{F}_p^{n,m}$ *and* $B \leftarrow^{\$} \mathbb{F}_p^{n,m}$ *such that* $\mathsf{rk}(B) = 1$*, then*

$$\left|\Pr\left[\mathcal{A}(\mathbb{G}, g, p, [A]) = 1\right] - \Pr\left[\mathcal{A}(\mathbb{G}, g, p, [B]) = 1\right]\right| \leq \varepsilon.$$

Note that exponential DDH and exponential Matrix-DDH are equivalent [35].

2.7 Fully Unique NIZK

A fully unique non-interactive zero-knowledge argument (FUNIZK) is a non-interactive proof system for an NP language $\mathcal{L}$ with adaptive computational soundness and perfect zero knowledge. A FUNIZK has the additional feature that for every statement x there exists a *unique* accepting proof π. The constructions of [38,40,39] are (succinct) FUNIZKs for all NP in the standard model and rely on iO. We refer to Remark 1 for a discussion on FUNIZK and a comparison with other similar notions.

Definition 13 (Fully Unique NIZK Argument). *Let* $\mathcal{L}$ *be an* NP *language with an associated relation* $\mathcal{R}$*. A FUNIZK for* $\mathcal{L}$ *is a tuple of PPT algorithms* $\mathsf{FUNIZK} = (\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ *with the following syntax:*

- $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)$*: given the security parameter* $\lambda \in \mathbb{N}$*, outputs a common reference string* crs*;*
- $\pi := \mathsf{Prove}(\mathsf{crs}, x, w)$*: given a common reference string* crs*, a statement* x*, and a witness* w*, outputs a* unique *proof* π*;*
- $b := \mathsf{Verify}(\mathsf{crs}, x, \pi)$*: given a common reference string* crs*, a statement* x*, and a proof* π*, outputs a bit* $b \in \{0, 1\}$*;*

satisfying the following properties:

- ***Completeness:*** *for every* $(x, w) \in \mathcal{R}$*:*

$$\Pr\left[\mathsf{Verify}(\mathsf{crs}, x, \pi) = 1 \;\middle|\; \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi := \mathsf{Prove}(\mathsf{crs}, x, w) \end{array}\right] = 1.$$

- ***Adaptive Computational Soundness:*** *for every PPT adversary* $\mathcal{A}$*:*

$$\Pr\left[x \notin \mathcal{L}, \mathsf{Verify}(\mathsf{crs}, x, \pi) = 1 \;\middle|\; \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ x, \pi \leftarrow \mathcal{A}(\mathsf{crs}) \end{array}\right] \leq \mathsf{negl}(\lambda).$$

- ***Perfect Zero-Knowledge:*** *there exists a polynomial time simulator* $S = (S_0, S_1)$ *such that for all* $(x, w) \in \mathcal{R}$ *the following two distributions are identical:*

$$\left\{(\mathsf{crs}, x, \pi) \;\middle|\; \begin{array}{l} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi := \mathsf{Prove}(\mathsf{crs}, x, w) \end{array}\right\} \equiv \left\{(\mathsf{crs}, x, \pi) \;\middle|\; \begin{array}{l} (\mathsf{crs}, \mathsf{st}) \leftarrow S_0(1^\lambda) \\ \pi := S_1(\mathsf{st}, x) \end{array}\right\}.$$

- ***Full Uniqueness:*** *for all* $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)$*, and any* $x \in \{0, 1\}^*$*, there exists a* unique *proof* π *such that* $\mathsf{Verify}(\mathsf{crs}, x, \pi) = 1$*.*

Note that uniqueness implies that the Prove algorithm is necessarily deterministic, and so is the simulator S_1.

Remark 1. Proof systems with fully unique proofs seem fundamentally connected to iO and have been shown [5,15,20] to imply witness encryption [21]. A related but different notion is the one of unique NIZK (UNIZK) [28,34]. In UNIZK the prover is allowed to sample a non-unique public key, and can use the corresponding secret key to compute proofs. Once this public key is fixed, it is required, via an extraction experiment, that different proofs for the same statement require the knowledge of different witnesses. This gives computationally unique proofs if it is hard to find two distinct witnesses. UNIZK are not known to

imply witness encryption and can be constructed without iO. However, this uniqueness flavor is not enough in our setting. Even though it would be hard for the prover to produce different proofs, given the secret proving key one can extract the full witness from the proof, which may contain a lot of anamorphic bits.

Several works have already proposed definitions similar to ours with the most similar being unique witness maps [15,16]. However, they do not have zero-knowledge and are only selectively sound. Therefore, we give our definition of FUNIZK which is satisfied by [38,40]. The constructions of [38,40] are (succinct) FUNIZKs for all NP assuming sub-exponentially-secure indistinguishability obfuscation, sub- exponentially-secure one-way functions, and various algebraic assumptions. The work of [39] removes the need of algebraic assumptions.

2.8 Min-entropy

The min-entropy is a measure used to quantify the amount of randomness of a probability distribution.

Definition 14. *Let X, Y be discrete random variables with support $\mathcal{X}, \mathcal{Y}$. The min-entropy of X and the average min-entropy of X given Y, are respectively defined as:*

$$\begin{aligned} \mathrm{H}_\infty(X) &= -\log\left(\max_{x_0 \in \mathcal{X}} \Pr\left[X = x_0\right]\right), \\ \widetilde{\mathrm{H}}_\infty(X \mid Y) &= -\log\left(\sum_{y_0 \in \mathcal{Y}} \Pr\left[Y = y_0\right] \cdot \max_{x_0 \in \mathcal{X}} \Pr\left[X = x_0 \mid Y = y_0\right]\right). \end{aligned}$$

We will furthermore make use of min-entropy and average min-entropy conditioned on an event.

Definition 15. *Let X, Y be as in Definition 14, and E an event. Then the min-entropy of X conditioned on E (resp. average min-entropy of X given Y conditioned on E) is defined as:*

$$\begin{aligned} \mathrm{H}_\infty(X \mid E) &= -\log\left(\max_{x \in \mathcal{X}} \Pr\left[X = x \mid E\right]\right), \\ \widetilde{\mathrm{H}}_\infty(X \mid Y; E) &= -\log\left(\sum_{y \in \mathcal{Y}} \Pr\left[Y = y \mid E\right] \cdot \max_{x \in \mathcal{X}} \Pr\left[X = x \mid Y = y, E\right]\right). \end{aligned}$$

Lemma 3. *Given X, Y, Z discrete random variables with support $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ respectively, and E an event, then*

1. $\mathrm{H}_\infty(X \mid E) \leq \log_2|\mathcal{X}_E|$ *where* $\mathcal{X}_E = \{x_0 \in \mathcal{X} : \Pr\left[X = x_0 | E\right] > 0\}$.
2. $\mathrm{H}_\infty(X \mid Y; E) \geq \mathrm{H}_\infty(X \mid Y, Z; E)$.
3. $\mathrm{H}_\infty(X, Y \mid Z; E) \geq \mathrm{H}_\infty(X \mid Z; E)$.
4. $\mathrm{H}_\infty(X \mid Y, Z; E) \geq \mathrm{H}_\infty(X, Y \mid Z; E) - \log\left(\max_{z_0 \in \mathcal{Z}} |\mathcal{Y}_{z_0,E}|\right)$ *where*

$$\mathcal{Y}_{z_0,E} = \{y_0 \in \mathcal{Y} : \Pr\left[Y = y_0 \mid Z = z_0,\ E\right] > 0\}.$$

   In particular,

$$\mathrm{H}_\infty(X \mid Y, Z; E) \geq \mathrm{H}_\infty(X \mid Z; E) - \log|\mathcal{Y}_E|,$$

   where $\mathcal{Y}_E = \{y_0 \in \mathcal{Y} : \Pr\left[Y = y_0 \mid E\right] > 0\}$.
5. $\mathrm{H}_\infty(X \mid Y; E) = \mathrm{H}_\infty(X \mid E)$ *if X, Y are mutually independent given E.*
6. $\mathrm{H}_\infty(X \mid Y; E) \geq \mathrm{H}_\infty(X \mid Y) + \log \Pr\left[E\right]$.

Proof. Items 1 to 5 are essentially a rephrased version of [19, Lemma 2.2] for conditional distributions. For Item 6, by the chain rule we have that

$$\Pr\left[X = x, E \mid Y = y\right] = \Pr\left[X = x \mid E, Y = y\right] \cdot \Pr\left[E \mid Y = y\right].$$

Therefore,

$$\begin{aligned}
2^{-\mathrm{H}_\infty(X\,|\,Y)} &= \sum_{y\in\mathcal{Y}} \Pr[Y=y] \max_x \Pr[X=x\,|\,Y=y] \\
&\geq \sum_{y\in\mathcal{Y}} \Pr[Y=y] \max_x \Pr[X=x, E\,|\,Y=y] \\
&= \sum_{y\in\mathcal{Y}} \Pr[Y=y] \max_x \Pr[X=x\;|\;E, Y=y] \cdot \Pr[E\;|\;Y=y] \\
&= \sum_{y\in\mathcal{Y}} \underbrace{\Pr[Y=y]\Pr[E\;|\;Y=y]}_{\Pr[E,Y=y]} \max_x \Pr[X=x\;|\;Y=y,E]\cdot \\
&= \Pr[E]\cdot \sum_{y\in\mathcal{Y}} \Pr[Y=y\;|\;E] \max_x \Pr[X=x\;|\;Y=y,E]\cdot \\
&= \Pr[E]\cdot 2^{-\mathrm{H}_\infty(X\,|\,Y;E)}.
\end{aligned}$$

The claim simply follows by taking the logarithms of both sides of the inequality. □

Intuitively, Item 6 states that further conditioning to the event E decreases the min-entropy by at most the information content of E. Note that, due to the logarithmic dependence on $\Pr[E]$, the bound becomes meaningless when E has very low probability.

Lemma 4 (Guessing Lemma). *Let X, Y, E be as in Lemma 3 and ϕ a probabilistic function with support in $\{f\colon \mathcal{Y}\to\mathcal{X}\}$. If ϕ and (X,Y) are mutually independent relative to E, then*

$$\Pr[\phi(Y)=X\,|\,E] \;\leq\; 2^{-\mathrm{H}_\infty(X\,|\,Y;E)}.$$

Proof. It is well-known that for a deterministic ϕ we have $\Pr[\phi(Y)=X] \leq 2^{-\mathrm{H}_\infty(X\,|\,Y)}$. For an arbitrary event E but ϕ still deterministic, it follows directly from the definition by taking the conditional distributions. Finally, for independent probabilistic ϕ, we condition on ϕ and use the law of total probability. □

2.9 Linear Algebra

We recall, without a proof, some definitions and lemmas from linear algebra. $\mathbb{F}_p^{n,m}$ is the set of $n\times m$ matrices with entries in $\mathbb{F}_p$. We use $\mathbb{F}_p^{n,m;k}$ to denote the subset of rank-k matrices.

Lemma 5. *Let $n\leq m$ and $A \leftarrow^{\$} \mathbb{F}_p^{m,n}$ and $B \leftarrow^{\$} \mathbb{F}_p^{m,n;n}$. Then $\Delta(A,B) \leq 1/p^{m-n}$.*

Lemma 6. *Let $k\leq n\leq m$. Given $A \leftarrow^{\$} \mathbb{F}_p^{m,k}$, $B \leftarrow^{\$} \mathbb{F}_p^{m,n}$ and $M \leftarrow^{\$} \mathbb{F}_p^{m,k;k}$ then $\Delta(A, BM) = 0$.*

Next, we denote with $G_p(n,m)$ the Grassmannian, consisting of all the m-dimensional subspaces of $\mathbb{F}_p^m$. Formally $G_p(n,m) = \{V \leq \mathbb{F}_p^n \;:\; \dim V = m\}$. We will need the following fact:

Lemma 7. *Let $k\leq m$. Then the Grassmannian $G_p(m,k)$ has size*

$$|G_p(m,k)| \;=\; \frac{[m!]_p}{[k!]_p\cdot[(m-k)!]_p},$$

where $[n!]_p = \prod_{t=1}^{n}(1+p+\ldots+p^{t-1})$.

3 Anamorphic Resistant Encryption From FUNIZK

We now present a compiler that turns any IND-CPA secure encryption scheme E^* into an *anamorphic-resistant* IND-CPA secure encryption scheme E, meaning that anamorphic message space of E has polynomial size. Interestingly, this construction, unlike the one we will present in Section 4, does not require the dictator to access the trapdoor of the public parameters.

3.1 Construction

The construction (Figure 4) is based on the following building blocks:

- a perfectly correct IND-CPA-secure encryption scheme $(\mathsf{E}^*.\mathsf{Gen}, \mathsf{E}^*.\mathsf{Enc}, \mathsf{E}^*.\mathsf{Dec})$ with randomness space $\{0,1\}^\lambda$;
- an Extremely Lossy Function $\mathsf{ELF.Gen}$ that we instantiate with input length 3λ and output length $\ell = \mathsf{poly}(\lambda)$;
- a Universal Hash Family $\mathcal{H}$ of functions of type $\{0,1\}^\ell \to \{0,1\}^\lambda$;
- a Fully Unique NIZK $\mathsf{FUNIZK} = (\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ for the relation

$$\mathcal{R} = \left\{ ((e, \mathsf{pk}), (m, r)) \;\middle|\; \begin{array}{l} \mathsf{pk} = (\mathsf{pk}^*, f, h, _, _), \\ e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m; h \circ f(r)) \end{array} \right\}.$$

$\mathsf{E.Init}(1^\lambda)$

1 : $f \leftarrow \mathsf{ELF.Gen}(2^{3\lambda}, 2^{3\lambda})$
2 : $h \leftarrow^{\$} \mathcal{H}$
3 : $\mathsf{crs} \leftarrow \mathsf{FUNIZK.Setup}(1^\lambda)$
4 : $\mathsf{pp}^* \leftarrow^{\$} \mathsf{E}^*.\mathsf{Init}(\lambda)$
5 : **return** $\mathsf{pp} = (f, h, \mathsf{crs}, \mathsf{pp}^*)$

$\mathsf{E.Gen}(\mathsf{pp})$

1 : Parse $\mathsf{pp} = (f, h, \mathsf{crs}, \mathsf{pp}^*)$
2 : $(\mathsf{pk}^*, \mathsf{sk}^*) \leftarrow \mathsf{E}^*.\mathsf{Gen}(\mathsf{pp}^*)$
3 : $\mathsf{pk} := (\mathsf{pk}^*, \mathsf{pp})$, $\mathsf{sk} := \mathsf{sk}^*$
4 : **return** $(\mathsf{pk}, \mathsf{sk})$

$\mathsf{E.Enc}(\mathsf{pk}, m)$

1 : Parse $\mathsf{pk} = (\mathsf{pk}^*, f, h, \mathsf{crs}, _)$
2 : $r \leftarrow^{\$} \{0,1\}^{3\lambda}$
3 : $e := \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m; h \circ f(r))$
4 : $\pi := \mathsf{FUNIZK.Prove}(\mathsf{crs}, (e, \mathsf{pk}), (m, r))$
5 : **return** $\mathsf{ct} = (e, \pi)$

$\mathsf{E.Dec}(\mathsf{sk}, \mathsf{ct})$

1 : Parse $\mathsf{ct} = (e, \pi)$
2 : $m := \mathsf{E}^*.\mathsf{Dec}(\mathsf{sk}, e)$
3 : **return** m

Fig. 4. Anamorphic resistant encryption scheme from ELF and FUNIZK.

3.2 Security

Correctness of E easily follows from the ones of the underlying building blocks and can be verified by inspection. Security and anamorphic resistance are established by the following theorems.

Theorem 1. *Let* $\mathsf{ELF.Gen}$ *be an extremely lossy function,* $\mathcal{H}$ *be a fully universal hash function,* FUNIZK *be a fully unique NIZK argument, and* E^* *be an* IND-CPA*-secure public key encryption scheme. Then the scheme of Figure 4 is* IND-CPA*-secure.*

Proof of Theorem 1. We proceed through an hybrids sequence $\mathcal{H}_0^b, \mathcal{H}_1^b, \mathcal{H}_2^b$ progressively modifying the IND-CPA security game. m_0, m_1 denotes the challenge messages queried by a given PPT adversary $\mathcal{A}$ and ct^* is the challenge ciphertext encrypting m_b, with $b \in \{0,1\}$ being the challenge bit. We further denote S_0, S_1 the FUNIZK simulators.

$\mathcal{H}_0^b$: This hybrid coincides with the real IND-CPA game, where the adversary receives the pair $\mathsf{ct}^* = (e, \pi)$ s.t. $e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}, m_b; h \circ f(r))$ and π is computed using $\pi := \mathsf{FUNIZK.Prove}(\mathsf{crs}, (e, \mathsf{pk}), (m_b, r))$.

$\mathcal{H}_1^b$**:** It is identical to $\mathcal{H}_0^b$ except that the NIZK is simulated, i.e., the crs is generated as $(\mathsf{crs}, \mathsf{st}) \leftarrow S_0(1^\lambda)$ and the proof $\pi := S_1(\mathsf{st}, (e, \mathsf{pk}))$.

$\mathcal{H}_2^b$**:** It is identical to $\mathcal{H}_1^b$ except that e is computed as $e := \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}, m_b; s)$ for a uniformly sampled $s \leftarrow^{\$} \{0,1\}^\lambda$.

We then prove that $\mathcal{H}_0^0 \approx_c \mathcal{H}_0^1$ using the above sequence as follows:

$\mathcal{H}_0^b \equiv \mathcal{H}_1^b$**:** Follows directly from the Perfect Zero-Knowledge property of FUNIZK.

$\mathcal{H}_1^b \approx_s \mathcal{H}_2^b$**:** Since $r \leftarrow^{\$} \{0,1\}^{3\lambda}$ and f is in injective mode, $\mathrm{H}_\infty(f(r) \,|\, f) = \mathrm{H}_\infty(r) = 3\lambda$. Since h is a UHF with output length λ, the Leftover Hash Lemma (Lemma 1) implies

$$\Delta((h, f, h \circ f(r)),\ (h, f, s)) \leq 2^{-\lambda}$$

for $s \leftarrow^{\$} \{0,1\}^\lambda$. As the hybrids' output is a function of the above distributions (and independently sampled random coins), we conclude $\mathcal{H}_1^b \approx_s \mathcal{H}_2^b$ by Lemma 2.

$\mathcal{H}_2^0 \approx_c \mathcal{H}_2^1$**:** Follows directly from the IND-CPA-security of E^*.

□

Remark 2. One can analogously prove that the scheme E is IND-CCA-secure if E^* is IND-CCA-secure.

Theorem 2. *Let* $\mathsf{ELF.Gen}$ *be an extremely lossy function,* $\mathcal{H}$ *be a universal hash function,* FUNIZK *be a fully unique NIZK argument and* E^* *a perfectly correct PKE. Suppose that* $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *is an* ε*-correct on average, semi-adaptive anamorphic triplet (cf. Definitions 3 and 5) for the scheme* E *of Figure 4 with anamorphic message space* $\widehat{M}$*. Then* $|\widehat{M}| = \mathsf{poly}(\lambda)$.

Proof of Theorem 2. Consider the distinguisher $\mathcal{D}$ of Figure 5 for the ELF mode. We recall that both M and $\widehat{M}$ are efficiently sampleable and decidable. Thus, by construction, $\mathcal{D}$ runs in polynomial time.

$\mathcal{D}(f)$

```
1 :  // Generate parameters using f
2 :  h ←$ H, crs ← FUNIZK.Setup(1^λ)
3 :  (pp*, _) ← E*.Init(1^λ)
4 :  pp := (f, h, crs, pp*)
5 :  // Encrypt and decrypt a random message
6 :  m ←$ M, m̂ ←$ M̂
7 :  (apk, ask, dk) ← AT.Gen(pp)
8 :  ct = (e, π) ← AT.Enc(apk, dk, m, m̂)
9 :  m̃ := AT.Dec(ask, dk, ct)
10 : // Check consistency and correctness
11 : NV := FUNIZK.Verify(crs, (e, apk), π)   // NIZK Verification
12 : DC := (m == E.Dec(sk, ct))              // Decryption Correctness
13 : KC := ((apk, ask) ∈ Supp(E.Gen(pp)))    // Key Correctness
14 : AC := (m̃ == m̂)                          // Anamorphic Correctness
15 : return (NV ∧ DC ∧ KC ∧ AC)
```

Fig. 5. Distinguisher algorithm breaking the security of the ELF.

The four bits of lines 11–14 stand for "NIZK Verification", "Decryption Correctness", "Key Correctness" and "Anamorphic Correctness", respectively. The intuition behind the design of $\mathcal{D}$ is as follows. Since FUNIZK

has *unique* proofs, the anamorphic encryption scheme can only attempt to embed the anamorphic message in the randomness used to generate the ciphertext e, that is, in the image of $h \circ f$. Indeed, the soundness of FUNIZK guarantees that the ciphertext e is correctly generated. Therefore, if the scheme has a large anamorphic message space, then many messages will collide while encrypting when f is instantiated in lossy mode, making anamorphic correctness information-theoretically hard. While, if f is instantiated in injective mode, we can exploit the correctness of the given AT. All of this leads to a distinguisher against the ELF security. We now proceed with the formal proof.

Proof outline. In the following, we denote $f_{\mathsf{inj}} \leftarrow \mathsf{ELF.Gen}(2^{3\lambda}, 2^{2\lambda})$ and $f_{\mathsf{lossy}} \leftarrow \mathsf{ELF.Gen}(2^{3\lambda}, R)$ respectively injective and lossy mode ELFs, for an arbitrary $R = \mathsf{poly}(\lambda)$ which we specify later.

Injective mode. First of all, we study the probability $\mathcal{D}(f_{\mathsf{inj}}) = 1$. We do so by claiming that all four bits $\mathsf{NV}, \ldots, \mathsf{AC}$ each equal 1 with overwhelming probability. For the first three this follows by semi-adaptive security as each of them is the result of a predicate that the dictator can check, and that is always true for the PKE in Figure 4. Regarding AC, this follows by ε-correctness on average of AT. To prove this case we use the following claims, which we first state and only later prove:

Claim 1. *For $\mathcal{D}(f_{\mathsf{inj}})$ it holds that* $\Pr[\mathsf{NV} \wedge \mathsf{DC} \wedge \mathsf{KC} = 1] \geq 1 - \mathsf{negl}(\lambda)$. *(We refer to Section 3.3 for the proof.)*

Claim 2. *For $\mathcal{D}(f_{\mathsf{inj}})$ it holds that* $\Pr[\mathsf{AC} = 1] \geq 1 - \mathsf{negl}(\lambda)$. *(We refer to Section 3.3 for the proof.)*

Applying the union bound, we conclude that

$$\Pr[\mathcal{D}(f_{\mathsf{inj}}) = 1] \;\geq\; 1 - \Pr[\mathsf{NV} \wedge \mathsf{DC} \wedge \mathsf{KC} = 0] - \Pr[\mathsf{AC} = 0] \;\geq\; 1 - \mathsf{negl}(\lambda).$$

Lossy mode. Next we study the probability that $\mathcal{D}(f_{\mathsf{lossy}}) = 1$. In what follows we denote $\mathsf{E}^m_{\mathsf{apk}}$ the set of valid encryption of m under key apk with respect to E, formally defined as

$$\mathsf{E}^m_{\mathsf{apk}} \;=\; \left\{ \mathsf{E.Enc}(\mathsf{apk}, m; r) \;\middle|\; r \in \{0,1\}^{3\lambda} \right\}.$$

We claim that the probability that ct is not valid while π is a valid proof, $m = \mathsf{E.Dec}(\mathsf{ask}, c)$ and $(\mathsf{apk}, \mathsf{ask})$ is a valid key pair, is negligible, crucially using the FUNIZK soundness and the PKE's perfect correctness:

Claim 3. *For $\mathcal{D}(f_{\mathsf{lossy}})$ it holds that* $\Pr\left[\mathsf{NV}, \mathsf{DC}, \mathsf{KC}, \mathsf{ct} \notin \mathsf{E}^m_{\mathsf{apk}}\right] \leq \mathsf{negl}(\lambda)$. *(We refer to Section 3.3 for the proof.)*

Next, assuming ct to be valid, we lower bound the min-entropy of $\widehat{m}$ conditioned on the random variables $\mathsf{AT.Dec}$ takes as input:

Claim 4. *For $\mathcal{D}(f_{\mathsf{lossy}})$ it holds that*

$$\mathrm{H}_\infty\left(\widehat{m} \;\middle|\; \mathsf{ask}, \mathsf{dk}, \mathsf{ct}; \mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right) \geq \log|\widehat{M}| - \log R + \log\left(\Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right]\right),$$

where R is the range parameter used to generate f_{lossy}. (We refer to Section 3.3 for the proof.)

Using both Claims 3, 4 we can eventually upper bound the probability that $\mathcal{D}(f_{\mathsf{lossy}})$ returns 1, as an application of Lemma 4. To finish the lossy mode case we claim that:

Claim 5. $\Pr[\mathcal{D}(f_{\mathsf{lossy}}) = 1] \leq R/|\widehat{M}| + \mathsf{negl}(\lambda)$. *(We refer to Section 3.3 for the proof.)*

Conclusion. By the above results, the following bound on the advantage of $\mathcal{D}$ can be derived when f_{lossy} is instantiated with image size at most R

$$\mathsf{Adv}_{\mathcal{D}} \;\geq\; \left|\Pr\left[\mathcal{D}(f_{\mathsf{inj}})=1\right]-\Pr\left[\mathcal{D}(f_{\mathsf{lossy}})=1\right]\right| \geq\; 1-R/|\widehat{M}|-\mathsf{negl}(\lambda).$$

Finally, let t be an upper bound on the running time of $\mathcal{D}$. By ELF security there exists an $R=\mathsf{poly}(\lambda)$ such that any t-time adversary has advantage at most $1/2$ in distinguishing f_{inj} from f_{lossy}. This in particular implies

$$1/2 \;\geq\; \mathsf{Adv}_{\mathcal{D}} \;\geq\; 1-R/|\widehat{M}|-\mathsf{negl}(\lambda).$$

By rearranging, $|\widehat{M}| \leq 2R+\mathsf{negl}(\lambda)=\mathsf{poly}(\lambda)$ which proves Theorem 2. □

3.3 Deferred proofs

We now proceed to prove the remaining claims stated in Section 3.2. Throughout the proof, for a bit-valued random variable B we will abuse the notation and write B for the event $\{B=1\}$.

Proof of Claim 1

Claim. For $\mathcal{D}(f_{\mathsf{inj}})$ it holds that $\Pr\left[\mathsf{NV}\wedge\mathsf{DC}\wedge\mathsf{KC}=1\right]\geq 1-\mathsf{negl}(\lambda)$.

Proof. Intuitively, if this were not the case, one could distinguish between real and anamorphic mode by testing all three properties associated to $\mathsf{NV},\mathsf{DC},\mathsf{KC}$. More precisely let $\mathcal{A}$ be the adversary for semi-adaptive security (Definition 3) of (E,AT) described in Figure 6.

$\mathcal{A}^{\mathcal{O}(\cdot,\cdot),\mathcal{O}_{\mathsf{key}}}(\mathsf{pp},\mathsf{pk})$

1 : Parse $\mathsf{pp}=(f,h,\mathsf{crs},\mathsf{pp}^*)$
2 : $m \leftarrow^{\$} M,\ \widehat{m} \leftarrow^{\$} \widehat{M}$
3 : $\mathsf{ct}=(e,\pi) \leftarrow \mathcal{O}(m,\widehat{m})$
4 : $(\mathsf{sk},\mathsf{td}) := \mathcal{O}_{\mathsf{key}}$
5 : // Compute the three bits NV, DC, KC
6 : $\mathsf{NV} := \mathsf{FUNIZK.Verify}(\mathsf{crs},(e,\mathsf{pk}),\pi)$ // NIZK Verification
7 : $\mathsf{DC} := (m == \mathsf{E.Dec}(\mathsf{sk},\mathsf{ct}))$ // Decryption Correctness
8 : $\mathsf{KC} := ((\mathsf{pk},\mathsf{sk}) \in \mathsf{Supp}(\mathsf{E.Gen}(\mathsf{pp})))$ // Key Correctness
9 : **return** $(\mathsf{NV}\wedge\mathsf{DC}\wedge\mathsf{KC})$

Fig. 6. Adversary $\mathcal{A}$ breaking the semi-adaptive security of (E,AT). $\mathcal{O}\in\{\mathcal{O}_{\mathsf{anam}},\mathcal{O}_{\mathsf{real}}\}$ is the encryption oracle. (cf. Definition 3)

In the real game all bits are always equal to 1. Regarding NV, it follows from the perfect completeness of FUNIZK. For DC it is a consequence of perfect correctness, while KC follows from the construction, as $(\mathsf{pk},\mathsf{sk})$ are actually generated by $\mathsf{E.Gen}(\mathsf{pp})$. On the other hand, in the anamorphic game, the public parameters pp received by $\mathcal{A}$ and the ones generated by $\mathcal{D}(f_{\mathsf{inj}})$ are identically distributed, as f_{inj} is generated in the injective mode in both cases. Since $\mathcal{O}=\mathcal{O}_{\mathsf{anam}}$, then $\mathsf{NV},\mathsf{DC},\mathsf{KC}$ are computed as the same (probabilistic) function of pp both in $\mathcal{A}(\mathsf{pp},\mathsf{apk})$ and $\mathcal{D}(f_{\mathsf{inj}})$. We then conclude that

$$\begin{aligned}\mathsf{negl}(\lambda) \;&\geq\; \left|\Pr\left[\mathcal{A}^{\mathcal{O}_{\mathsf{real}},\mathcal{O}_{\mathsf{key}}}(\mathsf{pp},\mathsf{pk})=1\right]-\Pr\left[\mathcal{A}^{\mathcal{O}_{\mathsf{anam}},\mathcal{O}_{\mathsf{key}}}(\mathsf{pp},\mathsf{apk})=1\right]\right| \\ &=\; \Pr\left[\overline{\mathsf{NV}}\vee\overline{\mathsf{DC}}\vee\overline{\mathsf{KC}}\right]=1-\Pr\left[\mathsf{NV}\wedge\mathsf{DC}\wedge\mathsf{KC}\right].\end{aligned}$$

□

Proof of Claim 2

Claim. For $\mathcal{D}(f_{\mathsf{inj}})$ it holds that $\Pr[\mathsf{AC}=1] \geq 1-\mathsf{negl}(\lambda)$.

Proof. By ε-correctness on average, since m is sampled uniformly in M and pp is correctly distributed due to f_{inj} being injective, we have that $\Pr[\mathsf{AC}=1] = \Pr[\widetilde{m}=\widehat{m}] \geq 1-\mathsf{negl}(\lambda)$. □

Proof of Claim 3

Claim. For $\mathcal{D}(f_{\mathsf{lossy}})$ it holds that $\Pr\left[\mathsf{NV},\mathsf{DC},\mathsf{KC},\mathsf{ct} \notin \mathsf{E}^m_{\mathsf{apk}}\right] \leq \mathsf{negl}(\lambda)$.

Proof. We begin observing that the clauses $\mathsf{DC}=\mathsf{KC}=1$ and $\mathsf{ct} \notin \mathsf{E}^m_{\mathsf{apk}}$ imply that $(e,\mathsf{apk}) \notin \mathcal{L}_{\mathcal{R}}$. Indeed, assume by contradiction that $(e,\mathsf{apk}) \in \mathcal{L}_{\mathcal{R}}$. Then $\mathsf{apk} = (\mathsf{pk}^*, f, h, \mathsf{crs}, \mathsf{pp}^*)$ and there exists (m', r') such that $e = \mathsf{E}.\mathsf{Enc}(\mathsf{pk}^*, m'; h \circ f(r'))$. $\mathsf{KC}=1$ implies that $(\mathsf{apk},\mathsf{ask})$ is in the support of $\mathsf{E}.\mathsf{Gen}(\mathsf{pp})$, which by construction implies $(\mathsf{pk}^*,\mathsf{ask})$ is in the support of $\mathsf{E}^*.\mathsf{Gen}(\mathsf{pp})$. The perfect correctness of E^* together with $\mathsf{DC}=1$ implies $\mathsf{E}^*.\mathsf{Dec}(\mathsf{ask},e) = m$, and so $m = m'$. All in all, this would imply $\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}$, yielding a contradiction. Thus,

$$(\mathsf{DC}=\mathsf{KC}=1) \wedge \mathsf{ct} \notin \mathsf{E}^m_{\mathsf{apk}} \;\Rightarrow\; (e,\mathsf{apk}) \notin \mathcal{L}_{\mathcal{R}}.$$

The Claim then follows by the following chain of inequalities:

$$\begin{aligned}
&\Pr\left[\mathsf{NV},\mathsf{DC},\mathsf{KC},\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] \leq \\
&\leq \Pr\left[\mathsf{NV},\ (e,\mathsf{apk}) \notin \mathcal{L}_{\mathcal{R}}\right] \\
&= \Pr\left[\mathsf{FUNIZK}.\mathsf{Verify}(\mathsf{crs},(e,\mathsf{apk}),\pi)=1,\ (e,\mathsf{apk}) \notin \mathcal{L}_{\mathcal{R}}\right] \\
&\leq \mathsf{negl}(\lambda),
\end{aligned}$$

where the last inequality result from the FUNIZK soundness. This is the case as $\mathcal{D}$ can be easily adapted into an adversary breaking soundness, who receives crs, extends it to pp, and eventually returns $((e,\mathsf{apk}),\pi)$. □

Proof of Claim 4

Claim. For $\mathcal{D}(f_{\mathsf{lossy}})$ it holds that

$$\mathrm{H}_\infty\left(\widehat{m} \;\middle|\; \mathsf{ask},\mathsf{dk},\mathsf{ct};\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right) \geq \log|\widehat{M}| - \log R + \log\left(\Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right]\right),$$

where R is the range parameter used to generate f_{lossy}.

Proof. The elements of $\mathsf{E}^m_{\mathsf{apk}}$ are pairs (e,π) and by the uniqueness property of FUNIZK, π is uniquely determined by e, and apk. Therefore,

$$|\mathsf{E}^m_{\mathsf{apk}}| = \left|\left\{\mathsf{E}^*.\mathsf{Enc}(\mathsf{apk},m;h\circ f(r)) \;\middle|\; r \in \{0,1\}^{3\lambda}\right\}\right| \leq R,$$

where R is the range parameter of the ELF (cf. Definition 8). Lemma 3 implies that:

$$\begin{aligned}
&\mathrm{H}_\infty\left(\widehat{m} \;\middle|\; \mathsf{ask},\mathsf{dk},\mathsf{ct};\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right) \;\geq \\
&\geq \mathrm{H}_\infty\left(\widehat{m} \;\middle|\; \mathsf{ask},\mathsf{dk},m,\mathsf{apk},\mathsf{ct};\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right) \\
&\geq \mathrm{H}_\infty\left(\widehat{m} \;\middle|\; \mathsf{ask},\mathsf{dk},m,\mathsf{apk};\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right) - \log|\mathsf{E}^m_{\mathsf{apk}}| \\
&\geq \mathrm{H}_\infty\left(\widehat{m} \mid \mathsf{ask},\mathsf{dk},m,\mathsf{apk}\right) + \log\Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] - \log|\mathsf{E}^m_{\mathsf{apk}}| \\
&\geq \mathrm{H}_\infty\left(\widehat{m}\right) + \log\Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] - \log|\mathsf{E}^m_{\mathsf{apk}}| \\
&\geq \log|\widehat{M}| + \log\Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] - \log R.
\end{aligned}$$

The first inequality is obtained further conditioning on (m,apk). The second is by Lemma 3, property 4. The third is again Lemma 3, property 6. The fourth follows as by construction $\widehat{m}$ is distributed independently of $(\mathsf{ask},\mathsf{dk},m,\mathsf{apk})$. The last one holds since $\widehat{m}$ is uniform over $\widehat{M}$. □

Proof of Claim 5

Claim. $\Pr[\mathcal{D}(f_{\mathsf{lossy}}) = 1] \leq R/|\widehat{M}| + \mathsf{negl}(\lambda)$.

Proof. Writing down the probabilities explicitly:

$$\begin{aligned}
\Pr[\mathcal{D}(f_{\mathsf{lossy}}) = 1] &= \Pr[\mathsf{NV}, \mathsf{DC}, \mathsf{KC}, \mathsf{AC}] \\
&\leq \Pr\left[\mathsf{NV},\ \mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] + \Pr\left[\mathsf{NV}, \mathsf{DC}, \mathsf{KC}, \mathsf{ct} \notin \mathsf{E}^m_{\mathsf{apk}}\right] && (1)\\
&\leq \Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] \cdot \Pr\left[\widetilde{m} = \widehat{m} \,\middle|\, \mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] + \mathsf{negl}(\lambda) && (2)\\
&\leq \Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] \cdot 2^{-\mathrm{H}_\infty\left(\widehat{m} \,\middle|\, \mathsf{ask},\mathsf{dk},\mathsf{ct};\mathsf{ct}\in\mathsf{E}^m_{\mathsf{apk}}\right)} + \mathsf{negl}(\lambda) && (3)\\
&\leq \Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right] \cdot \frac{R}{|\widehat{M}| \cdot \Pr\left[\mathsf{ct} \in \mathsf{E}^m_{\mathsf{apk}}\right]} + \mathsf{negl}(\lambda) && (4)\\
&= R/|\widehat{M}| + \mathsf{negl}(\lambda).
\end{aligned}$$

Where (1) follows by total probability and removing clauses in each term, (2) is due to Claim 3, (3) follows by applying Lemma 4 and (4) is true by Lemma 4. □

Given that semi-adaptive AE is a weaker notion of anamorphism, we immediately obtain the following:

Corollary 1 *Let* $\mathsf{ELF.Gen}$ *be an ELF,* $\mathcal{H}$ *be a UHF,* FUNIZK *be a FUNIZK. Suppose that* $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *is an* ε*-correct on average anamorphic triplet (cf. Definitions 4 and 5) for the scheme* E *of Figure 4 with anamorphic message space* $\widehat{M}$*. Then* $|\widehat{M}| = \mathsf{poly}(\lambda)$*.*

4 Anamorphic Resistant Encryption without iO

In this section, we remove the need of iO and rely on concretely efficient asymmetric primitives. In a nutshell, we modify the previous construction so that the dictator can directly check that the randomness used to prepare the ciphertext was generated using the ELF without relying on generic NIZK. The crucial technical tool is an ELF with a special kind of trapdoor, introduced below.

4.1 Revisiting Trapdoor ELFs

We first recall the trapdoor ELF construction in [43]. The full scheme is formally given in Figure 7. The main idea is to compose a sequence of trapdoor lossy functions as in [30], i.e. of the form $\mathbf{x} \mapsto [A\mathbf{x}]$ with A either full rank or rank 1. Each function is defined over its own group $\mathbb{G}_i$ of polynomial size. However, the size of each $\mathbb{G}_i$ is set to grow double-exponentially in i. In this way it is always possible for any polynomial time t and inverse-polynomial advantage ε to find an i such that matrix-DDH is ε-hard in $\mathbb{G}_i$ against t-time adversaries, thus proving the ELF security.

In order to preserve the trapdoor as in [30], no compressing step is applied between function applications (as opposed to [42]). However, this causes the output bit-length after each step to increase by a factor $\mathsf{poly}(2^i)$ needed to represent elements in $\mathbb{G}_i$. The final bit length can however be still polynomial in $\log M$ if only $\tau = \sqrt{\log \log M}$ many steps are taken. Note in particular that the scheme is efficient when sampling matrices $A_i \in \mathbb{F}_{p_i}^{m_i, n_i}$ with $m_i = c \cdot n_i$ for any constant c, as this expand the final output length only by a factor $c^{\sqrt{\log \log M}} = \mathsf{poly}(\log M)$. While in the original paper $c = 2$ is suggested so that A_i is full rank with overwhelming probability, we will need $c = 3$.

$\mathsf{TELF.GenInj}(M)$

$n_1 = \log M,\ \tau = \sqrt{\log \log M}$
for $i \in \{1, \ldots, \tau\}$:
  $(\mathbb{G}_i, g_i, p_i) \leftarrow \mathsf{GRP.Gen}(2^i)$
  $m_i \coloneqq 3 \cdot n_i$
  $A_i \leftarrow \mathbb{F}_{p_i}^{m_i, n_i} \ : \ \mathsf{rk}(A_i) = n_i$
  Find $n_{i+1} \ : \ \mathbb{G}_i^{m_i} \subseteq \{0,1\}^{n_{i+1}}$
$f \coloneqq (\mathbb{G}_i, g_i, p_i, [A_i]_i)_{i=1}^{\tau}$
$\mathsf{td} \coloneqq (A_i)_{i=1}^{\tau}$
return (f, td)

$\mathsf{TELF.Inv}(\mathsf{td}, y)$

Parse $\mathsf{td} = (A_i)_{i=1}^{\tau}$
Set $\mathbf{x}_{\tau+1} = y$
for $i \in \{\tau, \ldots, 1\}$:
  $\mathbf{Y}_i \coloneqq \varphi_i^{-1}(\mathbf{x}_{i+1})$
  Find $L_i \in \mathbb{F}_{p_i}^{n_i, m_i}$ left inverse of A_i
  $\mathbf{X}_i \coloneqq \mathbf{Y}_{i+1}^{L_i}$
  Find $\mathbf{x}_i \in \{0,1\}^{n_i} \ : \ \mathbf{X}_i = [\mathbf{x}_i]_i$
return $\mathbf{x}_1$

$\mathsf{TELF.Eval}(f, x)$

Parse $f = (\mathbb{G}_i, g_i, p_i, [A_i]_i)$
Set $\mathbf{x}_1 = x \in \{0,1\}^{n_1}$
for $i \in \{1, \ldots, \tau\}$:
  $\mathbf{x}_{i+1} = \varphi_i\left([A_i \mathbf{x}_i]_i\right)$
return $\mathbf{x}_{\tau+1}$

$\mathsf{TELF.GenLos}(M, R)$

$(f, \mathsf{td}) \leftarrow \mathsf{TELF.GenInj}(M)$
Parse $f = (\mathbb{G}_i, g_i, p_i, [A_i]_i)_{i=1}^{\tau}$
Set $j = \max\{i : p_i \leq R\}$
Replace $A_j \leftarrow \mathbb{F}_{p_j}^{m_j, n_j} \ : \ \mathsf{rk}(A_j) = 1$
return $f = (\mathbb{G}_i, g_i, p_i, [A_i]_i)_{i=1}^{\tau}$

Fig. 7. Trapdoor ELF from [43] parametrized by $c \in \mathbb{N}$. The notation $[a]_i = g_i^a$ is extended entry-wise to matrices. $\varphi_i : \mathbb{G}_i^{m_i} \to \{0,1\}^{n_{i+1}}$ maps group elements to their representation entry-wise. $\mathsf{ELF.Inv}$ is implicitly assumed to return $\perp$ if it does not find the discrete logarithm of some group element to be in $\{0,1\}$.

4.2 Construction

We are finally ready to present our second compiler. The main components are:

1. Any IND-CPA public key encryption scheme $\mathsf{E}^* = (\mathsf{E}^*.\mathsf{Gen}, \mathsf{E}^*.\mathsf{Enc}, \mathsf{E}^*.\mathsf{Dec})$;
2. A Lossy Trapdoor Function $\mathsf{LTF} = (\mathsf{GenInj}, \mathsf{GenLos}, \mathsf{Inv})$ with lossy image of size at most $2^{\mu(\lambda)}$;
3. The Trapdoor ELF of Figure 7, which we will instantiate with input length $\eta = \mu + 3\lambda$ and output length $\ell = \mathsf{poly}(\eta) = \mathsf{poly}(\lambda)$;
4. A Universal Hash Family $\mathcal{H}$ of functions of type $\{0,1\}^{\ell} \to \{0,1\}^{\lambda}$.

As in the previous section, we wish to encrypt a message m as $\mathsf{E.Enc}(\mathsf{pk}, m; h \circ f(r))$ with h a universal hash, and f a public ELF in injective mode. This time, however, in order to prove that c was computed correctly, we rely on the existence of a trapdoor for f. Ignoring for the moment the IND-CPA security, a straw-man idea would be to append $f(r)$ to the ciphertext above, let the dictator invert f to recover r, and finally check $f(r)$ was used to encrypt m.

This is, however, insufficient to prove anamorphic resistance. Recall the proof strategy is arguing that correctness of an anamorphic triplet holds when f is injective, but it is information-theoretically hard if f were to be lossy – thus yielding a distinguisher for the ELF. The second step critically requires a way to test membership in $\mathrm{Im}\, f$ even when f is in lossy mode, something the trapdoor alone does not allow. Moreover, if membership in $\mathrm{Im}\, f$ can only be tested with a trapdoor, we must also ensure that ELF security holds even when such trapdoor is given.

To solve these issue we will provide a direct reduction to exponential Matrix-DDH using the concrete Trapdoor ELF in [43] (see Section 4.1). This will enable us to provide the distinguisher with a tailored trapdoor to test membership in an approximation[17] of $\mathrm{Im}\, f$ without affecting ELF security.

[17] I.e., a set $S \supseteq \mathrm{Im}\, f$, whose size is polynomial in $|\mathrm{Im}\, f|$.

Finally, to achieve the IND-CPA-security, $f(r)$ has to be hidden from the IND-CPA adversary, while still allowing the dictator to recover it. Simply encrypting $f(r)$ does not seem to work, as this introduce the need for extra randomness, in which the anamorphic message could be hidden. Instead, we opt to rely on a TLF F and attach $F(f(r))$. This does indeed allow recovering $f(r)$ given a trapdoor for F. Moreover to prove IND-CPA, in the lossy mode of the TLR F only few bits of $f(r)$, say $\mu(\lambda)$, are leaked. By assuming that r has length $\mu + 3\lambda$ we can still apply the Leftover Hash Lemma to conclude that $h \circ f(r)$ is close to uniform, even when $F \circ f(r)$ is leaked.

$\mathsf{E.Init}(\lambda)$

1 : $(f, \mathsf{td}_1) \leftarrow \mathsf{TELF.GenInj}(2^{\mu+3\lambda})$
2 : Let $f : \{0,1\}^{\mu+3\lambda} \to \{0,1\}^{\ell}$
3 : $(F, \mathsf{td}_2) \leftarrow \mathsf{LTF.GenInj}(1^{\lambda}, 1^{\ell})$
4 : $h \leftarrow \mathcal{H}$ such that $h : \{0,1\}^{\ell} \to \{0,1\}^{\lambda}$
5 : $\mathsf{pp} \coloneqq (h, f, F)$, $\mathsf{td} \coloneqq (\mathsf{td}_1, \mathsf{td}_2)$
6 : **return** $(\mathsf{pp}, \mathsf{td})$

$\mathsf{E.Gen}(\mathsf{pp})$

1 : $(\mathsf{pk}^*, \mathsf{sk}^*) \leftarrow \mathsf{E}^*.\mathsf{Gen}(1^{\lambda})$
2 : $\mathsf{pk} \coloneqq (\mathsf{pk}^*, \mathsf{pp})$, $\mathsf{sk} \coloneqq \mathsf{sk}^*$
3 : **return** $(\mathsf{pk}, \mathsf{sk})$

$\mathsf{E.Enc}(\mathsf{pk}, m; r)$

1 : Parse $\mathsf{pk} = (\mathsf{pk}^*, h, f, F)$
2 : $e \coloneqq \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m; h \circ f(r))$
3 : $v \coloneqq F \circ f(r)$
4 : **return** $c = (e, v)$

$\mathsf{E.Dec}(\mathsf{sk}, c)$

1 : Parse $c = (e, v)$
2 : $m \coloneqq \mathsf{E}^*.\mathsf{Dec}(\mathsf{sk}, e)$
3 : **return** m

Fig. 8. Anamorphic resistant encryption scheme from Trapdoor ELFs.

4.3 Security

Theorem 3. *If* TLF *is a secure trapdoor lossy function,* $\mathcal{H}$ *a universal hash function and* E^* *is an* IND-CPA *secure scheme, then the scheme* E *of Figure 8 is* IND-CPA *secure.*

Proof of Theorem 3. We proceed through a hybrids sequence $\mathcal{H}_0, \ldots, \mathcal{H}_3$ progressively modifying the IND-CPA security game. m_0, m_1 denotes the challenge messages queried by a given PPT adversary $\mathcal{A}$ and c^* is the challenge ciphertext encrypting m_b, with $b \in \{0,1\}$ being the challenge bit.

$\mathcal{H}_0^b$: Real IND-CPA game with $c^* = (\mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m_b; h \circ f(r),\ F(f(r)))$.

$\mathcal{H}_1^b$: As $\mathcal{H}_0^b$ but $\mathsf{E.Init}$ samples $F \leftarrow \mathsf{LTF.GenLos}(1^{\lambda}, 1^{\ell})$.

$\mathcal{H}_2^b$: As $\mathcal{H}_1^b$ but $c^* = (\mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m_b; s),\ F(f(r)))$ with $s \leftarrow \{0,1\}^{\lambda}$.

We then prove that $\mathcal{H}_0^0 \approx_c \mathcal{H}_0^1$ using the above sequence as follows:

$\mathcal{H}_0^b \approx_c \mathcal{H}_1^b$: Follows directly from the security properties of TLFs, see Section 2.4.

$\mathcal{H}_1^b \approx_s \mathcal{H}_2^b$: Note that h and (f, r) are independent random variables, with $h : \{0,1\}^{\mu+3\lambda} \to \{0,1\}^{\lambda}$ a universal hash function. Since the image of F contains at most 2^{μ} elements, we have that

$$\begin{aligned} \mathrm{H}_{\infty}(f(r) \mid f, F(f(r))) &\geq \mathrm{H}_{\infty}(f(r) \mid f) - \log |\mathrm{Im}\, F| \\ &\geq \mathrm{H}_{\infty}(r) - \mu \;=\; 3\lambda, \end{aligned}$$

were in the second inequality we use the fact that f is guaranteed to be injective, thus preserving the min-entropy of r. By the Generalized Leftover Hash Lemma (Lemma 1), we have that:

$$\Delta\left((h \circ f(r), h, f, F \circ f(r)),\ (s, h, f, F \circ f(r))\right) \ \leq\ 2^{-\lambda}$$

for $s \leftarrow^{\$} \{0,1\}^{\lambda}$. The adversary's view in $\mathcal{H}_1, \mathcal{H}_2$ is a function of the two terms above (and independently distributed random coins). Hence, the statistical distance between these views is smaller than $2^{-\lambda}$ by Lemma 2.

$\mathcal{H}_2^0 \approx_c \mathcal{H}_2^1$**:** Follows directly from the IND-CPA security of E^*.

□

Remark 3. One can analogously prove that the scheme E is IND-CCA-secure if E^* is IND-CCA-secure.

Theorem 4. *If the exponential-DDH assumption holds for* $\mathsf{GRP.Gen}$ *used in the TELF presented in Figure 7, then for any anamorphic triplet for the PKE of Figure 8 with anamorphic message space* $\widehat{M}$*, that is simultaneously semi-adaptively-secure and* ε*-correct on average, we have that* $|\widehat{M}| \leq \mathsf{poly}(\lambda)$.

Proof of Theorem 4. Let $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ be an ε-correct on average and semi-adaptively secure triplet for the PKE in Figure 8. We build an adversary $\mathcal{A}$ breaking the exponential Matrix-DDH assumption. It internally uses its challenge matrix $[C]$ to instantiate a trapdoor ELF f as in Figure 7. Then it samples h, F to compute $\mathsf{pp} = (h, f, F)$, uses pp to generate anamorphic keys and then encrypt/decrypt a random message pair $(m, \widehat{m})$. Eventually, it checks whether the decrypted anamorphic message $\widetilde{m}$ is the same as the original anamorphic message $\widehat{m}$.

We claim that when f is injective[18], the correctness always holds. However, when f is lossy[19], it is information-theoretically hard to achieve correctness. The first step will follow from the anamorphic security and correctness. For the second one, we claim that $c = (e, v)$ returned by $\mathsf{AT.Enc}$ is such that $\rho = F^{-1}(v)$ lies in a polynomially small set. We do as follows.

First, as per Figure 7, recall that $f = \varphi_\tau \circ f_\tau \circ \ldots \varphi_1 \circ f_1$ with

$$f_i : \mathbb{F}_{p_i}^{n_i} \to \mathbb{G}_i^{m_i}\ :\ f_i(\mathbf{x}) = [A_i \mathbf{x}]_i$$

where $[a]_i = g_i^a$ is the entry-wise exponentiation by $g_i \in \mathbb{G}_i$ and $\varphi_i \colon \mathbb{G}_i^{m_i} \to \{0,1\}^{n_{i+1}}$ is an invertible function representing group elements as fixed-length strings. The adversary $\mathcal{A}$ will appropriately choose the index $j \in [\tau]$ and "program" $[A_j]_j$ with the challenge matrix (we later explain how). The first idea is that if $\mathcal{A}$ generates the remaining A_i for $i > j$, it can also invert φ_i, f_i for $i > j$. The problem is now to test membership in $\operatorname{Im} f_j$.

This is done by letting $\mathcal{A}$ program a *partial trapdoor* in A_j. Specifically, assume $\mathcal{A}$ receives a matrix $[C]_j \in \mathbb{G}_j^{2n_i, n_i}$ either uniform or rank 1. It then samples $B \leftarrow^{\$} \mathbb{F}_{p_j}^{3n_j, 2n_j}$ uniformly and set $[A_j]_j = [BC]_j$[20]. Then, knowing B, $\mathcal{A}$ can easily test membership in $[\operatorname{Im} B]_j$. Note that when C is rank-1, $\mathsf{AT.Enc}$ will receive a matrix $[BC]_i$ which only leaks a linear subspace of dimension 1 of $\operatorname{Im} B$. Using this we can prove that "guessing" a point in $\operatorname{Im} B \setminus \operatorname{Im} BC$ is statically hard. Thus testing membership in $\operatorname{Im} B$ essentially suffices to ensure membership in $\operatorname{Im} BC$. Finally note that if $\mathsf{rk}(C) = 1$, then $|\operatorname{Im} BC| = |\,[\operatorname{Im} BC]_j\,| = |\mathbb{G}_j| = \mathsf{poly}(\lambda)$. A detailed description of $\mathcal{A}$ is given in Figure 9. We recall that both M and $\widehat{M}$ are efficiently sampleable and decidable. Thus, by construction, $\mathcal{A}$ runs in polynomial time.

Formally, we will study the probability that $\mathcal{A}$ returns 1 when $(\mathbb{G}, g, p)$ was generated by $\mathsf{GRP.Gen}(2^j)$ for some $j \in \{1, \ldots, \tau\}$ and $C \in \mathbb{F}_p^{2n_j, n_j}$, so that the condition of Line 5 will never be satisfied. Recall that by the TELF construction $2^\lambda \leq p_j^{n_j}$. Let b be the challenger's bit, i.e., $\mathsf{rk}(C) = 1$ whenever $b = 0$, and C is uniformly sampled whenever $b = 1$.

[18] I.e., when $\mathcal{A}$'s matrix is uniformly random.
[19] I.e., when $\mathcal{A}$'s matrix is rank-1.
[20] Note that $[BC]_j$ can be computed given only B and $[C]_j$.

$\mathcal{A}(\mathbb{G}, g, p, [C])$

1 : // Extract parameters j and λ
2 : Find $j \in \mathbb{N}$ such that $2^{2^j} \leq p < 2 \cdot 2^{2^j}$
3 : Let $n \in \mathbb{N}$ be such that $[C] \in \mathbb{G}^{2n,n}$
4 : Find $\lambda \in \mathbb{N}$ such that $n = n_j(3\lambda)$
5 : **if** any of the above steps failed: **return** 0
6 : // Build pp from $[C]$
7 : Sample $(\mathbb{G}_i, g_i, p_i, [A_i]_i, A_i)$ as in $\mathsf{TELF.GenInj}(2^{3\lambda})$ for $i \in \{1, \ldots, \tau\} \setminus \{j\}$
8 : Set $(\mathbb{G}_j, g_j, p_j) \coloneqq (\mathbb{G}, g, p)$, sample $B \leftarrow \mathbb{F}_{p_j}^{3n_j, 2n_j}$ and set $[A]_j \coloneqq [BC]_j$
9 : $f \coloneqq (\mathbb{G}_i, g_i, p_i, [A]_i)_{i=1}^{\tau},\ h \leftarrow \mathcal{H},\ (F, \mathsf{td}_2) \leftarrow \mathsf{LTF.GenInj}(1^\lambda, 1^\ell),\ \mathsf{pp} \coloneqq (h, f, F)$
10 : // Encrypt and decrypt a random message
11 : Sample $m \leftarrow^{\$} M$ and $\widehat{m} \leftarrow^{\$} \widehat{M}$
12 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}) \leftarrow \mathsf{AT.Gen}(\mathsf{pp})$ with $\mathsf{apk} = (\mathsf{apk}^*, _)$
13 : $c \leftarrow \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$ with $c = (e, v)$
14 : $\widetilde{m} \coloneqq \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{dk}, c)$
15 : // Check the validity of c
16 : Using the A_i, find $u \in \mathbb{G}_j^{m_j} : F \circ \varphi_\tau \circ f_\tau \circ \ldots \circ f_{j+1} \circ \varphi_j(u) = v$
17 : Using td_2, find $\rho : F(\rho) = v$
18 : $\mathsf{AC} \coloneqq (\widetilde{m} = \widehat{m})$ // Anamorphic Correctness
19 : $\mathsf{RC} \coloneqq (u \in [\operatorname{Im} B]_j)$ // Randomness Correctness
20 : $\mathsf{EC} \coloneqq (e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{apk}^*, m; h(\rho)))$ // Encryption Correctness
21 : **return** $(\mathsf{AC} \wedge \mathsf{RC} \wedge \mathsf{EC})$

Fig. 9. Adversary $\mathcal{A}$ for exponential matrix-DDH. We denote $f_i(\mathbf{x}_i) = [A_i \mathbf{x}_i]_i$ which can be efficiently inverted given A_i, and $\varphi_i : \mathbb{G}_i^{m_i} \to \{0,1\}^{n_{i+1}}$ an efficiently invertible map representing group elements as fixed-length strings. $n_j(\lambda)$ is the input-size of f_j when setting up an ELF with $\mathsf{ELF.Gen}(2^\lambda, 2^\lambda)$. Note $n_j(\lambda) = \mathsf{poly}(\lambda)$.

High rank case. When $b = 1$, by Lemma 5, C is full rank except with probability $p_j^{-n_j} \leq 2^{-\lambda}$. In this case, by Lemma 6, $A_j = B \cdot C$ is a uniformly distributed matrix in $\mathbb{F}_{p_j}^{3n_j, n_j}$. In particular, when C is full rank, the parameters pp generated by $\mathcal{A}$ in line 9 are distributed as the ones generated by $\mathsf{E.Init}$. By ε-correctness on average we then have that,

$$\begin{aligned}
\Pr[\neg\mathsf{AC} \mid b = 1] &\leq \Pr[\neg\mathsf{AC} \mid \mathsf{rk}(C) = n_j,\ b = 1] + \Pr[\mathsf{rk}(C) < n_j \mid b = 1] \\
&\leq \Pr[\widetilde{m} \neq \widehat{m} \mid \mathsf{rk}(C) = n_j,\ b = 1] + \Pr[\mathsf{rk}(C) < n_j \mid b = 1] \\
&\leq \varepsilon(\lambda) + p_j^{-n_j} \ = \ \mathsf{negl}(\lambda).
\end{aligned}$$

The remaining tests succeed with overwhelming probability by the following claim. The following claim holds due to semi-adaptive security, since the adversary can compute both RC and EC and use them to distinguish the modes.

Claim 6. $\Pr[\mathsf{RC}, \mathsf{EC} \mid b = 1] \geq 1 - \mathsf{negl}(\lambda)$. *(The proof is in Section 4.4.)*

The probability that $\mathcal{A}$ returns 1 is then readily bounded through a union bound

$$\begin{aligned}
&\Pr\left[\mathcal{A}(\mathbb{G}_j, g_j, p_j, [C]_j) = 1 \,\middle|\, b = 1\right] \\
&\geq\ 1 - \Pr[\neg\mathsf{AC} \mid b = 1] - \Pr[\neg(\mathsf{RC}, \mathsf{EC}) \mid b = 1] \\
&\geq\ 1 - \mathsf{negl}(\lambda).
\end{aligned}$$

Low-rank case. We will use the following information-theoretical lemma, which formalizes the intuition that, for a low-rank D, it is difficult to guess an element of $(\operatorname{Im} B) \setminus (\operatorname{Im} BD)$ given only BD. Note that, in our setting, $2/p_j^{-n_j} \leq \mathsf{negl}(\lambda)$.

Claim 7. *Let $B \leftarrow \mathbb{F}_p^{3n,2n}$, $D \leftarrow \mathbb{F}_p^{2n,n;1}$ and ϕ be a function-valued random variable with values in $\{f : \mathbb{F}_p^{3n,2n} \to \mathbb{F}_p^{3n}\}$, such that ϕ and (B, D) are independent. Then*

$$\Pr\left[\phi(BD) \in (\operatorname{Im} B) \setminus (\operatorname{Im} BD)\right] \;\leq\; 2 \cdot p^{-n}.$$

(We refer to Section 4.4 for the proof).

Next we define the following sets, respectively approximating the set of correctly-derived random coins, and of valid normal-mode ciphertexts encrypting m:

$$\begin{aligned} S_{\mathsf{pp}} \;&=\; \{\operatorname{Im}\left(\varphi_{\tau+1} \circ f_\tau \circ \ldots \varphi_j \circ f_j\right)\}, \\ \mathsf{E}^m_{\mathsf{pp},\mathsf{apk}^*} \;&=\; \{(e, v) \;:\; e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{apk}^*, m; h(\rho)),\; v = F(\rho),\; \rho \in S_{\mathsf{pp}}\}. \end{aligned}$$

Moreover, consider the event $\mathsf{Good} = \{c \in \mathsf{E}^m_{\mathsf{pp},\mathsf{apk}^*}\}$. Note that, for any pp generated by $\mathcal{A}$ in line 9 when $b = 0$, we have that $\mathsf{rk}(A_j) \leq 1$. Thus, $p_j \geq |\operatorname{Im} f_j| \geq |S_{\mathsf{pp}}|$. This, in particular, implies that for any $\mathsf{pp}, \mathsf{apk}^*$ and m generated by $\mathcal{A}$ when $b = 0$ we have that $|\mathsf{E}^m_{\mathsf{pp},\mathsf{apk}^*}| \leq p_j$. We can now upper-bound the probability that $\mathcal{A}$ incorrectly believes it is in the high-rank-mode as follows:

$$\begin{aligned} &\Pr\left[\mathcal{A}(\mathbb{G}_j, g_j, p_j, [C]) = 1 \,|\, b = 0\right] \\ =\; &\Pr\left[\widetilde{m} = \widehat{m},\; u \in [\operatorname{Im} B]_j,\; e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{apk}^*, m; h(\rho)) \,\middle|\, b = 0\right] \\ \leq\; &\Pr\left[\widetilde{m} = \widehat{m},\; u \in [\operatorname{Im} BC]_j,\; e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{apk}, m; h(\rho)) \,\middle|\, b = 0\right] \\ &+ \Pr\left[u \in [\operatorname{Im} B \setminus \operatorname{Im} BC]_j \,\middle|\, b = 0\right] \\ =\; &\Pr\left[\widetilde{m} = \widehat{m},\; \rho \in S_{\mathsf{pp}},\; e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{apk}^*, m; h(\rho)) \,|\, b = 0\right] + \mathsf{negl}(\lambda) \\ =\; &\Pr\left[\widetilde{m} = \widehat{m},\; \mathsf{Good} \,|\, b = 0\right] + \mathsf{negl}(\lambda). \end{aligned}$$

In order to bound the remaining term we observe that $\widehat{m}$ is the output of $\mathsf{AT}.\mathsf{Dec}(\mathsf{ask}, \mathsf{dk}, c)$. We thus study the average min-entropy of $\widehat{m}$ conditioned on those input variables, and the events $(\mathsf{Good},\; b = 0)$.

Claim 8. *With the above notation, it holds that*

$$\mathrm{H}_\infty(\widehat{m} \,|\, \mathsf{ask}, \mathsf{dk}, c; \mathsf{Good}, b = 0) \;\leq\; \log|\widehat{M}| + \log\Pr\left[\mathsf{Good} \,|\, b = 0\right] - \log p_j.$$

(We refer to Section 4.4 for the proof).

Using Lemma 4 we can continue to bound the accepting probability:

$$\begin{aligned} &\Pr\left[\mathcal{A}(\mathbb{G}_j, g_j, p_j, [C]) = 1 \,|\, b = 0\right] \\ \leq &\Pr\left[\widetilde{m} = \widehat{m},\; \mathsf{Good} \,|\, b = 0\right] + \mathsf{negl}(\lambda) \\ \leq &\Pr\left[\mathsf{Good} \,|\, b = 0\right] \cdot 2^{-\mathrm{H}_\infty(\widehat{m} \,|\, \mathsf{ask}, \mathsf{dk}, c; \mathsf{Good}, b=0)} + \mathsf{negl}(\lambda) \\ \leq\; &\Pr\left[\mathsf{Good} \,|\, b = 0\right] \cdot \frac{p_j}{\Pr\left[\mathsf{Good} \,|\, b = 0\right] \cdot |\widehat{M}|} + \mathsf{negl}(\lambda) \\ = &\; p_j \cdot |\widehat{M}|^{-1} + \mathsf{negl}(\lambda). \end{aligned}$$

Conclusion. We showed so far that running $\mathcal{A}$ with input $(\mathbb{G}_j, g_j, p_j, [D])$ with the group being generated with $\mathsf{GRP}.\mathsf{Gen}(2^j)$ and $D \in \mathbb{F}_{p_j}^{m,n}$ with $m \geq 2n_j$ and $n \geq n_j$ then the advantage of $\mathcal{A}$ is bounded by

$$\mathsf{Adv}_{\mathcal{A}} \;\geq\; 1 - \mathsf{negl}(\lambda) - \left(\frac{p_j}{|\widehat{M}|} - \mathsf{negl}(\lambda)\right) \;=\; 1 - \frac{p_j}{|\widehat{M}|} - \mathsf{negl}(\lambda).$$

Let $t = \mathsf{poly}(\lambda)$ an upper bound on the execution of $\mathcal{A}$ for any $j \in \{1, \ldots, \tau\}$. Note that by the construction in Figure 7, when the input has length $\mu + 3\lambda$, then the group operations are efficient in λ and n_j, m_j are polynomials in λ for all j. Let Q be the polynomial whose existence is guaranteed be exponential matrix-DDH (see Definition 12)[21]. Without the loss of generality we will assume Q to be nondecreasing when restricted to any of its coordinates[22].

Then choose j, such that $2^j \geq \log Q(t, n_\tau, m_\tau, 2) > 2^{j-1}$. Since we assume Q to be nondecreasing coordinate-wise and the n_j, m_j are increasing w.r.t. j by construction, we obtain $2^j \geq \log Q(t, n_j, m_j, 2)$, which implies

$$1/2 \geq \mathsf{Adv}_{\mathcal{A}} \Rightarrow |\widehat{M}| \leq 2p_j + \mathsf{negl}(\lambda).$$

By construction of the TELF we have that $p_j \leq 2 \cdot 2^{2^j}$, and so, by our choice of j, $p_j \leq 2 \cdot Q(t, n_\tau, m_\tau, 2)^2 = \mathsf{poly}(\lambda)$. We can therefore conclude $|\widehat{M}| \leq \mathsf{poly}(\lambda)$. □

$\mathcal{D}^{\mathcal{O}(\cdot,\cdot),\mathcal{O}_{\mathsf{key}}}(\mathsf{pp}, \mathsf{pk})$

1 : // Obtain a random encryption
2 : Parse $\mathsf{pp} = (f, h, \mathsf{crs}, \mathsf{pp}^*)$
3 : Parse $\mathsf{pk} = (\mathsf{pk}^*, _, _, _)$
4 : $m \leftarrow^{\$} M$, $\widehat{m} \leftarrow^{\$} \widehat{M}$
5 : $c = (e, \pi) \leftarrow \mathcal{O}(m, \widehat{m})$
6 : // Use the trapdoors
7 : $(\mathsf{sk}, \mathsf{td}) \leftarrow \mathcal{O}_{\mathsf{key}}$
8 : Parse $\mathsf{td} = (\mathsf{td}_1, \mathsf{td}_2)$
9 : $r := \mathsf{TELF.Inv}(\mathsf{td}_1, \rho)$
10 : $\rho := \mathsf{LTF.Inv}(\mathsf{td}_2, v)$
11 : // Compute the checks
12 : $\mathsf{RC}^* = (F \circ f(r) = v)$ // Randomness Correctness
13 : $\mathsf{EC}^* = (e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m; h(\rho)))$ // Encryption Correctness
14 : **return** $(\mathsf{RC}^* \wedge \mathsf{EC}^*)$

Fig. 10. Distinguisher $\mathcal{D}$ breaking the semi-adaptive security of $(\mathsf{E}, \mathsf{AT})$. $\mathcal{O} \in \{\mathcal{O}_{\mathsf{anam}}, \mathcal{O}_{\mathsf{real}}\}$ is the encryption oracle. (cf. Definition 3)

4.4 Deferred proofs

Proof of Claim 6

Claim. $\Pr[\mathsf{RC}, \mathsf{EC} \mid b = 1] \geq 1 - \mathsf{negl}(\lambda)$.

Proof. We construct an adversary $\mathcal{D}$ for the semi-adaptive security of AT, whose pseudocode is presented on Figure 10. Initially, $\mathcal{D}(\mathsf{pp}, \mathsf{pk})$ queries for the encryption of $(m, \widehat{m}) \leftarrow^{\$} M \times \widehat{M}$ and obtains $c = (e, v)$. Then it requests $(\mathsf{sk}, \mathsf{td})$, extracts td_1 to invert f and td_2 to invert F, and computes $\rho := \mathsf{LTF.Inv}(\mathsf{td}_2, v)$ and $r := \mathsf{TELF.Inv}(\mathsf{td}_1, \rho)$. Finally, it performs two checks:

[21] That is, for any $\eta \leq \log Q(t, n, m, 1/\varepsilon)$, any t-time adversary cannot solve an $n \times m$ sized instance over $(\mathbb{G}, g, p) \leftarrow \mathsf{GRP.Gen}(\eta)$ with advantage greater than ε.

[22] This is always possible up to upper bound $Q'(\mathbf{x}) > Q(\mathbf{x})$ with Q' nondecreasing in each entry. A way to do so is to take $z = \|\mathbf{x}\|_2^2$, $f(z) = Q(z, \ldots, z)$, observe that for some constants c, n we have $cz^n \geq f(z)$ and finally set $Q'(\mathbf{x}) = c \cdot \|\mathbf{x}\|_2^{2n}$.

– $\mathsf{RC}^* = (F \circ f(r) = v)$;
– $\mathsf{EC}^* = (e = \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m; h(\rho)))$;

where pk^* is extracted as the first entry of apk and h is the UHF in pp. Eventually $\mathcal{D}$ returns 0 is any of the above checks fail.

Note that $\mathcal{D}$ always returns 1 in case of the real PKE, due to the correctness of the inversion algorithm for F and f. Consequently, considering the above events in the anamorphic game:

$$\mathsf{negl}(\lambda) \;\geq\; \mathsf{Adv}_{\mathcal{D}}(\lambda) \;\geq\; \Pr\left[\neg(\mathsf{RC}^*,\, \mathsf{EC}^*)\right].$$

Consider the *image check* event $\mathsf{IC}^* := \{v \in \operatorname{Im} F \circ f\}$. Then $\mathsf{RC}^* \;\Rightarrow\; \mathsf{IC}^*$, and so $\Pr\left[\neg(\mathsf{IC}^*, \mathsf{EC}^*)\right] \leq \Pr\left[\neg(\mathsf{RC}^*, \mathsf{EC}^*)\right] \leq \mathsf{negl}(\lambda)$. Moreover, noticed that $\mathsf{IC}^*, \mathsf{EC}^*$ are both functions of $(\mathsf{pp}, \mathsf{td}_1, m)$. Since the tuples $(\mathsf{pp}, \mathsf{td}_1, m)$ produced by $\mathcal{A}$ and $\mathcal{D}$, respectively, have statistical distance at most $p_j^{-n_j} \leq 2^{-\lambda}$, the events IC^* and EC^* also occur with negligible probability for $\mathcal{D}$. Finally, $\mathsf{EC}^* = \mathsf{EC}$ (cf. Figure 9), while, letting RK be the event $\{\mathsf{rk}(C) = n_j\}$, we have the following chain of implications:

$$\begin{aligned}(\neg\mathsf{RC}) \wedge \mathsf{RK} \;&\Leftrightarrow\; u \notin [\operatorname{Im} B]_j \wedge \mathsf{rk}(C) = n_j \;\Rightarrow\; u \notin [\operatorname{Im} BC]_j \\ &\Rightarrow\; u \notin \operatorname{Im} f_j \;\Rightarrow\; \rho \notin \operatorname{Im} f \\ &\Rightarrow\; v \notin \operatorname{Im} F \circ f \;\Rightarrow\; \neg\mathsf{IC}^*.\end{aligned}$$

Where we used the fact that f and F are injective. This concludes the proof, since

$$\begin{aligned}\Pr\left[\neg(\mathsf{RC},\, \mathsf{EC})\right] \;&\leq\; \Pr\left[(\neg\mathsf{RC}),\, \mathsf{RK}\right] + \Pr\left[\neg\mathsf{RK}\right] + \Pr\left[\neg(\mathsf{EC})\right] \\ &\leq\; \Pr\left[\neg\mathsf{IC}^*\right] + p_j^{-n_j} + \mathsf{negl}(\lambda) \;\leq\; \mathsf{negl}(\lambda).\end{aligned}$$

□

Proof of Claim 7

Claim. Let $B \leftarrow \mathbb{F}_p^{3n,2n}$, $D \leftarrow \mathbb{F}_p^{2n,n;1}$ and ϕ be a function-valued random variable with values in $\{f : \mathbb{F}_p^{3n,2n} \to \mathbb{F}_p^{3n}\}$, such that ϕ and (B, D) are independent. Then

$$\Pr\left[\phi(BD) \in (\operatorname{Im} B) \setminus (\operatorname{Im} BD)\right] \;\leq\; 2 \cdot p^{-n}.$$

Proof. For any matrix $S \in \mathbb{F}_p^{3n,n}$ we can associate as a consequence of the base extension theorem a *parity-check matrix* $L_S \in \mathbb{F}_p^{3n-d,3n}$, such that:

– $d = \mathsf{rk}(S) = \dim(\operatorname{Im} S)$;
– $\mathsf{rk}(L_S) = 3n - d$;
– $L_S \cdot S = \Omega$ the zero matrix.

We can then define the following three hybrid distributions:

1. $(BD, \operatorname{Im}(L_{BD} \cdot B))$ where $B \leftarrow^{\$} \mathbb{F}_p^{3n,2n}$ and $D \leftarrow^{\$} \mathbb{F}_p^{2n,n;1}$;
2. $(CD, \operatorname{Im}(L_{CD} \cdot C))$ where $C \leftarrow^{\$} \mathbb{F}_p^{3n,2n;2n}$ and $D \leftarrow^{\$} \mathbb{F}_p^{2n,n;1}$;
3. (CD, V) where C, D are as above and $V \leftarrow^{\$} G_p(3n-1, 2n-1)$.

The first two distribution have statistical distance smaller than p^{-n}, since $\Delta(B, C) \leq p^{-n}$ by Lemma 5. To show the second and the third distributions are the same, we study the distance of their second component conditioning on $CD = A_0$ for every $A_0 \in \mathbb{F}_p^{3n,n;1}$.

Under this condition, (C, D) are uniformly distributed over the set

$$S(A_0) = \{(C_0, D_0) \in \mathbb{F}_p^{3n,2n;2n} \times \mathbb{F}_p^{2n,n;1} \;:\; A_0 = C_0 D_0\}.$$

We need to show that $\mathrm{Im}\,(L_{CD} \cdot C)$ conditioned on $CD = A_0$ is uniform in $G_p(3n-1, 2n-1)$. To do so, we show that the map $\phi\colon S(A_0) \to G_p(3n-1, 2n-1)$ sending $(C, D) \mapsto \mathrm{Im}\,(L_{CD} \cdot C)$ is surjective and balanced[23].

Let $V = \phi(C_0, D_0)$ and let $V' \in G_q(3n-1, 2n-1)$. Let $U = \mathrm{Im}\, C_0$ and let $U' = L_{A_0}^{-1}(V')$. Since L_{A_0} matrix is full rank we have that $\dim U' = 2n$ and $\mathrm{Ker}\, L_{A_0} = W \subseteq U'$. Moreover, C_0 being full rank implies that $\dim U = 2n$ and $W \subseteq U$. Consequently, there exists a matrix $T \in \mathbb{F}_p^{3n,3n}$ such that:

- T is invertible.
- $TA_0 = A_0$ (i.e., W contains only eigenvectors of eigenvalue 1).
- $T \cdot U = U'$.

This implies that $\mathrm{Im}\,(L_{A_0} T C_0) = L_{A_0} \cdot T \cdot \mathrm{Im}\,(C_0) = L_{A_0} T \cdot U = L_{A_0} U' = V'$, and so $V' \in \mathrm{Im}\,\psi$, and ϕ is surjective. Furthermore, as T is invertible, the map $(C, D) \mapsto (TC, D)$ is a bijection between $\psi^{-1}(V)$ and $\psi^{-1}(V')$, and so ϕ is balanced.

We can thus conclude that the second and third distributions are identical. Finally, we prove the claim:

$$\begin{aligned}
&\Pr\left[\phi(BD) \in \mathrm{Im}\, B \setminus \mathrm{Im}\, BD\right] = \\
&= p^{-n} + \Pr\left[L_{BD} \cdot \phi(BD) \in \mathrm{Im}\,(L_{BD} B) \setminus \{0\}\right] \\
&\le p^{-n} + \Pr\left[L_{CD} \cdot \phi(CD) \in V \setminus \{0\}\right] \\
&= p^{-n} + \sum_{y_0 \neq 0} \Pr\left[y_0 \in V\right] \Pr\left[L_{CD} \cdot \phi(CD) = y_0\right] \\
&\le p^{-n} + \sum_{y_0 \neq 0} p^{-n} \Pr\left[L_{CD} \phi(CD) = y_0\right] \\
&= p^{-n} + p^{-n} \Pr\left[L_{CD} \phi(CD) = y_0\right] \;\le\; 2p^{-n}.
\end{aligned}$$

The first equality follows as $v \in \mathrm{Im}\, B \setminus \mathrm{Im}\, BD$ iff its projection is a non-zero vector in $\mathrm{Im}\, L_{BD} B$. The inequality is a consequence of $(BD, \mathrm{Im}\,(L_{BD} B))$ and (CD, V) having statistical distance smaller than p^{-n}. The second equality follows as V is statistically independent from ϕ, C, D. The last inequality follows as $y_0 \neq 0$ and V is uniform in $G_p(3n-1, 2n-1)$, which implies that, by Lemma 7:

$$\Pr\left[y_0 \in V\right] \;=\; \frac{|G_p(3n-2, 2n-2)|}{|G_p(3n-1, 2n-1)|} \;=\; \frac{p^{2n-1}-1}{p^{3n-1}-1} \;\le\; \frac{1}{p^n}.$$

□

Proof of Claim 8

Claim. With the above notation, it holds that

$$\mathrm{H}_\infty(\widehat{m} \mid \mathsf{ask}, \mathsf{dk}, c; \mathsf{Good}, b = 0) \;\le\; \log|\widehat{M}| + \log \Pr\left[\mathsf{Good} \mid b = 0\right] - \log p_j.$$

Proof.

$$\begin{aligned}
&\mathrm{H}_\infty(\widehat{m} \mid \mathsf{ask}, \mathsf{dk}, c; \mathsf{Good}, b = 0) \\
&\ge \mathrm{H}_\infty(\widehat{m} \mid \mathsf{pp}, \mathsf{apk}, m, \mathsf{ask}, \mathsf{dk}, c; \mathsf{Good}, b = 0) && (5) \\
&\ge \mathrm{H}_\infty(\widehat{m} \mid \mathsf{pp}, \mathsf{apk}, m, \mathsf{ask}, \mathsf{dk}; \mathsf{Good}, b = 0) - \log p_j && (6) \\
&\ge \mathrm{H}_\infty(\widehat{m} \mid \mathsf{pp}, \mathsf{apk}, m, \mathsf{ask}, \mathsf{dk}; b = 0) + \log \Pr\left[\mathsf{Good} \mid b = 0\right] - \log p_j && (7) \\
&= \mathrm{H}_\infty(\widehat{m} \mid b = 0) + \log \Pr\left[\mathsf{Good} \mid b = 0\right] - \log p_j && (8) \\
&= \log|\widehat{M}| + \log \Pr\left[\mathsf{Good} \mid b = 0\right] - \log p_j, && (9)
\end{aligned}$$

[23] That is, the preimage of any two element in $\mathrm{Im}\,\phi$ have the same size

where (5) follows by further conditioning on $(\mathsf{pp}, \mathsf{apk}, m)$. For (6), we notice that for any $(\mathsf{pp}, \mathsf{apk}, m)$ in their support and conditioned on Good, we have that $c \in \mathsf{E}^m_{\mathsf{pp},\mathsf{apk}^*}$ with $|\mathsf{E}^m_{\mathsf{pp},\mathsf{apk}^*}| \leq p_j$, and use Lemma 3, subitem 4. (7) follows by Lemma 3, subitem 6. In turn, (8) follows from the fact that, by construction, $\widehat{m}$ and $(\mathsf{pp}, \mathsf{apk}, m, \mathsf{ask}, \mathsf{dk})$ are mutually independent, even when conditioned on $b = 0$. Finally, (9) holds due to $\widehat{m} \leftarrow^{\$} \widehat{M}$. □

5 The "Worst" of Both Worlds

In Sections 3 and 4 we have shown two PKEs for which any ε-correct on average anamorphic triplet yielding semi-adaptive AE can send at most a logarithmic number of anamorphic bits, i.e., the size of its anamorphic message is $|\widehat{M}| = \mathsf{poly}(\lambda)$. Moreover, [10] showed how to construct PKEs which do not admit any ε-correct on average anamorphic triplet yielding adaptive AE.

Theorem 5 ([10], Informal). *There exists a compiler that, given as an input any* IND-CPA *(resp.* IND-CCA*) secure PKE scheme* E'*, produces an* IND-CPA *(resp.* IND-CCA*) secure PKE* E^* *for which no* ε*-correct on average anamorphic triplet can yield adaptive AE (i.e., Definition 4).*

Since ε-correctness on average is a key requirement for an anamorphic triplet, we essentially have that:

1. Our constructions in Sections 3 and 4 tell us that we can build a PKE E_1 where the anamorphic message space of any semi-adaptive AE (Definition 3) is polynomially bounded;
2. The compiler in [10] shows us how to construct a PKE E_2 where no anamorphic triplet can yield adaptive AE (Definition 4).

Nevertheless, these two results tell us nothing about the existence of a PKE E_3 that *simultaneously* has polynomially-bounded anamorphic message space when considering semi-adaptive AE, and prevents anamorphic encryption altogether when considering the notion of adaptive AE. In this section, we show how to construct such a scheme E_3.

Since semi-adaptive AE is a weaker notion than adaptive AE, a PKE admitting only semi-adaptive AE with small anamorphic message space will only admit adaptive AE with an *equally small* anamorphic message space. Nevertheless, we seek a stronger limitations for adaptive AE, namely, the impossibility of transmitting *even a single* anamorphic bit.

We achieve our goal by showing that when our compiler in Figure 4 takes as input a PKE E^* for which no ε-correct on average anamorphic triplet can yield adaptive AE (Definition 4), the same holds for the resulting PKE E.

Hence, we can use the PKE of [10][24] for which adaptive AE is impossible and feed it to our compiler of Figure 4. In more details, given any PKE scheme E', we can first pass it through the compiler of [10] obtaining the PKE scheme E^* and then, give E^* in input to the compiler of Figure 4, yielding a PKE scheme E which is the worst possible PKE from the users' point of view – and the best one from dictators' point of view. Formally, we prove the following theorem:

Theorem 6. *Let* E^* *be a PKE in the Public Parameters model for which no* ε*-correct on average adaptive AE exists. Then, the same holds for the PKE scheme* E *obtained by applying the compiler of Figure 4 on input* E^**.*

Proof. We prove the theorem by contradiction. Namely, suppose that there exists an anamorphic triplet AT for the PKE scheme E providing adaptive AE, and that it is ε-correct on average. Then we can construct an anamorphic triplet AT^* for the PKE E^* (the one given in input to the compiler) providing adaptive AE and that is ε-correct on average. The triplet AT^* is given in Figure 11. The intuition is straightforward: since E and E^* differ by the presence of the FUNIZK proof, the reduction can simply simulate the missing proof. Moreover, the extra elements in the public parameters of E can be stored in the double key dk^*.

[24] We remark that the compiler of [10] preserves the perfect correctness of the underlying PKE.

$\mathsf{AT}^*.\mathsf{Gen}(\mathsf{pp}^*)$

1 : $f \leftarrow \mathsf{ELF.Gen}(2^{3\lambda}, 2^{3\lambda})$
2 : $h \leftarrow^{\$} \mathcal{H}$
3 : $(\mathsf{crs}, \mathsf{st}) \leftarrow S_0(\lambda)$
4 : $\mathsf{pp} \coloneqq (f, h, \mathsf{crs}, \mathsf{pp}^*)$
5 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}) \leftarrow \mathsf{AT.Gen}(\mathsf{pp})$
6 : Parse $\mathsf{apk} \coloneqq (\mathsf{apk}', f, h, \mathsf{crs}, \mathsf{pp}^*)$
7 : $\mathsf{apk}^* \coloneqq (\mathsf{apk}', \mathsf{pp}^*)$
8 : $\mathsf{dk}^* \coloneqq (\mathsf{dk}, f, h, \mathsf{st}, \mathsf{crs})$
9 : **return** $(\mathsf{apk}^*, \mathsf{ask}, \mathsf{dk}^*)$

$\mathsf{AT}^*.\mathsf{Enc}(\mathsf{apk}^*, \mathsf{dk}^*, m, \widehat{m})$

1 : Parse $\mathsf{apk}^* \coloneqq (\mathsf{apk}', \mathsf{pp}^*)$
2 : Parse $\mathsf{dk}^* \coloneqq (\mathsf{dk}, f, h, \mathsf{crs})$
3 : $\mathsf{apk} \coloneqq (\mathsf{apk}', f, h, \mathsf{crs}, \mathsf{pp}^*)$
4 : $c = (e, \pi) \leftarrow \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$
5 : **return** e

$\mathsf{AT}^*.\mathsf{Dec}(\mathsf{ask}, \mathsf{dk}^*, e)$

1 : Parse $\mathsf{dk}^* \coloneqq (\mathsf{dk}, f, h, \mathsf{crs})$
2 : $\pi' \coloneqq S_1(\mathsf{st}, (e, \mathsf{pk}, h, f))$
3 : $c' \coloneqq (e, \pi')$
4 : **return** $\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{dk}, c')$

Fig. 11. Anamorphic triplet AT^* for E^* constructed from AT for E. $\mathsf{ELF.Gen}$ is an ELF (Definition 8), $\mathcal{H}$ a family of hash functions $\{0,1\}^{3\lambda} \to \{0,1\}^{\lambda}$ and $S = (S_0, S_1)$ the simulator of a FUNIZK (Definition 13) for the same relation R defined for Figure 4.

Claim 9. *If the anamorphic triplet* AT *is* ε*-correct on average, then the anamorphic triplet* AT^* *is also* ε*-correct on average.*

Proof. We show that the ε-correctness on average of AT^* simply follows from the one of AT. To prove this, it suffices to show that the anamorphic triplet $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ is run on inputs that are identically distributed to when the triplet is run on top of an honest execution of E.

1. Regarding the input of $\mathsf{AT.Gen}$, the public parameters pp given in input to $\mathsf{AT.Gen}$ are generated with the same distribution of $\mathsf{E.Init}$, as required by Definition 5. Indeed, the ELF f and the UHF h are sampled with the same distribution of $\mathsf{E.Init}$ by construction, while crs is identically distributed in both cases thanks to the perfect zero-knowledge of FUNIZK.
2. Regarding the inputs of $\mathsf{AT.Enc}$, $\mathsf{AT}^*.\mathsf{Enc}$ can reassemble apk as it was generated by $\mathsf{AT.Gen}$ thanks to f, h, crs that are available in dk^* [25]. Thus, the inputs to $\mathsf{AT.Enc}$ are identically distributed to the ones of an honest execution of the triplet. In particular, this implies that the ciphertext $c = (e, \pi)$ obtained from $\mathsf{AT.Enc}$ is computed properly. The proof π is discarded and only e is given in output[26].
3. Regarding the inputs of $\mathsf{AT.Dec}$, $\mathsf{AT}^*.\mathsf{Dec}$ can recreate the proof π that was discarded by $\mathsf{AT}^*.\mathsf{Enc}$ using the simulator of FUNIZK. In this step, it is crucial that FUNIZK is perfect zero-knowledge with fully unique proofs, as this implies that Prove and S_1 produce the *same* proof. In particular, this allows to supply $\mathsf{AT.Dec}$ with a ciphertext $c = (e, \pi)$ that has the same distribution of $\{\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})\}$ even though $\mathsf{AT}^*.\mathsf{Dec}$ does not know m and $\widehat{m}$. Hence, $\mathsf{AT.Dec}^*$ recomputes $\widehat{m}$ with the same negligible error ε.

□

Claim 10. *If the anamorphic triplet* AT *yields adaptive AE for* E*, then the anamorphic triplet* AT^* *yields adaptive AE for* E^**.*

Proof. Consider any adversary $\mathcal{A}^*$ against the anamorphic security of $(\mathsf{E}^*, \mathsf{AT}^*)$. We construct an adversary $\mathcal{A}$ against the anamorphic security of $(\mathsf{E}, \mathsf{AT})$. The adversary $\mathcal{A}$ is given in Figure 12.

We claim that the advantage of $\mathcal{A}^*$ in distinguishing the anamorphic game from the real one is at most equal to $\mathcal{A}$'s advantage plus a negligible term, i.e.,

$$\mathsf{Adv}^{\mathsf{Anam}}_{\mathcal{A}^*, \mathsf{E}^*, \mathsf{AT}^*}(1^\lambda) \leq \mathsf{Adv}^{\mathsf{Anam}}_{\mathcal{A}, \mathsf{E}, \mathsf{AT}}(1^\lambda) + \mathsf{negl}(\lambda). \tag{10}$$

[25] These elements cannot be stored in apk^* as it would clearly compromise AE security.
[26] Again, this is done to guarantee AE security.

$\mathcal{A}^{\mathcal{O}}(\mathsf{pp}, \mathsf{pk}, \mathsf{sk})$

1 : Parse $\mathsf{pp} \coloneqq (f, h, \mathsf{crs}, \mathsf{pp}^*)$
2 : Run $\mathcal{A}^*(\mathsf{pp}^*, \mathsf{pk}, \mathsf{sk})$
3 : Whenever $\mathcal{A}^*$ makes a query $(m, \widehat{m})$:
4 : $\quad c \leftarrow \mathcal{O}(m, \widehat{m})$
5 : $\quad$ Parse $c \coloneqq (e, \pi)$
6 : $\quad$ Give e to $\mathcal{A}^*$
7 : **return** $\mathcal{A}^*$'s output

Fig. 12. Adversary $\mathcal{A}$ breaking the anamorphic security of $(\mathsf{E}, \mathsf{AT})$. $\mathcal{O} \in \{\mathcal{O}_{\mathsf{real}}, \mathcal{O}_{\mathsf{anam}}\}$ is the encryption oracle of Figure 3 returning the output of either $\mathsf{E}.\mathsf{Enc}$ or $\mathsf{AT}.\mathsf{Enc}$ (in $\mathsf{AnamorphicG}_{\mathsf{AT}}$ and $\mathsf{RealG}_{\mathsf{E}}$).

We will show that $\mathcal{A}$ almost perfectly simulates both games for $\mathcal{A}^*$. We claim that the arguments that $\mathcal{A}$ and $\mathcal{A}^*$ receive in $\mathsf{AnamorphicG}_{\mathsf{AT}}$ and $\mathsf{AnamorphicG}_{\mathsf{AT}^*}$, respectively, (cf. Figure 2) are identically distributed. In case of pp this follows directly from the construction of E^* (cf. Figure 4), in case of pk and sk this follows directly from construction of AT^* (cf. Figure 11). Moreover, the output of $\mathsf{AT}^*.\mathsf{Enc}$ equals, by construction, the first element of the output of $\mathsf{AT}.\mathsf{Enc}$ (cf. Figure 11), so the value e provided to $\mathcal{A}^*$ is the same in both games. Hence, we conclude that:

$$\Pr\left[\mathsf{AnamorphicG}_{\mathsf{AT}}(1^\lambda, \mathcal{A}) = 1\right] = \Pr\left[\mathsf{AnamorphicG}_{\mathsf{AT}^*}(1^\lambda, \mathcal{A}^*) = 1\right].$$

We now analyze the behavior of $\mathcal{A}$ and $\mathcal{A}^*$ in $\mathsf{RealG}_{\mathsf{E}}$ and $\mathsf{RealG}_{\mathsf{E}^*}$ respectively. In $\mathsf{RealG}_{\mathsf{E}^*}$, the adversary $\mathcal{A}^*$ will be given the output of $\mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}, m; r)$ for $r \leftarrow^{\$} \{0,1\}^\lambda$. On the other hand, in $\mathsf{RealG}_{\mathsf{E}}$, the adversary $\mathcal{A}$ will be given the output of $\mathsf{E}^*\mathsf{Enc}(\mathsf{pk}, m; h \circ f(r))$ for $r \leftarrow^{\$} \{0,1\}^{3\lambda}$ (cf. Figure 4). We now claim that the two distributions are statistically close, which was proved in $\mathcal{H}_1 \approx_s \mathcal{H}_2$ of Section 3.2 [27]. Given the last claim, it follows that $\mathcal{A}$ simulates the anamorphic game to $\mathcal{A}^*$ perfectly except with negligible probability, i.e.,

$$\Pr\left[\mathsf{RealG}_{\mathsf{E}}(1^\lambda, \mathcal{A}) = 1\right] = \Pr\left[\mathsf{RealG}_{\mathsf{E}^*}(1^\lambda, \mathcal{A}^*) = 1\right] + \mathsf{negl}(\lambda).$$

Subtracting and invoking the triangle inequality, we obtain Equation (10). □

Since E^* does not admit any anamorphic triplet that yields adaptive AE and that is ε-correct on average, we obtain a contradiction. Therefore, E cannot admit any anamorphic triplet that is ε-correct on average yielding adaptive AE. □

Combining Theorem 6 with the results of [10] yields a PKE E ensuring that semi-adaptive anamorphic triplets can send at most $O(\log \lambda)$ anamorphic bits and that does not admit (adaptive) anamorphic encryption. This is formally stated in the following corollary.

Corollary 1. *There exists a PKE scheme* E *such that:*

- *No ε-correct on average anamorphic triplet yields adaptive AE;*
- *For any ε-correct on average anamorphic triplet with anamorphic message space $\widehat{M}$ yielding semi-adaptive AE it holds that $|\widehat{M}| = \mathsf{poly}(\lambda)$.*

Remark 4. Unfortunately, the strategy used to prove Theorem 6 does not work if we use our compiler of Section 4 instead of the one in Section 3. To see why, consider the final ciphertext produced by the compiler of Section 4 which is of the form $c = (e, v)$, where $e \coloneqq \mathsf{E}^*.\mathsf{Enc}(\mathsf{pk}^*, m; h \circ f(r))$ and $v \coloneqq F \circ f(r)$. Intuitively, we are aiming for a ciphertext that cannot transmit *even one* anamorphic bit but, for any given e, there are multiple admissible values of v. Therefore, v could directly be used (e.g., via rejection sampling) to possibly encode a (short) anamorphic message.

[27] We point out that the fact that $\mathcal{A}^*$ knows the secret key of the PKE does not have any impact on his view of the produced ciphertext, as it will decrypt to the same regular message that has been queried.

6 Acknowledgements

This work is part of the grant JDC2023-050791-I, funded by MCIN/AEI/10.13039/501100011033 and the ESF+. This result is part of projects that have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under project PICOCRYPT (grant agreement No. 101001283), and from the Spanish Government under projects PRODIGY (TED2021-132464B-I00) and ESPADA (PID2022-142290OB-I00). The last two projects are co-funded by European Union EIE, and NextGenerationEU/PRTR funds. This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 innovation program (grant PROCONTRA-885666). This work was also supported by the project PARTHENON (B53D23013000006), under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU. This work has also been supported by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme in the scope of the CONFIDENTIAL6G project under Grant Agreement `101096435`. The contents of this publication are the sole responsibility of the authors and do not in any way reflect the views of the EU.

References

1. Fabio Banfi, Konstantin Gegier, Martin Hirt, Ueli Maurer, and Guilherme Rito. Anamorphic encryption, revisited. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 3–32, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.
2. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2), May 2012.
3. Mihir Bellare, Alexandra Boldyreva, and Adriana Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 171–188, Interlaken, Switzerland, May 2–6, 2004. Springer Berlin Heidelberg, Germany.
4. Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway. Security of symmetric encryption against mass surveillance. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 1–19, Santa Barbara, CA, USA, August 17–21, 2014. Springer Berlin Heidelberg, Germany.
5. Nir Bitansky and Arka Rai Choudhuri. Characterizing deterministic-prover zero knowledge. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 535–566, Durham, NC, USA, November 16–19, 2020. Springer, Cham, Switzerland.
6. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112, Chicago, IL, USA, May 2–4, 1988. ACM Press.
7. Pedro Branco, Nico Döttling, and Jesko Dujmovic. Rate-1 incompressible encryption from standard assumptions. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part II*, volume 13748 of *LNCS*, pages 33–69, Chicago, IL, USA, November 7–10, 2022. Springer, Cham, Switzerland.
8. Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218, Dallas, TX, USA, May 23–26, 1998. ACM Press.
9. Ran Canetti, Oded Goldreich, and Shai Halevi. On the random-oracle methodology as applied to length-restricted signature schemes. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 40–57, Cambridge, MA, USA, February 19–21, 2004. Springer Berlin Heidelberg, Germany.
10. Davide Carnemolla, Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Anamorphic resistant encryption: the good, the bad and the ugly. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 472–503, Cham, 2025. Springer Nature Switzerland.
11. J.Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979.
12. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Anamorphic encryption: New constructions and homomorphic realizations. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 33–62, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.
13. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Limits of black-box anamorphic encryption. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part II*, volume 14921 of *LNCS*, pages 352–383, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.
14. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Generic anamorphic encryption, revisited: New limitations and constructions. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part II*, volume 15602 of *LNCS*, pages 275–303, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.
15. Suvradip Chakraborty, Manoj Prabhakaran, and Daniel Wichs. Witness maps and applications. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 220–246, Edinburgh, UK, May 4–7, 2020. Springer, Cham, Switzerland.
16. Suvradip Chakraborty, Manoj Prabhakaran, and Daniel Wichs. A map of witness maps: New definitions and connections. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part II*, volume 13941 of *LNCS*, pages 635–662, Atlanta, GA, USA, May 7–10, 2023. Springer, Cham, Switzerland.
17. Xuan Thanh Do, Giuseppe Persiano, Duong Hieu Phan, and Moti Yung. Anamorphism beyond one-to-one messaging: Public-key with anamorphic broadcast mode. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 429–455, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.
18. Yevgeniy Dodis and Eli Goldin. Anamorphic-resistant encryption; or why the encryption debate is still alive. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 440–471, Cham, 2025. Springer Nature Switzerland.
19. Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.

20. Antonio Faonio, Jesper Buus Nielsen, and Daniele Venturi. Predictable arguments of knowledge. In Serge Fehr, editor, *PKC 2017, Part I*, volume 10174 of *LNCS*, pages 121–150, Amsterdam, The Netherlands, March 28–31, 2017. Springer Berlin Heidelberg, Germany.
21. Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.
22. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press.
23. Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *44th FOCS*, pages 102–115, Cambridge, MA, USA, October 11–14, 2003. IEEE Computer Society Press.
24. Joseph Jaeger and Roy Stracovsky. Dictators? Friends? Forgers. - breaking and fixing unforgeability definitions for anamorphic signature schemes. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part II*, volume 15485 of *LNCS*, pages 105–137, Kolkata, India, December 9–13, 2024. Springer, Singapore, Singapore.
25. Dmitry Khovratovich, Ron D. Rothblum, and Lev Soukhanov. How to prove false statements: Practical attacks on fiat-shamir. Cryptology ePrint Archive, Paper 2025/118, 2025.
26. Miroslaw Kutylowski, Giuseppe Persiano, Duong Hieu Phan, Moti Yung, and Marcin Zawada. Anamorphic signatures: Secrecy from a dictator who only permits authentication! In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 759–790, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.
27. Miroslaw Kutylowski, Giuseppe Persiano, Duong Hieu Phan, Moti Yung, and Marcin Zawada. The self-anti-censorship nature of encryption: On the prevalence of anamorphic cryptography. *Proc. Priv. Enhancing Technol.*, 2023(4):170–183, 2023.
28. Matt Lepinski, Silvio Micali, and abhi shelat. Fair-zero knowledge. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 245–263, Cambridge, MA, USA, February 10–12, 2005. Springer Berlin Heidelberg, Germany.
29. Jesper Buus Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 111–126, Santa Barbara, CA, USA, August 18–22, 2002. Springer Berlin Heidelberg, Germany.
30. Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571, Santa Barbara, CA, USA, August 17–21, 2008. Springer Berlin Heidelberg, Germany.
31. Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 187–196, Victoria, BC, Canada, May 17–20, 2008. ACM Press.
32. Giuseppe Persiano, Duong Hieu Phan, and Moti Yung. Anamorphic encryption: Private communication against a dictator. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 34–63, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland.
33. Giuseppe Persiano, Duong Hieu Phan, and Moti Yung. Public-key anamorphism in (CCA-secure) public-key encryption and beyond. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part II*, volume 14921 of *LNCS*, pages 422–455, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.
34. Willy Quach, LaKyah Tyner, and Daniel Wichs. Unique NIZKs and steganography detection. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*, pages 34–63, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland.
35. Jorge Luis Villar. Optimal reductions of some decisional problems to the rank problem. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 80–97, Beijing, China, December 2–6, 2012. Springer Berlin Heidelberg, Germany.
36. Weihao Wang, Shuai Han, and Shengli Liu. Anamorphic authenticated key exchange: Double key distribution under surveillance. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part V*, volume 15488 of *LNCS*, pages 168–200, Kolkata, India, December 9–13, 2024. Springer, Singapore, Singapore.
37. Yi Wang, Rongmao Chen, Xinyi Huang, and Moti Yung. Sender-anamorphic encryption reformulated: Achieving robust and generic constructions. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VI*, volume 14443 of *LNCS*, pages 135–167, Guangzhou, China, December 4–8, 2023. Springer, Singapore, Singapore.
38. Brent Waters and David J. Wu. Adaptively-sound succinct arguments for NP from indistinguishability obfuscation. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 387–398, Vancouver, BC, Canada, June 24–28, 2024. ACM Press.
39. Brent Waters and David J. Wu. A pure indistinguishability obfuscation approach to adaptively-sound SNARGs for NP. Cryptology ePrint Archive, Report 2024/933, 2024.
40. Brent Waters and Mark Zhandry. Adaptive security in SNARGs via iO and lossy functions. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 72–104, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

41. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167, Toronto, Ontario, Canada, October 27–29, 1986. IEEE Computer Society Press.
42. Mark Zhandry. The magic of ELFs. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 479–508, Santa Barbara, CA, USA, August 14–18, 2016. Springer Berlin Heidelberg, Germany.
43. Mark Zhandry. On ELFs, deterministic encryption, and correlated-input security. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 3–32, Darmstadt, Germany, May 19–23, 2019. Springer, Cham, Switzerland.