Limits of Black-Box Anamorphic Encryption

Dario Catalano[1], Emanuele Giunta[2,3], and Francesco Migliaro[1]

[1] Dipartimento di Matematica e Informatica, Università di Catania, Italy.
catalano@dmi.unict.it, francesco.migliaro@phd.unict.it
[2] IMDEA Software Institute, Madrid, Spain.
emanuele.giunta@imdea.org
[3] Universidad Politecnica de Madrid, Spain.

Abstract. (Receiver) Anamorphic encryption, introduced by Persiano *et al.* at Eurocrypt 2022, considers the question of achieving private communication in a world where secret decryption keys are under the control of a dictator. The challenge here is to be able to establish a secret communication channel to exchange covert (i.e. anamorphic) messages on top of some already deployed public key encryption scheme.

Over the last few years several works addressed this challenge by showing new constructions, refined notions and extensions. Most of these constructions, however, are either ad hoc, in the sense that they build upon specific properties of the underlying PKE, or impose severe restrictions on the size of the underlying anamorphic message space.

In this paper we consider the question of whether it is possible to have realizations of the primitive that are both generic and allow for large anamorphic message spaces. We give strong indications that, unfortunately, this is not the case.

Our first result shows that *any black-box realization* of the primitive, i.e. any realization that accesses the underlying PKE only via oracle calls, *must* have an anamorphic message space of size at most $\mathsf{poly}(\lambda)$ (λ security parameter).

Even worse, if one aims at stronger variants of the primitive (and, specifically, the notion of asymmetric anamorphic encryption, recently proposed by Catalano *et al.*) we show that such black-box realizations are plainly impossible, i.e. no matter how small the anamorphic message space is.

Finally, we show that our impossibility results are rather tight: indeed, by making more specific assumptions on the underlying PKE, it becomes possible to build generic AE where the anamorphic message space is of size $\Omega(2^\lambda)$.

Table of Contents

1 Introduction

Anamorphic encryption [PPY22] (AE, for short) is a novel paradigm to allow private communication in a world where a dictator has the power to control every user to some extent. This includes knowing the secret keys corresponding to public encryption key (violating *receiver privacy*) and limiting user's freedom to choose the message sent (violating *sender freedom*). Such capabilities are indeed plausible in dictatorships, where, for instance, citizens may be subject to censorship measures.

In [PPY22] Persiano *et al.* proposed two flavors of Anamorphic Encryption depending on whether sender freedom or receiver privacy are violated. The two notions (called *sender anamorphic* and *receiver anamorphic* encryption, respectively) build from similar ideas. For the case of receiver anamorphic encryption, which is the focus of this work, AE can be deployed in one (of two) possible modes: regular and anamorphic. In regular mode, the encryption scheme works exactly as an ordinary public key scheme. In anamorphic mode, a public key (apk) is produced along with *two* secret keys: a regular looking one (ask) and a second one called *double key* (dk). Bob shares dk privately with Alice and uses apk as his public key. When forced to reveal his secret key though, he only hands over ask.

To avoid suspicion ($\mathsf{apk}, \mathsf{ask}$) have to be compatible with the regular mode scheme. However, Alice can further use dk as a symmetric key to embed an extra message into the ciphertext that remains hidden even when ask is leaked. More specifically, when employed in anamorphic mode, the scheme permits the encryption of two messages: a regular-looking m, meant to be observed by the dictator, and a covert $\widehat{m}$. The resulting anamorphic ciphertext reveals either m, when decrypted with ask, or $\widehat{m}$ when decrypted anamorphically using dk. A crucial requirement then is that regular ciphertexts should be indistinguishable from anamorphically created ones.

In [PPY22] Persiano et al. argued that, to address privacy needs in the presence of a dictator, coming up with new schemes might be useless: the dictator can simply ban them as illegal and prevent their usage. Therefore, the challenge lies in showing that *existing*, possibly currently deployed, constructions are (or can be easily made) anamorphic.

Over the last two years several works [PPY22, KPP$^+$23b, BGH$^+$24, WCHY23, CGM24] took on this challenge, resulting in new constructions surprisingly covering a large class of known encryption schemes[4]. Most of them, however, exploit rather *specific* properties of the underlying encryption scheme. In particular a dictator can always ban PKE schemes achieving those properties known to yield anamorphic encryption. For this reason generic constructions, that do not depend on the underlying encryption scheme, appear more appealing.

The first such black-box construction was proposed in [PPY22]. We briefly recall it here. Given any pseudo-random function F and *any* PKE chosen by the dictator, anamorphic mode is set as follows. Public and secret keys ($\mathsf{apk}, \mathsf{ask}$) are

[4] [KPP$^+$23b] refers to this phenomenon as the prevalence of anamorphic cryptography.

generated from the PKE regular key generation procedure, while the double key dk is a random seed k for F. To encrypt a regular message m and a covert *bit* $\widehat{m}$ Alice produces an encryption c of m such that $F(k, c) = \widehat{m}$ through rejection sampling. Bob then retrieves m by decrypting c using sk and $\widehat{m}$ as $F(k, c)$.

Although elegant and generic, this construction does not support *large* anamorphic messages. Indeed it is possible to convey at most $O(\log \lambda)$ bit long anamorphic messages per ciphertext while keeping the sender polynomial time[5]. The same limitation affects the generic construction in [BGH+24]. Although the issue might be mitigated through other means, such as sending multiple ciphertexts [WCHY23], these may dangerously increase the risk of detection (e.g. through traffic analysis). Moreover, it is unclear why users not interested in sending covert messages should adopt such a risky behavior.

This state of affairs thus leads to the (quite natural) question of whether such limitation is inherent. More precisely, we ask:

Question: *Is it possible to build a compiler that, given any (standard) PKE, turns it into an anamorphic scheme with large anamorphic message space, i.e. of size $\Omega(2^\lambda)$, with security parameter λ?*

We argue this to be a fundamental question as a positive answer would provide a viable solution regardless of the imposed encryption standard. A negative one instead would imply that crafting *ad hoc* AE for currently used PKE (which may however be eventually banned) is the only way to ensure efficient secure communication in this setting. Furthermore, such a negative answer would leave open the possibility of practically adoptable PKE not admitting anamorphic counterparts with large message space[6]. This would significantly affect the main *raison d'être* of anamorphic encryption as cryptographic primitive: a (powerful enough) dictator can simply decide to allow the usage of that single scheme and declare illegal all other ones.

1.1 Our Contributions

In this paper we (partially) answer the question above in the negative. We do so by focusing only on generic constructions with black-box access to the underlying PKE chosen by the dictator. More precisely we prove that:

- Any black-box construction of AE, i.e. that only accesses the underlying PKE through oracle key-generation, encryption and decryption calls, *must* have anamorphic message space of size at most $|\widehat{M}| = \mathsf{poly}(\lambda)$ with λ the security parameter.

[5] Finding through rejection sampling c such that $F(k, c) = \widehat{m}$ requires on expectation $O(2^{|\widehat{m}|})$ iterations, with $|\widehat{m}|$ the bit-length of $\widehat{m}$, thus $|\widehat{m}| = O(\log \lambda)$.

[6] Here by practically adoptable we mean an efficient PKE leading to an AE with large (anamorphic) message space, without resorting to the multiple ciphertexts strategy mentioned above.

– Black-box constructions of (*weak*) asymmetric AE, an enhanced notion recently introduced in [CGM24] and discussed below, are impossible, regardless of the anamorphic message space size.

The notion of asymmetric AE referred to in our second result can be informally thought as a public-key variant of AE-mode. In this setting the receiver generates two extra keys dk, tk (as opposed to one). dk is shared with the sender(s) to encrypt anamorphic messages, while decryption is performed with tk. Slightly more precisely, dk acts as an asymmetric key that allows to encrypt anamorphic messages bot not to decrypt anamorphic ciphertexts. *Weak* asymmetric security then requires that *anamorphic messages* remains indistinguishable given the regular public key apk and the double key dk. The technical challenges of these two results are exposed in details in the next section. First, however, we clarify how they should be interpreted, and possibly circumvented.

Implications. A direct consequence of our results is that the "rejection sampling" construction [PPY22] discussed before is optimal from different perspectives. The first one is indeed that communicating $O(\log \lambda)$ bits per ciphertext is the best a black-box construction can achieve. Moreover, the usage of dk as a symmetric key to encrypt and decrypt anamorphic messages cannot be avoided. Note also that both results readily extends to stronger primitives such as black-box *robust* AE [BGH$^+$24] and *anamorphic extensions* [BGH$^+$24] in the respective setting. Similarly, as our second result holds for *weak* asymmetric AE, it further extends to all stronger notions (e.g. fully asymmetric AE [CGM24]).

Limitations. We stress that our results concern generic AE constructions interacting with the PKE only through oracle call (we call these black-box constructions). This clearly excludes all constructions that, while generic, are non black-box. Thus, both our negative results might be bypassed via some explicit usage of a circuit[7] evaluating the underlying PKE. This could be achieved, for instance, by relying on garbled circuits [Yao86] or iO [BGI$^+$01]. We remark, however, that such techniques are unlikely to yield practically efficient solutions.

Towards practical solutions, a more promising approach consists in assuming that the underlying PKE satisfies some extra properties. This is for instance the case for constructions in [BGH$^+$24, KPP$^+$23b, CGM24].

In this sense, inspired by our proof techniques, we provide a novel "semi-generic" construction that is asymmetric (in the sense discussed above) and supports *exponential* anamorphic message space. More precisely, we require the underlying PKE to have polynomially small message space and to be *dense* (i.e. any random string in the ciphertext space is a valid encryption with significant probability).

Since our results are quite technical, in what follows we provide some intuition about the main ideas they build upon.

[7] As in the case of Identity Based Encryption [Sha84, BF01] which admits no black-box construction from pairing-free groups [PRV12, Zha22] and yet can be *generically* realized from them through garbled circuits [DG17].

1.2 Technical Overview

As a starting point assume $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ is a generic construction that turns *any* PKE into an anamorphic encryption scheme. We restrict to AEs that access the underlying PKE algorithms only through oracle queries. As customary in the black-box separations literature [IR89], we then study their behavior when interacting with an *ideal* PKE $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$. The latter scheme is similar to the one proposed by Gertner *et al.* in [GKM$^+$00] and by Zhandry and Zhang in [ZZ20], and is based on two truly random permutations specifying the key generation and encryption behavior. Decryption instead consists in (inefficiently) inverting the encryption permutation.

Ciphertext Selection Lemma. Our first step towards both results is to prove a fundamental property of the anamorphic encryption procedure $\mathsf{AT.Enc}$. Namely that, up to negligible probability, it can only return one of the ciphertexts it obtains from oracle calls to $\mathsf{E.Enc}$. First notice that, being $\mathsf{AT.Enc}$ anamorphic, its produced ciphertexts have to be indistinguishable from regular-mode ones. As security is assumed to hold for any PKE, this has to be the case also for the ideal PKE mentioned above. In this latter case, however, there is essentially no way to meaningfully manipulate ciphertexts. Thus, the only way for $\mathsf{AT.Enc}$ to return a valid ciphertext (i.e. encrypting the intended regular message m), has to be to simply choose it among the obtained ones[8].

Limits of black-box anamorphic conversion. To prove our first lower bound we start from an information-theoretic game where a sender $\mathcal{S}$ wishes to communicate a message m to a receiver $\mathcal{R}$. The rules of this game are that a random oracle H is available to both, and all $\mathcal{S}$ can do is choose one of the outputs y it received from H and send it to $\mathcal{R}$. The goal of $\mathcal{R}$ is to get back m from y with overwhelming probability. Finally $\mathcal{S}$ and $\mathcal{R}$ are allowed to have shared randomness. We call this setting a *Random Oracle Channel*. Assuming both procedures can access H only $\mathsf{poly}(\lambda)$ many times, we prove that the message space size $|M|$ has to be polynomially bounded. Intuitively, this should be the case as $\mathcal{S}$'s choice can bias at most $\log(\lambda)$ many bits of y, while $\mathcal{R}$'s queries seem useful only when it finds a preimage to y.

Our final step consists in building a Random Oracle Channel from a black-box AE scheme. The basic idea is that $\mathcal{S}^{\mathsf{H}}$ computes and sends the anamorphic encryption of a message $\widehat{m}$, while $\mathcal{R}^{\mathsf{H}}$ decrypts it. Crucially, both parties use H to answer encryption queries performed respectively by $\mathsf{AT.Enc}$ and $\mathsf{AT.Dec}$[9], which results in a good approximation of the ideal PKE scheme. Next, we use the ciphertext selection lemma to argue that $\mathsf{AT.Enc}$ can only "choose" one of the ciphertexts it observed. Thus, the anamorphic ciphertext that $\mathcal{S}$ forwards to

[8] Actually, another possible way is by decrypting a (random) ciphertext with $\mathsf{E.Dec}$ hoping it returns m. For carefully chosen ideal PKE's parameters however this strategy only succeeds with negligible probability.

[9] In this technical overview we deliberately ignore the significant technical challenges related to dealing with decryption queries. See Section 5 for details

$\mathcal{R}$ is a value it got from H, which, in turn, means that $(\mathcal{S}, \mathcal{R})$ defines a random oracle channel. As a consequence, its associated (anamorphic) message space has to be polynomially bounded.

Impossibility of Asymmetric AE. Another application of the ciphertext selection lemma is that (weak) asymmetric AE as discussed above cannot be realized black-box. An intuitive reason for this is that $\mathsf{AT.Enc}$, in order to *correctly* choose a ciphertext encrypting the anamorphic message $\widehat{m}$ it wish to send, must somehow distinguish those that encrypt $\widehat{m}$ from those that do not.

This suggests the following proof strategy. An (efficient) adversary $\mathcal{A}$ refuting the weak asymmetric property can initially query its challenger to get c^*, either the anamorphic encryption of $(m, \widehat{m}_0)$ or $(m, \widehat{m}_1)$, where m here denotes the regular message, whereas $\widehat{m}_0, \widehat{m}_1$ are anamorphic ones. Then it locally runs $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}_0)$ with both apk and dk being provided to $\mathcal{A}$ at the beginning. When $\mathsf{AT.Enc}$ calls the underling PKE encryption procedure, $\mathcal{A}$ replies with the correct ciphertext for all but a randomly chosen query. For this latter query it replies with c^*. Finally, when $\mathsf{AT.Enc}$ returns c', $\mathcal{A}$ outputs 1 if $c^* = c'$ and 0 otherwise.

Oversimplifying the analysis, if c^* is an encryption of $\widehat{m}_0$, then $\mathsf{AT.Enc}$ should choose it with significant probability ($\approx 1/q$ with q the total number of encryption queries). If c^* encrypts $\widehat{m}_1$, on the other hand, correctness of encryption dictates that $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}_0)$ can output it only with negligible probability.

This simple strategy fails for a variety of technical reasons, some of which are not discussed here. The most challenging one though, is that c^* may be incorrectly distributed. More specifically, during its execution $\mathsf{AT.Enc}$ expects *regular* ciphertexts as answers to its encryption calls. Yet, c^* is an *anamorphic* one. Although the security definition from [PPY22] guarantees that regular ciphertexts are indistinguishable from anamorphic ones, this only holds when given $\mathsf{apk}, \mathsf{ask}$ but *not* dk. As $\mathsf{AT.Enc}$ gets dk it may easily distinguish c^* and potentially abort, thus preventing our proof to go through.

To address this issue we analyze in depth an abstract object, that we call *symmetric choice functions*. Such an object is meant to describe $\mathsf{AT.Enc}$'s behavior but, we believe, could be of independent interest.

Informally, a choice function is any (probabilistic) function that outputs one of its arguments, without modifying it in any way. If it does not depend on the order of its input, we further call it *symmetric*.

We prove that symmetric choice functions have the very interesting property of being *consistent* in their choices. Specifically, imagine that on (uniformly distributed) input $x_1, \ldots, x_k$ the choice function f outputs one of them (and let us call z such a value). Interestingly, on input $(z, u_2, \ldots, u_k)$, for uniformly distributed $u_2, \ldots, u_k$, f will output back z with probability at least $1/k - \varepsilon$. This may seem obvious at first, as the inputs look uniformly distributed in both cases. Notice however, that while z is chosen from uniformly distributed inputs,

its distribution is (or at least might be) biased by f and, thus, it might not be uniform anymore[10].

The previously unjustified step in our (simplified) analysis is then fixed by showing that $\mathsf{AT.Enc}$ essentially behaves like a symmetric choice function. Thus, when receiving from $\mathcal{A}$ a c^* that (anamorphically) encrypts $(m, \widehat{m}_0)$, along with $q-1$ almost uniformly random ciphertexts[11], it will choose the same c^* again with probability at least $1/q - \varepsilon$.

A semi-generic Realization. Both our impossibility results crucially rely on the ciphertext selection lemma discussed above. Interesting this lemma requires the ideal PKE to satisfy certain conditions. In particular, its proof does not go through for the special case of PKE with *small* message space and *dense* ciphertext space[12]. We show that this is no coincidence and, in fact, we prove such restrictions to be sufficient to achieve efficient asymmetric anamorphic conversions with large (anamorphic) message spaces. Specifically, we prove that if one starts with a PKE with the two properties above, that also guarantees a mild pseudorandom property on the produced ciphertexts (see section 6 for details about this), then there exists a simple black-box asymmetric AE with exponential anamorphic message space. This construction can be seen as the dual of the rejection sampling scheme from [PPY22] when swapping the role of regular and anamorphic messages. At a (very) high level, one starts with a PKE $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ satisfying the conditions above, together with a PKE $\Pi^{\mathsf{pr}} = (\Pi^{\mathsf{pr}}.\mathsf{Gen}, \Pi^{\mathsf{pr}}.\mathsf{Enc}, \Pi^{\mathsf{pr}}.\mathsf{Dec})$, whose ciphertexts are indistinguishable from (uniformly) distributed ones in the ciphertext space of Π. We remark that, as we show in Appendix A.1, it is easy to construct such a Π^{pr} from standard PKE and pseudorandom permutations. Equipped with Π and Π^{pr}, the construction is as follows. To encrypt a (regular) message m and a covert one $\widehat{m}$, one keeps encrypting $\widehat{m}$ with $\Pi^{\mathsf{pr}}.\mathsf{Enc}$ until is found a ciphertext c such that decrypting c with $\mathsf{E.Dec}$ outputs m. Notice that, since Π has dense ciphertexts and small message space, this procedure is expected to end in polynomial time.

1.3 Related Works

Anamorphic Encryption is a notion similar to other ones studied in the past by cryptographers, such as key-escrow (e.g. [Mic93, Bla94, FY95]), deniable encryption (e.g. [CDNO97]), kleptography (e.g. [YY96, YY97]) and public key steganography (e.g. [vH04]), but at the same time it differs from all of these. We refer to the work of Persiano *et al.* [PPY22] for a comparison of these notions.

In [KPP$^+$23b, CGM24] the notion of receiver AE has been refined requiring privacy for the normal and covert messages even against the holders of dk. In

[10] Think, for instance, to the case when f is the minimum function: this satisfies our notion of symmetric choice function but, even when executed on inputs uniformly distributed in some finite set X, its output is hardly uniform in X.

[11] As is the case when the underlying PKE is the ideal one.

[12] By dense, we mean that a significant fraction of the strings in the ciphertexts space are actually valid ciphertexts.

[BGH+24] the notion of *robust* AE and Anamorphic Extension have been introduced. Later on in [WCHY23] the notion of robustness has been extended and adapted also to sender AE. In [KPP+23a] the notion of Anamorphic Signatures has been introduced in order to face a more extreme scenario in which all communications must pass through a central authority controlled by the dictator. In this context, the usage of encryption channels becomes even more problematic and to face this problem they rely on authenticated channels (i.e., using digital signatures) to be able to send covert messages between parties.

The study of black-box separations instead started from the seminal work of Impagliazzo and Rudich [IR89], from which stemmed a long standing and active area of research [Sim98, KST99, GT00, GKM+00, GMR01, GGK03]. All these works however share similar limitations with ours. Indeed they can only rule out black-box constructions using the underlying primitive as an oracle, but not *any* possible construction. Concrete examples are the case of Identity Based Encryption and Non-Interactive Zero-Knowledge proofs from pairing-free prime order groups. In both case negative results are known, [PRV12, Zha22] for the former and [Giu23] for the latter. Yet, generic *non-black-box* constructions were given respectively in [DG17] and [JJ21].

Concurrent Works. A concurrent work by Dodis and Goldin also investigates the limitations of anamorphic encryption. In particular they study a mildly different model where the dictator generates (backdoored) public parameters, and show that in such model relative to a random oracle there exists an *anamorphic resistant scheme*, i.e. a PKE such that any anamorphic triplet for it must have polynomially small message space. They further introduce a new notion they call unforgeability which strengthen robustness. Although related, their results are ultimately incomparable with ours: Our lower-bound for black-box constructions does not imply their result, because they describe an explicit anamorphic resistant PKE. On the other hand, their result does not directly imply ours due to their different model, which grants more power to the dictator and makes the construction of AE harder.

2 Preliminaries

2.1 Notation

$\lambda \in \mathbb{N}$ denotes the security parameter. A function $f : \mathbb{N} \to \mathbb{R}^+$ is *negligible* if it vanishes faster than the inverse of any polynomial. $\mathsf{negl}(\lambda)$ denotes a generic negligible function. In logic propositions, we use commas to denote logical AND. We also use "wildcard notation" in place of existential quantifier, e.g. $(a, \cdot) \in A \times B$ means $\exists b : (a, b) \in A \times B$.

Given a probabilistic Turing Machine $\mathcal{A}$ we denote $y \leftarrow \mathcal{A}(x; r)$ its output on input x and random tape r. The notation $y \leftarrow^{\$} \mathcal{A}(x)$ is short for $y \leftarrow \mathcal{A}(x; r)$ with r being a uniformly sampled tape. With PPT we denote probabilistic polynomial time. Given a set S we denote by $x \leftarrow^{\$} S$ the uniformly random sampling of an element x from the set S. We further write $x \sim U(S)$ to indicate that x is a

uniformly distributed random variable over S. Unless otherwise specified, we assume *adversaries* in security definitions to be *stateful*, and procedures in a given scheme (e.g. a PKE) to be *stateless*.

2.2 Anamorphic Encryption

The notion of (receiver) anamorphic encryption was first introduced in [PPY22] to model private communication in the presence of a dictator who controls the PKE scheme in use and knows each user's secret key. In this paper, we use a more general definition proposed in [CGM24] which contains [PPY22] as a special case. To achieve the above seemingly impossible goal, the receiver is allowed to generate its own public and secret key $\mathsf{apk}, \mathsf{ask}$ in *anamorphic mode*, exchange secretly with the sender a *double key* dk, and locally storing a *trapdoor key* tk to decrypt anamorphic messages from the sender.

Definition 1 (Anamorphic Triplet). *Formally, an anamorphic triplet* $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *is a triplet of efficient algorithms such that*

- $\mathsf{AT.Gen}(\lambda) \xrightarrow{\$} (\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$ *with* $\mathsf{apk}, \mathsf{ask}$ *being the anamorphic public and secret keys while* dk, tk *are the double and (a possibly empty) trapdoor key.*
- $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}) \xrightarrow{\$} c$*, with* $m \in M$ *and* $\widehat{m} \in \widehat{M}$ *being respectively the standard and anamorphic messages encrypted in* c*.*
- $\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c) \to \widehat{m}/\perp$*, with* $\widehat{m}$ *being the anamorphic message encrypted in* c*.*

In the definition above we do not explicitly provide $\mathsf{apk}, \mathsf{dk}$ as part of $\mathsf{AT.Dec}$ input, as we implicitly assume them to be contained in ask and tk respectively.

Definition 2 (Anamorphic Encryption). *A PKE* $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *is an Anamorphic Encryption scheme if it is* IND-CPA *secure and there exists an anamorphic triplet* $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *such that any* PPT *adversary* $\mathcal{A}$ *has negligible advantage, defined as*

$$\mathsf{Adv}^{\mathsf{Anam}}_{\mathcal{A},\Pi,\Sigma}(\lambda) \;:=\; \left|\Pr\left[\mathsf{RealG}_{\Pi}(\lambda, \mathcal{A}) = 1\right] - \Pr\left[\mathsf{AnamorphicG}_{\Sigma}(\lambda, \mathcal{A}) = 1\right]\right|$$

where RealG_{Π} *and* $\mathsf{AnamorphicG}_{\Sigma}$ *are described in Figure 1.*

Finally, regarding correctness we follow the game-based definition provided [BGH+24], provided in the Appendix, Section B.1. For the sake of generality however we will only use a weaker notion, holding only for uniformly sampled messages (and correct keys). Formally, given $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \xleftarrow{\$} \mathsf{AT.Gen}(\lambda)$ and $m, \widehat{m}$ uniformly sampled messages, then

$$\Pr\left[\widetilde{m} \neq \widehat{m} \;\middle|\; \widetilde{m} \leftarrow \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c),\; c \xleftarrow{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})\right] \;\leq\; \mathsf{negl}(\lambda).$$

$\mathsf{RealG}_{\Pi}(\lambda, \mathcal{A})$

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

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

1 : Sample a random r
2 : **return** $\mathsf{E.Enc}(\mathsf{pk}, m; r)$

$\mathsf{AnamorphicG}_{\Sigma}(\lambda, \mathcal{A})$

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

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

1 : Sample a random r
2 : **return** $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}; r)$

Fig. 1. Anamorphic Encryption security game.

2.3 Asymmetric Anamorphic Encryption

The notion of *Asymmetric Anamorphic Encryption* [CGM24], intuitively, requires that the Anamorphic Triplet Σ realizes an asymmetric scheme for covert messages. The notion is formalized through the game in Figure 2, where $\mathcal{D}$ is a PPT adversary, $b \in \{0, 1\}$ and $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ is an Anamorphic Triplet. The advantage of a given distinguisher $\mathcal{D}$ is defined as

$$\mathsf{Adv}^{\mathsf{Asy\text{-}Anam}}_{\mathcal{D},\Sigma}(\lambda) := \left|\Pr\left[\mathsf{AsyAnam\text{-}IND\text{-}CPA}^{0}_{\Sigma}(\lambda, \mathcal{D}) = 1\right] - \Pr\left[\mathsf{AsyAnam\text{-}IND\text{-}CPA}^{1}_{\Sigma}(\lambda, \mathcal{D}) = 1\right]\right|.$$

$\mathsf{AsyAnam\text{-}IND\text{-}CPA}^{b}_{\Sigma}(\lambda, \mathcal{D})$

1 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda)$
2 : $(m, \widehat{m}_0, \widehat{m}_1) \leftarrow^{\$} \mathcal{D}(\mathsf{apk}, \mathsf{ask}, \mathsf{dk})$
3 : $c \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}_b)$
4 : **return** $\mathcal{D}(c)$

Fig. 2. Asymmetric Anamorphic Encryption security game.

Definition 3 (Asymmetric Anamorphic Encryption). *An Anamorphic Encryption scheme* Π *equppied with an anamorphic triplet* Σ *is an Asymmetric Anamorphic Encryption scheme if for every* PPT *distinguisher* $\mathcal{D}$*,*

$$\mathsf{Adv}^{\mathsf{Asy\text{-}Anam}}_{\mathcal{D},\Sigma}(\lambda) \leq \mathsf{negl}(\lambda).$$

In this paper we define a weaker notion, called *Weak Asymmetric Anamorphic Encryption*. We weaken the previous definition requiring that the adversary in the security game has no access to ask. More precisely, let $\mathcal{D}$ be a PPT adversary, $b \in \{0, 1\}$ and $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ be an Anamorphic Triplet. The

Weak Asymmetric AE security game is then detailed in Figure 3. The advantage of a distinguisher $\mathcal{D}$ for such game is defined as

$$\mathsf{Adv}^{\mathsf{Weak\text{-}Asy\text{-}Anam}}_{\mathcal{D},\Sigma}(\lambda) \coloneqq \left|\Pr\left[\mathsf{Weak\text{-}AsyAnam\text{-}IND\text{-}CPA}^{0}_{\Sigma}(\lambda,\mathcal{D})=1\right] - \Pr\left[\mathsf{Weak\text{-}AsyAnam\text{-}IND\text{-}CPA}^{1}_{\Sigma}(\lambda,\mathcal{D})=1\right]\right|.$$

$\underline{\mathsf{Weak\text{-}AsyAnam\text{-}IND\text{-}CPA}^{b}_{\Sigma}(\lambda,\mathcal{D})}$

1 : $(\mathsf{apk},\mathsf{ask},\mathsf{dk},\mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda)$

2 : $(m,\widehat{m}_0,\widehat{m}_1) \leftarrow^{\$} \mathcal{D}(\mathsf{apk},\mathsf{dk})$

3 : $c \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk},\mathsf{dk},m,\widehat{m}_b)$

4 : **return** $\mathcal{D}(c)$

Fig. 3. Weak Asymmetric Anamorphic Encryption security game.

Definition 4 (Weak Asymmetric Anamorphic Encryption). *An Anamorphic Encryption scheme* Π *equipped with an anamorphic triplet* Σ *is a Weak Asymmetric Anamorphic Encryption scheme if for every* PPT *distinguisher* $\mathcal{D}$

$$\mathsf{Adv}^{\mathsf{Weak\text{-}Asy\text{-}Anam}}_{\mathcal{D},\Sigma}(\lambda) \leq \mathsf{negl}(\lambda).$$

2.4 Other flavors of AE

Robustness of Anamorphic Encryption Robustness for receiver anamorphic encryption has been introduced in [BGH+24]. Informally, it requires that it should be difficult to find a message m that, when encrypted normally (i.e., using $\mathsf{E.Enc}$) and then *anamorphically* decrypted (i.e. using $\mathsf{AT.Dec}$) results in some $\widehat{m} \neq \perp$. Later, in [WCHY23], the notion has been extended to sender anamorphic encryption, requiring in addition to the previous property, also that there exists a negligible probability of decrypting $\widehat{m} \neq \perp$ using a different secret key from the one corresponding to the public key used to anamorphically encrypt $\widehat{m}$. A formal definition of robustness for (receiver) Anamorphic Encryption can be found in Appendix B.2. One can verify that Robust AE implies AE.

Fully Asymmetric AE Fully Asymmetric (receiver) AE (fasy-AE, for short) has been introduced in [CGM24], it is a notion reminiscent of *Single-Receiver* AE from [KPP+23b], indeed it takes the latter notion and makes one step further. Informally, a fasy-AE guarantees the privacy of both the regular and the anamorphic messages with respect to users having access *also* to dk (but not to ask and tk of course). In [CGM24] the relation between this notions has been explored. A formal definition of fasy-AE can be found in Appendix B.3. One can verify that fasy-AE implies Weak Asymmetric AE.

3 Anamorphic Encryption from Black-Box PKE

3.1 Ideal PKE

In this section we model an idealized (and inefficient) PKE scheme, inspired by the one presented in [GKM+00, ZZ20], accessible through three oracles $\mathsf{E.Gen}$, $\mathsf{E.Enc}, \mathsf{E.Dec}$. Internally the scheme is defined by two random functions ϕ and ψ tracking respectively the relation between public/secret keys, and the one between messages/ciphertexts. More in detail SK, PK are the secret and public keys sets while $\{0,1\}^\mu$, $\{0,1\}^\rho$, $\{0,1\}^\ell$ are respectively the messages, randomness (for encryption) and ciphertexts spaces. Then ϕ, ψ are sampled so that

- $\phi : \mathsf{SK} \to \mathsf{PK}$ is a uniformly random bijection.
- $\psi : \mathsf{PK} \times \{0,1\}^\mu \times \{0,1\}^\rho \to \{0,1\}^\ell$ random function s.t. $\psi(\mathsf{pk}, \cdot, \cdot)$ is injective.

Note that at this stage we do not constrain μ, ρ, ℓ, that are respectively the bit-size of messages, randomness and ciphertexts. Some later results will however only apply for certain parameters choice.

$\mathsf{E.Gen}(\lambda; \mathsf{sk})$

1 : $\mathsf{pk} \leftarrow \phi(\mathsf{sk})$
2 : **return** $(\mathsf{pk}, \mathsf{sk})$

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

1 : $c \leftarrow \psi(\mathsf{pk}, m, r)$
2 : **return** c

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

1 : $\mathsf{pk} \leftarrow \phi(\mathsf{sk})$
2 : **for** $(m, r) \in \{0,1\}^\mu \times \{0,1\}^\rho$
3 : **if** $\psi(\mathsf{pk}, m, r) = c$: **return** m
4 : **return** $\perp$.

Fig. 4. Ideal PKE with $\phi : \mathsf{SK} \to \mathsf{PK}$ and $\psi : \mathsf{PK} \times \{0,1\}^\mu \times \{0,1\}^\rho \to \{0,1\}^\ell$ as above.

It is easy to observe that this scheme achieves semantic security (IND-CPA) if $\rho = \Omega(\lambda)$ and $|\mathsf{SK}| = \Omega(2^\lambda)$ as ciphertexts are random strings, and distinguishing the encryptions of two different messages requires a number of queries to $\mathsf{E.Enc}$ exponential in ρ. For completeness a proof appears in the Appendix, Section C.1.

3.2 Black-Box Anamorphic Encryption

Definition 5 (Black-Box Anamorphic Triplet). *A triplet* $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *is said to be a black-box anamorphic triplet (for any PKE* Π*) if every algorithm in* Σ *can access the procedures in* Π ***only*** *through oracle access, i.e. providing input and random coins to these procedures and obtaining* ***only*** *the output of such procedures call in return.*

We remark that we may occasionally and informally refer to an Black-Box Anamorphic Triplet as a Black-Box *Anamorphic Encryption.*

3.3 General Properties

Assume there exists a generic compiler $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ turning any IND-CPA secure PKE into an anamorphic encryption scheme, accessing the underlying PKE algorithms only through oracle queries. We can then study the behavior of such construction when applied to the ideal PKE $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ defined in Figure 4. A first property it has to satisfy is that, up to negligible probability, the public and secret anamorphic keys have to be a valid key pair for the underlying PKE.

Lemma 1. *If $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ is an anamorphic triplet for the ideal PKE Π, then there exists a negligible ε such that*

$$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda) \quad \Rightarrow \quad \Pr[\phi(\mathsf{ask}) \neq \mathsf{apk}] \leq \varepsilon(\lambda).$$

Proof. Let $\mathcal{A}$ be a PPT adversary playing the game in Definition 2 which on input pk, sk, runs the key generation algorithm $(\mathsf{pk}', \mathsf{sk}) \leftarrow \mathsf{E.Gen}(\lambda; \mathsf{sk})$ and returns 1 if $\mathsf{pk} = \mathsf{pk}'$ and 0 otherwise. From the definition of $\mathsf{E.Gen}$ in Figure 4, the secret key coincides with the random tape of $\mathsf{E.Gen}$. Thus in the real game $\mathsf{pk}' = \mathsf{pk}$ occurs always. Conversely in the anamorphic game, the adversary receives $\mathsf{apk}, \mathsf{ask}$ generated through $\mathsf{AT.Gen}$. Again by construction $\mathsf{pk}' = \phi(\mathsf{ask})$, meaning $\mathcal{A}$ returns 1 if and only if $\mathsf{apk} = \phi(\mathsf{ask})$. In conclusion

$$\mathsf{Adv}(\mathcal{A}) = |1 - \Pr[\phi(\mathsf{ask}) = \mathsf{apk}]| = \Pr[\phi(\mathsf{ask}) \neq \mathsf{apk}]$$

which is negligible as we assumed Σ to be an anamorphic triplet for the ideal PKE. □

The next property we study informally states that ciphertexts have to be unpredictable enough. While this could be stated in terms of (pseudo) min-entropy, for our purpose the following less general formulation will suffice.

Lemma 2. *Given $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ a black-box anamorphic triplet and uniformly sampled s, r and messages $m, \widehat{m}$, let*

$$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow \mathsf{AT.Gen}(\lambda; s), \qquad c \leftarrow \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}; r).$$

For any set S independent from r, with $|S| \leq \mathsf{poly}(\lambda)$ then $\Pr[c \in S] \leq \mathsf{negl}(\lambda)$.

Proof. Consider the following adversary $\mathcal{A}$ against the anamorphic security game in Definition 2 instantiated when Σ is combined with the ideal PKE with $\rho = \Omega(2^{\lambda})$. Its attack consists in encrypting twice a random message pair, and checking if the resulting ciphertexts are the same, see Figure 5.

If $c \in S$ with significant probability, as this set has polynomially bounded size, two ciphertexts sampled independently from it will collide with noticeable probability, allowing $\mathcal{A}$ to distinguish the two games.

More formally, in the real game $c_1 = c_2$ only if the random coins used to produce both ciphertexts are the same, which occurs with probability $2^{-\rho}$. To analyze the anamorphic game let

$$V_{\delta} = \{(m_0, \widehat{m}_0, s_0) : \Pr[c \in S \mid m = m_0, \ \widehat{m} = \widehat{m}_0, \ s = s_0] \geq \delta\}.$$

Using a variant of Markov inequality we can then prove that

$\mathcal{A}^{\mathcal{O}}(\mathsf{apk}, \mathsf{ask})$:

1 : Sample $m \leftarrow^{\$} \{0,1\}^{\mu}$ and $\widehat{m} \leftarrow^{\$} \widehat{M}$
2 : $c_1 \leftarrow \mathcal{O}(m, \widehat{m})$
3 : $c_2 \leftarrow \mathcal{O}(m, \widehat{m})$
4 : **return** $c_1 == c_2$

Fig. 5. Adversary against the security game in Figure 1. $\mathcal{O}$ is the encryption oracle provided in both RealG and AnamorphicG.

Claim 1. $\delta = 1/2 \cdot \Pr[c \in S]$ *implies that* $\Pr[(m, \widehat{m}, s) \in V_\delta] \geq \delta$.

A proof appear in the Appendix, Section C.2. Calling for notational simplicity $\mathbf{v} = (m, \widehat{m}, s)$, it can now be shown that for all $\mathbf{v}_0 \in V_\delta$, $\Pr[c_1 = c_2 \mid \mathbf{v} = \mathbf{v}_0] =$

$$\begin{aligned}
&= \Pr[c_1 = c_2 \mid c_1, c_2 \in S,\ \mathbf{v} = \mathbf{v}_0] \cdot \Pr[c_1 \in S,\ c_2 \in S \mid \mathbf{v} = \mathbf{v}_0] \\
&\geq |S|^{-1} \cdot \Pr[c_1 \in S,\ c_2 \in S \mid \mathbf{v} = \mathbf{v}_0] \\
&= |S|^{-1} \cdot \Pr[c_1 \in S \mid \mathbf{v} = \mathbf{v}_0] \cdot \Pr[c_2 \in S \mid \mathbf{v} = \mathbf{v}_0] \\
&\geq |S|^{-1} \cdot \delta^2
\end{aligned}$$

where the second equality follows as c_1, c_2 are mutually independent conditioned on $\mathbf{v} = \mathbf{v}_0$, as in that case they are only a function of the (independently sampled) random coins used to compute them, and the random subset S is distributed independently from them. As a consequence $\Pr[c_1 = c_2 \mid \mathbf{v} \in V_\delta] \geq |S|^{-1} \cdot \delta^2$, which allow us to lower bound the probability $\mathcal{A}$ finds a collision in the anamorphic game as, fixing $\delta = 1/2 \cdot \Pr[c \in S]$,

$$\Pr[c_1 = c_2] \ \geq\ \Pr[c_1 = c_2 \mid \mathbf{v} \in V_\delta] \cdot \Pr[\mathbf{v} \in V_\delta] \ \geq\ \delta^3 \cdot |S|^{-1}.$$

Combining this with the bound on the collision probability in the real game, the advantage of $\mathcal{A}$ is then bounded by $\mathsf{Adv}(\mathcal{A}) \geq \delta^3 \cdot |S|^{-1} - 2^{-\rho}$. Having set $\delta = 1/2 \cdot \Pr[c \in S]$ we conclude the proof as we assumed $\rho = \Omega(\lambda)$, $|S|$ polynomially bounded and the black-box anamorphic triplet to be secure. □

A consequence of the above result is that AT.Enc almost never returns a ciphertext that was observed by AT.Gen. To formally state this, we first define this set of ciphertexts.

Definition 6. *Given a black-box anamorphic triplet* Σ *we define* $E^{\mathsf{Gen}}_{\mathsf{in}}$, $E^{\mathsf{Enc}}_{\mathsf{in}}$ *the sets of tuples* (pk, m, r, c) *such that respectively* AT.Gen *and* AT.Enc *on input* in *eventually query* $c = \mathsf{E.Enc}(\mathsf{pk}, m; r)$. *Analogously,* $D^{\mathsf{Gen}}_{\mathsf{in}}$, $D^{\mathsf{Enc}}_{\mathsf{in}}$ *are the sets of tuples* (sk, c, m) *such that respectively* AT.Gen *and* AT.Enc *on input* in *computes* $m = \mathsf{E.Dec}(\mathsf{sk}, c)$.

Definition 7. *Given a black-box anamorphic triplet* Σ *we define the set of ciphertexts observed by* AT.Gen *on input* s *as*

$$C^{\mathsf{Gen}}_s \ :=\ \{c\ :\ (\cdot,\cdot,\cdot,c) \in E^{\mathsf{Gen}}_s \ \ \vee\ \ (\cdot,c,\cdot) \in D^{\mathsf{Gen}}_s\}.$$

Corollary 1. *With the same notation of Lemma 2,* $\Pr\left[c \in C_s^{\mathsf{Gen}}\right] \leq \mathsf{negl}(\lambda)$.

3.4 Ciphertext Selection Lemma

The core technical result of this section is a characterization of the encryption procedure for a black-box anamorphic triplet. Informally, our result states that such procedure can only obtain *valid* ciphertexts through encryption queries to $\mathsf{E.Enc}$ and then return one of them. This is perhaps not surprising as there is no assumption on the underlying PKE scheme. Thus, no meaningful manipulation of ciphertexts after their generation is possible. This intuition is captured by the following *ciphertext selection lemma*. First, we formally define the set of valid ciphertexts queried by $\mathsf{AT.Enc}$.

Definition 8. *Given input* $\mathsf{in} = (\mathsf{apk}, \mathsf{ask}, m, \widehat{m}, r)$ *the set of valid ciphertexts queried by* $\mathsf{AT.Enc}$ *is* $C_{\mathsf{in}}^{\mathsf{Enc}} = \{c : (\mathsf{apk}, m, \cdot, c) \in E_{\mathsf{in}}^{\mathsf{Enc}}\}$.

We recall that our ideal PKE is parametrized by μ, ρ, ℓ, respectively the message, random coins and ciphertext bit-length. Notably, the following result requires $\ell - \rho = \Omega(\lambda)$ to hold. This means the lemma cannot be specialized to black-box anamorphic schemes where the underlying PKE is assumed to have *small* message space $\mu = O(\log \lambda)$ and *dense* ciphertext space $\ell = \rho + \mu + O(\log \lambda)$, i.e. such that a noticeable fraction of strings with length ℓ are valid ciphertexts. We will later prove in Section 6 this to be no coincidence as in this case efficient "semi-generic" constructions do exist.

Lemma 3. *Given* $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *a black-box anamorphic triplet, let* r, s *be uniform random coins and* $m, \widehat{m}$ *uniformly sampled messages. Setting*

$$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow \mathsf{AT.Gen}(\lambda; s), \quad \mathsf{in} = (\mathsf{apk}, \mathsf{dk}, m, \widehat{m}, r), \quad c \leftarrow \mathsf{AT.Enc}(\mathsf{in}),$$

if $\rho = \Omega(\lambda)$ *and* $\ell - \rho = \Omega(\lambda)$, *then* $\Pr\left[c \notin C_{\mathsf{in}}^{\mathsf{Enc}}\right] \leq \mathsf{negl}(\lambda)$.

Proof. To prove the lemma let $\mathcal{A}$ be an adversary against the anamorphic security definition as described in Figure 6. Given $(\mathsf{apk}, \mathsf{ask})$ it requests the encryption c of a random message m and locally decrypts it computing $m' = \mathsf{E.Dec}(\mathsf{ask}, c)$. It returns 1 if and only if $m \neq m'$.

$\mathcal{A}^{\mathcal{O}}(\mathsf{apk}, \mathsf{ask})$:

1 : Sample $m \leftarrow^{\$} \{0,1\}^{\mu}$ and $\widehat{m} \leftarrow^{\$} \widehat{M}$
2 : $c \leftarrow \mathcal{O}(m, \widehat{m})$
3 : $m' \leftarrow \mathsf{E.Dec}(\mathsf{ask}, c)$
4 : **return** 1 if $m \neq m'$

Fig. 6. Adversary for the anamorphism game (Fig. 1). $\mathcal{O}$ is the encryption oracle.

Since the ideal PKE scheme achieves perfect correctness $\mathcal{A}$ never returns 1 when executed in the real game. To study the anamorphic game, let s be the random tape of $\mathsf{AT.Gen}$, so that $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow \mathsf{AT.Gen}(\lambda; s)$, and r the one of $\mathsf{AT.Enc}$ when executed to answer $\mathcal{A}$'s only query. For notational convenience $\mathsf{in} = (\mathsf{apk}, \mathsf{dk}, m, \widehat{m}, r)$ so that $c = \mathsf{AT.Enc}(\mathsf{in})$. We then define the two events

$$\mathsf{Bad} \ : \ \phi(\mathsf{ask}) \neq \mathsf{apk} \ \vee \ c \in C_s^{\mathsf{Gen}} \qquad \mathsf{Good} \ : \ c \in C_{\mathsf{in}}^{\mathsf{Enc}}.$$

Lemma 1 and Corollary 1 together imply that $\Pr[\mathsf{Bad}] \leq \mathsf{negl}(\lambda)$. Next we claim that the following probability is also negligible.

Claim 2. $\Pr[m = m', \neg\mathsf{Bad}, \neg\mathsf{Good}] \leq \mathsf{negl}(\lambda)$.

These two inequalities immediately imply the thesis as, through a union bound

$$\begin{aligned} \Pr[m = m'] \ &\leq \ \Pr[m = m', \neg\mathsf{Bad}, \neg\mathsf{Good}] + \Pr[\mathsf{Bad}] + \Pr[\mathsf{Good}] \\ &\leq \ \Pr[\mathsf{Good}] + \mathsf{negl}(\lambda). \end{aligned}$$

By our initial observation $\mathsf{Adv}^{\mathsf{anam}}(\mathcal{A}) = \Pr[m \neq m']$ with m' distributed as in the anamorphic game. As a consequence $\Pr[\neg\mathsf{Good}] \ \leq \ \mathsf{Adv}^{\mathsf{anam}}(\mathcal{A}) + \mathsf{negl}(\lambda)$, that is negligible.

Proof of Claim 2. Let $C = C_s^{\mathsf{Gen}} \cup C_{\mathsf{in}}^{\mathsf{Enc}}$. We denote V_m the set of ciphertexts encrypting m under apk, that is $V_m = \{\psi(\mathsf{apk}, m, r) \ : \ r \in \{0,1\}^\rho\}$. The claim can then be translated in terms of C and V_m. Indeed, if the studied event occurs then $c \notin C$. Similarly $m = m'$ and $\neg\mathsf{Bad}$ both implies that $m = \mathsf{E.Dec}(\mathsf{ask}, c) \Rightarrow \psi(\phi(\mathsf{ask}), m, r) = c \ \Rightarrow \ \psi(\mathsf{apk}, m, r) = c$ for some r, which means $c \in V_m$. Therefore

$$\begin{aligned} &(m = m', \neg\mathsf{Bad}, \neg\mathsf{Good}) \ \Rightarrow \ c \in V_m \setminus C \quad \Rightarrow \\ \Rightarrow \quad &\Pr[m = m', \neg\mathsf{Bad}, \neg\mathsf{Good}] \ \leq \ \Pr[c \in V_m \setminus C]. \end{aligned}$$

To prove the latter probability to be negligible, let q be a bound on the total queries of $\mathsf{AT.Gen}$ and $\mathsf{AT.Enc}$. Let $c_1, \ldots, c_d$ be the (ordered) ciphertexts $\mathsf{AT.Enc}$ queries to $\mathsf{E.Dec}(\mathsf{ask}, \cdot)$ and for notational convenience we name $c_{d+1} \coloneqq c$. Let C_i be the set of ciphertext either returned by $\mathsf{E.Enc}(\mathsf{apk}, \cdot, \cdot)$ or queried to $\mathsf{E.Dec}(\mathsf{ask}, \cdot)$ by either $\mathsf{AT.Gen}(\lambda; s)$ or $\mathsf{AT.Enc}(\mathsf{in})$ *before* the latter queries $\mathsf{E.Dec}(\mathsf{apk}, c_i)$. Note this means $C_i \subseteq C \cup \{c_1, \ldots, c_{i-1}\}$. Crucially, given only this information, the set of ciphertexts $V_m \setminus C_i$ is uniformly distributed over $\{0,1\}^\ell \setminus C_i$. Once again the event above can be decomposed through a chain of implications:

$$\begin{aligned} c \in V_m \setminus C \ &\Rightarrow \ \bigvee_{i=1}^{d+1} (c_i \in V_m \setminus C \ \wedge \ \{c_1, \ldots, c_{i-1}\} \cap V_m \setminus C = \varnothing) \\ &\Rightarrow \ \bigvee_{i=1}^{d+1} (c_i \in V_m \setminus (C \cup \{c_1, \ldots, c_{i-1}\})) \\ &\Rightarrow \ \bigvee_{i=1}^{d+1} (c_i \in V_m \setminus C_i). \end{aligned}$$

Using a union bound, along with the fact that $V_m \setminus C_i$ is a uniformly distributed subset of $\{0,1\}^\ell \setminus C_i$ and independent from c_i, we can conclude that

$$\begin{aligned}\Pr\left[c \in V_m \setminus C\right] &\leq \sum_{i=1}^{d+1} \Pr\left[c_i \in V_m \setminus C_i\right] \\ &\leq \sum_{i=1}^{d+1} \frac{|V_m \setminus C_i|}{|\{0,1\}^\ell \setminus C_i|} \leq (d+1) \cdot \frac{2^\rho}{2^\ell - q}\end{aligned}$$

with the last quantity being negligible as we assumed $\ell - \rho = \Omega(\lambda)$ while d, q are polynomially bounded. □

□

Remark 1. Lemma 3 holds only for *stateless* anamorphic triplets. If stateful encryption/decryption is allowed, then we can only prove a slightly weaker result. Specifically c has to lie, with overwhelming probability, in the set of valid ciphertexts observed by $\mathsf{AT.Enc}$ *and* $\mathsf{AT.Gen}$ (as opposed to only $\mathsf{AT.Enc}$). We stress this to be sufficient for a slightly weaker version of Theorem 2 (See Remark 2) to hold true. The proof is analogous up to the fact that Corollary 1 cannot be applied anymore.

3.5 Symmetric Choice Functions

Thanks to the Ciphertext Selection Lemma, the encryption procedure of any black-box anamorphic triplet can be abstracted as a process observing a list of ciphertexts and eventually choosing one of them. We will call such a function returning one of its arguments a *choice function*. In this section we show this class of functions satisfies interesting properties, which will be useful in the proof of Theorem 3, Section 5.2. First we provide a formal definition of choice functions and in particular *symmetric* ones, which do not depend on the order of their arguments.

Definition 9. *Given a finite set X, a random function $f \sim \{g : X^k \to X\}$ is a choice function if $f(x_1, \ldots, x_k) \in \{x_1, \ldots, x_k\}$ for all $x_1, \ldots, x_k \in X$. Furthermore, a choice function is called symmetric if for any permutation π we have $f(x_1, \ldots, x_k) = f(x_{\pi(1)}, \ldots, x_{\pi(k)})$.*

A rather non-trivial property of symmetric choice functions is that they are *consistent* with their choices. More specifically, assume that on random inputs $u_1, \ldots, u_k$ the function $f(u_1, \ldots, u_k)$ chose z among them. Then given more random inputs $v_2, \ldots, v_k$, the function $f(z, v_2, \ldots, v_k)$ will chose z again with probability at least $\approx 1/k$. At first sight this might seem trivial, as z could appear to be random and f unable to distinguish it from the other elements. However this reasoning is incorrect. Indeed, although z is chosen from uniformly sampled variables, this choice can bias its distribution. The above intuition is therefore wrong, but we nevertheless prove this lower bound with the following Lemma. A full proof appears in the Appendix, Section C.3.

Lemma 4. *Let $f \sim \{g : X^k \to X\}$ be a symmetric choice function. Given $\mathbf{u} \sim U(X^k)$, $\mathbf{v} \sim U(X^{k-1})$ uniformly distributed, let $z = f(\mathbf{u})$. Then*

$$\Pr\left[f(z, \mathbf{v}) = z\right] \;\geq\; \frac{1}{k} - O\left(\frac{1}{|X|}\right).$$

4 Random Oracle Channels

4.1 Definition

In order to provide lower bounds for black-box Anamorphic Encryption, we first study a simpler scenario where a *sender* $\mathcal{S}$ has to communicate a message $m \in M$ to a *receiver* $\mathcal{R}$ under some constraints. In particular, both parties have access to a random oracle H and $\mathcal{S}$, which obtains values $y_1, \ldots, y_k$ during its interaction with H, can only chose one of them and send it to $\mathcal{R}$, who eventually has to recover the original message. We will call this setting a *Random Oracle Channel*.

Definition 10. *A RO-channel is a tuple $(\mathcal{S}, \mathcal{R}, M, k, h)$ with $\mathcal{S}, \mathcal{R}$ Probabilistic Turing Machines (not necessarily PPT), $M \subseteq \{0,1\}^*$ and $k, h = \mathsf{poly}(\lambda)$ such that*

1. *$\mathcal{S}, \mathcal{R}$ make respectively at most k and h queries to H.*
2. *$\forall m \in M$, calling $y_j = \mathsf{H}(x_j)$ with $j \in \{1, \ldots, k\}$ the queries $\mathcal{S}^{\mathsf{H}}(m)$ performs, then $\mathcal{S}^{\mathsf{H}}(m) \to y_i$ for some $i \in \{1, \ldots, k\}$.*
3. *There exists a negligible $\varepsilon(\lambda)$ such that $\forall m \in M$ and uniformly sampled common random tape s*

$$\Pr\left[m \neq m' \,\middle|\, y \leftarrow \mathcal{S}^{\mathsf{H}}(m; s),\ m' \leftarrow \mathcal{R}^{\mathsf{H}}(y; s)\right] \leq \varepsilon(\lambda).$$

The main problem about RO channels is determining how large can $|M|$ be as a function of k, h. Intuitively, due to the high limitations imposed on $\mathcal{S}, \mathcal{R}$, we expect $|M|$ to be small, and indeed our results eventually implies that $|M| = \mathsf{poly}(\lambda)$ or that, equivalently, in this setting it is possible to communicate at most $O(\log \lambda)$ bits.

4.2 Bound for RO-Channel

Theorem 1. *For any RO-Channel $(\mathcal{S}, \mathcal{R}, M, k, h)$ we have that asymptotically $|M| \leq 2(h+k)^2$. In particular $|M| = \mathsf{poly}(\lambda)$.*

Proof. The result is proven by showing that any RO-channel can be compiled into two unbounded $\mathcal{S}^*$, $\mathcal{R}^*$ with shared randomness that reliably communicate a message $m \in M$ by only sending $\ell = O(\log \lambda)$ bits. More specifically the shared randomness is of the form (F, G, s) with $F : \{0,1\}^{\mathsf{poly}(\lambda)} \to \{0,1\}^{\ell}$ and $G : \{0,1\}^{\ell} \to \{0,1\}^{\lambda}$ random functions, and s the random tape used by $\mathcal{S}, \mathcal{R}$.

$\mathcal{S}^*$ on input m executes $\mathcal{S}(m; s)$ and simulates the RO through the function $G \circ F$. More formally, when $\mathcal{S}$ queries the RO on input x_i, it returns $y_i = G(F(x_i))$

$\mathcal{S}^*(m; F, G, s)$:	$\mathcal{R}^*(z; F, G, s)$:
1 : Run $\mathcal{S}(m; s)$	1 : Run $\mathcal{R}(G(z); s)$
2 : **when** $\mathcal{S}$ queries x_i:	2 : **when** $\mathcal{R}$ queries x_i:
3 : $z_i \leftarrow F(x_i),\ y_i \leftarrow G(z_i)$	3 : $y_i \leftarrow G(F(x_i))$
4 : **reply** with $\mathcal{S} \leftarrow y_i$	4 : **reply** with $\mathcal{R} \leftarrow y_i$
5 : **when** $\mathcal{S}$ returns y_i:	5 : **when** $\mathcal{R}$ returns m:
6 : **return** z_i	6 : **return** m

Fig. 7. Unbounded $\mathcal{S}^*, \mathcal{R}^*$ using $(\mathcal{S}, \mathcal{R})$ to communicate m by only sending ℓ bits.

and locally stores $z_i = F(x_i)$. Finally, once $\mathcal{S}$ chooses its output y_i, $\mathcal{S}^*$ returns $z_i \in \{0,1\}^\ell$. In order to recover m, $\mathcal{R}^*$ internally executes $\mathcal{R}$ simulating the RO as before. A full description of $\mathcal{S}^*, \mathcal{R}^*$ is provided in Figure 7.

Let δ be the probability that $\mathcal{S}^*$ and $\mathcal{R}^*$ fail to communicate correctly, i.e.

$$\delta \ \coloneqq\ \Pr\left[m \neq m' \,|\, m' \leftarrow \mathcal{R}^*(z; F, G, s),\ z \leftarrow \mathcal{S}^*(m; F, G, s)\right].$$

Then, the success probability $1 - \delta$ is bounded by the conditional min-entropy of m given z. This implies that

$$\begin{aligned} \mathrm{H}_\infty(m \,|\, z) \ \geq\ \mathrm{H}_\infty(m) - \ell \ =\ \log_2 |M| - \ell \quad &\Rightarrow \quad (1-\delta) \ \leq\ 2^{-\mathrm{H}_\infty(m \,|\, z)} \ =\ \frac{2^\ell}{|M|} \\ &\Rightarrow \quad |M| \leq \frac{2^\ell}{1-\delta}. \end{aligned}$$

Where the first inequality follows from the fact that $z \in \{0,1\}^\ell$ [DRS04, Lemma 2.2]. Next we study the success probability for the specific case of $\mathcal{S}^*, \mathcal{R}^*$ and a suitable choice of ℓ. Let X be the set of queries that, given $m \sim U(M)$ and a random tape s, the initial algorithms $\mathcal{S}$, $\mathcal{R}$ jointly performs to the RO. Calling Coll the event that two such points collides with respect to F, since $|X| \leq h + k$

$$\Pr[\mathsf{Coll}] \leq \frac{(h+k)^2}{2} \cdot \frac{1}{2^\ell}.$$

Next we observe that, as $G : \{0,1\}^\ell \to \{0,1\}^\lambda$ is a random function, if $\neg\mathsf{Coll}$, then $\mathcal{S}^*$, $\mathcal{R}^*$ perfectly simulate the RO. In particular, calling ε the error probability of the given RO-channel, i.e., $\Pr[m \neq m' \,|\, \neg\mathsf{Coll}]$, we have that

$$\begin{aligned} \delta \ &=\ \Pr[m \neq m' \,|\, \mathsf{Coll}] \Pr[\mathsf{Coll}] + \Pr[m \neq m' \,|\, \neg\mathsf{Coll}] \Pr[\neg\mathsf{Coll}] \\ &\leq\ \Pr[\mathsf{Coll}] + \Pr[m \neq m' \,|\, \neg\mathsf{Coll}] \\ &\leq\ \frac{(h+k)^2}{2 \cdot 2^\ell} + \varepsilon. \end{aligned}$$

Setting $\ell = 2\log(h+k)$ we obtain $1 - \delta \geq 1/2 - \varepsilon$ and in particular

$$|M| \ \leq\ \frac{2^{2\log(h+k)}}{1/2 - \varepsilon} \ =\ 2 \cdot (h+k)^2 + \mathsf{negl}(\lambda) \quad \Rightarrow \quad |M| \leq 2 \cdot (h+k)^2$$

where the equality holds because $\frac{1}{1-2\varepsilon} = 1 + \mathsf{negl}(\lambda)$ and last inequality holds asymptotically in λ as $|M|$ is an integer and $\mathsf{negl}(\lambda)$ is eventually less than 1. $\square$

5 Lower bounds and Impossibility

5.1 Communication Rate Bound

In this section we answer our question on black-box anamorphic encryption proving that its anamorphic message space must be polynomially bounded, or equivalently that it is impossible to communicate more than $O(\log \lambda)$ bits per ciphertext. The main technique, as described in the introduction, is to combine the information-theoretic lower bound for RO-channel with the ciphertext-selection lemma. The latter indeed informally implies that communication using black-box anamorphic encryption scheme happens almost as in a RO-channel: the sender can only perform certain queries to $\mathsf{E.Enc}(\mathsf{apk}, m, \cdot)$ and eventually return one of the replies. Similarly, the receiver is allowed to query $\mathsf{E.Enc}(\mathsf{apk}, m, \cdot)$ to extract information about the sender's hidden message. We can thus present our first result.

Theorem 2. *Let* $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *be a black-box anamorphic triplet with anamorphic message space* $\widehat{M}$*. Then* $|\widehat{M}| = \mathsf{poly}(\lambda)$*. More precisely, calling* q_e *and* q_d *the queries performed to* $\mathsf{E.Enc}$ *respectively by* $\mathsf{AT.Enc}$ *and* $\mathsf{AT.Dec}$*, then* $|\widehat{M}| \leq 2(q_e + q_d)^2$*.*

Proof. Applying the above black-box anamorphic triplet scheme to the ideal PKE $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ defined in Section 3.1, we describe a RO-channel with anamorphic message space $\widehat{M}$. A detailed presentation of $\mathcal{S}, \mathcal{R}$ appears in Figure 8. Initially both procedures hold shared randomness used to setup the anamorphic encryption parameters, and later simulate the ideal PKE. This is of the form $(s^*, r^*, m^*, \phi^*, \psi^*, \xi^*)$ with

- (s^*, r^*): random tapes for $\mathsf{AT.Gen}$ and $\mathsf{AT.Enc}$.
- m^*: random regular (i.e. non anamorphic) message in $M = \{0,1\}^\mu$.
- ϕ^*: random bijection from SK to PK, as in the ideal PKE.
- ψ^*: random function mapping (pk, m, r) to ciphertexts in $\{0,1\}^\ell$.
- ξ^*: biased random function mapping $\mathsf{SK} \times \{0,1\}^\ell$ to $M \cup \{\perp\}$, such that $\xi^*(\mathsf{sk}, c) = m_0$ with probability $2^{\ell-\rho}$ for all $m_0 \in M$.

Given the above shared randomness $\mathcal{S}$, $\mathcal{R}$ proceed as follows:

1. Key Generation. Initially they both setup the Anamorphic Encryption parameters $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$ running $\mathsf{AT.Gen}(\lambda; s^*)$ (lines 1-6). In this phase, each time the key generation queries $\mathsf{E.Gen}(\lambda; \mathsf{sk})$, they use ϕ^* to reply with $(\phi^*(\mathsf{sk}), \mathsf{sk})$. When it queries an encryption $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ they both reply with $\psi^*(\mathsf{pk}, m, r)$. When it queries a decryption $\mathsf{E.Dec}(\mathsf{sk}, c)$, if c was previously obtained as the encryption of some m they reply with m. Else, they reply with $\xi^*(\mathsf{sk}, c)$.

2. Encryption. $\mathcal{S}^{\mathsf{H}}(\widehat{m})$ proceeds computing c^*, the anamorphic encryption of $(m^*, \widehat{m})$ with keys $(\mathsf{apk}, \mathsf{dk})$ and randomness r^* (lines 9-14). During this computation, each time $\mathsf{AT.Enc}$ queries $\mathsf{E.Gen}(\lambda; \mathsf{sk})$ it replies as above using ϕ^*. When it queries an encryption $\mathsf{E.Enc}(\mathsf{pk}, m; r)$, if the same request was performed by $\mathsf{AT.Gen}$ it replies consistently, i.e. with $\psi^*(\mathsf{pk}, m; r)$. Otherwise it invokes its RO, replying with $c = \mathsf{H}(\mathsf{pk}, m, r)$. Decryption queries are handled as before. Finally it returns c^*.

3. Decryption. $\mathcal{R}$ on input c^* finally computes $\widetilde{m} \leftarrow \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c^*)$ (lines 9-14, right procedure). During this execution, each time $\mathsf{AT.Dec}$ queries $\mathsf{E.Gen}(\lambda; \mathsf{sk})$, it replies as above using ϕ^*. When it queries $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ it replies with $\psi^*(\mathsf{pk}, m, r)$ if the same query was performed by $\mathsf{AT.Gen}$, or with $\mathsf{H}(\mathsf{pk}, m, r)$ otherwise. Finally, queries to $\mathsf{E.Dec}(\mathsf{sk}, c)$ are handled as before, with the exception that to $\mathsf{E.Dec}(\mathsf{ask}, c^*)$ it always replies with m^* (see line 7). Eventually it returns $\widetilde{m}$.

$\mathcal{S}^{\mathsf{H}}(\widehat{m}\,; (s^*, r^*, m^*, \phi^*, \psi^*, \xi^*))$:		$\mathcal{R}^{\mathsf{H}}(c^*\,; (s^*, r^*, m^*, \phi^*, \psi^*, \xi^*))$:	
1 :	$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow \mathsf{AT.Gen}(\lambda; s^*)$	1 :	$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow \mathsf{AT.Gen}(\lambda; s^*)$
2 :	**when** queried $\mathsf{E.Enc}(\mathsf{pk}, m; r)$:	2 :	**when** queried $\mathsf{E.Enc}(\mathsf{pk}, m; r)$:
3 :	Get $c \leftarrow \psi^*(\mathsf{pk}, m, r)$	3 :	Get $c \leftarrow \psi^*(\mathsf{pk}, m, r)$
4 :	Set $\xi^*(\mathsf{sk}, c) \leftarrow m \ : \ \mathsf{pk} = \phi^*(\mathsf{sk})$	4 :	Set $\xi^*(\mathsf{sk}, c) \leftarrow m \ : \ \mathsf{pk} = \phi^*(\mathsf{sk})$
5 :	Set $\mathsf{H}(\mathsf{pk}, m, r) \leftarrow c$	5 :	Set $\mathsf{H}(\mathsf{pk}, m, r) \leftarrow c$
6 :	**reply** c	6 :	**reply** c
7 :		7 :	Set $\xi^*(\mathsf{ask}, c^*) \leftarrow m^*$
8 :	// Get the Anamorphic Encryption	8 :	// Decrypt the Anamorphic Ciphertext
9 :	Run $c^* \leftarrow \mathsf{AT.Enc}(\mathsf{apk}, m; r)$	9 :	Run $\widetilde{m} \leftarrow \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c^*)$
10 :	**when** queried $\mathsf{E.Enc}(\mathsf{pk}, m, r)$:	10 :	**when** queried $\mathsf{E.Enc}(\mathsf{pk}, m, r)$:
11 :	Get $c \leftarrow \mathsf{H}(\mathsf{pk}, m, r)$	11 :	Get $c \leftarrow \mathsf{H}(\mathsf{pk}, m, r)$
12 :	Set $\xi^*(\mathsf{sk}, c) \leftarrow m \ : \ \mathsf{pk} = \phi^*(\mathsf{sk})$	12 :	Set $\xi^*(\mathsf{sk}, c) \leftarrow m \ : \ \mathsf{pk} = \phi^*(\mathsf{sk})$
13 :	**reply** c	13 :	**reply** c
14 :	**return** c^*	14 :	**return** $\widetilde{m}$
15 :	// Key Gen. and Decryption query	15 :	// Key Gen. and Decryption query
16 :	**when** queried $\mathsf{E.Gen}(\lambda; \mathsf{sk})$:	16 :	**when** queried $\mathsf{E.Gen}(\lambda; \mathsf{sk})$:
17 :	**reply** $(\phi^*(\mathsf{sk}), \mathsf{sk})$	17 :	**reply** $(\phi^*(\mathsf{sk}), \mathsf{sk})$
18 :	**when** queried $\mathsf{E.Dec}(\mathsf{sk}, c)$:	18 :	**when** queried $\mathsf{E.Dec}(\mathsf{sk}, c)$:
19 :	**reply** $\xi^*(\mathsf{sk}, c)$	19 :	**reply** $\xi^*(\mathsf{sk}, c)$

Fig. 8. RO-Channel based on black-box Anamorphic Encryption. The notation $\mathsf{H}(\mathsf{pk}, m, r) \leftarrow c$ denotes that future calls to H on (pk, m, r) return c without calling H.

Given the above description of $\mathcal{S}, \mathcal{R}$ we proceed illustrating immediate properties they satisfy. First of all $\mathcal{S}$ returns up to negligible probability a value it received from the RO. This follows from the Ciphertext Selection Lemma (Lemma 3) and Lemma 1. Indeed, they imply $\mathsf{AT.Enc}$ will almost always return a ciphertext c it obtained from $\mathsf{E.Enc}$ and which was not observed by $\mathsf{AT.Gen}$, meaning that c is evaluated from H (as opposed to ψ^* to keep consistency with $\mathsf{AT.Gen}$'s view). Another immediate observation is that $\mathcal{S}$ and $\mathcal{R}$ respectively performs q_e and q_d RO calls, i.e. the number of queries to $\mathsf{E.Enc}$ respectively from $\mathsf{AT.Enc}$ and $\mathsf{AT.Dec}$. This follows as the RO may be called at most once for each such query.

To conclude that $(\mathcal{S}, \mathcal{R}, \widehat{M}, q_e, q_d)$ is a RO-Channel we only need to establish correctness. To do so we rely on the anamorphic encryption scheme's correctness, Section 2.2: given correctly generated keys and messages $(m, \widehat{m})$

$$\Pr\left[\widetilde{m} \neq \widehat{m} \,\middle|\, \widetilde{m} \leftarrow \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c),\ c \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})\right] \ \leq\ \mathsf{negl}(\lambda).$$

Note this holds only when all queries the anamorphic encryption scheme performs to the underlying PKE are answered correctly. Our last step is then to prove $\mathcal{S}, \mathcal{R}$ simulate the ideal PKE correctly. Let $\mathsf{View}^{\mathsf{real}}$ be the sequence of oracle replies $\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec}$ (in this order) would observe when executed with the correct PKE, and $\mathsf{View}^{\mathsf{sim}}$ the sequence of values they get with $\mathcal{S}, \mathcal{R}$. We claim them to be statistically close, implying that $\Pr[\widetilde{m} \neq \widehat{m}] \leq \mathsf{negl}(\lambda)$.

Claim 3. $\Delta(\mathsf{View}^{\mathsf{real}}, \mathsf{View}^{\mathsf{sim}}) \leq \mathsf{negl}(\lambda)$.

A proof of this Claim is presented in the Appendix, Section C.4. Finally, applying Theorem 1, we conclude that $|\widehat{M}| \leq 2(q_e + q_d)^2$. □

Remark 2. Again, this lower bound holds for *stateless* black-box triplets. If *stateful* anamorphic encryption/decryption is allowed, Lemma 3 only guarantees that c is a valid ciphertext observed by $\mathsf{AT.Enc}$ or $\mathsf{AT.Gen}$ (see Remark 1). This worsen the final bound to $\widehat{M} \leq 2(q_e + q_d + 2q_g)^2$ with q_g the total queries of $\mathsf{AT.Gen}$. The proof is readily adapted by replacing ψ^* with H calls both in $\mathcal{S}$ and $\mathcal{R}$.

5.2 Impossibility of Asymmetric AE

The bounds provided in the previous section applies to any black-box anamorphic triplet. Although our bound can be achieved asymptotically, see [PPY22], the only known constructions encrypt anamorphic messages in a *symmetric* fashion. That is, sender and receiver must have exchanged a secret key in advance. The lack of black-box *asymmetric* anamorphic scheme is however no coincidence. In this section we will indeed prove that such constructions are impossible.

More precisely, we will prove that any black-box anamorphic triplet scheme satisfying Definition 5, must be insecure with respect to the Weak Asymmetric security notion (Definition 4) when instantiated for the ideal PKE scheme.

Theorem 3. *For any black-box anamorphic triplet* $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$*, when applied to the ideal PKE* $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *(Section 3.1) there exists* $\mathcal{A}$ PPT *such that,*

$$\mathsf{Adv}^{\mathsf{Weak\text{-}Asy\text{-}Anam}}_{\mathcal{A},\Sigma}(\lambda) \;\geq\; \frac{1}{\mathsf{poly}(\lambda)}.$$

Proof. At a high level the strategy of $\mathcal{A}$, fully described in Figure 9, is as follows. First it gets a challenge ciphertext c^* encrypting either $(m^*, \widehat{m}_0)$ or $(m^*, \widehat{m}_1)$ random messages of its choice. Next it locally runs $\mathsf{AT.Enc}$ to encrypt $\widehat{m}_0$ and during its execution replaces the response of a randomly chosen query to $\mathsf{E.Enc}$ with c^*. If c^* encrypts $\widehat{m}_0$, $\mathsf{AT.Enc}$ should return it with significant probability, whereas if it encrypts $\widehat{m}_1$, this should only happen with negligible probability.

This simple approach however faces a number of technical challenges. First, we need to ensure $\mathcal{A}$ is unlikely to *overwrite* an encryption query that was previously performed by $\mathsf{AT.Gen}$, as this will create detectable inconsistencies. Next, the query c^* may not follow the expected distribution *given* dk. This may be the case since anamorphic security only guarantees c^* to be indistinguishable from any other ciphertext given $\mathsf{ask}, \mathsf{apk}$ but not dk. Thus c^* is *not* hard to distinguish and creates a non-negligible change in the view of $\mathsf{AT.Enc}$.

To address the first issue we rely on a preprocessing phase (Lines 1-5): $\mathcal{A}$ initially runs $\mathsf{AT.Enc}$ for ϑ many times (we fix ϑ later) and stores the randomness used in encryption queries of the form $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$. The idea is that if $\mathsf{AT.Gen}$ performs a query of this kind, either it is easily observed in the preprocessing or $\mathsf{AT.Enc}$ queries it with sufficiently low probability for our argument to go through. After this phase, the attack is executed as mentioned above (lines 6-13), choosing the query to program randomly among those of the form $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$ where r was not observed in the preprocessing phase.

Regarding the second issue, we will use the fact that $\mathsf{AT.Enc}$ can be roughly treated as a symmetric choice function (see Section 3.5). This will help us conclude that, when c^* is the encryption of $\widehat{m}_0$, the probability of choosing it again is significant.

Let $q = \mathsf{poly}(\lambda)$ be the number of queries made by $\mathsf{AT.Enc}$ to $\mathsf{E.Enc}$. Our first step is to show that although c^* is biased, this can only increase the probability of certain (bad) events by a factor of $\approx q$, plus a non-negligible term accounting for the probability that $\mathcal{A}$ overwrites a query previously asked by $\mathsf{AT.Gen}$. To be more precise we call Bias the joint view of $\mathsf{AT.Gen}$, which generates $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$, $\mathsf{AT.Enc}$ executed as in line 9, and $\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c')$. Similarly, let Real be the same view, with the exception that at line 11 $\mathcal{A}$ returns the correct ciphertext $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$. Then we can claim the following bound.

Claim 4. *For any predicate* p

$$\Pr\left[\, p(\mathsf{Bias}) = 1\right] \leq q \cdot \Pr\left[\, p(\mathsf{Real}) = 1\right] + \frac{q^2}{\vartheta + 1} + \mathsf{negl}(\lambda).$$

The proof of this claim appears in the Appendix, Section C.5. Next we proceed studying the probability that $\mathcal{A}$ returns 1 when c^* is an encryption of $\widehat{m}_b$ for $b \in \{0, 1\}$ separately.

$\mathcal{A}(\mathsf{apk}, \mathsf{dk})$:

1 : // Preprocessing phase
2 : Set $R \leftarrow \varnothing$, sample $m^* \leftarrow^{\$} M$ and $\widehat{m}_0, \widehat{m}_1 \leftarrow^{\$} \widehat{M}$
3 : **for** ϑ times:
4 : Run $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$
5 : **when** it queries $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$: Store $R \leftarrow R \cup \{r\}$
6 : // Attack phase
7 : Sample a random $i \leftarrow^{\$} \{1, \ldots, q\}$
8 : Give $(m^*, \widehat{m}_0, \widehat{m}_1)$ to the challenger and obtain c^*
9 : Run $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$
10 : **when** it queries the i-th time a new $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$ with $r \notin R$:
11 : **reply** with c^*
12 : **when** it returns c':
13 : **return** $c^* == c'$

Fig. 9. Adversary for the Weak Asymmetric AE game, where $\vartheta = \mathsf{poly}(\lambda)$ and $q = \mathsf{poly}(\lambda)$ is the number of queries made by $\mathsf{AT.Enc}$ to $\mathsf{E.Enc}$.

Encryption of $\widehat{m}_1$. In this case let Err be the event $\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c') \neq \widehat{m}_0$. From correctness of the anamorphic encryption scheme, if $\mathcal{A}$ replies with the correct ciphertext at line 11, this event occurs only with negligible probability. Using Claim 4 we have then that

$$\begin{aligned} \Pr[c' = c^* \mid b = 1] \;&\leq\; \Pr[\mathsf{Err}] + \mathsf{negl}(\lambda) \;\leq\; q \cdot \mathsf{negl}(\lambda) + \frac{q^2}{\vartheta + 1} + \mathsf{negl}(\lambda) \\ &=\; \frac{q^2}{\vartheta + 1} + \mathsf{negl}(\lambda) \end{aligned}$$

where the first inequality follows as c^* is the encryption of $\widehat{m}_1$, and therefore, up to negligible probability $\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c^*) = \widehat{m}_1 \neq \widehat{m}_0$.

Encryption of $\widehat{m}_0$. We start by fixing some notation. We will call S^*, S the sets of randomness r so that the query $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$ was respectively performed by $\mathsf{AT.Enc}$ inside the challenger call in line 8 or $\mathsf{AT.Enc}$ executed in line 9. As a direct consequence of the Ciphertext Selection Lemma and Lemma 2 we then claim that

Claim 5. *Calling* $\mathsf{BadChoice} : (\nexists r' \in S \setminus R \,:\, c' = \mathsf{E.Enc}(\mathsf{apk}, m^*; r'))$ *and analogously* $\mathsf{BadChoice}^* : (\nexists r^* \in S^* \setminus R \,:\, c^* = \mathsf{E.Enc}(\mathsf{apk}, m^*; r^*))$ *then*

$$\Pr[\mathsf{BadChoice}^*] \;\leq\; \mathsf{negl}(\lambda), \qquad \Pr[\mathsf{BadChoice}] \;\leq\; \frac{q^2}{\vartheta + 1} + \mathsf{negl}(\lambda).$$

A proof appears in the Appendix, Section C.5. Next, our goal is to argue that $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$ is *close* to a symmetric choice function, taking as

input the ciphertexts it requests through encryption calls and returning one of them. Conditioning on $\neg\mathsf{BadChoice}$ guarantees that this is a choice function. To argue it is also almost symmetric we use a sequence of hybrid adversaries where we replace $\mathsf{E.Enc}$ with an actual symmetric choice function $\mathcal{F}$, described in Figure 10.

- $\mathcal{A}_1$: The adversary described in Figure 9, when the challenger encrypts $\widehat{m}_0$.
- $\mathcal{A}_2$: As $\mathcal{A}_1$, but to compute c^* it samples $c_1, \ldots, c_q \leftarrow^{\$} \{0,1\}^\ell$ and evaluates the function $\mathcal{F}$, described in Figure 10, setting $c^* = \mathcal{F}(c_1, \ldots, c_q)$.
- $\mathcal{A}_3$: As $\mathcal{A}_2$, but to compute c' it samples $c_2, \ldots, c_q \leftarrow^{\$} \{0,1\}^\ell$ and evaluates the function $\mathcal{F}$, described in Figure 10, setting $c' = \mathcal{F}(c^*, c_2, \ldots, c_q)$.

$\mathcal{F}(c_1, \ldots, c_q)$:

1 : Sample a random permutation $\pi : \{1, \ldots, q\} \to \{1, \ldots, q\}$.
2 : Run $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$
3 : **when** it queries a new $\mathsf{E.Enc}(\mathsf{apk}, m^*; r)$ with $r \notin R$ the i-th time:
4 : **reply** $c_{\pi(i)}$
5 : **when** it queries $\mathsf{E.Dec}(\mathsf{ask}, c)$ with $c \in \{c_1, \ldots, c_q\}$:
6 : **reply** m^*.
7 : **when** it returns c_{out}
8 : **if** $c_{\mathsf{out}} \in \{c_1, \ldots, c_q\}$: **return** c_{out}
9 : **else** : **return** a random $c_{\mathsf{out}} \leftarrow^{\$} \{c_1, \ldots, c_q\}$

Fig. 10. Symmetric choice function used to replace $\mathsf{E.Enc}$ in $\mathcal{A}_1, \mathcal{A}_2$. Note this is implicitly parametrized by $\mathsf{apk}, \mathsf{dk}$ and R. Equality to ask can be checked querying $\mathsf{E.Gen}$.

For notational convenience we will call c_i^*, c_i' the ciphertexts generated by $\mathcal{A}_i$. Then we can claim that $\mathcal{F}$ is a symmetric choice function and that the statistical distance between the ciphertexts generated by these adversaries is small.

Claim 6. *$\mathcal{F}$ is a symmetric choice function (see Definition 9).*

Claim 7. $\Delta((c_1^*, c_1'), (c_2^*, c_2')) \leq \mathsf{negl}(\lambda)$.

Claim 8. $\Delta((c_2^*, c_2'), (c_3^*, c_3')) \leq \frac{2q^2}{1+\vartheta} + \mathsf{negl}(\lambda)$.

All three Claims are proven in the Appendix, Section C.5. Combining them with Lemma 4 we have that $\Pr[c_3^* = c_3'] \geq q^{-1} - \mathsf{negl}(\lambda)$ and in particular

$$\begin{aligned}\Pr[c^* = c' \mid b = 0] \;=\; \Pr[c_1^* = c_1'] \;&\geq\; \Pr[c_3^* = c_3'] - \frac{2q^2}{\vartheta + 1} - \mathsf{negl}(\lambda) \\ &\geq\; \frac{1}{q} - \frac{2q^2}{\vartheta + 1} - \mathsf{negl}(\lambda).\end{aligned}$$

Advantage Bound. Combining both intermediate results, a bound on the advantage of $\mathcal{A}$ can be derived as

$$\begin{aligned}\mathsf{Adv}(\mathcal{A}) &= |\Pr[c^* = c' \mid b = 0] - \Pr[c^* = c' \mid b = 1]| \\ &\geq \left(\frac{1}{q} - \frac{2q^2}{\vartheta + 1} - \mathsf{negl}(\lambda)\right) - \left(\frac{q^2}{\vartheta + 1} + \mathsf{negl}(\lambda)\right) \\ &\geq \frac{1}{q} - \frac{3q^2}{\vartheta + 1} - \mathsf{negl}(\lambda).\end{aligned}$$

Setting $\vartheta = 6q^3 - 1$ we get that the advantage is negligibly close to $1/2q$. As $q = \mathsf{poly}(\lambda)$ the Theorem is proven. □

Remark 3. As done previously, the Theorem only refers to a *stateless* anamorphic triplet. In this case however we choose not to discuss about stateful variants as, even in anamorphic mode, the scheme is *asymmetric*, with potentially many senders holding the same dk. Thus keeping state in such case does not appear meaningful.

6 Dual Construction

As stated in Section 3.4, the Ciphertext Selection Lemma holds only for a certain parameters choice of the ideal PKE. Thus it would not apply to black-box AE for the *specific* class of PKE with *small* message space and *dense* ciphertext space. In this section we prove the above restriction[13] to be necessary. We do so by showing that for this class of PKE (further satisfying a technical condition explained below) there exists a simple compiler to a black-box asymmetric AE with exponential anamorphic message space. We call this the "Dual Construction" as it is reminiscent of the black-box solution in [PPY22], but swapping the role of regular and anamorphic messages. Let $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ be the PKE with small message space and dense ciphertext space, having $(\mathsf{apk}, \mathsf{ask})$ as a pair of keys. At high level the idea is to have another PKE scheme, call it $\Pi^{\mathsf{pr}} = (\Pi^{\mathsf{pr}}.\mathsf{Gen}, \Pi^{\mathsf{pr}}.\mathsf{Enc}, \Pi^{\mathsf{pr}}.\mathsf{Dec})$, with a corresponding pair of keys $(\mathsf{dpk}, \mathsf{dsk})$, and to set $\mathsf{dk} = (\mathsf{dpk}, \mathsf{ask})$. In order to encrypt a normal message m and a covert message $\widehat{m}$ in a normal looking ciphertext c, $\widehat{m}$ is encrypted with $\Pi^{\mathsf{pr}}.\mathsf{Enc}(\mathsf{dpk}, \widehat{m})$ until it obtains through rejection sampling a c such that $\mathsf{E.Dec}(\mathsf{ask}, c) = m$. A detailed description of the anamorphic triplet $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ for Π appears in Figure 11.

For this to work we need the PKE Π^{pr} to produce ciphertexts that look "uniformly distributed" over the ciphertext space of Π. This can be achieved through a (weak) *pseudorandom ciphertexts* PKE scheme [vH04, Möl04], see Appendix, Section A.1. Next we formalize the conditions that $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ has to satisfy, in order to apply this generic compiler to it and obtain an AE scheme. Those are

[13] Which does not affect the generality our result, but only prevents it to be extended to such specific case.

1. Given $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$, there exists $p = 1/\mathsf{poly}(\lambda)$ such that, for all $m \in M$ and c uniform over the ciphertext space $\Pr[\mathsf{E.Dec}(\mathsf{sk}, c) = m] \geq p$.
2. Given $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$, and $m \in M$, then $c \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{pk}, m)$ implies that c is uniformly distributed over the ciphertexts that decrypt to m, i.e.

$$c \sim U(\{c_0 \ : \ \mathsf{E.Dec}(\mathsf{sk}, c_0) = m\}) .$$

The first means at the same time that the plaintext space is polynomially small and the ciphertext space dense. The second one instead is introduced for technical reasons, to ensure that a ciphertext for Π obtained through rejection sampling has the same distribution of a fresh encryption computed with $\mathsf{E.Enc}$, which is needed to prove the final construction to be anamorphic. Finally, note that the ideal PKE scheme defined in Section 3.1 satisfies both conditions when $\ell - \rho = O(\log \lambda)$.

$\mathsf{AT.Gen}(\lambda)$

1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$
2 : $(\mathsf{dpk}, \mathsf{dsk}) \leftarrow^{\$} \Pi^{\mathsf{pr}}.\mathsf{Gen}(\lambda)$
3 : $\mathsf{apk} \leftarrow \mathsf{pk},\ \mathsf{ask} \leftarrow \mathsf{sk}$
4 : $\mathsf{dk} \leftarrow (\mathsf{dpk}, \mathsf{ask}),\ \mathsf{tk} \leftarrow \mathsf{dsk}$
5 : **return** $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$

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

1 : Parse dk as $(\mathsf{dpk}, \mathsf{sk})$
2 : **do** // Rejection Sampling
3 : $\quad c \leftarrow^{\$} \Pi^{\mathsf{pr}}.\mathsf{Enc}(\mathsf{dpk}, \widehat{m})$
4 : **while** $m \neq \mathsf{E.Dec}(\mathsf{sk}, c)$
5 : **return** c

$\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c)$

1 : **return** $\Pi^{\mathsf{pr}}.\mathsf{Dec}(\mathsf{tk}, c)$

Fig. 11. Black-Box Anamorphic Triplet Σ^{bb}. Note $\mathsf{AT.Enc}$ runs in expected polynomial time $O(1/p) = \mathsf{poly}(\lambda)$. This can be turned into PPT by limiting the *while* loop to λ/p iterations, making however $\mathsf{AT.Enc}$'s usage of Π non-uniform.

6.1 Anamorphism

Theorem 4. *If Π^{pr} is a PKE with weak pseudorandom ciphertext (see Appendix, Section A.1) and Π is an* IND-CPA *secure PKE satisfying the two conditions in Section 6, then Π equppied with Σ^{bb} defined in Figure 11 is a Black-Box Anamorphic Encryption scheme.*

Proof. Let $\mathcal{D}$ be an adversary distinguishing RealG_{Π} from $\mathsf{AnamorphicG}_{\Sigma^{\mathsf{bb}}}$. We reduce it to an adversary $\mathcal{A}$ against the weak pseudorandom-ciphertext property of Π^{pr}, described in Figure 12. Precisely, $\mathcal{A}$ plays the game $\mathsf{W\text{-}PRCtG}^{b}_{\Pi^{\mathsf{pr}},\mathcal{A}}$ with access to $\mathcal{O}$ which, on input $\widehat{m}$, returns either a random string s when $b = 0$ or the result of $\Pi^{\mathsf{pr}}.\mathsf{Enc}(\mathsf{pk}, \widehat{m})$, with pk chosen by the challenger, when $b = 1$. Its strategy is to run $\mathcal{D}$, answering its encryption queries via $\mathcal{O}$. It does so through

rejection sampling as done in $\mathsf{AT.Enc}$, performing ϑ attempts each time before giving up (we specify a suitable ϑ later in the proof).

$\mathcal{A}^{\mathcal{O}}(\lambda)$:

Sample $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$ and run $\mathcal{D}(\mathsf{pk}, \mathsf{sk})$
when $\mathcal{D}$ queries $(m_i, \widehat{m}_i)$ the i-th time:
  for ϑ times: // Rejection Sampling with ϑ attempts
    Get $c \leftarrow \mathcal{O}(\widehat{m}_i)$ from the PRC encryption oracle
    if $m_i = \mathsf{E.Dec}(\mathsf{sk}, c)$: **reply** c to $\mathcal{D}$ and **break**
  if no reply was given to $\mathcal{D}$ in the previous loop:
    return $\perp$ // i.e. abort
when $\mathcal{D}$ returns b': **return** b'

Fig. 12. Adversary $\mathcal{A}$ parametrized by ϑ reducing $\mathcal{D}$ for $\mathsf{Anamorphism}$ to $\mathsf{W\text{-}PRCtG}$.

Formally let q be an upper bound on the total queries performed by $\mathcal{D}$ and recall p to be a lower bound on the probability that a random ciphertext for Π decrypts to a given message (by the hypothesis on Π, $p \geq 1/\mathsf{poly}(\lambda)$). We call Abort_i the event where $\mathcal{A}$ aborts after the i-th query of $\mathcal{D}$, and $\mathsf{Abort} = \bigvee_{i=1}^{q} \mathsf{Abort}_i$. First we claim that for a sufficiently large ϑ, this occurs with negligible probability.

Claim 9. *If* $\vartheta \geq \log_2(q) \cdot \lambda / p$ *then* $\Pr[\mathsf{Abort}] \leq \mathsf{negl}(\lambda)$.

Then, up to negligible probability, it suffices to study the advantage of $\mathcal{A}$ conditioning on $\neg\mathsf{Abort}$. If $b = 0$, then $\mathcal{A}$ obtains random strings from $\mathcal{O}$ in the ciphertext space. In particular its replies to the i-th query $(m_i, \widehat{m}_i)$ is replied with c uniformly distributed over the ciphertext such that $\mathsf{E.Dec}(\mathsf{sk}, c) = m_i$. Our second condition on Π implies follows the same distribution of $\mathsf{E.Enc}(\mathsf{pk}, m_i)$, and so $\mathcal{A}$ perfectly simulates the real game in Figure 1.

Conversely if $b = 1$, its behavior is identical to $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m_i, \widehat{m}_i)$ (up to the negligible failing probability). Thus conditioning on $\neg\mathsf{Abort}$ it perfectly simulates the view of $\mathcal{D}$ in the anamorphic game. We thus conclude that

$$\mathsf{Adv}^{\mathsf{Anam}}_{\mathcal{D},\Pi,\Sigma^{\mathsf{bb}}}(\lambda) \;\leq\; \mathsf{Adv}^{\mathsf{W\text{-}PRCtG}}_{\mathcal{A},\Pi^{\mathsf{pr}}}(\lambda) + \mathsf{negl}(\lambda). \qquad \square$$

Proof of Claim 9. Let $\mathsf{Abort}_{i,j}$ be the event in which, while replying to the i-th query, $\mathcal{A}$ gets a ciphertext c such that $\mathsf{E.Dec}(\mathsf{sk}, c) \neq m_i$ in the j-th repetition of

the loop.[14] As these events are all mutually independent, though a union bound

$$\begin{aligned}\Pr[\mathsf{Abort}] \;&\leq\; \sum_{i=1}^{q} \Pr[\mathsf{Abort}_i] \;=\; \sum_{i=1}^{q} \prod_{j=1}^{\vartheta} \Pr[\mathsf{Abort}_{i,j}] \\ &\leq\; q(1-p)^{\vartheta} \;\leq\; q \cdot 2^{-\vartheta p} \;\leq\; 2^{-\lambda}\end{aligned}$$

where the first inequality follows as $\Pr[\mathsf{Abort}_{i,j}]$ is smaller than $1-p$, while the second one follows as $(1-p)^{1/p} \leq 1/2$ for all $p \in [0,1]$. □

6.2 Weak Asymmetric

Theorem 5. *If Π^{pr} is an* IND-CPA *secure PKE and Π is a PKE satisfying the two conditions in Section 6, then, Π equipped with Σ^{bb} defined in Figure 11 is a Weak Asymmetric Anamorphic Encryption scheme.*

Proof. Let $\mathcal{D}$ be a distinguisher for the weak asymmetric anamorphic security game (See Section 2.3). We use it to construct an adversary $\mathcal{A}$ for the IND-CPA security of Π^{pr} fully described in Figure 13. The reduction simply generates the public parameters of Π and runs $\mathcal{D}$. To reply to $\mathcal{D}$'s encryption query, $\mathcal{A}$ adopts the same strategy as in the proof of Theorem 4, i.e. it performs rejection sampling on the ciphertexts generated by the IND-CPA oracle for Π^{pr}.

$\mathcal{A}(\mathsf{dpk})$:

1 : Get $(\mathsf{pk},\mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$ and set the double key $\mathsf{dk} \leftarrow (\mathsf{dpk},\mathsf{sk})$
2 : Run $\mathcal{D}(\mathsf{pk},\mathsf{dk})$ until it queries $(m, \widehat{m}_0, \widehat{m}_1)$
3 : **for** λ/p times: // Rejection sampling as AT.Enc
4 : Send $(\widehat{m}_0, \widehat{m}_1)$ to the encryption oracle and get c
5 : **if** $m = \mathsf{E.Dec}(\mathsf{sk}, c)$: **reply** c to $\mathcal{D}$ and **break**
6 : **if** no reply was given to $\mathcal{D}$ in the previous loop:
7 : **reply** $\perp$ to $\mathcal{D}$
8 : **when** $\mathcal{D}$ returns b': **return** b'

Fig. 13. $\mathcal{A}$ reducing weak asymmetric anamorphic security of Σ^{bb} to IND-CPA of Π^{pr}.

Calling b the challenge bit for $\mathcal{A}$, it is immediate to observe that $\mathcal{A}$ perfectly emulates the behavior of $\mathsf{AT.Enc}(\mathsf{apk},\mathsf{dk},m,\widehat{m}_b)$, including the (small) error probability. We can thus conclude that

$$\mathsf{Adv}^{\mathsf{Weak\text{-}Asy\text{-}Anam}}_{\mathcal{D},\Sigma^{\mathsf{bb}}}(\lambda) \;\leq\; \mathsf{Adv}^{\text{IND-CPA}}_{\mathcal{A},\Pi^{\mathsf{pr}}}(\lambda) + \mathsf{negl}(\lambda). \qquad \square$$

[14] This is technically not well-defined as $\mathcal{A}$ may break the loop before the j-th iteration. This can be fixed re-defining $\mathcal{A}^*$ to (pointlessly) continue the loop execution ϑ times and observe $\mathcal{A}$ and $\mathcal{A}^*$ are functionally equivalent. We nevertheless omit such details.

Remark 4. One can verify that both properties of weak pseudorandom ciphertexts and IND-CPA security are implied by the regular pseudorandom ciphertexts property. So, if Π^{pr} has pseudorandom ciphertexts it satisfies both conditions for anamorphism and weak asymmetric anamorphism.

Acknowledgments

The authors would like to thank Chris Brzuska for pointing out a flaw in an early version of Theorem 1's proof, discussed during TCC'23. This study has been supported by the project "PrepAring cRypTograpHy for privacy-awarE blockchaiN applicatiONs (PARTHENON)" – PRIN 2022 - Finanziato dall'Unione europea - Next Generation EU – CUP: E53D2300799 0006. and partially supported by PRODIGY Project (TED2021-132464B-I00) funded by MCIN/AEI/10.13039/501100011033/ and the European Union NextGenerationEU/PRTR.

References

BF01. Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer, Heidelberg, August 2001.

BGH+24. Fabio Banfi, Konstantin Gegier, Martin Hirt, Ueli Maurer, and Guilherme Rito. Anamorphic encryption, revisited. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024*, pages 3–32, Cham, 2024. Springer Nature Switzerland.

BGI+01. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001.

Bla94. Matt Blaze. Protocol failure in the escrowed encryption standard. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, and Ravi S. Sandhu, editors, *ACM CCS 94*, pages 59–67. ACM Press, November 1994.

CDNO97. Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In Burton S. Kaliski Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 90–104. Springer, Heidelberg, August 1997.

CGM24. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Anamorphic encryption: New constructions and homomorphic realizations. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024*, pages 33–62, Cham, 2024. Springer Nature Switzerland.

DG17. Nico Döttling and Sanjam Garg. Identity-based encryption from the Diffie-Hellman assumption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 537–569. Springer, Heidelberg, August 2017.

DRS04. Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 523–540. Springer, Heidelberg, May 2004.

FY95. Yair Frankel and Moti Yung. Escrow encryption systems visited: Attacks, analysis and designs. In Don Coppersmith, editor, *CRYPTO'95*, volume 963 of *LNCS*, pages 222–235. Springer, Heidelberg, August 1995.

GGK03. Rosario Gennaro, Yael Gertner, and Jonathan Katz. Lower bounds on the efficiency of encryption and digital signature schemes. In *35th ACM STOC*, pages 417–425. ACM Press, June 2003.

Giu23. Emanuele Giunta. On the impossibility of algebraic NIZK in pairing-free groups. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 702–730. Springer, Heidelberg, August 2023.

GKM+00. Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The relationship between public key encryption and oblivious transfer. In *41st FOCS*, pages 325–335. IEEE Computer Society Press, November 2000.

GMR01. Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *42nd FOCS*, pages 126–135. IEEE Computer Society Press, October 2001.

GT00. Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *41st FOCS*, pages 305–313. IEEE Computer Society Press, November 2000.

IR89. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st ACM STOC*, pages 44–61. ACM Press, May 1989.

JJ21. Abhishek Jain and Zhengzhong Jin. Non-interactive zero knowledge from sub-exponential DDH. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 3–32. Springer, Heidelberg, October 2021.

KPP+23a. 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, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part II*, volume 14082 of *Lecture Notes in Computer Science*, pages 759–790. Springer, 2023.

KPP+23b. 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.

KST99. Jeong Han Kim, Daniel R. Simon, and Prasad Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *40th FOCS*, pages 535–542. IEEE Computer Society Press, October 1999.

Mic93. Silvio Micali. Fair public-key cryptosystems. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 113–138. Springer, Heidelberg, August 1993.

Möl04. Bodo Möller. A public-key encryption scheme with pseudo-random ciphertexts. In Pierangela Samarati, Peter Y. A. Ryan, Dieter Gollmann, and Refik Molva, editors, *ESORICS 2004*, volume 3193 of *LNCS*, pages 335–351. Springer, Heidelberg, September 2004.

PPY22. 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. Springer, Heidelberg, May / June 2022.

PRV12. Periklis A. Papakonstantinou, Charles Rackoff, and Yevgeniy Vahlis. How powerful are the DDH hard groups? *Electron. Colloquium Comput. Complex.*, page 167, 2012.

Sha84. Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 47–53. Springer, Heidelberg, August 1984.

Sim98. Daniel R. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In Kaisa Nyberg, editor, *EUROCRYPT'98*, volume 1403 of *LNCS*, pages 334–345. Springer, Heidelberg, May / June 1998.

vH04. Luis von Ahn and Nicholas J. Hopper. Public-key steganography. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 323–341. Springer, Heidelberg, May 2004.

WCHY23. 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, *Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4-8, 2023, Proceedings, Part VI*, volume 14443 of *Lecture Notes in Computer Science*, pages 135–167. Springer, 2023.

Yao86. Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.

YY96. Adam Young and Moti Yung. The dark side of "black-box" cryptography, or: Should we trust capstone? In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 89–103. Springer, Heidelberg, August 1996.

YY97. Adam Young and Moti Yung. The prevalence of kleptographic attacks on discrete-log based cryptosystems. In Burton S. Kaliski Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 264–276. Springer, Heidelberg, August 1997.

Zha22. Mark Zhandry. To label, or not to label (in generic groups). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part III*, volume 13509 of *LNCS*, pages 66–96. Springer, Heidelberg, August 2022.

ZZ20. Mark Zhandry and Cong Zhang. Indifferentiability for public key cryptosystems. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 63–93. Springer, Heidelberg, August 2020.

A Primitives

A.1 Asymmetric Encryption with pseudorandom ciphertexts

We recall the notion of asymmetric encryption pseudorandom ciphertext from [vH04, Möl04]. Let $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ be an asymmetric encryption scheme with message space M and ciphertext space C. We define the game $\mathsf{PRCtG}^b_{\Pi,\mathcal{D}}(\lambda)$, for $b \in \{0,1\}$, as in Fig. 14 and call, for any PPT adversary $\mathcal{D}$ its

advantage in distinguish between the two as

$$\mathsf{Adv}^{\mathsf{PRCtG}}_{\mathcal{D},\Pi}(\lambda) := \left|\Pr\left[\mathsf{PRCtG}^0_{\Pi,\mathcal{D}}(\lambda) = 1\right] - \Pr\left[\mathsf{PRCtG}^1_{\Pi,\mathcal{D}}(\lambda) = 1\right]\right|.$$

$\mathsf{PRCtG}^b_{\Pi,\mathcal{D}}(\lambda)$

1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$
2 : **return** $\mathcal{D}^{\mathcal{O}^b_{\mathsf{pr}}(\mathsf{pk},\cdot)}(\lambda, \mathsf{pk})$ where
3 : $\quad\mathcal{O}^0_{\mathsf{pr}}(\mathsf{pk}, m)$ returns a random string in C
4 : $\quad\mathcal{O}^1_{\mathsf{pr}}(\mathsf{pk}, m) = \mathsf{E.Enc}(\mathsf{pk}, m)$

Fig. 14. The pseudorandom ciphertext game for asymmetric encryption Π.

Definition 11. *Let Π be an asymmetric encryption scheme. Π has pseudorandom ciphertexts if for every* PPT *adversary $\mathcal{D}$ we have*

$$\mathsf{Adv}^{\mathsf{PRCtG}}_{\mathcal{D},\Pi}(\lambda) \leq \mathsf{negl}(\lambda).$$

We also define a weak variant of an asymmetric encryption with pseudorandom ciphertexts in which the distinguisher is not provided with the public key of the scheme. As above, let $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ be an asymmetric encryption scheme with message space M and ciphertext space C We define the game $\mathsf{W\text{-}PRCtG}^b_{\Pi,\mathcal{D}}(\lambda)$, for $b \in \{0, 1\}$, as in Figure A.1. Then, as above, we define the advantage of any adversary $\mathcal{D}$ distinguishing between the two as

$$\mathsf{Adv}^{\mathsf{W\text{-}PRCtG}}_{\mathcal{D},\Pi}(\lambda) := \left|\Pr\left[\mathsf{W\text{-}PRCtG}^0_{\Pi,\mathcal{D}}(\lambda) = 1\right] - \Pr\left[\mathsf{W\text{-}PRCtG}^1_{\Pi,\mathcal{D}}(\lambda) = 1\right]\right|.$$

$\mathsf{W\text{-}PRCtG}^b_{\Pi,\mathcal{D}}(\lambda)$

1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$
2 : **return** $\mathcal{D}^{\mathcal{O}^b_{\mathsf{pr}}(\mathsf{pk},\cdot)}(\lambda)$ where
3 : $\quad\mathcal{O}^0_{\mathsf{pr}}(\mathsf{pk}, m)$ returns a random string in C
4 : $\quad\mathcal{O}^1_{\mathsf{pr}}(\mathsf{pk}, m) = \mathsf{E.Enc}(\mathsf{pk}, m)$

Fig. 15. The *weak* pseudorandom-ciphertext game for asymmetric encryption Π.

Definition 12. *Let Π be an asymmetric encryption scheme. Π has weak pseudorandom ciphertexts if for every* PPT *adversary $\mathcal{D}$*

$$\mathsf{Adv}^{\mathsf{W\text{-}PRCtG}}_{\mathcal{D},\Pi}(\lambda) \leq \mathsf{negl}(\lambda).$$

It is easy to observe that the property of asymmetric pseudorandom ciphertexts implies the weak variant defined above.

Next, we give a generic compiler to turn any PKE scheme $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ with message and ciphertext space respectively M and C, into a PKE $\Pi' = (\mathsf{E.Gen}', \mathsf{E.Enc}', \mathsf{E.Dec}')$ with the same message and ciphertext space that has weak pseudorandom ciphertexts. The idea is to shuffle the ciphertexts produced by the encryption algorithm of Π using a pseudorandom permutation (PRP) $F : \mathcal{K} \times C \to C$, which key $k \in \mathcal{K}$ is stored in the public key pk' of Π'. Note that as the adversary in W-PRCtG game is not allowed to see pk', then he can't see k. Since F is a PRP, for the adversary is computationally hard to distinguish between ciphertexts produced with the "tweaked" PKE Π' from the output of a random permutation, i.e., truly-random strings in C.
The construction of Π' is given in Figure A.1.

$\mathsf{E.Gen}'(\lambda)$

1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$
2 : $k \leftarrow^{\$} \mathcal{K}$
3 : $\mathsf{pk}' = (\mathsf{pk}, k), \mathsf{sk}' = \mathsf{sk}$
4 : **return** $(\mathsf{pk}', \mathsf{sk}')$

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

1 : Parse pk' as (pk, k)
2 : $c \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{pk}, m)$
3 : $c' = F(k, c)$
4 : **return** c'

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

1 : $c = F^{-1}(k, c)$
2 : **return** $\mathsf{E.Dec}(\mathsf{sk}', c)$

Fig. 16. PKE Π' with weak pseudorandom ciphertexts.

Theorem 6. *If F is a PRP and Π is an asymmetric encryption scheme, then Π' defined in Figure A.1 is a PKE with weak pseudorandom ciphertexts.*

Proof. Let $\mathcal{D}$ be a distinguisher for $\mathsf{W\text{-}PRCtG}^0_{\Pi',\mathcal{D}}$ from $\mathsf{W\text{-}PRCtG}^1_{\Pi',\mathcal{D}}$, then we can construct a distinguisher $\mathcal{A}$ for the PRP game of F. The pseudocode of $\mathcal{A}$ is given in Figure A.1.

Namely, if $\mathcal{A}$ is playing the game PRP^0, then the oracle $\mathcal{O}$ given to $\mathcal{A}$ is a truly random permutation f, while if it is playing the game PRP^1, then $\mathcal{O}$ answer the query of $\mathcal{A}$ with the output of a keyed function F. The strategy of $\mathcal{A}$ consists in answer the queries of $\mathcal{D}$ using $\mathcal{O}$ and emulating $\mathsf{E.Enc}'$. Now, if $\mathcal{A}$ is playing the game PRP^0 then the answers which is giving to $\mathcal{D}$ are the outputs of a random permutation applied on ciphertexts produced by $\mathsf{E.Enc}$, i.e., a random string in C, like in $\mathsf{W\text{-}PRCtG}^0_{\Pi',\mathcal{D}}$. If $\mathcal{A}$ is playing the game PRP^1 then the answers which is giving to $\mathcal{D}$ are the outputs of a keyed function F applied on ciphertexts produced by $\mathsf{E.Enc}$, i.e., its behavior is exactly the one of $\mathsf{E.Enc}'$, like in $\mathsf{W\text{-}PRCtG}^1_{\Pi',\mathcal{D}}$. We can conclude that

$$\mathsf{Adv}^{\mathsf{W\text{-}PRCtG}}_{\mathcal{D},\Pi'}(\lambda) \leq \mathsf{Adv}^{\mathsf{PRP}}_{\mathcal{A}}(\lambda). \qquad \square$$

$\mathcal{A}^{\mathcal{O}}(\lambda)$

1 : $(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$
2 : Run $\mathcal{D}$
3 : **when** $\mathcal{D}$ queries m_i:
4 : $\quad c \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{pk}, m_i)$
5 : $\quad c' = \mathcal{O}(c)$
6 : $\quad$ **reply** to $\mathcal{D}$ with c'
7 : **when** $\mathcal{D}$ returns b': **return** b'

Fig. 17. Distinguher $\mathcal{A}$ for PRP reducing a distinguisher $\mathcal{D}$ for W-PRCtG.

B Supplementary Definitions

B.1 Correctness of Anamorphic Encryption

Let Π be a PKE scheme equipped with an Anamorphic Triplet $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$. The correctness game, for $b \in \{0, 1\}$ and $\mathcal{A}$ a PPT adversary, is defined in Figure 18.

$\mathsf{Cor}^{b}_{\Pi,\Sigma,m}(\mathcal{A})$

1 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda)$
2 : **return** $\mathcal{A}^{\mathcal{O}^b(\mathsf{apk},\mathsf{ask},\mathsf{dk},\mathsf{tk},\cdot)}(\mathsf{apk}, \mathsf{ask})$ where
3 : $\quad \mathcal{O}^0(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}, \widehat{m}) = \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}))$
4 : $\quad \mathcal{O}^1(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}, \widehat{m}) = \widehat{m}$

Fig. 18. Anamorphic Encryption correctness game.

And we define the advantage of an adversary $\mathcal{A}$ in breaking the correctness property as

$$\mathsf{Adv}^{\mathsf{cor}}_{\mathcal{A},\Pi,\Sigma,m}(\lambda) = \left|\Pr\left[\mathsf{Cor}^{0}_{\Pi,\Sigma,m}(\mathcal{A}) = 1\right] - \Pr\left[\mathsf{Cor}^{1}_{\Pi,\Sigma,m}(\mathcal{A}) = 1\right]\right|.$$

Definition 13 (δ-Correctness). *An Anamorphic Encryption scheme Π equipped with Anamorphic Triplet Σ is said to be δ-correct for a negligible $\delta(\lambda)$ if for an arbitrary $m \in M$ and for all PPT adversary $\mathcal{A}$ it holds that*

$$\mathsf{Adv}^{\mathsf{cor}}_{\mathcal{A},\Pi,\Sigma,m}(\lambda) \leq \delta(\lambda).$$

B.2 Robustness of Anamorphic Encryption

Let $\Pi = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ be a PKE scheme equipped with an Anamorphic Triplet $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$. The robustness game, for $b \in \{0, 1\}$ and $\mathcal{A}$ a PPT adversary, is defined in Figure 19.

$\mathsf{Robust}^{b}_{\Pi,\Sigma}(\mathcal{A})$

1 : $((\mathsf{apk},\mathsf{ask}),\mathsf{dk},\mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda)$

2 : **return** $\mathcal{A}^{\mathcal{O}^{b}(\mathsf{apk},\mathsf{ask},\mathsf{tk},\cdot)}(\mathsf{apk},\mathsf{ask})$ where

3 : $\quad \mathcal{O}^{0}(\mathsf{apk},\mathsf{ask},\mathsf{tk},m) = \mathsf{AT.Dec}(\mathsf{ask},\mathsf{tk},\mathsf{E.Enc}(\mathsf{apk},m))$

4 : $\quad \mathcal{O}^{1}(\mathsf{apk},\mathsf{ask},\mathsf{tk},m) = \perp$

Fig. 19. Anamorphic Encryption robustness game.

And we define the advantage of an adversary $\mathcal{A}$ in breaking the robustness property as

$$\mathsf{Adv}^{\mathsf{rob}}_{\mathcal{A},\Pi,\Sigma}(\lambda) = \left|\Pr\left[\mathsf{Robust}^{0}_{\Pi,\Sigma}(\mathcal{A}) = 1\right] - \Pr\left[\mathsf{Robust}^{1}_{\Pi,\Sigma}(\mathcal{A}) = 1\right]\right|.$$

Definition 14 (Robustness). *An Anamorphic Encryption scheme* Π *equppied with Anamorphic Triplet* Σ *is said to be robust if for all* PPT *adversary* $\mathcal{A}$ *it holds that*

$$\mathsf{Adv}^{\mathsf{rob}}_{\mathcal{A},\Pi,\Sigma}(\lambda) \leq \mathsf{negl}(\lambda).$$

B.3 Fully Asymmetric Anamorphic Encryption

Let Π be a PKE scheme equipped with an Anamorphic Triplet $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$. The Fully Asymmetric game, for $b \in \{0,1\}$ and $\mathcal{A}$ a PPT adversary, is defined in Figure 20.

$\mathsf{FAsyAnam\text{-}IND\text{-}CPA}^{b}_{\Sigma}(\mathcal{A})$

1 : $(\mathsf{apk},\mathsf{ask},\mathsf{dk},\mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda)$

2 : $(m_0, m_1, \widehat{m}_0, \widehat{m}_1) \leftarrow^{\$} \mathcal{A}(\mathsf{apk},\mathsf{dk})$

3 : $c \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk},\mathsf{dk}, m_b, \widehat{m}_b)$

4 : **return** $\mathcal{A}(c)$

Fig. 20. Fully Asymmetric Anamorphic Encryption game.

We define the advantage of an adversary $\mathcal{A}$ in breaking the Fully Asymmetric property as

$$\begin{aligned}\mathsf{Adv}^{\mathsf{FAsy\text{-}Anam}}_{\mathcal{A},\Sigma}(\lambda) = \Big|&\Pr\left[\mathsf{FAsyAnam\text{-}IND\text{-}CPA}^{0}_{\Sigma}(\mathcal{A}) = 1\right] \\ &- \Pr\left[\mathsf{FAsyAnam\text{-}IND\text{-}CPA}^{1}_{\Sigma}(\mathcal{A}) = 1\right]\Big|.\end{aligned}$$

Notice that the adversary does not receive any (additional) encryption oracle as having both apk and dk it can create both regular and anamorphic ciphertexts on its own.

Definition 15 (Fully Asymmetric AE). *An Anamorphic Encryption scheme* Π *equppied with Anamorphic Triplet* Σ *is said to be Fully Asymmetric if for every* PPT *adversary* $\mathcal{A}$ *it holds that*

$$\mathsf{Adv}^{\mathsf{FAsy\text{-}Anam}}_{\mathcal{A},\Sigma}(\lambda) \leq \mathsf{negl}(\lambda).$$

C Postponed Proofs

C.1 IND-CPA of Ideal PKE

In this section we prove that the ideal PKE of Section 3.1 is IND-CPA secure.

Theorem 7. *If* $\rho = \Omega(\lambda)$ *and* $|\mathsf{SK}| = \Omega(2^\lambda)$ *then the ideal PKE scheme in Fig. 4 is IND-CPA secure.*

Proof. Given a PPT adversary $\mathcal{A}$, let pk be the chosen public key, m_0, m_1 the plaintexts $\mathcal{A}$ sends to the challenger, and c^* be the challenge ciphertext, i.e. such that $c^* = \psi(\mathsf{pk}, m_b, r^*)$ for $b \sim U(\{0,1\})$ and $r^* \sim U(\{0,1\}^\rho)$. Recall $\mathcal{A}$ can only access the ideal PKE through oracle queries. We define two bad events. BadSK in which $\mathcal{A}$ queries at any point $\mathsf{E.Dec}(\mathsf{sk}, \cdot)$ or $\mathsf{E.Gen}(\lambda; \mathsf{sk})$, i.e. it guesses the secret key correctly. BadRnd in which $\mathcal{A}$ queries at any point $\mathsf{E.Enc}(\mathsf{pk}, \cdot\,; r^*)$, i.e. it guesses the randomness correctly. Calling q the (polynomially bounded) number of total PKE queries performed by $\mathcal{A}$, the following bounds hold for the events above:

Claim 10. *With the previous notation*

$$\Pr[\mathsf{BadSK}] \leq \frac{q}{|\mathsf{SK}| - q} = \mathsf{negl}(\lambda), \qquad \Pr[\mathsf{BadRnd}] \leq \frac{q}{2^\rho - q} = \mathsf{negl}(\lambda).$$

Conditioning on those events not occurring, we show $\mathcal{A}$ has *almost* no information on b, i.e., conditioning on $\neg\mathsf{BadSK} \wedge \neg\mathsf{BadRnd}$ then b is almost uniformly distributed from the point of view of $\mathcal{A}$. The idea is that it might still have queried many encryptions of one messages, and if none of those collided with c^* then he may guess the encrypted message to be the other one. Formally, let View be the view[15] of $\mathcal{A}$ when it halts and $\neg\mathsf{BadSK}$ and $\neg\mathsf{BadRnd}$ occur (excluding c^* from the view). Further call R_0 the set of random coins such that $\mathsf{E.Enc}(\mathsf{pk}, m_0; r)$ was not queried by $\mathcal{A}$ and R_1 the same set but with respect to m_1. Finally, for ease of notation, let us call $f_b(\cdot) = \psi(\mathsf{pk}, m_b; \cdot)$ for $b \in \{0,1\}$. Then conditioning on the view, c^* is uniform over $f_0(R_0) \cup f_1(R_1)$ and $b = 0$ iff $c^* \in f_0(R_0)$. Thus

$$\Pr[b = 0 \,|\, \mathsf{View}] = \Pr[c^* \in f_0(R_0) \,|\, \mathsf{View}] = \frac{|f_0(R_0)|}{|f_0(R_0) \cup f_1(R_1)|} = \frac{|R_0|}{|R_0| + |R_1|}$$

[15] i.e. the joint distribution of $\mathcal{A}$'s input, random coins and oracle replies. Note, oracle *queries* are a deterministic function of the view, and thus need not to be included.

Finally, as $2^\rho \geq |R_b| \geq 2^\rho - q$, we have that

$$\frac{1}{2} - \frac{q}{2^{\rho+1}} \;\leq\; \frac{|R_0|}{|R_0| + |R_1|} \;\leq\; \frac{1}{2} + \frac{q}{2^{\rho+2} - 2q}.$$

The same bounds then applies to the conditional probability that $b = 1$. We can thus conclude that, calling b' the final bit guessed by $\mathcal{A}$

$$\begin{aligned}
\frac{1}{2} \cdot \mathsf{Adv}(\mathcal{A}) \;&=\; \left| \Pr\left[b = b'\right] - \frac{1}{2} \right| \\
&\leq\; \left| \Pr\left[b = b', \neg\mathsf{BadSK}, \neg\mathsf{BadRnd}\right] - \frac{1}{2} \right| + \Pr\left[\mathsf{BadSK}\right] + \Pr\left[\mathsf{BadRnd}\right] \\
&\leq\; \mathsf{negl}(\lambda) + \mathsf{negl}(\lambda) + \mathsf{negl}(\lambda). \qquad \square
\end{aligned}$$

C.2 Markov Lower Bound

Proof of Claim 1. We first show a Markov-type lower bound, that is, given a discrete variable X with support $\Omega \subseteq [0, 1]$ and expectation μ, then for all $\delta \in [0, 1]$ we have

$$\Pr\left[X \geq \delta\right] \geq \mu - \delta.$$

Indeed, dividing Ω in $\Omega^- = \{x \;:\; x < \delta\}$ and $\Omega^+ = \Omega \setminus \Omega^-$, by definition of expectation

$$\begin{aligned}
\mu \;=\; \sum_{x_0 \in \Omega} x_0 \Pr\left[X = x_0\right] \;&=\; \sum_{x_0 \in \Omega^-} x_0 \Pr\left[X = x_0\right] + \sum_{x_0 \in \Omega^+} x_0 \Pr\left[X = x_0\right] \\
&\leq\; \delta \Pr\left[X < \delta\right] + \Pr\left[X \geq \delta\right] \;\leq\; \delta + \Pr\left[X \geq \delta\right]
\end{aligned}$$

where the first inequality follows upper bounding $x_0 \in \Omega^-$ with δ and $x_0 \in \Omega^+$ with 1.

Next we use this Markov-type inequality to prove the claim. In our case the random variable X is such that $X = \Pr\left[c \in S \,|\, m = m_0, \widehat{m} = \widehat{m}_0, s = s_0\right]$ with probability $\Pr\left[m = m_0, \widehat{m} = \widehat{m}_0, s = s_0\right]$ for all $m_0, \widehat{m}_0, s_0$. Then is easy to see that X has average $\Pr\left[c \in S\right]$ and that it is contained in $[0, 1]$. Moreover $\Pr\left[(m, \widehat{m}, s) \in V_\delta\right] = \Pr\left[X \geq \delta\right]$. We thus conclude that

$$\Pr\left[(m, \widehat{m}, s) \in V_\delta\right] \;=\; \Pr\left[X \geq \delta\right] \;\geq\; \mu - \delta \;=\; \frac{1}{2} \cdot \Pr\left[c \in S\right]. \qquad \square$$

C.3 Symmetric Choice Functions

Proof of Lemma 4. Let $n = |X|$ and $P(x_1, \ldots, x_k) = \Pr\left[f(x_1, \ldots, x_k) = x_1\right]$. By definition of choice function f has to return one of its arguments, meaning that for $x_1, \ldots, x_k$ all distinct

$$P(x_1, \ldots, x_k) + P(x_2, \ldots, x_k, x_1) + \ldots + P(x_k, x_1, \ldots, x_{k-1}) \;=\; 1.$$

As a first step we state some properties of P.

Claim 11. *The following bounds for the sum of P over X^k holds:*

$$\sum\nolimits_{\mathbf{x}} P(\mathbf{x}) \leq n^k, \qquad\qquad \sum\nolimits_{\mathbf{x}} P(\mathbf{x}) \geq \frac{n^k}{k} - kn^{k-1}.$$

Next we study the distribution of $z = f(\mathbf{u})$.

Claim 12. *For all $a \in X$, $\Pr[z = a] \geq \left(\frac{k}{n^k} \sum_{\mathbf{x}} P(a, \mathbf{x})\right) - \frac{k^3}{n^2}$.*

Using both claim, the theorem's proof follows as

$$\begin{aligned}
\Pr[f(z, \mathbf{v}) = z] &= \sum\nolimits_{\mathbf{y}} \frac{1}{n^{k-1}} \cdot \Pr[f(z, \mathbf{y}) = z] \\
&= \frac{1}{n^{k-1}} \sum\nolimits_{a,\mathbf{y}} \Pr[z = a] \Pr[f(a, \mathbf{y}) = a] \\
&\geq \frac{1}{n^{k-1}} \sum\nolimits_{a,\mathbf{y}} \left(\sum\nolimits_{\mathbf{x}} \frac{k}{n^k} P(a, \mathbf{x}) - \frac{k^3}{n^2} \right) P(a, \mathbf{y}) \\
&= \frac{k}{n^{2k-1}} \sum\nolimits_{a,\mathbf{y},\mathbf{x}} P(a, \mathbf{x}) P(a, \mathbf{y}) - \frac{k^3}{n^{k+1}} \sum\nolimits_{a,\mathbf{y}} P(a, \mathbf{y}) \\
&\geq \frac{k}{n^{2k-1}} \sum\nolimits_{a} \left(\sum\nolimits_{\mathbf{x}} P(a, \mathbf{x}) \right)^2 - \frac{k^3}{n} \\
&\geq \frac{k}{n^{2k-1}} \cdot \frac{1}{n} \left(\frac{n^k}{k} - k \cdot n^{k-1} \right)^2 - O(n^{-1}) \\
&= \frac{k}{n^{2k}} \cdot \left(\frac{n^{2k}}{k^2} + (n^{k-1}k)^2 - 2n^{2k-1} \right) - O(n^{-1}) \\
&= \frac{k}{n^{2k}} \cdot \frac{n^{2k}}{k^2} - O(n^{-1}) = \frac{1}{k} - O(n^{-1}).
\end{aligned}$$

Where the first inequality follows by Claim 12, the second one applying Claim 11 on the second term. The third inequality follows from AM-QM where, calling $s(a) = \sum_{\mathbf{x}} P(a, \mathbf{x})$, the sum of $s(a)$ coincides with the sum of P over X^k, and is therefore lower bounded as per Claim 11.

Proof of Claim 11. The first part is trivial as $P(\mathbf{x}) \leq 1$. For the second part let $S = \{(x_1, \ldots, x_k) \in X^k \ : \ \forall i, j (x_i \neq x_j)\}$. The size of $X^k \setminus S$ is smaller than $\binom{k}{2} \cdot n^{k-1}$, as it is a union of the $\binom{k}{2}$ sets $D_{i,j}$ containing all vectors $\mathbf{x}$ with $x_i = x_j$ (so that $|D_{i,j}| = n^{k-1}$). As a consequence then $|S| \geq n^k - \binom{k}{2} n^{k-1}$.

Next we can partition S into a collection $\mathcal{P}$ of $|S|/k$ classes of size k, each containing the cyclic shift of a vector $\mathbf{x} \in S$. Formally

$$[(x_1, \ldots, x_k)] \coloneqq \{(x_{1+i}, \ldots, x_{k+i}) \ : \ i \in \mathbb{Z}/k\mathbb{Z}\}$$

note that the vectors in S have entries that are all distinct, so each such cyclic shift produces a different vector. Moreover, as observed previously, the sum of $P(\mathbf{x})$ for $\mathbf{x} \in [\mathbf{x}]$ equals 1, as the choice function must return one of its entries. We thus conclude that

$$\sum\nolimits_{\mathbf{x} \in X^k} P(\mathbf{x}) \ \geq \ \sum\nolimits_{\mathbf{x} \in S} P(\mathbf{x}) \ = \ \frac{|S|}{k} \ \geq \ \frac{n^k}{k} - \binom{k}{2} \frac{n^{k-1}}{k} \ \geq \ \frac{n^k}{k} - kn^{k-1}. \ \square$$

Proof of Claim 12. Let $S = \{(x_2, \ldots, x_k) \in X^{k-1} : \forall i, j (x_i \neq a,\ x_i \neq x_j)\}$. To lower bound its size let D_i the set of points in X^{k-1} with i-th coordinate equal to a and $D_{i,j}$ the subset of X^{k-1} with $x_i = x_j$. Then[16]

$$|X^{k-1} \setminus S| = \left| \bigcup_{i=2}^{k} D_i \cup \bigcup_{i<j} D_{i,j} \right| \leq kn^{k-2} + \binom{k-1}{2} n^{k-2} \leq k^2 n^{k-2}.$$

Thus $|S| \geq n^{k-1} - k^2 n^{k-2}$. We can finally lower bound the probability that $z = a$ as

$$\begin{aligned} \Pr[z = a] \ &\geq\ k \sum_{\mathbf{x} \in S} P(a, \mathbf{x}) \frac{1}{n^k} \ \geq\ \frac{k}{n^k} \sum_{\mathbf{x} \in X^{k-1}} P(a, \mathbf{x}) \ -\ \frac{k^3}{n^k} \sum_{\mathbf{x} \in X^{k-1} \setminus S} P(a, \mathbf{x}) \\ &\geq\ \frac{k}{n^k} \sum_{\mathbf{x} \in X^{k-1}} P(a, \mathbf{x}) \ -\ \frac{k^3}{n^2}. \end{aligned}$$

The first bound follows by restricting all components of $\mathbf{u}$ to be different, lower bounding the probability of this not happening with 0, and later, as $z = a \Rightarrow a \in \{u_1, \ldots, u_k\}$, grouping all vectors shifting the (only) entry equal to a in the first position (meaning that each term $P(a, \mathbf{x})$ is repeated k times). □

□

C.4 Communication Bound Proof

Proof of Claim 3. We prove the claim through a sequence of hybrid distributions $V_0, \ldots, V_4$. Recall $\xi^* : \mathsf{SK} \times \{0,1\}^\ell \to M \cup \{\perp\}$ is a biased random function such that $\xi^*(\mathsf{sk}, c) = m_0$ with probability $2^{\rho-\ell}$ for all $m_0 \in M$. Moreover $\psi^* : \mathsf{PK} \times M \times \{0,1\}^\rho \to \{0,1\}^\ell$ is a truly random function.

V_0: The real view $\mathsf{View}^{\mathsf{real}}$.

V_1: As V_0 but queries to $\mathsf{E.Dec}(\mathsf{sk}, c)$ are replied with m if $c = \mathsf{E.Enc}(\mathsf{pk}, m; r)$ was previously obtained where $\mathsf{pk} = \phi(\mathsf{sk})$, or with $\xi^*(\mathsf{sk}, c)$ otherwise. Moreover queries to $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ are replied with $\psi^*(\mathsf{pk}, m, r)$.

V_2: As V_1, but during the execution of $\mathsf{AT.Dec}$, the query $\mathsf{E.Dec}(\mathsf{ask}, c^*)$ always returns m^*.

V_3: As V_2, but while executing $\mathsf{AT.Dec}$, the query $\mathsf{E.Dec}(\mathsf{sk}, c)$ is answered with
- c^* if $(\mathsf{sk}, c) = (\mathsf{ask}, c^*)$.
- m if $\mathsf{AT.Gen}$ or $\mathsf{AT.Dec}$ already got $c = \mathsf{E.Enc}(\mathsf{pk}, m; r)$ with $\mathsf{pk} = \phi(\mathsf{sk})$.
- $\xi^*(\mathsf{sk}, c)$ otherwise.

V_4: The simulated view $\mathsf{View}^{\mathsf{sim}}$.

The proof will follow showing the statistical distance between every two consecutive distributions is negligible (denoted with $V_i \approx V_{i+1}$). To fix notation $V_{i,n}$ represents the first n replies observed in V_i while q denote the maximum number of queries, so that $V_i = V_{i,q}$.

[16] Here we assume $\binom{n}{m} = 0$ when $n < m$.

$V_0 \approx V_1$. We prove by induction that

$$\Delta(V_{0,n}, V_{1,n}) \leq 2n \cdot \frac{q}{2^\rho}.$$

The base case is trivial. Assuming this to hold for n, we study the $(n+1)$-th query in both distributions, conditioning on $V_{0,n} = v = V_{1,n}$ a given view. If this query is $\mathsf{E.Gen}(\lambda; \mathsf{sk})$, the reply is identically distributed in both executions.

If next query is $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ and this was already asked the reply remains consistent. Else, let C be the set of ciphertexts either obtained through encryption queries or appearing in decryption ones with $\mathsf{E.Dec}(\mathsf{sk}, c) \neq m$. D is the set of ciphertexts such that $\mathsf{E.Dec}(\mathsf{sk}, c) = m$ was previously observed. R is the set of randomness r such that $\mathsf{E.Enc}(\mathsf{pk}, m, r)$ was ask before. Let c, c' be the replies in V_0, V_1 respectively. We study the probability of $\Pr[c = c_0 \mid V_{0,n} = v]$:

- If $c_0 \in C$ then $\Pr[c = c_0] = 0$, as we assumed the query to be different from previous ones.
- If $c_0 \in D$ then $c = c_0$ if the queried randomness matches the one such that $c_0 = \psi(\mathsf{pk}, m, r_0)$. Due to the distribution of ψ, such r_0 is uniform over $\{0,1\}^\rho \setminus R$, therefore

$$\Pr[c = c_0 \mid V_{0,n} = v] \;=\; \frac{1}{|\{0,1\}^\rho \setminus R|} \;\leq\; \frac{1}{2^\rho - q}.$$

  Where the inequality follows as $|R| \leq q$, as each query increases the size of R by at most one.
- If $c_0 \notin (C \cup D)$, since conditioning on $c \notin D$ implies that c is uniform over $\{0,1\}^\ell \setminus (D \cup C)$, we have that

$$\begin{aligned}\Pr[c = c_0 \mid V_{0,n} = v] \;&=\; \Pr[c = c_0 \mid c \notin D,\ V_{0,n} = v] \cdot \Pr[c \notin D,\ V_{0,n} = v] \\ &=\; \frac{1}{|\{0,1\}^\ell \setminus (C \cup D)|} \cdot \left(1 - \frac{|D|}{|\{0,1\}^\rho \setminus R|}\right)\end{aligned}$$

  Using again the fact that the size of $(C \cup D)$ and R is at most q it can then be easily shown that

$$\begin{aligned}\Pr[c = c_0 \mid V_{0,n} = v] \;&\leq\; \frac{1}{2^\ell - q} \;=\; \frac{1}{2^\ell} + \frac{q}{2^\ell(2^\ell - q)} \\ \Pr[c = c_0 \mid V_{0,n} = v] \;&\geq\; \frac{1}{2^\ell} \cdot \left(1 - \frac{q}{2^\rho - q}\right) \;=\; \frac{1}{2^\ell} - \frac{q}{2^\ell(2^\rho - q)}.\end{aligned}$$

Finally, $c' = \psi^*(\mathsf{pk}, m, r)$ is uniform over $\{0,1\}^\ell$. Thus the statistical distance of c, c' conditioning on $V_{0,n} = v = V_{1,n}$ can be bounded as:

$$\begin{aligned}
\Delta(c_{|V_{0,n}=v},\ c'_{|V_{1,n}=v}) \ &= \ \frac{1}{2}\sum_{c_0} \left|\Pr\left[c = c_0 \,|\, V_{0,n} = v\right] - \Pr\left[c' = c_0 \,|\, V_{1,n} = v\right]\right| \\
&= \ \frac{1}{2}\sum_{c_0} \left|\Pr\left[c = c_0 \,|\, V_{0,n} = v\right] - \frac{1}{2^\ell}\right| \\
&\leq \ \frac{1}{2}\sum_{c_0 \in C} \frac{1}{2^\ell} \ + \ \frac{1}{2}\sum_{c_0 \in D} \frac{1}{2^\rho - q} \ + \ \frac{1}{2}\sum_{c_0 \notin C \cup D} \frac{q}{2^\ell(2^\rho - q)} \\
&\leq \ \frac{1}{2}\left(\frac{q}{2^\ell} + \frac{q}{2^\rho - q} + \frac{q}{2^\rho - q}\right) \ \leq \ 2 \cdot \frac{q}{2^\rho}.
\end{aligned}$$

Where the first inequality follows as $c \notin C$ for the first term, because $1/(2^\rho - q)$ is always greater than $2^{-\ell}$ for the second term, and as the distance between the conditional probability of $c = c_0$ from $2^{-\ell}$ when $c_0 \notin C$ was previously upper-bounded by $1/(2^\ell(2^\rho - q))$ for the third term. The second inequality again uses the fact that $C \cup D$ has size at most q, and the last one holds asymptotically given q polynomially bounded, and $\ell - \rho = \Omega(\lambda)$. This suffices to prove the inductive step for the encryption query case.

Lastly, if next query is $\mathsf{E.Dec}(\mathsf{sk}, c)$, if this was previously queried or $c = \mathsf{E.Enc}(\mathsf{pk}, m; r)$ was previously observed, the reply is identical in both distributions. Otherwise, let m, m' be the replies in V_0, V_1 respectively. By the definition of $\xi^*(\mathsf{sk}, c)$, for all $m_0 \in M$

$$\Pr\left[m' = m_0\right] \ = \ \frac{2^\rho}{2^\ell}.$$

Regarding m, for each m_0 let C_{pk} be the set of ciphertext computed with pk or involved in a decryption query with $\mathsf{sk} = \phi^{-1}(\mathsf{pk})$. Further let $C(m_0)$ the set of valid encryption of m_0 under pk, i.e. $C(m_0) = \{\psi(\mathsf{pk}, m_0, r) \ : \ r \in \{0,1\}^\rho\}$. Conditioning on previous queries, $C(m_0) \setminus C_{\mathsf{pk}}$ is uniform over $\{0,1\}^\ell \setminus C_{\mathsf{pk}}$, thus

$$\Pr\left[m = m_0 \,|\, V_{0,n} = v\right] \ = \ \Pr\left[c \in C(m_0)\right] \ = \ \frac{|C(m_0) \setminus C_{\mathsf{pk}}|}{|\{0,1\}^\ell \setminus C_{\mathsf{pk}}|}.$$

From this expression, using the fact that C_{pk} has size smaller than q and $|C(m_0)| = 2^\rho$, we can bound the distance of the above probability from $2^{\ell-\rho}$ in absolute value:

$$\begin{aligned}
\Pr\left[m = m_0 \,|\, V_{0,n} = n\right] \ &\leq \ \frac{2^\rho}{2^\ell - q} \ \leq \ \frac{2^\rho}{2^\ell} + \frac{2^\rho}{2^\ell - q} \cdot \frac{q}{2^\ell} \\
\Pr\left[m = m_0 \,|\, V_{0,n} = n\right] \ &\geq \ \frac{2^\rho - q}{2^\ell} \ \geq \ \frac{2^\rho}{2^\ell} - \frac{q}{2^\ell}.
\end{aligned}$$

This implies as noted that the distance from the same event in V_1 is bounded by $q/2^\ell$, i.e.

$$\left|\Pr\left[m = m_0 \,|\, V_{0,n} = v\right] - \Pr\left[m' = m_0 \,|\, V_{1,n} = v\right]\right| \ \leq \ \frac{q}{2^\ell}.$$

where the inequality hold asymptotically if q is polynomially bounded and $\ell-\rho = \Omega(\lambda)$. The same bound can be shown for the remaining case $m = \perp$. Indeed V_1 returns a decryption error with probability $1 - 2^{\rho+\mu-\ell}$. In V_0 instead, let C_{pk} be as before and $C(\perp)$ be the set of invalid ciphertext under key pk. Then as before $C(\perp) \setminus C_{\mathsf{pk}}$ is uniform over $\{0,1\}^\ell \setminus C_{\mathsf{pk}}$. Therefore

$$\Pr\left[m = \perp \mid V_{0,n} = v\right] \;=\; \frac{|C(\perp) \setminus C_{\mathsf{pk}}|}{|\{0,1\}^\ell \setminus C_{\mathsf{pk}}|}.$$

Using the fact that $|C(\perp)| = 2^\ell - 2^{\rho+\mu}$, the distance of the probability above from the one measured in V_1 we can bound as

$$\begin{aligned}
\Pr\left[m = \perp \mid V_{0,n} = v\right] \;&\leq\; \frac{2^\ell - 2^{\rho+\mu}}{2^\ell - q} \;\leq\; \left(1 - \frac{2^{\rho+\mu}}{2^\ell}\right) + \frac{q}{2^\ell} \cdot \frac{2^\ell - 2^{\rho+\mu}}{2^\ell - q} \\
\Pr\left[m = \perp \mid V_{0,n} = v\right] \;&\geq\; \frac{2^\ell - 2^{\rho+\mu} - q}{2^\ell} \;=\; \left(1 - \frac{2^{\rho+\mu}}{2^\ell}\right) + \frac{q}{2^\ell}.
\end{aligned}$$

Hence the probability of the events $m = \perp$ and $m' = \perp$ given the previous queries have distance smaller that $q \cdot 2^\ell$. Combining the provided inequalities yields a bound on the conditional statistical distance

$$\begin{aligned}
&\Delta(m_{|V_{0,n}=v}, m'_{|V_{1,n}=v}) \;= \\
&= \frac{1}{2} \sum_{m_0 \in M \cup \{\perp\}} \left|\Pr\left[m = m_0 \mid V_{0,n} = v\right] - \Pr\left[m' = m_0 \mid V_{1,n} = v\right]\right| \\
&\leq \frac{1}{2} \sum_{m_0 \in M \cup \{\perp\}} \frac{q}{2^\ell} \;=\; \frac{1}{2} \cdot \frac{(2^\mu + 1)q}{2^\ell} \;\leq\; \frac{q}{2^\rho}.
\end{aligned}$$

where the last inequality holds asymptotically as $\ell \geq \mu + \rho$, q is polynomially bounded. This suffices to imply the inductive case and, as we exhausted the three query types, it also conclude the proof for $\Delta(V_0, V_1) \leq \mathsf{negl}(\lambda)$.

$V_1 \approx V_2$: The only difference in the two worlds is the reply to $\mathsf{E.Dec}(\mathsf{ask}, c^*)$ provided during the execution of $\mathsf{AT.Dec}$. To prove V_1, V_2 have low statistical distance, it suffices to show that the event $\mathsf{E.Dec}(\mathsf{ask}, c^*) \neq m^*$ occurs only with negligible probability. This is true due in V_0, where queries to the underlying PKE are answered correctly, due to the security notion for anamorphic encryption.

Indeed, one can define an adversary $\mathcal{A}(\mathsf{apk}, \mathsf{ask})$ which initially samples a random messages pair $(m', \widehat{m})$, queries its encryption $c' \leftarrow \mathcal{O}(m', \widehat{m})$ and checks that $m' = \mathsf{E.Dec}(\mathsf{ask}, c')$. If $\mathcal{O}$ produced c' with $\mathsf{E.Enc}$ then the condition $\mathcal{A}$ checks is always verified. Hence in V_0

$$\Pr\left[\mathsf{E.Dec}(\mathsf{ask}, c') \neq m'\right] \;=\; \mathsf{Adv}^{\mathsf{anam}}(\mathcal{A}) \;=\; \varepsilon(\lambda)$$

for a negligible $\varepsilon(\lambda)$. Calling Bad the event $\mathsf{E.Dec}(\mathsf{ask}, c^*) \neq m^*$, then $\Pr[\mathsf{Bad}] \leq \varepsilon(\lambda) + \Delta(V_0, V_1)$ and in particular

$$\begin{aligned}\Delta(V_1, V_2) &\leq \Pr[\mathsf{Bad}]\,\Delta(V_{1|\mathsf{Bad}}, V_{2|\mathsf{Bad}}) + \Pr[\neg\mathsf{Bad}]\,\Delta(V_{1|\neg\mathsf{Bad}}, V_{2|\neg\mathsf{Bad}}) \\ &\leq \Pr[\mathsf{Bad}] + \Delta(V_{1|\neg\mathsf{Bad}}, V_{2|\neg\mathsf{Bad}}) \\ &\leq \varepsilon(\lambda) + \Delta(V_0, V_1).\end{aligned}$$

Where in the last inequality we used the fact that, conditioning on $\neg\mathsf{Bad}$ the two distributions are identical.

$V_2 \approx V_3$: . The main difference between V_2 and V_3 is that in the latter, decryption call $\mathsf{E.Dec}(\mathsf{sk}, c)$ do not depends on encryption queries of $\mathsf{AT.Enc}$. In particular, if $\mathsf{AT.Dec}$ were to query the decryption of a ciphertext only computed by $\mathsf{AT.Enc}$, the reply in V_2 would by construction return the encrypted message, while in V_3 it would be $\xi^*(\mathsf{sk}, c)$. To show $V_2 \approx V_3$ we prove the above event occurs with negligible probability in both distributions.

Let $c_1, \ldots, c_h$ be the ciphertext obtained by $\mathsf{AT.Enc}$, W_i the replies to PKE queries performed only by $\mathsf{AT.Gen}$ and $\mathsf{AT.Dec}$ in V_i (for $i \in \{1, ..., 4\}$), and $W_{i,n}$ the same subsequence of $V_{1,n}$. Finally let Coll_n the event that in the n-th query, $\mathsf{AT.Dec}$ queries $\mathsf{E.Dec}(\cdot, c)$ such that $c \in \{c_1, \ldots, c_h\} \setminus \{c^*\}$.

First we determine the random variables c is a function of $\mathsf{AT.Dec}$'s query only depend on its input $(\mathsf{ask}, \mathsf{tk}, c^*)$ and its view, and in turns $(\mathsf{ask}, \mathsf{tk})$ is a deterministic function of s^*, i.e. $\mathsf{AT.Gen}$'s random tape, and $\mathsf{AT.Gen}$'s view. Thus c is a function of $W_{i,n}, s^*, c^*$.

Next, for all j such that $c_j \neq c^*$, we study the min-entropy of c^*. Both in V_2 and V_3, as c_j was by definition not obtained from encryption queries performed by $\mathsf{AT.Gen}$ and $\mathsf{AT.Dec}$, c_j is uniformly random and independent from W_n, s^*. It may however share mutual information with c^*, which by the Ciphertext Selection Lemma (Lemma 3), is chosen among $c_1, \ldots, c_h$. Let $I \sim \{1, \ldots, h\}$ be a random variable denoting the index of such choice, i.e. such that $c^* = c_I$. Then the min-entropy of $c_j \neq c^*$ given $\mathsf{AT.Dec}$'s information can be bounded as

$$\begin{aligned}\mathrm{H}_\infty(c_j \,|\, W_{i,n}, s^*, c^*) &= \mathrm{H}_\infty(c_j \,|\, W_{i,n}, s^*, c_I) \\ &= \mathrm{H}_\infty(c_j \,|\, W_{i,n}, s^*, (c_1, \ldots, c_{j-1}, c_{j+1}, \ldots, c_n), I) \\ &\geq \mathrm{H}_\infty(c_j \,|\, W_{i,n}, s^*, (c_1, \ldots, c_{j-1}, c_{j+1}, \ldots, c_n)) - \log_2 h \\ &\geq \mathrm{H}_\infty(c_j) - \log_2 q \;=\; \ell - \log_2 q.\end{aligned}$$

Where the first inequality follows as $I \in \{1, \ldots, h\}$ and the second one as $h \leq q$. Hence $\Pr[c = c_j] \leq q \cdot 2^{-\ell}$, and, by a union bound, $\Pr[\mathsf{Coll}_n] \leq hq \cdot 2^{-\ell} \leq q^2 \cdot 2^{-\ell}$.

Finally, as $\mathsf{AT.Gen}$ performs at most q decryption queries, the probability that $\exists n : \mathsf{Coll}_n$ is, again from a union bound, smaller than $q^3 \cdot 2^{-\ell}$. This concludes the proof as, conditioning on $\nexists n : \mathsf{Coll}_n$, the two distributions V_2, V_3 are identical.

$V_3 = V_4$: Follows by inspection as:

– Encryption queries for $\mathsf{AT.Gen}$ are replied with ψ^*, while for $\mathsf{AT.Enc}, \mathsf{AT.Dec}$, the RO H is used, keeping however consistency with previous queries performed by $\mathsf{AT.Gen}$. Hence every new query is always uniformly distributed over $\{0,1\}^\ell$ – as specified in V_1.
– $\mathcal{R}$ programs $\xi^*(\mathsf{ask}, c^*) = m^*$, see Line 7 of Figure 8. In particular for $\mathsf{AT.Dec}$, the query $\mathsf{E.Dec}(\mathsf{ask}, c^*)$ always returns m^*, as specified in V_2.
– When $\mathsf{AT.Gen}$ queries $\mathsf{E.Dec}(\mathsf{sk}, c)$, with $(\mathsf{sk}, c) \neq (\mathsf{ask}, c^*)$, then the output is m if $\mathsf{E.Enc}(\phi^*(\mathsf{sk}), m, r) = c$ was previously obtained by $\mathsf{AT.Gen}$ or $\mathsf{AT.Dec}$, and the unprogrammed value of $\xi^*(\mathsf{sk}, c)$ otherwise. In particular the output does not depend on $\mathsf{E.Enc}$'s queries, as specified in V_3.

The proof of Claim 3 is therefore completed. □

C.5 Impossibility of Weak Asymmetric AE Proof

Proof of Claim 4. The proof is divided in two parts. First we show that "programming" a ciphertext previously queried by $\mathsf{AT.Gen}$ is unlikely, and then prove the bound studying the distribution of c^* and $\widetilde{c}$, with $\widetilde{c}$ being the correct ciphertext returned in Real.

Rewriting Probability. To fix some notation let $r_1, \ldots, r_q$ be the randomness used by $\mathsf{AT.Gen}$ in queries of the form $\mathsf{E.Enc}(\mathsf{apk}, m^*; r_i)$. S_j is the same set relative to the queries of $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$ in the preprocessing phase, while S is again the same set for the last execution of $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$ (assuming though that $\mathcal{A}$ replies with the correct ciphertext instead of c^*). With this notation then $R = S_1 \cup \ldots \cup S_\vartheta$, as in the definition of $\mathcal{A}$. The bad event we wish to bound the probability of is $\mathsf{Rew} = \exists r_i \in S \setminus R$. Note that upon conditioning on the input $\mathsf{in} = (\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0) = \mathsf{in}_0$ we have that $S_1, \ldots, S_\vartheta, S$ are independent and equally distributed. Finally, for each r_i we call $p_i(\mathsf{in}_0) = \Pr[r_i \in S_i | \mathsf{in} = \mathsf{in}_0]$. Then

$$
\begin{aligned}
\Pr[\mathsf{Rew}] &= \sum\nolimits_{\mathsf{in}_0} \Pr[\mathsf{Rew} \mid \mathsf{in} = \mathsf{in}_0] \Pr[\mathsf{in} = \mathsf{in}_0] \\
&= \sum\nolimits_{\mathsf{in}_0} \Pr[\exists r_i \in S \setminus (S_1 \cup \ldots \cup S_\vartheta) \mid \mathsf{in} = \mathsf{in}_0] \Pr[\mathsf{in} = \mathsf{in}_0] \\
&\leq \sum\nolimits_{\mathsf{in}_0} \sum\nolimits_{i=1}^{q} \Pr[r_i \in S \setminus (S_1 \cup \ldots \cup S_\vartheta) \mid \mathsf{in} = \mathsf{in}_0] \Pr[\mathsf{in} = \mathsf{in}_0] \\
&= \sum\nolimits_{\mathsf{in}_0} \sum\nolimits_{i=1}^{q} p_i(\mathsf{in}_0) \left(1 - p_i(\mathsf{in}_0)\right)^\vartheta \Pr[\mathsf{in} = \mathsf{in}_0] \\
&\leq \sum\nolimits_{\mathsf{in}_0} \sum\nolimits_{i=1}^{q} \frac{1}{\vartheta + 1} \Pr[\mathsf{in} = \mathsf{in}_0] \;\leq\; \frac{q}{\vartheta + 1}.
\end{aligned}
$$

Where the first inequality is a union bound and the second one follows as $p_i(\mathsf{in}_0) \in [0,1]$.

Predicate Probability. Let $\widetilde{c}$ be the correct reply $\mathcal{A}$ should have given to $\mathsf{AT.Enc}$ on Figure 9, line 11, i.e. $\widetilde{c} = \mathsf{E.Enc}(\mathsf{apk}, m^*; r)$. Further call vb and vr be the

vectors obtained removing c^* and $\widetilde{c}$ respectively from Bias and Real. Then, up to rearranging, $\mathsf{Bias} = (\mathsf{vb}, c^*)$ and $\mathsf{Real} = (\mathsf{vr}, \widetilde{c})$.

We begin studying $\widetilde{c}$. For any (partial) view $\mathsf{rv} = v$, let us call C_v the set of ciphertext observed in the given view. Then, calling $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ the query $\mathsf{AT.Enc}$ performed to get $\widetilde{c}$, either r was queried by $\mathsf{AT.Gen}$ or the query is performed for the first time (or else $\mathcal{A}$ would not "try" to program this query). Note that conditioning on $\neg\mathsf{Rew}$ the first events never occurs. In the second case instead, we can prove as in the proof of Claim 3 that $\Pr[c \in C_v \mid \mathsf{rv} = v,\ \neg\mathsf{Rew}] \leq q/(2^\rho - q)$. Furthermore, conditioning again on $\mathsf{rv} = v$ and $\neg\mathsf{Rew}$, the ciphertext $\widetilde{c} \sim U(\{0,1\}^\ell \setminus C_v)$. Thus, for all $c_0 \notin C_v$

$$\begin{aligned}
&\Pr[\widetilde{c} = c_0 \mid \neg\mathsf{Rew}, \mathsf{rv} = v] \ = \\
&= \Pr[\widetilde{c} = c_0 \mid \widetilde{c} \notin C, \neg\mathsf{Rew}, \mathsf{rv} = v] \cdot \Pr[\widetilde{c} \notin C \mid \neg\mathsf{Rew}, \mathsf{rv} = v] \\
&\geq \frac{1}{2^\ell} \cdot \left(1 - \frac{q}{2^\rho - q}\right) \ \geq \ \frac{1}{2^\ell} - \frac{q}{2^\ell(2^\rho - q)}.
\end{aligned}$$

In particular then the probability of getting $\widetilde{c} = c_0$ given view v is larger than

$$\begin{aligned}
\Pr[\widetilde{c} = c_0 \mid \mathsf{rv} = v] \ &\geq \ \Pr[\widetilde{c} = c_0 \mid \mathsf{rv} = v, \neg\mathsf{Rew}] - \Pr[\widetilde{c} = c_0, \mathsf{Rew} \mid \mathsf{rv} = v] \\
&\geq \ \frac{1}{2^\ell} - \frac{q}{2^\ell(2^\rho - q)} - \Pr[\widetilde{c} = c_0, \mathsf{Rew} \mid \mathsf{rv} = v].
\end{aligned}$$

Next we focus on c^*. To study its distribution, let $c_1, \ldots, c_q$ be the ciphertext queried by $\mathsf{AT.Enc}$ whose output is c^*. Then by the ciphertext selection lemma, the event $c^* \notin \{c_1, \ldots, c_q\}$ occurs with negligible probability. Hence $\Pr[p(\mathsf{Bias})] =$

$$\begin{aligned}
&= \sum_{v_0} \sum_{c_0} \Pr[p(v_0, c_0)] \Pr[c^* = c_0,\ \mathsf{bv} = v_0] \\
&\leq \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \Pr[p(v_0, c_0)] \Pr[c^* = c_0,\ \mathsf{bv} = v_0] + \Pr[c^* \in C_{v_0}] \\
&\leq \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \Pr[p(v_0, c_0)] \Pr\left[\begin{matrix} c^* = c_0,\ \mathsf{bv} = v_0 \\ c^* \in \{c_i\}_{i=1}^q \end{matrix}\right] + \Pr[c^* \notin \{c_i\}_{i=1}^q] + \mathsf{negl}(\lambda) \\
&\leq \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \Pr[p(v_0, c_0)] \Pr[c_0 \in \{c_i\}_{i=1}^q,\ \mathsf{bv} = v_0] \ + \ \mathsf{negl}(\lambda)
\end{aligned}$$

Where the first inequality follows removing the terms with $c_0 \in C_{v_0}$, the second one as $\Pr[c^* \notin C_{v_0}]$ is negligible by Lemma 2, the third follows from the Ciphertext Selection Lemma for the second term and because $c^* = c_0$ and $c^* \in \{c_i\}_{i=1}^q$ implies $c_0 \in \{c_i\}_{i=1}^q$ for the first term. We then continue the chain of inequalities

with a union bound:

$$
\begin{aligned}
&\leq \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \sum_{i=1}^{q} \Pr\left[p(v_0, c_0)\right] \Pr\left[c_i = c_0 \mid \mathsf{bv} = v_0\right] \Pr\left[\mathsf{bv} = v_0\right] \; + \; \mathsf{negl}(\lambda) \\
&\leq \; q \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \Pr\left[p(v_0, c_0)\right] \frac{1}{2^\ell - q} \Pr\left[\mathsf{bv} = v_0\right] + \mathsf{negl}(\lambda) \\
&\leq \; q \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \Pr\left[p(v_0, c_0)\right] \frac{1}{2^\ell} \Pr\left[\mathsf{bv} = v_0\right] + \frac{q}{2^\ell - q} + \mathsf{negl}(\lambda) \\
&\leq \; q \sum_{v_0} \sum_{c_0 \notin C_{v_0}} \Pr\left[p(v_0, c_0)\right] \Pr\left[\widetilde{c} = c_0, \mathsf{rv} = v_0\right] + \frac{q^2}{2^\rho - q} + q \Pr\left[\mathsf{Rew}\right] + \mathsf{negl}(\lambda) \\
&\leq \; q \sum_{v_0} \sum_{c_0} \Pr\left[p(v_0, c_0)\right] \Pr\left[\widetilde{c} = c_0, \mathsf{rv} = v_0\right] + q \Pr\left[\mathsf{Rew}\right] + \mathsf{negl}(\lambda) \\
&\leq \; q \Pr\left[p(\mathsf{Real})\right] + \frac{q^2}{\vartheta + 1} + \mathsf{negl}(\lambda)
\end{aligned}
$$

Where the second inequality follows as, conditioning on $c_i \notin C_{v_0}$ we have that $c_i \sim U(\{0,1\}^\ell \setminus C_{v_0})$ given the view, where $|C_{v_0}| \leq q$. The fourth follows observing that bv and rv are identically distributed, and from the bound we previously found on $\Pr\left[\widetilde{c} = c_0 \mid \mathsf{rv} = v_0\right]$ applied on $2^{-\ell}$. The fifth by summing over a domain non-negative terms. The claim is therefore proven. □

Proof of Claim 5. The event $\nexists r^* \in S \setminus R : c^* = \mathsf{E.Enc}(\mathsf{apk}, m^*; r')$ is equivalent to requiring that either $c^* \notin C^{\mathsf{Enc}}_{\mathsf{in}}$, with in being $\mathsf{AT.Enc}$'s input, or $c^* \in \{\mathsf{E.Enc}(\mathsf{apk}, m^*; r) \; : \; r \in R\} = C_R$. Note C_R has polynomially bounded size (in particular $|C_R| \leq q\vartheta$) and its distribution is independent from the random coins used to generate c^*. We can thus use Lemma 3 and Lemma 2 we conclude that

$$\Pr\left[\mathsf{BadChoice}^*\right] \; \leq \; \Pr\left[c' \notin C^{\mathsf{Enc}}_{\mathsf{in}}\right] + \Pr\left[c' \in C_R\right] \; \leq \; \mathsf{negl}(\lambda).$$

The result is analogous for c' up to using Claim 4. □

Proof of Claim 6. $\mathcal{F}$, described in Figure 10 is a choice function since, if $c_{\mathsf{out}} \in \{c_1, \ldots, c_q\}$ it returns c_{out} while otherwise its output is a random element from its input.

It is also symmetric since its execution of $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$ depends on a random permutation of its input. Thus for any $\eta : \{1, \ldots, q\} \to \{1, \ldots, q\}$ permutation we have that $\left(c_{\pi(\eta(i))}\right)_{i=1}^{q}$ follows the same distribution of $\left(c_{\pi(i)}\right)_{i=1}^{q}$, meaning that c_{out} also does not depend on the input order. □

Proof of Claim 7. Since c'_1 and c'_2 are computed in the same way given c^*_1, c^*_2, it suffice two prove $\Delta(c^*_1, c^*_2) \leq \mathsf{negl}(\lambda)$. Let v_1, v_2 be the view of $\mathsf{AT.Gen}$ and $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, \widehat{m}_0)$ executed by $\mathcal{A}_1, \mathcal{A}_2$ and computing c^*_1 and c^*_2 respectively. Then we have that, using notation from figure 10, c^* and c_{out} are deterministic functions of v_1 and v_2 respectively. Thus $\Delta(c^*_1, c_{\mathsf{out}}) \leq \Delta(v_1, v_2)$. Note

however that c_{out} may differ from the actual output of $\mathcal{F}$. In particular $c_{\mathsf{out}} = c_2^*$ only when $\neg\mathsf{BadChoince}^*$. We can therefore bound, using Claim 5

$$\begin{aligned}
\Delta(c_1^*, c_2^*) &\leq \Delta(c_1^*, c_{2|\neg\mathsf{BadChoice}^*}^*) + \Pr[\mathsf{BadChoice}^*] \\
&= \Delta(c_1^*, c_{\mathsf{out}|\neg\mathsf{BadChoice}^*}) + \Pr[\mathsf{BadChoice}^*] \\
&\leq \frac{1}{1 - \Pr[\mathsf{BadChoice}^*]} \cdot \Delta(c_1^*, c_{\mathsf{out}}) + \Pr[\mathsf{BadChoice}^*] \\
&\leq \Delta(c_1^*, c_{\mathsf{out}}) + 2\Pr[\mathsf{BadChoice}^*] \\
&\leq \Delta(v_1, v_2) + \mathsf{negl}(\lambda).
\end{aligned}$$

To prove the latter statistical distance is also negligible, let $v_{b,n}$ be the vector consisting of the first n queries in v_b. Then we will show by induction that $\Delta(v_{0,n}, v_{1,n}) \leq n \cdot \frac{2q}{2^\rho}$.

The base step is trivial. Moreover for n smaller than the first query of $\mathsf{AT.Enc}$, the two distributions are identical by construction. Assuming the thesis for n, we study the $(n+1)$-th query of $\mathsf{AT.Enc}$ according to its type, conditioning on $v_{1,n} = v = v_{2,n}$.

Key Generation: Queries to $\mathsf{E.Gen}$ are answered identically in both worlds, thus the statistical distance does not increase after performing such queries.

Encryption: When querying $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ this query is answered identically in both worlds, except when $\mathsf{pk} = \mathsf{apk}$ and $m = m^*$. In this case, if the query was already performed before, the answer is consistent. Otherwise, let c_1, c_2 be the replies returned by $\mathcal{A}_1$ and $\mathcal{A}_2$. By construction $c_2 \sim U(\{0,1\}^\ell)$ is uniformly random, even upon conditioning on the view so far.

Conversely to study c_1, let C the set of ciphertext observed so far. Then it can be shown as done in the proof of Claim 3 that

$$\Pr[c_1 \in C \mid v_{1,n} = v] \leq \frac{q}{2^\rho - q}$$

and that, conditioning on $c_1 \notin C$, then $c_1 \sim U(\{0,1\}^\ell \setminus C)$. Then for all $c_0 \notin C$

$$\begin{aligned}
\Pr[c_1 = c_0 \mid v_{1,n} = v] &= \Pr[c_1 = c_0 \mid c_1 \notin C,\ v_{1,n} = v] \cdot \Pr[c_1 \notin C \mid v_{1,n} = v] \\
&\in \left[\frac{1}{2^\ell} - \frac{q}{2^\ell(2^\rho - \rho)}\,;\ \frac{1}{2^\ell} + \frac{q}{2^\ell(2^\ell - q)}\right].
\end{aligned}$$

where the lower bound follows lower-bounding the first factor with $1/2^\ell$ and the second one with $(1 - q/(2^\rho - q))$. Conversely the upper bound follows upper-bounding the first factor with $1/(2^\ell - q)$ and the second one with 1. We eventually

get that

$$\begin{aligned}\Delta(c_{1|v_{1,n}=v}, c_{2|v_{2,n}}) &= \frac{1}{2}\sum_{c_0} |\Pr[c_1 = c_0 \mid v_{1,n} = v] - \Pr[c_2 = c_0 \mid v_{2,n} = v]| \\ &\leq \frac{1}{2}\Pr[c_1 \in C] + \frac{1}{2}\Pr[c_2 \in C] + \frac{1}{2}\sum_{c_0 \notin C} \frac{q}{2^\ell(2^\rho - q)} \\ &\leq \frac{1}{2}\left(\frac{q}{2^\rho - q} + \frac{q}{2^\ell} + \frac{q}{2^\rho - q}\right) \leq \frac{2q}{2^\rho}.\end{aligned}$$

where the second inequality follows as the distance between the two probability for $c_0 \notin C$ is smaller than $q/(2^\ell(2^\rho - q))$, and the last one holds asymptotically as q is polynomially bounded and $\rho = \Omega(\lambda)$. It immediately follows that $\Delta(v_{1,n+1}, v_{2,n+1}) \leq (n+1) \cdot 2q \cdot 2^{-\rho}$ from the inductive hypothesis.

Decryption: If the $(n+1)$-th query is $\mathsf{E.Dec}(\mathsf{sk}, c)$, let C be the set of ciphertext observed so far (which a function of the current view v). If this query was performed before or either $\mathsf{sk} \neq \mathsf{ask}$ or $c \notin \{c_1, \ldots, c_q\} \setminus C$ then the query is replied identically in the two distributions. To conclude it thus suffices to show that in the second view the event $\mathsf{Bad} : \mathsf{sk} = \mathsf{ask} \wedge c \in \{c_1, \ldots, c_q\} \setminus C$ occurs only with negligible probability. This is true as each $c_i \in \{c_1, \ldots, c_q\} \setminus C$, even conditioned on the view, is uniform over $\{0,1\}^\ell$. Thus, by a union bound $\Pr[\mathsf{Bad} \mid v_{2,n} = v] \leq 2^{-\ell}$. Calling m_1, m_2 the replies in the two distributions we thus get

$$\begin{aligned}\Delta(m_{1|v_{1,n}=v},\ m_{2|v_{2,n}=v}) &\leq \Delta(m_{1|v_{1,n}=v},\ m_{2|\neg\mathsf{Bad},\, v_{2,n}=v}) + \Pr[\mathsf{Bad}] \\ &\leq \Pr[\mathsf{Bad}] \leq \frac{q}{2^\ell}.\end{aligned}$$

Given this, the inductive step easily follows as before. □

Proof of Claim 8. Analogous to proof of Claim 7, up to noticing that this time it suffices to prove the bound for $\Delta(c'_2, c'_3)$. The proof is identical up to the fact that in this case $\Pr[\mathsf{BadChoice}] \leq \frac{q^2}{\vartheta+1} + \mathsf{negl}(\lambda)$, which introduces the non-negligible term in the final result. □