Generic Anamorphic Encryption, Revisited: New Limitations and Constructions

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. The notion of Anamorphic Encryption (Persiano *et al.* Eurocrypt 2022) aims at establishing private communication against an adversary who can access secret decryption keys and influence the chosen messages. Persiano *et al.* gave a simple, black-box, rejection sampling-based technique to send anamorphic *bits* using any IND-CPA secure scheme as underlying PKE.

In this paper however we provide evidence that their solution is not as general as claimed: indeed there exists a (contrived yet secure) PKE which lead to insecure anamorphic instantiations. Actually, our result implies that such stateless black-box realizations of AE are impossible to achieve, unless weaker notions are targeted or extra assumptions are made on the PKE. Even worse, this holds true even if one resorts to powerful non-black-box techniques, such as NIZKs, iO or garbling.

From a constructive perspective, we shed light those required assumptions. Specifically, we show that one could bypass (to some extent) our impossibility by either considering a weaker (but meaningful) notion of AE or by assuming the underlying PKE to (always) produce high min-entropy ciphertexts.

Finally, we prove that, for the case of *Fully-Asymmetric* AE, iO *can* actually be used to overcome existing impossibility barriers. We show how to use iO to build Fully-Asymmetric AE (with small anamorphic message space) generically from any IND-CPA secure PKE with sufficiently high min-entropy ciphertexts. Put together our results provide a clearer picture of what black-box constructions can and cannot achieve.

Table of Contents

1 Introduction **4**
- 1.1 Our contributions . . . 5
- 1.2 Technical Overview . . . 6
- 1.3 Other related work . . . 9

2 Preliminaries **10**
- 2.1 Notation . . . 10
- 2.2 Public Key Encryption . . . 10
- 2.3 Anamorphic Encryption . . . 10
- 2.4 Fully-Asymmetric AE . . . 12

3 Impossibility of Stateless Black-Box AE **13**
- 3.1 Counterexample to Rejection Sampling . . . 13
- 3.2 Ideal Weak PKE . . . 15
- 3.3 Impossibility Result . . . 16

4 Positive Results for Stateless Black-Box Triplets **18**
- 4.1 Semi-Adaptive AE . . . 18
- 4.2 Rejection-Sampling Security . . . 19
- 4.3 Extension of Negative Results . . . 19

5 Extensions to Non-Black-Box Techniques **21**
- 5.1 Verifiable Obfuscation . . . 21
- 5.2 Compiling Out Verifiable Obfuscation . . . 22

6 Generic Fully-Asymmetric AE **23**
- 6.1 From Obfuscation . . . 23
- 6.2 From Obfuscation and Perfect Correctness . . . 24

A Supplementary Definitions **28**
- A.1 Pseudorandom Permutations . . . 28
- A.2 Correctness of Anamorphic Encryption . . . 29
- A.3 Weak Asymmetric Anamorphic Encryption . . . 29
- A.4 Indistinguishability Obfuscator and Puncturable PRFs . . . 30

B Supplementary Lemmas **31**
- B.1 Statistical Distance . . . 31
- B.2 Rejection Sampling . . . 31

C Postponed Proof **32**
C.1 IND-CPA of the Counterexample PKE 32
C.2 Counterexample to Rejection Sampling 33
C.3 IND-CPA of the Ideal PKE 33
C.4 Impossibility of Stateless Black-Box AE 35
C.5 Rejection-Sampling is Semi-Adaptive AE 39
C.6 Overcoming impossibility 41
C.7 Verifiable Obfuscation implications 42
C.8 Compiling out Verifiable Obfuscation 44
C.9 First Construction from Obfuscation 45
C.10 Second Construction from Obfuscation 50

1 Introduction

The recently proposed paradigm of Anamorphic Encryption (AE) aims at enabling confidential communication in scenarios where an adversary wields strong control over users. This includes accessing user's private encryption keys (violating so-called *receiver privacy*) and restricting the messages users are allowed to transmit (infringing *sender freedom*). Such capabilities are conceivable in autocratic regimes, where citizens might face strong censorship measures.

In their work, Persiano *et al.* [PPY22] introduced two variations of Anamorphic Encryption, depending on whether sender freedom or receiver privacy are compromised. Here we discuss only the second one, called receiver Anamorphic Encryption, as it is the focus of this work. Informally, receiver AE can be deployed in one of two modes: regular and anamorphic. In the regular mode, the encryption scheme functions as a standard public key scheme. In anamorphic mode, a public key (apk) is generated along with *two* secret keys: a conventional one (ask) and an additional key, referred to as the "double key" (dk). Bob shares dk privately with Alice and uses apk as his public key. When required to surrender his secret key to the adversary, he reveals only ask.

Clearly $(\mathsf{apk}, \mathsf{ask})$ are expected to be compatible with the regular scheme. Yet, Alice can use dk as a symmetric key to embed an extra message into her ciphertext, which remains concealed even knowing ask. More in detail, when operating in anamorphic mode, the scheme enables the encryption of two messages: a regular-looking one m, intended to be seen by the adversary, and a covert one $\widehat{m}$. The resulting anamorphic ciphertext reveals either m, when decrypted with ask, or $\widehat{m}$ when anamorphically decrypted via dk. The main security requirement is for anamorphic ciphertexts to be indistinguishable from regular ones.

In [PPY22], Persiano et al. observed that devising new schemes supporting anamorphic mode might prove futile. An influential enough adversary could indeed push for adopting new PKE schemes as standard and ban the anamorphic ones. Thus, the real challenge here is proving that existing, possibly deployed in practice, constructions have an anamorphic nature. This was tackled by several recent works proposing novel realizations [PPY22, KPP+23b, BGH+24, WCHY23, CGM24a, CGM24b] and refining security notions [BGH+24, KPP+23b, CGM24a]. Most of them, however, build upon rather specific properties of the underlying PKE. Exceptions are the rejection sampling scheme from [PPY22] and the robust one in [BGH+24, Section 4.1], both claimed to work for any IND-CPA secure encryption scheme. In what follows we recall the former construction, as it plays a pivotal role in our work.

Starting from a pseudorandom function f and any PKE, anamorphic mode is constructed as follows. Public and secret keys $(\mathsf{apk}, \mathsf{ask})$ are produced according to the given PKE, whereas the double key dk is a random seed k for f. To encrypt a regular message m and an anamorphic bit $\widehat{m}$, one uses rejection sampling to produce a ciphertext c for m such that $f_k(c) = \widehat{m}$. Regular decryption works as expected, while $\widehat{m}$ is retrieved as $f_k(c)$. In spite of its elegance, this solution only supports up to $O(\log \lambda)$-bit long anamorphic messages, with λ security parameter.

This was recently shown to be optimal by Catalano et al. [CGM24b]. Specifically, they prove secure *black-box* constructions can convey at most $O(\log n)$ covert bit per ciphertext, and cannot achieve stronger notions, such as Fully-Asymmetric security [CGM24a]. In their work, *black-box* refers to generic construction accessing the underlying PKE only through *oracle* calls. This is for instance the case for the rejection sampling scheme above.

1.1 Our contributions

In this paper we revisit the question of studying *generic* constructions of Anamorphic Encryption scheme from PKE and show that the answer is more convoluted than anticipated by previous works. We make progress on this question along several directions. First, in the context of black-box constructions (Sections 3-4):

1. We show the rejection sampling scheme (RS) is actually insecure when applied to a (admittedly contrived, but still ind-cpa) PKE. Thus, RS *does not* generically realize AE.
2. More generally, we prove that *stateless* black-box anamorphic encryption is impossible. Stateless here means that sender and receiver do not keep a synced state. In particular, the usage of synced state in [BGH$^+$24, Section 4.1] is necessary.
3. We introduce a weaker security notion for AE called *Semi-Adaptive* (see below for a discussion on this). We show that RS achieves either semi-adaptivity for any PKE, or the original notion from [PPY22] but only for PKEs with high min-entropy ciphertexts.
4. We extend the message space size bound and impossibility from [CGM24b] to Semi-Adaptive security and high min-entropy PKEs. This shows that [PPY22] is again optimal among semi-adaptively secure black-box constructions.

Next, we ask whether obfuscation, garbling or NIZKs could help bypass these limitations. Towards this, we then allow constructions to obfuscate (or prove statements about) circuits with gates evaluating the underlying PKE procedures. The PKE is otherwise accessed through oracle calls as before. In Sections 5-6:

5. We show that any secure AE in this model can be compiled into a secure black-box AE not using $\mathsf{iO}/\mathsf{NIZK}$. In particular, all negative results above relative to (plain) AE extends to this setting.
6. We realize Fully Asymmetric AE [CGM24a] with semi-adaptive security from iO. We remark that this is implied to be impossible in the black-box setting (without iO) by point 4.

We provide a summary of our results in Table 1. The notion of fully asymmetric AE [CGM24a] mentioned above informally considers an asymmetric variant of the original AE definition. Alice now generates *two* covert keys $(\mathsf{dk}, \mathsf{tk})$ (as opposed to dk only). dk is shared with the sender(s) and acts as the encryption key for anamorphic messages. tk is instead kept private and used for decryption.

	Black-Box		Black-Box + $\mathsf{iO}/\mathsf{NIZK}$	
	Possible?	$\lvert M\rvert$	Possible?	$\lvert M\rvert$
Stateless AE	✗	–	✗	–
Stateful AE	✓	$\leq \mathsf{poly}(\lambda)$	✓	$\leq \mathsf{poly}(\lambda)$
SA-AE	✓	$\leq \mathsf{poly}(\lambda)$	✓	$\leq \mathsf{poly}(\lambda)$
Fully-Asym. + AE	✗	–	✗	–
Fully-Asym. + SA-AE	✗	–	✓	$\leq \mathsf{poly}(\lambda)$

Table 1. *SA-AE* is Semi-Adaptive AE, *Fully-Asym.* is short for Fully Asymmetric. Black-Box refers to triplets using the underlying PKE only through oracle calls. Black-Box + $\mathsf{iO}/\mathsf{NIKZ}$ allows obfuscating and proving statements about circuits with PKE operation gates. $|M|$ is the anamorphic message space size. λ is the security parameter.

1.2 Technical Overview

In what follows we discuss the intuitions and technical challenges underlying our results, simplifying where necessary to aid intuition. Throughout this Section, $\Sigma = (\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ will be an anamorphic triplet turning *any* PKE into an AE.

Revisiting rejection sampling AE. Our first step is to construct an artificial PKE which, in spite of being IND-CPA and correct, does not give rise to a secure AE when RS (presented earlier) is applied to it. The main idea is to introduce an hard to find weak message, with few associated ciphertexts. We start with a PKE $(\mathsf{E.Gen}, \mathsf{E.Dec}, \mathsf{E.Dec})$ with exponential message space M, an injective OWF $F : M \to \{0,1\}^*$ and a small set B disjoint from the PKE ciphertext space.

Our weakened PKE, $(\mathsf{E.Gen}^*, \mathsf{E.Enc}^*, \mathsf{E.Dec}^*)$ works as follows. $\mathsf{E.Gen}^*$ first runs $\mathsf{E.Gen}$ to get pk, sk, it samples a random message m^* from M and sets $y^* \leftarrow F(m^*)$. The public key pk^* is (pk, y^*) and the secret key sk^* is (sk, m^*). $\mathsf{E.Enc}^*$ is as $\mathsf{E.Enc}$ for all messages m except if $F(m) = y^*$, in which case it outputs a random string in B. Finally, $\mathsf{E.Dec}^*$ runs as $\mathsf{E.Dec}$ for all ciphertexts not in B, while in this latter case it outputs m^*.

It is easy to show that, being M exponentially large, the scheme is IND-CPA if so is the underlying PKE. However regular and anamorphic mode are easily distinguished. Indeed, an adversary holding ask could query c_0, c_1 encryptions of $(m^*, 0)$ and $(m^*, 1)$. In regular mode, both ciphertext will collide with probability $1/|B|$, which is significant. In anamorphic mode instead, collisions almost never happen due to correctness, as $f_k(c_0) = 0$ and $f_k(c_1) = 1$.

Impossibility of stateless black-box AE. Building from the counterexample illustrated above we prove that black-box Anamorphic Encryption is impossible to realize. This improves upon a recent result by Catalano *et al.* [CGM24b] that shows that any such conversion can at best produce an AE with small (anamorphic) message space.

Our proof follows the same general approach of [CGM24b]: we start by describing an *ideal* public key encryption $\mathsf{\Pi} = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$, based on truly random permutations specifying the key generation and encryption/decryption behavior. In our case, this is further augmented with a mechanism to (artificially) introduce weak messages given the secret key, i.e. with few associated ciphertext as before. The resulting scheme is provably IND-CPA. Therefore, a black-box AE has to be secure when applied to it.

To reach a contradiction then, it suffices to provide an attack against the resulting scheme. We proceed as before. Given a "weak" message m^*, the attacker asks (several) encryptions for $(m^*, 0)$ and $(m^*, 1)$. As before, these have a significant chance of colliding when using the regular encryption scheme. In anamorphic mode, on the other hand, correctnessof $\mathsf{AT.Enc}$ and the fact that it is stateless, implies that a collision occurs with significantly lower probability.

In order for this simple argument to go through, however, one has to make sure that the anamorphic encryption procedure does not realize m^* to be weak[4]. A crucial step in our proof consists in showing that, when there are sufficiently many (but still polynomially many) ciphertexts associated to m^*, $\mathsf{AT.Enc}$ cannot distinguish *too often* weak messages from regular ones.

Finally, note the above attack only works against *stateless* anamorphic schemes. In such cases indeed correctness should prevent encryptions of $(m^*, 0)$ and $(m^*, 1)$ to collide. This is remarkably not the case for *stateful* constructions. Indeed in that case the two ciphertexts would be allowed to collide, as they will later be decrypted with different states. This is the reason why the generic construction in [BGH+24] does not contradicts our result.

Achievable security for stateless black-box AE. Having established that (stateless) AE cannot be realized generically, the natural question becomes either what security notion *can* be achieved, or what class of PKEs do we need to exclude to circumvent the above barrier.

Regarding the latter, we show a sufficient condition to be *high min-entropy* ciphertexts. That is, for any valid key and message, each ciphertexts has $\Omega(\lambda)$ bits of min-entropy. In this case we can prove RS to be secure as all produced ciphertexts c are distinct up to negligible probability and the bits $f_k(c)$ are computationally close to uniformly and independently distributed.

About the former, on the other hand, we propose a new definition called *semi-adaptive* AE. Informally, this modifies the original notion by letting the adversary access the secret key only *after* all the encryption queries are made[5]. Even though

[4] In principle, $\mathsf{AT.Enc}$ could try to encrypt m^* several times looking for collisions. If this occurs, it could then ignore the covert message and simply output a (regular) encryption of m^*. Such a behavior, while affecting correctness, would fool our distinguisher.

[5] The semi-adaptive name comes from the fact that encryption queries can be asked adaptively after having seen the public key but cannot depend on explicit knowledge of the secret key. This is reminiscent of semi-adaptive security for functional encryption [CW14] where the adversary is allowed to ask the challenge query after having seen the public key but before making key derivation queries.

we don't have any compelling case use for semi-adaptive AE we believe it could be used to model security in contexts where an adversary/dictator having the power to forces users to surrender their secret key still cannot check their behavior before some point in time (e.g. before her/his rise to power).

Extensions to non-black box techniques. Next we consider the question of whether powerful non-black-box techniques such as NIZKs, garbling or iO can be used to overcome our results so far.

Our first answer for (semi-adaptive) AE is negative. We show that a large class of general non-black-box techniques would not be useful here. Towards this goal we begin by targeting a very powerful primitive, called *Verifiable Virtual Black Box Obfuscation* (VO), which is an extension of verifiable obfuscation from [BGJS16] and subsumes all the above techniques. Informally this, along with regular obfuscation, further allows verifying a given predicate P of the obfuscated circuit C, with P chosen by the obfuscator.

Next, we study anamorphic triplet defined relative to PKE oracles and to an *ideal* VO oracles. We take this route because, informally, we cannot "obfuscate the PKE oracles". In other words, obfuscation does not relativize. Our ideal VO, instead, can take as input circuits with PKE gates, obfuscate them by simply assigning random labels, and later evaluate them through the PKE oracles. This is a well-known approach, an example of which can be found in [GHMM18, Section 4] to model garbling relative to an ideal OWF.

Finally, we show that relative to those PKE and VO oracles, any AE triplet can be compiled into one that never accesses VO while preserving (semi-adaptive) security. This is done by letting sender and receiver (relative to the PKE only!) share a PRP key k and simulate the obfuscator with $f_k(C)$. Among themselves they can easily evaluate and verify by just inverting f_k. Given an adversary $\mathcal{A}$ relative to PKE it can be lifted to one relative to the PKE and VO by simply not making any VO query. The result follows by proving that in the two worlds (i.e. with the ideal VO or with the simulated one) the views are computationally close. Thus obfuscation, as well as NIZK and garbling, is of no help here.

Fully-Asymmetric AE from obfuscation. A interesting aspect of the compiler discussed above is that it requires sender and receiver to cooperate. This is acceptable for the standard notion of AE, where dk is treated as a symmetric key. It is however not acceptable for stronger notions such as fully-asymmetric AE, where the receiver has private key information that wishes not to share with the sender. Hence the above result does not extend to the case fully-asymmetric AE.

Interestingly, we show that this is no coincidence and indeed we prove that using iO it is possible to build Fully-Asymmetric AEs (with small anamorphic message space) generically from any IND-CPA secure PKE. The usage of iO thus allows to bypass the (extended) result[6] from [CGM24b] which proved fully-asymmetric black-box (semi-adaptive) AE to be impossible. We give two such

[6] Here by extended we mean reinterpreted in light of the results in this paper.

constructions, both building upon the Sahai-Waters [SW14] realization of public key encryption from iO.

The basic idea is to interpret the rejection sampling scheme from [PPY22] as a secret-key encryption scheme and turn it into an asymmetric one exactly as done in [SW14]. Our first construction closely follows Sahai-Waters and inherits their exponential security loss arising from their PRG usage. Recall, that the Sahai-Waters scheme uses a PRG G, that takes a seed of size $\lambda/2$, to produce the random coins needed to encrypt. Typically such a loss is acceptable as it only means that larger λ have to be chosen in case of need. For the case of Anamorphic Encryption however this might be problematic as the concrete value for λ might be fixed by the adversary so that breaking the PKE is unfeasible, but distinguishing regular from anamorphic ciphertexts becomes doable.

Our second construction avoids this issue by removing the PRG altogether but assuming perfect correctness of the underlying PKE instead. Very informally, the idea is as follows. We modify the obfuscated circuit used to encrypt by adding an "unreachable" condition for which a fixed output is returned. Specifically, the condition is that, on input (m, r), one checks whether $m = m_1^*$ and $\mathsf{E.Enc}(\mathsf{pk}, m, r) = c^*$ where m_1^*, c^* are hard-coded in the circuit and c^* is an encryption of a message $m_0^* \neq m_1^*$. Here is where perfect correctness comes into play: it allows to rule out the possibility that c^* could be obtained as the encryption of an a $m \neq m_1^*$, making such condition unreachable. Later, using the IND-CPA security, we set c^* as the encryption of m_1^*, thus making the condition reachable.

As a final note, we remark that, as we adapt the rejection sampling construction, both scheme *still* either requires the underlying PKE to produce high min-entropy ciphertexts, or only achieve semi-adaptive security. This is however in line with our previous results. Indeed, achieving plain AE in the black-box + iO model when the underlying PKE is assumed to guarantee only correctness and IND-CPA security was shown to be impossible.

1.3 Other related work

Anamorphic Encryption shares similarities with previously studied notions, 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]). We refer to the work of Persiano *et al.* [PPY22] for an in-depth comparison among these notions.

In [KPP$^+$23b, CGM24a] the notion of receiver AE has been further refined by requiring privacy for the normal and covert messages to hold even when knowing dk. In [BGH$^+$24] the notion of *robust* AE and Anamorphic Extension have been introduced. Later on in [WCHY23] the notion of robustness has been extended and adapted also to the case of sender AE. In [KPP$^+$23a] the notion of Anamorphic Signatures has been introduced in order to deal with a more extreme scenario where all communications must pass through a central authority under the adversary's control. In this context, the usage of encryption channels becomes even more complicated, thus, in order to get around it, they rely on authenticated channels (i.e., using digital signatures) to be able to establish secure communications between parties.

The study of black-box separations started from the seminal work of Impagliazzo and Rudich [IR89], which gave rise to a fritful and active area of research [Sim98, KST99, GT00, GKM+00, GMR01, GGK03]. All these works however only rule out black-box constructions that use the underlying primitive as an oracle (i.e. not *all* possible constructions).

2 Preliminaries

2.1 Notation

$[n]$ denotes the set $\{1, \ldots, n\}$. $\lambda \in \mathbb{N}$ is 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. 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. With $\approx$ and $\stackrel{\mathrm{p}}{=}$ we denote respectively the computationally and perfect indistinguishability. 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*. Also, we may omit the game in the adversary's advantage Adv when clear from context.

2.2 Public Key Encryption

We denote with $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ a PKE scheme with message space M. Along with the standard properties of correctness and IND-CPA, we consider the following one, requiring ciphertexts to have high min-entropy for any key and message choice.

Definition 1. *A PKE scheme has high min-entropy ciphertexts if, for any* $(\mathsf{pk}, \mathsf{sk})$ *in the range of* $\mathsf{E.Gen}$*, and for any message* $m \in M$

$$\mathrm{H}_\infty(\mathsf{E.Enc}(\mathsf{pk}, m)) = \Omega(\lambda).$$

2.3 Anamorphic Encryption

The definition of (receiver) Anamorphic Encryption that we use in this paper is the one from [CGM24a], which is a generalization of the original one by Persiano *et al.* [PPY22]. 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 storing a *trapdoor key* tk to decrypt anamorphic messages from the sender.

Definition 2 (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.
– $\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c) \to \widehat{m}/\bot$, *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. Moreover, we may omit tk when empty.

Definition 3 (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*

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

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

$\mathsf{RealG}_{\Pi}(\lambda, \mathcal{A})$	$\mathsf{AnamorphicG}_{\Sigma}(\lambda, \mathcal{A})$
$(\mathsf{pk}, \mathsf{sk}) \xleftarrow{\$} \mathsf{E.Gen}(\lambda)$	$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \xleftarrow{\$} \mathsf{AT.Gen}(\lambda)$
return $\mathcal{A}^{\mathcal{O}_{\mathsf{real}}}(\mathsf{pk}, \mathsf{sk})$	**return** $\mathcal{A}^{\mathcal{O}_{\mathsf{anam}}}(\mathsf{apk}, \mathsf{ask})$
$\mathcal{O}_{\mathsf{real}}(m, \widehat{m})$	$\mathcal{O}_{\mathsf{anam}}(m, \widehat{m})$
Sample a random r	Sample a random r
return $\mathsf{E.Enc}(\mathsf{pk}, m; r)$	**return** $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}; r)$

Fig. 1. Anamorphic Encryption security game.

Regarding correctness we recall the game-based definition provided [BGH+24] in the Appendix, Section A.2. For the sake of generality however we will mainly refer to a weaker notion, *correctness on average*, holding only for uniformly sampled messages (and correct keys).

Definition 4. *An anamorphic triplet is* ε*-correct on average if, for a negligible* ε, *sampling* $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \xleftarrow{\$} \mathsf{AT.Gen}(\lambda)$ *and a random message* $m \xleftarrow{\$} M$ *from the regular message space, then for all* $\widehat{m} \in \widehat{M}$

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

Finally, as the focus of our investigation is on black-box constructions, we proceed to provide a formally define them as done in [CGM24b].

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 informally refer to an Black-Box Anamorphic Triplet as a Black-Box *Anamorphic Encryption.*

2.4 Fully-Asymmetric AE

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 2.

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

1 : $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \xleftarrow{\$} \mathsf{AT.Gen}(\lambda)$
2 : $b \xleftarrow{\$} \{0,1\}$
3 : $(m_0, m_1, \widehat{m}_0, \widehat{m}_1) \xleftarrow{\$} \mathcal{A}(\mathsf{apk}, \mathsf{dk})$
4 : $c \xleftarrow{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m_b, \widehat{m}_b)$
5 : $b' \xleftarrow{\$} \mathcal{A}(c)$
6 : **return** $b == b'$

Fig. 2. Fully-Asymmetric Anamorphic Encryption game.

We define the advantage of $\mathcal{A}$ against the Fully-Asymmetric property as

$$\mathsf{Adv}^{\mathsf{fasy\text{-}anam}}_{\mathcal{A},\Sigma}(\lambda) \;=\; 2 \cdot \left|1/2 - \Pr\left[\mathsf{FAsyAnam}\text{-IND-CPA}_{\Sigma}(\lambda, \mathcal{A}) = 1\right]\right|.$$

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 6 (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*

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

3 Impossibility of Stateless Black-Box AE

3.1 Counterexample to Rejection Sampling

In [PPY22], along with the definition of Anamorphic Encryption, a supposedly generic stateless construction based on rejection sampling was proposed. In this section we recall their construction, and show it to be insecure when applied to an artificially weakened (but still IND-CPA) encryption scheme.

Given any PKE with public and secret keys $(\mathsf{pk}, \mathsf{sk})$, sender and receiver of [PPY22]'s AE initially exchange a PRF key k acting as the double key. To communicate a bit $\widehat{m}$ the sender produces many ciphertexts $c_1, \dots, c_\vartheta$ for the regular message m, and eventually sends the first c_i such that $f_k(c_i) = \widehat{m}$. This mildly deviate from the original, which does not prescribe an exit condition if a proper c is never found. In particular it only run (at best) in *expected polynomial time*[7]. Here instead we bound the attempts to ϑ and eventually send a new $c \xleftarrow{\$} \mathsf{E.Enc}(\mathsf{apk}, m)$ if no desired c_i was found, giving up on correctness. A full description of the triplet RS is given in Figure 3.

$\mathsf{RS.Gen}(\lambda)$	$\mathsf{RS.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$	$\mathsf{RS.Dec}(\mathsf{ask}, \mathsf{dk}, c)$
$(\mathsf{apk}, \mathsf{ask}) \xleftarrow{\$} \mathsf{E.Gen}(\lambda)$	**for** $i \in \{1, \dots, \vartheta\}$:	**return** $f_k(c)$
$k \xleftarrow{\$} \mathsf{PRF.Gen}(\lambda)$	$\quad c_i \xleftarrow{\$} \mathsf{E.Enc}(\mathsf{apk}, m)$	
$\mathsf{dk} \leftarrow k$	$\quad$ **if** $f_k(c_i) = \widehat{m}$: **return** c_i	
return $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk})$	**return** $\mathsf{E.Enc}(\mathsf{apk}, m)$	

Fig. 3. Anamorphic Triplet RS with $\vartheta = \mathsf{poly}(\lambda)$ repetitions.

A key requirement for RS to work is the existence of *many* distinct ciphertexts linked to m. In other words, $\mathsf{E.Enc}(\mathsf{pk}, m; r)$ needs to have high min-entropy given pk and m. To see why, assume that only $\mathsf{poly}(\lambda)$ ciphertexts can be obtained encrypting a given m. Then the probability that two *regular* encryptions of m collide is noticeable. However, two anamorphic ciphertexts of m with anamorphic message 0 and 1 collides with negligible probability due to anamorphic correctness. Hence the two modes would be readily distinguishable.

The issue above should not to occur when m is chosen by an adversary who only knows pk, as such m would allow breaking IND-CPA. However IND-CPA alone cannot prevent to find it given *both* pk and sk. This is exactly the setting of the anamorphic security game. A counterexample can therefore be built as follows: given any PKE with exponential message space, we artificially weaken a random message m^*. The public key is extended to contain $F(m^*)$ with F an injective one-way function, and sk is extended with m^*. Encryption is the same,

[7] Even worse, on some input, the encryption alogirhtm may never terminate. Looking ahead, setting $|B| = 1$ in our counterexample implies this to happen for some message pair $(m^*, \widehat{m})$.

except for m^* where a ciphertext is a random element from a polynomially small set B disjoint from the given PKE's ciphertext space. Decryption runs either the old decryption or, if $c \in B$, returns m^*. A detailed description is given in Figure 4, while proof for the next Proposition appears in Appendix C.1.

$\mathsf{E.Gen}^*(\lambda)$

```
1 :  pk, sk ←$ E.Gen(λ)
2 :  m* ←$ M,  y* ← F(m*)
3 :  pk* ← (pk, y*),  sk* ← (sk, m*)
4 :  return (pk*, sk*)
```

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

```
1 :  Parse pk* = (pk, y*)
2 :  if F(m) = y*:
3 :     return c ←$ B
4 :  else :
5 :     return c ←$ E.Enc(pk, m)
```

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

```
1 :  Parse sk* = (sk, m*)
2 :  if c ∈ B:   return m*
3 :  else :   return E.Dec(sk, c)
```

Fig. 4. Weakened PKE from any PKE $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ with message space M. $F : M \to \{0,1\}^*$ is an injective OWF and B a set of size $|B| = \mathsf{poly}(\lambda)$ disjoint from the given PKE's ciphertext space.

Proposition 1. *Given a correct and* IND-CPA *encryption* $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *with* $|M| = \Omega(2^\lambda)$ *and* F *injective OWF, then the scheme presented in figure 4 is correct and* IND-CPA *secure.*

Proposition 2. *The triplet* RS *defined in Figure 3 is not a secure anamorphic triplet with respect to the PKE described in Figure 4 when* $|B| \geq 4\vartheta$.

Proof of Proposition 2. We describe an adversary $\mathcal{A}$ breaking anamorphic security in Figure 5. Initially it extracts m^* from ask, which RS computes correctly by construction. Then uses m^* to produce two ciphertexts, supposedly encrypting the anamorphic bit 0 and 1. Finally, it returns 1 only if the two ciphertexts collide.

It is immediate to see that in the real game $\mathcal{A}$ returns 1 with probability $1/|B|$ as c_0, c_1 are uniformly and independently sampled from B. To study the anamorphic game, let $\mathsf{Fail}_0, \mathsf{Fail}_1$ the events in which line 4 is executed when $\mathsf{RS.Enc}$ encrypts respectively $(m^*, 0)$ and $(m^*, 1)$. We then claim those events to occur with probability far from 1. A proof appears in the Appendix, Section C.2.

Claim 1. $\Pr[\mathsf{Fail}] \leq 1/2 + \mathsf{negl}(\lambda)$, *where* $\mathsf{Fail} = \mathsf{Fail}_0 \vee \mathsf{Fail}_1$.

Next, if $\neg\mathsf{Fail}$, either c_0 or c_1 is a regular ciphertext, and therefore a collision occurs with probability $1/|B|$. Conversely, $f_k(c_0) = 0$ and $f_k(c_1) = 1$ implies

$$
\begin{array}{ll}
\multicolumn{2}{l}{\underline{\mathcal{A}^{\mathcal{O}}(\mathsf{apk}, \mathsf{ask}):}} \\
1: & \text{Parse } \mathsf{ask} = (\mathsf{sk}^*, m^*) \\
2: & \text{Query } c_0 \leftarrow^{\$} \mathcal{O}(m^*, 0) \text{ and } c_1 \leftarrow^{\$} \mathcal{O}(m^*, 1) \\
3: & \textbf{return } c_0 = c_1
\end{array}
$$

Fig. 5. Adversary breaking security of the RS triplet applied to the weak PKE in Figure 4. $\mathcal{O}$ is the encryption oracle provided in the anamorphic security game 1.

that no collision can occur. We then conclude that, calling c_0', c_1' the ciphertext obtained in the real game, the advantage of $\mathcal{A}$ is lower-bounded by

$$
\begin{aligned}
\mathsf{Adv}(\mathcal{A}) \;\geq\; & \Pr[c_0' = c_1'] - \Pr[c_0 = c_1] \;=\; \Pr[c_0' = c_1'] - \Pr[c_0 = c_1 \mid \mathsf{Fail}] \Pr[\mathsf{Fail}] \\
=\; & \frac{1}{4\vartheta} - \frac{1}{4\vartheta}\left(\frac{1}{2} + \mathsf{negl}(\lambda)\right) \;=\; \frac{1}{8\vartheta} - \mathsf{negl}(\lambda). \qquad \square
\end{aligned}
$$

Remark 1. Modifying the rejection sampling triplet to avoid this attack is trivial. We can define RS.Enc to behave as $\mathsf{E.Enc}^*$ when asked to encrypt $(m^*, \cdot)$[8]. Our goal indeed is not to show that the weak PKE above does not admit anamorphic triplets, but rather that the rejection sampling construction does not apply to *all* PKEs.

3.2 Ideal Weak PKE

The counterexample proposed against the rejection sampling triplet (Figure 3) can be generalized to show that black-box Anamorphic Encryption is not possible. Following the same general approach of [CGM24b], we begin describing an ideal public key encryption, but this time with artificially weakened messages. Then, we prove this ideal PKE, in spite of being IND-CPA secure and correct, cannot admit a secure *stateless* anamorphic triplet. Hence building stateless black-box triplets assuming the underlying PKE scheme to only be correct and IND-CPA secure is impossible.

Our PKE is informally defined by two random functions ϕ, ψ roughly describing the key generation and encryption. Moreover, to introduce *weak* messages, the scheme is further defined by $m_1^*, \ldots, m_\lambda^*$ random functions (taking in input elements from SK) and τ. The latter is a function acting on the encryption random coins that on a *good* message is the identity, whereas on a *weak* one is (extremely) compressing to ensure many collisions. More precisely, we denote PK, SK the public and secret key space, while $\{0,1\}^\mu$, $\{0,1\}^\rho$, $\{0,1\}^\ell$ are respectively the messages, encryption's coins, and ciphertexts space. Then ϕ, ψ, τ and m_i^* are sampled uniformly satisfying the following constraints:

1. $\phi : \mathsf{SK} \to \mathsf{PK}$ is a bijection.
2. $\psi : \mathsf{PK} \times \{0,1\}^\mu \times \{0,1\}^\rho \to \{0,1\}^\ell$ such that $\psi(\mathsf{pk}, \cdot, \cdot)$ is injective.

[8] Although correctness is unavoidably lost with respect to the anamorphic message.

3. $m_i^* : \mathsf{SK} \to \{0,1\}^\mu$.
4. $\tau(\phi(\mathsf{sk}), m, r) = r$ if m is not weak, i.e. $m \notin \{m_i^*(\mathsf{sk})\}_{i=1}^{\lambda}$.
5. $|\mathrm{Im}\ (\tau(\mathsf{pk}, m, \cdot))| \leq 2^i$ if m is the i-th weak message, i.e. $\mathsf{pk} = \phi(\mathsf{sk})$ and $m = m_i^*(\mathsf{sk})$.

Looking ahead, we impose m_i^* to have at most 2^i ciphertexts to later let our adversary choose the right i for its attack to succeed. For ease of notation we will denote $\psi_\tau(\mathsf{pk}, m, r) = \psi(\mathsf{pk}, m, \tau(\mathsf{pk}, m, r))$. Moreover, as in [CGM24b], we fix parameters so that $\rho = \Omega(\lambda)$ and $\ell - (\rho + \mu) = \Omega(\lambda)$. Next, given ϕ, ψ, τ, m_i^* distributed as above, our ideal weak PKE is presented in Figure 6.

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

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

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

1 : **return** $\psi_\tau(\mathsf{pk}, m, r)$

$\mathsf{E.Find}(\mathsf{sk}, i)$:

1 : **return** $m_i^*(\mathsf{sk})$

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

1 : **if** there exists (m, r) such that $c = \psi_\tau(\phi(\mathsf{sk}), m, r)$:

2 : **return** m

3 : **else** : **return** $\perp$

Fig. 6. Ideal Weak PKE. $\phi : \mathsf{SK} \to \mathsf{PK}$ and $\psi : \mathsf{PK} \times \{0,1\}^\mu \times \{0,1\}^\rho \to \{0,1\}^\ell$ are distributed as above. $\rho = \Omega(\lambda)$ and $\ell = \rho + \mu + \Omega(\lambda)$.

In order to claim that a black-box anamorphic triplet should be required to work for the above PKE, we first need to show it to be efficiently simulatable[9], correct and IND-CPA secure. This is addressed in the following Lemma, whose proof appears in Appendix C.3.

Lemma 1. *Relative to the ideal weak PKE* $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec}, \mathsf{E.Find})$ *presented in Figure 6, there exists a PKE defined by the triplet* $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *that is perfectly correct and* IND-CPA *secure. Moreover the ideal weak PKE can be simulated efficiently.*

3.3 Impossibility Result

Toward contradiction let $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ be a black-box *stateless* anamorphic tuple, i.e. which access the underlying PKE only through oracle calls. By definition, as long as the given PKE is correct and IND-CPA, such tuple is required to be secure according to the security notion in Definition 3. To show such a tuple cannot exists, in this section we provide an efficient adversary breaking

[9] This requirement is actually to avoid the PKE oracle to provide help in solving problems that would be hard in PPT time.

the anamorphism game when we apply the given tuple to the ideal weak PKE presented in Figure 6.

Our adversary is similar to the one presented for the rejection sampling triplet. Initially it finds a weak message m^* and then it queries (several) ciphertexts encrypting $(m^*, 0)$ and $(m^*, 1)$. These have a significant chance of colliding in the real game, whereas in anamorphic mode a collision should only occur with small probability due to correctness and the lack of state. As opposed to the rejection sampling case however, more care has to be taken in those arguments. Indeed, if $\mathsf{AT.Enc}$ understands m^* to be a weak message[10], it could give up any attempt to encrypt the anamorphic message and simply return a regular ciphertext. To avoid this, $\mathsf{AT.Enc}$'s view when asked to encrypt m^* has to be almost the same as with a random message.

Crucially, the latter is only possible as we study *black-box* anamorphic triplets. Recall these access the underlying PKE through oracle calls and have to be correct and secure relative to *any* PKE. In particular, relative to the four oracle $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec}, \mathsf{E.Find})$, a generic triplet for the PKE defined by the first three procedures cannot query $\mathsf{E.Find}$, as not *every* PKE admits such procedure. This will be the main reason why the underlying anamorphic triplet, in spite of having access to sk, is almost unable to distinguish weak messages from regular ones.

$\mathcal{A}_{\vartheta,\nu}(\mathsf{pk}, \mathsf{sk})$:

Get the weak message $m^* \leftarrow \mathsf{E.Find}(\mathsf{sk}, \log_2 \nu)$
for $i \in \{1, \ldots, \vartheta\}$:
  Query $c_{0,i} \xleftarrow{\$} \mathcal{O}(m^*, 0)$ and $c_{1,i} \xleftarrow{\$} \mathcal{O}(m^*, 1)$
if $\nexists i, j$ such that $c_{0,i} = c_{1,j}$:
  return 0   // The real PKE is likely to have collisions
else : **return** 1

Fig. 7. Adversary breaking a black-box anamorphic tuple $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ applied to the ideal weak PKE relative to oracles $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec}, \mathsf{E.Find})$. $\mathcal{A}$ is parametrized by $\vartheta, \nu = \mathsf{poly}(\lambda)$. $\mathcal{O}$ is the encryption oracle in the anamorphism game.

Theorem 1. *For any* $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *black-box anamorphic triplet* ε*-correct on average, where each procedures performs at most* $q = \mathsf{poly}(\lambda)$ *queries, when applied to the ideal PKE* $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec}, \mathsf{E.Find})$ *in Figure 6 there exists a* PPT *adversary* $\mathcal{A}_{\vartheta,\nu}$ *(Figure 7) such that*

$$\nu \geq \lambda^2 q^4, \quad \vartheta = \sqrt{\nu/2} \quad \Rightarrow \quad \mathsf{Adv}(\mathcal{A}_{\vartheta,\nu}) = \Omega(1).$$

[10] e.g. by finding a collision while producing many fresh encryptions of m^*, which for an average message should almost never occur.

A proof of the above theorem appears in Appenidx C.4.

Remark 2. Our result can actually be strengthened to show (stateless) black-box triplet with ε-correctness on average cannot exists, where $\varepsilon = o(1/q^2)$. We leave it as an intriguing open problem to understand whether secure constructions with polynomial error $\Omega(1/q^2)$ exists.

4 Positive Results for Stateless Black-Box Triplets

Having shown that no *stateless* black-box anamorphic triplet can be secure for all PKE schemes, in this section we consider the following two questions:

1. What (mildly) weaker security notion can still be satisfied?
2. Under what condition on the PKE can plain anamorphic security be achieved?

The first one is answered providing a relaxation of the definition in [PPY22] which we call *semi-adaptive security*. We answer the second one instead restricting to PKEs with high min-entropy ciphertexts (see Definition 1). This suffices to rule out pathological cases (e.g. Figure 4). Although these restrictions allow bypassing Theorem 1, we finally show that bounds and negative results in [CGM24b] extend to these settings.

4.1 Semi-Adaptive AE

The core issue exploited in the proof of Theorem 1 is that the adversary can access in the query phase both public and private key. To avoid such class of attacks, we now discuss a relaxation of Definition 3. The only difference we introduce is that ask is provided at the end of the query phase instead of the beginning. We call this new definition *semi-adaptive* AE. The name indeed is reminiscent of semi-adaptive security for Functional Encryption [CW14], where challenge queries are performed before observing (functional) secret keys.

Formally, 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 Semi-Adaptive Anamorphism game, for $\mathcal{A}$ a PPT adversary, is defined in Figure 8. We define the advantage of an adversary $\mathcal{A}$ in breaking the Semi-Adaptive property as

$$\mathsf{Adv}^{\mathsf{sa\text{-}anam}}_{\mathcal{A},\Pi,\Sigma}(\lambda) = \left|\Pr\left[\mathsf{SA\text{-}RealG}_{\Pi}(\lambda, \mathcal{A}) = 1\right] - \Pr\left[\mathsf{SA\text{-}AnamG}_{\Sigma}(\lambda, \mathcal{A}) = 1\right]\right|.$$

Definition 7 (Semi-Adaptive AE). *A PKE Π equppied with an Anamorphic Triplet Σ is said to be Semi-Adaptive Anamorphic if for every PPT adversary $\mathcal{A}$ it holds that*

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

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

$(\mathsf{pk}, \mathsf{sk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda)$
Run $\mathcal{A}(\mathsf{pk})$
for $i = 1, \ldots, \mathsf{poly}(\lambda)$:
  $(m_i, \widehat{m}_i) \leftarrow^{\$} \mathcal{A}$
  $c_i \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{pk}, m_i)$
  Give c_i to $\mathcal{A}$
Give sk to $\mathcal{A}$
return $\mathcal{A}$'s output

$\mathsf{SA\text{-}AnamG}_{\Sigma}(\lambda, \mathcal{A})$

$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow^{\$} \mathsf{AT.Gen}(\lambda)$
Run $\mathcal{A}(\mathsf{apk})$
for $i = 1, \ldots, \mathsf{poly}(\lambda)$:
  $(m_i, \widehat{m}_i) \leftarrow^{\$} \mathcal{A}$
  $c_i \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m_i, \widehat{m}_i)$
  Give c_i to $\mathcal{A}$
Give ask to $\mathcal{A}$
return $\mathcal{A}$'s output

Fig. 8. Semi-Adaptive Anamorphic Encryption game.

4.2 Rejection-Sampling Security

Having formally defined a weaker security notion for Anamorphic Encryption, our next step is proving RS to achieve it generically. This will provide an answer to the first question asked at the beginning of this section, as RS is stateless, black-box, and we show security to hold *for any* PKE. As mentioned, contrived schemes such as the counter-example in Section 3.1, should not affect the proof anymore. Indeed, according our new notion, an adversary can only query messages that depend on the public key, thus excluding adversaries such as the one in Figure 7. This formally leads to the following Theorem, proven in the Appendix, Section C.6.

Theorem 2. *The rejection sampling triplet* RS *described in Figure 3 when applied to an* IND-CPA *secure PKE, yields a black-box Semi-Adaptive Anamorphic Encryption scheme.*

Alternatively, in order to achieve plain anamorphic security with black-box construction, some restrictions have to be imposed on the class of PKEs the triplet is proven secure with. In this direction, answering our second question, we show RS to be secure if the underlying PKE has high min-entropy ciphertexts (Definition 1). This offers a trade-off between security levels and generality. The first theorem indeed captures all PKE but provides weaker guaranteed. The second one only applies to a (broad) class of PKEs, but guarantees stronger security. A proof for this second theorem appears in the Appenidx, Section C.6.

Theorem 3. *The triplet* RS *when applied to an* IND-CPA *PKE with high min-entropy ciphertexts yields a black-box Anamorphic Encryption scheme.*

4.3 Extension of Negative Results

To conclude this section, we show how the negative results in [CGM24b] can be extended to our new definition. More precisely, in [CGM24b] the RS tuple was claimed to be an optimal as no black-box AE can have super-polynomial message space, and stronger notions such as Fully-Asymmetric security [CGM24a] cannot be achieved altogether. Both results are technically surpassed by our new impossibility in Section 3, based on the observation that RS does not attain the claimed generality. As we shown before however, RS can still be proven secure either

- According to our weaker notion (Definition 7) for *any* PKE that is correct and IND-CPA.
- According to the original anamorphism notion (Definition 3) for those PKE that are correct, IND-CPA and have high min-entropy ciphertexts.

This leaves the open the question on whether RS is optimal in both contexts. We show this to be the case by extending the message-space upper bound and the impossibility of Fully-Asymmetric AE to both cases. This is formally is formally stated in the following Corollary.

Corollary 1. *Let* $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *be a black-box anamorphic triplet achieving Semi-Adaptive AE and ε-correctness on average, for the class of PKEs that are correct,* IND-CPA *and have high min-entropy ciphertexts. Then*

1. *Its message space M must satisfy* $|M| = \mathsf{poly}(\lambda)$.
2. *There exists a* PPT *adversary breaking weak asymmetric security [CGM24b] (see Appendix A.3).*

Proof of Corollary 1. Regarding the limitation to PKEs with high min-entropy ciphertexts it suffices to observe that the ideal encryption scheme proposed in [CGM24b, Sec. 3.1] has high min-entropy ciphertexts. This is true as, given a message m, a public key pk and fixing a PKE oracle, the encryption is defined as $\psi(\mathsf{pk}, m, r)$ for a random string $r \in \{0,1\}^\rho$ where $\rho = \Omega(\lambda)$ and ψ is a (fixed) injective function. Thus

$$\mathrm{H}_\infty(\mathsf{E.Enc}(\mathsf{pk}, m)) \;=\; \mathrm{H}_\infty(\psi(\mathsf{pk}, m, r)) \;=\; \mathrm{H}_\infty(r) \;=\; \rho \;=\; \Omega(\lambda).$$

Regarding the Semi-Adaptive notion we assume the triplet to achieve, we need to show that Lemmas 1, 2 and 3 (the latter being referred as the *ciphertext-selection lemma*) in [CGM24b] works even in this case.

- [CGM24b, Lemma 1] still applies as the adversary makes no encryption query, and is therefore also a valid adversary for the game in Definition 7.
- [CGM24b, Lemma 2] still applies because the adversary (see [CGM24b, Fig. 5]) makes no usage of the secret key. Thus it is also a valid reduction to Definition 7.
- [CGM24b, Lemma 3] still applies since the adversary (see [CGM24b, Fig. 6]) only uses the secret key after performing all its encryption queries. It is then a valid reduction also to Definition 7 up to syntactical adaptations.

In particular, given the ciphertext selection lemma, the message space upper bound follows through an information theoretic argument. Analogously, the adversary breaking weak-asymmetric anamorphic security's advantage is proven to be significant only through the ciphertext selection lemma and information-theoretic arguments. This conclude the proof of our Corollary. $\square$

5 Extensions to Non-Black-Box Techniques

In this and the following section we study whether known non-black-box tools could be used to bypass our negative results. Recall that for plain Anamorphic Encryption these include the impossibility in Theorem 1 and the bound in Corollary 1 (first part). Fur Fully-Asymmetric encryption instead only Corollary 1 (second part) applies. Regarding non-black-box techniques, we specifically focus on the usage of NIZKs [BFM88], garbling [Yao86] and obfuscation [BGI$^+$01].

This section is devoted to plain Anamorphic Encryption, providing evidence suggesting that those tools would not be helpful. Section 6 instead addresses the case of Fully-Asymmetric anamorphism. In particular we show how it can be *generically* realized (albeit with small message space) from obfuscation.

5.1 Verifiable Obfuscation

In order to address the above question, we begin introducing a (strong) primitive that subsumes NIZK, garbling and obfuscation. We target Virtual-Black-Box Verifiable Obfuscation (VO), a natural extension of the notion presented in [BGJS16]. Informally, VO enhance plain obfuscation by allowing to obfuscate a circuit C along with a (public) predicate P. Everyone can then later verify that $P(C) = 1$ given only P and an obfuscation of C.

Next, we need to adjust our model. Note that assuming powerful tools such as VO relative to a PKE oracle is not sufficient to bypass our negative results. The issue is that obfuscation-like techniques do not *relativize*, or informally, we cannot obfuscate oracles[11]. To address this, we study black-box constructions relative to the given PKE oracles and an *ideal* obfuscator. To obfuscate, it simply assign a random label and to evaluate, it retrieve the circuit associated to said label and evaluates it[12]. The advantage is that circuits with oracle-call gates to the PKE can now be obfuscated. More in details, our ideal obfuscator is defined by a length-preserving random permutation $\xi : \{0,1\}^* \to \{0,1\}^*$, i.e. such that $\xi : \{0,1\}^n \to \{0,1\}^n$ is a random permutation for all n. A full description is provided in Figure 9. Is easy to see ideal VO implies the aforementioned tools. This is formally stated in the following Lemma. A proof appears in the Appendix, Section C.7.

Lemma 2. *Relative to a PKE oracle and the ideal VO in Figure 9, there exist:*

- *NIZKs for all* NP *relations* $\mathcal{R}$ *relative to the given PKE oracles, i.e. such that* $\mathcal{R}$ *may depend on the PKE input/output relations.*
- *Virtual Black-Box emulation, and in particular indistinguishability obfuscation and garbling, for circuits* C *of polynomial size relative to the PKE oracles, i.e. which may contain PKE gates.*

[11] see for instance the discussion in [HJK$^+$16].
[12] this approach is not new, see for instance [GHMM18, Section 4] for ideal garbling.

$\mathsf{VO.Obf}(C, P)$:

Sample $r \xleftarrow{\$} \{0,1\}^\lambda$

$\widetilde{C} \leftarrow \xi(C, P, r)$

return $\widetilde{C}$

$\mathsf{VO.Eval}(\widetilde{C}, x)$:

$(C, P, r) = \xi^{-1}(\widetilde{C})$

return $C(x)$

$\mathsf{VO.Vfy}(\widetilde{C}, P)$:

$(C, P', r) = \xi^{-1}(\widetilde{C})$

if $P \neq P'$: **return** 0

else : **return** $P(C)$

Fig. 9. Ideal Verifiable Obfuscator. $\xi : \{0,1\}^* \to \{0,1\}^*$ is a length-preserving truly random permutation. A representation of C may contain oracle calls/gates to the PKE.

5.2 Compiling Out Verifiable Obfuscation

We now show that our negative results, as well as those presented in [CGM24b], regarding plain Anamorphic Encryption holds even relative to an ideal VO. We do so proving that any black-box anamorphic triplet defined relative to the PKE oracle and the ideal VO, can be compiled into a new triplet which does not make use of verifiable obfuscation, but is still secure.

The idea is that sender and receiver do not need to hide anything from each other. Hence the sender could safely share the random coins he used to generate the public parameters with the sender, rendering NIZK or obfuscation useless. More formally, assume $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ to be a black-box PKE relative to a verifiable obfuscation oracle (Figure 9). We then produce a new scheme $(\mathsf{AT.Gen}^*, \mathsf{AT.Enc}^*, \mathsf{AT.Dec}^*)$ that does not access the VO oracle and is as secure as the initial triplet. This is presented in Figure 10.

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

Sample a PRP key k

$(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}) \xleftarrow{\$} \mathsf{AT.Gen}^{\mathsf{VO}_k}(\lambda)$

$\mathsf{dk}^* \leftarrow (\mathsf{dk}, k)$

return $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}^*)$

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

Parse $\mathsf{dk}^* = (\mathsf{dk}, k)$

$c \xleftarrow{\$} \mathsf{AT.Enc}^{\mathsf{VO}_k}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$

return c

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

Parse $\mathsf{dk}^* = (\mathsf{dk}, k)$

$\widehat{m} \leftarrow \mathsf{AT.Dec}^{\mathsf{VO}_k}(\mathsf{ask}, \mathsf{dk}, c)$

return $\widehat{m}$

$\mathsf{VO.Obf}_k(C, P)$

Sample $r \xleftarrow{\$} \{0,1\}^\lambda$

$\widetilde{C} \leftarrow f_k(C, P, r)$

return $\widetilde{C}$

$\mathsf{VO.Eval}_k(\widetilde{C}, x)$

$(C, P, r) \leftarrow f_k^{-1}(\widetilde{C})$

return $C(x)$

$\mathsf{VO.Vfy}_k(\widetilde{C}, P')$

$(C, P, r) \leftarrow f_k^{-1}(\widetilde{C})$

return $(P = P') \wedge P(C)$

Fig. 10. Compiler from black-box AE relative to a verifiable obfuscation oracle. f_k is a length-preserving PRP. $\mathsf{VO}_k = (\mathsf{VO.Obf}_k, \mathsf{VO.Eval}_k, \mathsf{VO.Vfy}_k)$.

Theorem 4. *Let* $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ *be a [Semi-Adaptive] black-box anamorphic triplet relative to a verifiable obfuscation oracle. If* f *is a length-preserving strong PRP, then* $(\mathsf{AT.Gen}^*, \mathsf{AT.Enc}^*, \mathsf{AT.Dec}^*)$ *is a secure [Semi-Adaptive] black-box anamorphic triplet.*

A proof of the above theorem appears in Appendix C.8.

Remark 3. The compiler presented in Figure 10 only preserves anamorphic security (or weaker variants thereof). Stronger notions such as Fully-Asymmetric security are not preserved. In particular this does not violate negative results in [CGM24b] (and our extension in Corollary 1) regarding the plain impossibility of Fully-Asymmetric AE.

6 Generic Fully-Asymmetric AE

In this section we keep studying whether our impossibility results could be bypassed through non-black-box techniques, focusing now on *Fully-Asymmetric AE.* In such case, we show constructions are possible from obfuscation. Specifically, we prove two adaptations of the Sahai-Waters encryption to achieve generically[13]:

- *Semi-Adaptive* Anamorphic Security 7.
- Fully-Asymmetric Security 6.

The first one applies to any secure PKE, but suffers an exponential security loss (as in [SW14]). The second one instead only suffers a polynomial loss in its reductions, but requires the PKE to be perfectly correct. As RS, both are also anamorphic extensions [BGH+24].

6.1 From Obfuscation

Our first construction informally follows by interpreting the RS triplet as a *secret-key* encryption scheme, and turning it into a *public-key* one using the same strategy of [SW14]. In details, we modify RS (see Figure 3) as follows: First, the PRF is replaced with a puncturable PRF. Next, given a PRG G, we set the double key as $\widetilde{C}$, i.e. the obfuscation of a program that, on input m and a seed s, returns the evaluation of the PRF on the encryption of m with random coins $G(s)$. In this way, in order to encrypt $(m, \widehat{m})$ the sender looks for a seed such that $\widetilde{C}(m, s) = \widehat{m}$ and eventually returns an encryption of m with randomness $G(s)$. The PRF key k is instead kept as the trapdoor key, and decryption is performed computing $\widehat{m} = f_k(c)$. A full description of the circuit used for obfuscation is presented in Figure 11 while the resulting scheme is illustrated in Figure 12. For now on, we use κ to refer to the value $\mathrm{H}_\infty(\mathsf{E.Enc}(\mathsf{pk}, m))$.

[13] i.e. regardless of the underlying PKE.

$C_{\mathsf{pk},k}(m, s)$

Encrypt $c \leftarrow \mathsf{E.Enc}(\mathsf{pk}, m; G(s))$
Return $f_k(c)$

Fig. 11. Circuit used in the Anamorphic Encryption procedure.

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

Sample $\mathsf{apk}, \mathsf{ask} \xleftarrow{\$} \mathsf{E.Gen}(\lambda)$
Generate k a PRF key for f
Obfuscate $\widetilde{C} \xleftarrow{\$} \mathsf{iO}(C_{\mathsf{apk},k})$
$\mathsf{dk} \leftarrow \widetilde{C},\ \mathsf{tk} \leftarrow k$
return $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$

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

for ϑ times:
  Sample $s \xleftarrow{\$} \{0,1\}^\sigma$
  if $\widetilde{C}(m, s) = \widehat{m}$: // $\widetilde{C} = \mathsf{dk}$.
    return $c \leftarrow \mathsf{E.Enc}(\mathsf{apk}, m; G(s))$
// After ϑ failed attempts
return $\mathsf{E.Enc}(\mathsf{apk}, m)$

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

Parse $\mathsf{tk} = k$ the PRF key
return $f_k(c)$

Fig. 12. Fully-Asymmetric Anamorphic Encryption from iO. $G : \{0,1\}^\sigma \to \{0,1\}^\rho$ is a PRG with $\{0,1\}^\rho$ being the random coins space of $\mathsf{E.Enc}$.

Theorem 5. *If* $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *is an* IND-CPA *public key encryption satisfying Definition 1,* $G : \{0,1\}^\sigma \to \{0,1\}^\rho$ *is a PRG with* $\sigma = \kappa/2$*,* f *is a puncturable PRF, and* iO *is a secure obfuscator. Then the Anamorphic Triplet in Figure 12 yields a Fully-Asymmetric Anamorphic Encryption.*

A proof of the above theorem appears in Appenidx C.9. We remark that, as for the case of RS, the assumption on the PKE having high min-entropy ciphertexts could be removed, although in such case the scheme achieves only semi-adaptive anamorphic security.

6.2 From Obfuscation and Perfect Correctness

Our first construction, obtained adapting Sahai-Waters' scheme, inherits an exponential loss in the security parameter. While in general such a loss is acceptable, as it only means that a higher λ has to be chosen, in the context of Anamorphic Encryption this might not be the case. Indeed it could be possible to choose a concrete λ so that breaking the PKE is unfeasible, but distinguishing regular from anamorphic ciphertexts is not hard. For this reason we propose an alternative construction avoiding the above issue.

From a technical perspective, the security loss mentioned above comes from the PRG G usage. This is used to ensure that the set of ciphertexts reachable via $\mathsf{E.Enc}(\mathsf{pk}, m; G(s))$ is sparse in the set of all ciphertexts, which later means that puncturing f_k on the challenge ciphertext yields a functionally equivalent program. This is necessary to then rely on iO.

We address the issue removing G. For the proof we use a different strategy, assuming the PKE to achieve perfect correctness. First, we modify the obfuscated program adding an unsatisfiable branch in which a fixed output is returned. Such condition in our case is that on input (m, r), the obfuscated program $\widetilde{C}$ checks whether $m = m_1^*$ and $\mathsf{E.Enc}(\mathsf{apk}, m; r) = c^*$, where m_1^*, c^* are hard-coded in $\widetilde{C}$ and c^* is an encryption of m_0^*, i.e., a message different from m_1^*. Then, using IND-CPA of the PKE we make this branch reachable by setting c^* as an encryption of m_1^*. Note that perfect correctness is essential as otherwise it may be possible to find $m' \neq m$ and r such that $c^* = \mathsf{E.Enc}(\mathsf{apk}, m'; r)$.

Formally, the new scheme is obtained setting $C_{\mathsf{pk},k}(m, r)$ as the circuit returning $f_k(c)$ with $c \leftarrow \mathsf{E.Enc}(\mathsf{pk}, m; r)$, and modifying $\mathsf{AT.Enc}$ in Figure 12 by sampling $r \leftarrow^{\$} \{0,1\}^\rho$ and if $\widetilde{C}(m, r) = \widehat{m}$ return $\mathsf{E.Enc}(\mathsf{apk}, m; r)$.

Theorem 6. *If* $(\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ *is a perfectly correct* IND-CPA *secure encryption scheme with high min-entropy ciphertexts (Definition 1), f is a puncturable PRF and* iO *a secure obfuscator, then the Anamorphic Triplet described above yields a Fully-Asymmetric Anamorphic Encryption scheme. Namely, for any* PPT *distinguisher $\mathcal{D}_1$ that distinguishes* RealG *from* $\mathsf{AnamorphicG}$ *there exists an adversary $\mathcal{B}_1$ such that*

$$\mathsf{Adv}(\mathcal{D}_1) \leq \mathsf{Adv}^{\mathsf{PRF}}(\mathcal{B}_1) + q^2\vartheta^2 \cdot 2^{-\kappa}$$

and, for any PPT *adversary $\mathcal{D}_2$ that wins the game* $\mathsf{FAsyAnam}$-IND-CPA *there exist adversaries $\mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ such that*

$$\mathsf{Adv}(\mathcal{D}_2) \leq 2(\vartheta + 1)\mathsf{Adv}^{\text{IND-CPA}}(\mathcal{B}_2) + 2\mathsf{Adv}^{\mathsf{iO}}(\mathcal{B}_3) + \mathsf{Adv}^{\mathsf{PRF}}(\mathcal{B}_4) + 3\vartheta^2 \cdot 2^{-\kappa}.$$

Where $q = \mathsf{poly}(\lambda)$ is the number of queries asked by a distinguisher and $\vartheta = \mathsf{poly}(\lambda)$ is the number of attempts that $\mathsf{AT.Enc}$ *does to anamorphically encrypt.*

A proof of the above theorem appears in Appenidx C.10. Again removing the assumption on the PKE having high min-entropy ciphertexts still allows proving the scheme satisfies semi-adaptive anamorphic security, along with the regular Fully-Asymmetric notion.

References

BCC88. Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.*, 37(2):156–189, 1988.

BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988.

BGH+24. Fabio Banfi, Konstantin Gegier, Martin Hirt, Ueli Maurer, and Guilherme Rito. Anamorphic encryption, revisited. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 3–32. Springer, Heidelberg, May 2024.

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.

BGI+12. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, 2012.

BGI14. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.

BGJS16. Saikrishna Badrinarayanan, Vipul Goyal, Aayush Jain, and Amit Sahai. Verifiable functional encryption. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 557–587. Springer, Heidelberg, December 2016.

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.

BR99. Mihir Bellare and Phillip Rogaway. On the construction of variable-input-length ciphers. In Lars R. Knudsen, editor, *FSE'99*, volume 1636 of *LNCS*, pages 231–244. Springer, Heidelberg, March 1999.

BW13. Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.

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.

CGM24a. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Anamorphic encryption: New constructions and homomorphic realizations. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 33–62. Springer, Heidelberg, May 2024.

CGM24b. Dario Catalano, Emanuele Giunta, and Francesco Migliaro. Limits of black-box anamorphic encryption. In *Advances in Cryptology – CRYPTO 2024*, 2024.

CW14. Jie Chen and Hoeteck Wee. Semi-adaptive attribute-based encryption and improved delegation for Boolean formula. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 277–297. Springer, Heidelberg, September 2014.

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.

GHMM18. Sanjam Garg, Mohammad Hajiabadi, Mohammad Mahmoody, and Ameer Mohammed. Limits on the power of garbling techniques for public-key encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 335–364. Springer, Heidelberg, August 2018.

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.

HJK+16. Dennis Hofheinz, Tibor Jager, Dakshita Khurana, Amit Sahai, Brent Waters, and Mark Zhandry. How to generate and use universal samplers. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 715–744. Springer, Heidelberg, December 2016.

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.

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.

KPTZ13. Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 669–684. ACM Press, November 2013.

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.

LR88. Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.

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.

MMN16. Mohammad Mahmoody, Ameer Mohammed, and Soheil Nematihaji. On the impossibility of virtual black-box obfuscation in idealized models. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 18–48. Springer, Heidelberg, January 2016.

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.

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.

SW14. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.

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.

A Supplementary Definitions

A.1 Pseudorandom Permutations

Definition 8 (PRP). *Let $f : \{0,1\}^s \times \{0,1\}^n \to \{0,1\}^n$, where $s, n = \mathsf{poly}(\lambda)$, then f is a Pseudorandom Permutation (PRP) if for every* PPT *distinguisher $\mathcal{D}$*

$$\left|\Pr\left[\mathcal{D}^{f^*}(\lambda) \to 1\right] - \Pr\left[\mathcal{D}^{f_k}(\lambda) \to 1\right]\right| \le \mathsf{negl}(\lambda)$$

where f^* is a truly random permutation, and the key k is random and uniformly sampled from $\{0,1\}^s$.

If the above condition hold when $\mathcal{D}$ has access to both f_k and f_k^{-1}, we say f to be a *strong* PRP [LR88]. For PRP taking values over a set of variable length strings, the notion *length-preserving* PRF/PRP [BR99] will come in handy.

Definition 9 (Length-Preserving PRP). *Given $S \subseteq \{0,1\}^*$, a PRP $f : \{0,1\}^s \times S \to \{0,1\}^*$ is length-preserving if, for all $k \in \{0,1\}^s$ and for all $x \in S$, it holds that $|f_k(x)| = |x|$.*

If f is also a *strong* PRP then f is a *strong* length-preserving PRP.

A.2 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 13.

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

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. 13. 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 10 (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*

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

A.3 Weak Asymmetric Anamorphic Encryption

Let $\mathcal{D}$ be a PPT adversary, $b \in \{0,1\}$ and $\Sigma = (\mathsf{E.Gen}, \mathsf{E.Enc}, \mathsf{E.Dec})$ be an Anamorphic Triplet. The Weak Asymmetric AE security game is then detailed in Figure 14. The advantage of a distinguisher $\mathcal{D}$ for such game is defined as

$$\begin{aligned}\mathsf{Adv}^{\mathsf{weak\text{-}asy\text{-}anam}}_{\mathcal{D},\Sigma}(\lambda) := \Big|&\Pr\left[\mathsf{Weak\text{-}AsyAnam\text{-}}\mathrm{IND\text{-}CPA}^{0}_{\Sigma}(\lambda,\mathcal{D}) = 1\right] \\ &- \Pr\left[\mathsf{Weak\text{-}AsyAnam\text{-}}\mathrm{IND\text{-}CPA}^{1}_{\Sigma}(\lambda,\mathcal{D}) = 1\right]\Big|.\end{aligned}$$

Definition 11 (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 $\mathcal{D}$*

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

$$
\begin{array}{ll}
\multicolumn{2}{l}{\underline{\mathsf{Weak\text{-}AsyAnam}\text{-IND-CPA}^b_{\Sigma}(\lambda, \mathcal{D})}} \\
1: & (\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk}) \leftarrow^{\$} \mathsf{E.Gen}(\lambda) \\
2: & (m, \widehat{m}_0, \widehat{m}_1) \leftarrow^{\$} \mathcal{D}(\mathsf{apk}, \mathsf{dk}) \\
3: & c \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m}_b) \\
4: & \textbf{return } \mathcal{D}(c)
\end{array}
$$

Fig. 14. Weak Asymmetric Anamorphic Encryption security game.

A.4 Indistinguishability Obfuscator and Puncturable PRFs

We briefly recall the definitions of Indistinguishability Obfuscator [BGI+01] and Puncturable PRFs [BW13, KPTZ13, BGI14], taking notation from [SW14].

Definition 12 (Indistinguishability Obfuscator). *A uniform* PPT *algorithm* iO *is called an Indistinguishability Obfuscator for a circuit class* $\{\mathcal{C}_\lambda\}$ *if:*

– *For all* $\lambda \in \mathbb{N}$, *for all* $C \in \mathcal{C}_\lambda$, *for all inputs* x, *it holds that*

$$\Pr\left[C'(x) = C(x) : C' \leftarrow^{\$} \mathsf{iO}(\lambda, C)\right] = 1.$$

– *For any* PPT *adversaries* $\mathcal{S}, \mathcal{D}$, *there exists a negligible* ε *such that, given* $(C_0, C_1, \sigma) \leftarrow^{\$} \mathcal{S}(\lambda)$, *if* $\Pr[\forall x, C_0(x) = C_1(x)] > 1 - \varepsilon(\lambda)$, *then it holds that*

$$|\Pr[\mathcal{D}(\sigma, \mathsf{iO}(\lambda, C_0)) = 1] - \Pr[\mathcal{D}(\sigma, \mathsf{iO}(\lambda, C_1)) = 1]| \leq \varepsilon(\lambda).$$

Definition 13 (Puncturable PRF). *A triplet of algorithm* (PRF.Gen, PRF.Eval, PRF.Puncture) *is said to be a Puncturable PRF if, given* $n(\lambda), m(\lambda)$ *two computable functions, the two following requirements are satisfied:*

– *For every* PPT *adversary* $\mathcal{A}$ *such that* $\mathcal{A}(\lambda)$ *outputs a set* $S \subseteq \{0,1\}^n$, *then for all* $x \in \{0,1\}^n \setminus S$, *it holds that*

$$\Pr\Big[\mathsf{PRF.Eval}(k, x) = \mathsf{PRF.Eval}(k_S, x) : k \leftarrow^{\$} \mathsf{PRF.Gen}(\lambda), k_S \leftarrow \mathsf{PRF.Puncture}(k, S)\Big] = 1.$$

– *For every* PPT *adversary* $(\mathcal{A}_1, \mathcal{A}_2)$ *such that* $\mathcal{A}_1(\lambda)$ *outputs a set* $S \subseteq \{0,1\}^n$ *and a state* σ, *given* $k \leftarrow^{\$} \mathsf{PRF.Gen}(\lambda), k_S \leftarrow \mathsf{PRF.Puncture}(k, S)$, *it holds that*

$$\begin{aligned}&|\Pr[\mathcal{A}_2(\sigma, k_S, S, \mathsf{PRF.Eval}(k, S)) = 1] \\ &- \Pr[\mathcal{A}_2(\sigma, k_S, S, U(m(\lambda) \cdot |S|)) = 1]| = \mathsf{negl}(\lambda).\end{aligned}$$

Where $\mathsf{PRF.Eval}(k, S)$, *for* $S = \{x_1, \ldots, x_l\}$, *denotes the concatenation of* $\mathsf{PRF.Eval}(k, x_1), \ldots, \mathsf{PRF.Eval}(k, x_l)$ *and* $U(\ell)$ *denotes the uniform distribution over* ℓ *bits.*

B Supplementary Lemmas

B.1 Statistical Distance

Given two discrete random variables x, y distributed over a set S, we define their statistical distance (or *total-variation* or ℓ_1) as

$$\Delta(x,y) \;=\; \frac{1}{2}\sum_{a\in S} \Pr\left[x=a\right] - \Pr\left[y=a\right].$$

The following lemma will come in handy to inductively study the statistical distance of two tuples of random variables.

Lemma 3. *Given four random variables $x_1, x_2 \sim X$, $y_1, y_2 \sim Y$ and setting $X^+ = \{a \in X : \Pr\left[x_i = a\right] > 0,\ i \in [2]\}$, if there exists $A \subseteq X$ such that*

$$P(x_1 \in A) \leq \varepsilon_1, \quad \Delta(x_1, x_2) \leq \varepsilon_2, \quad \Delta(y_{1|x_1=x}, y_{2|x_2=x}) \leq \varepsilon_3 \quad \forall x \in X^+ \setminus A,$$

for positive real numbers $\varepsilon_1, \varepsilon_2, \varepsilon_3 \in \mathbb{R}^+$, then $\Delta((x_1, y_1), (x_2, y_2)) \leq \varepsilon_1 + \varepsilon_2 + \varepsilon_3$.

B.2 Rejection Sampling

The following Lemma shows that performing rejection sampling with a predicate that is independent from the candidate output does not alter the distribution. This represents a common step in the security proof of RS as well as our construction in Section 6.

Lemma 4. *Given probability density p over $\mathcal{X}$, $c_1, \ldots, c_{\vartheta+1}$ independently sampled from this distribution and $b_1, \ldots, b_\vartheta \sim \{0,1\}$, let c be equal to c_i for the smallest i such that $b_i = 1$, or $c_{\vartheta+1}$ if no such i exists.*

If $b_1, \ldots, b_\vartheta$ are distributed uniformly and independently from each others and from $c_1, \ldots, c_{\vartheta+1}$, then c is distributed over $\mathcal{X}$ with probability density p.

Proof. For any $c_0 \in \mathcal{X}$, we proceed computing $\Pr\left[c = c_0\right] =$

$$\begin{aligned}
&= \sum_{i=1}^{\vartheta} \Pr\left[c_i = c_0 \,\middle|\, \begin{matrix} b_1 = \ldots = b_{i-1} = 0 \\ b_i = 1 \end{matrix}\right] \cdot \Pr\left[\begin{matrix} b_1 = \ldots = b_{i-1} = 0 \\ b_i = 1 \end{matrix}\right] + \\
&\qquad + \Pr\left[c_{\vartheta+1} = c_0 \,|\, b_1 = \ldots = b_\vartheta = 0\right] \cdot \Pr\left[b_1 = \ldots = b_\vartheta = 0\right] \\
&= \sum_{i=1}^{\vartheta} p(c_0) \cdot \frac{1}{2^i} + p(c_0) \cdot \frac{1}{2^\vartheta} \;=\; p(c_0).
\end{aligned}$$

The second equality follows as the bits b_i are independently distributed and uniform over $\{0,1\}$, and the fact that $\Pr\left[c_i = c_0\right] = p(c_0)$ as we assumed c_i to follow the distribution defined by p. The thesis follows. $\square$

C Postponed Proof

C.1 IND-CPA of the Counterexample PKE

Proof of Proposition 1. Correctness follows as F is injective, B is disjoint from the original PKE's ciphertext space, and because of the initial PKE's correctness.

Regarding IND-CPA, let $\mathcal{A}$ be an adversary for the weakened scheme. We design $\mathcal{B}$ breaking the original PKE. Informally, on input pk, $\mathcal{B}$ samples a random message m^*, computes $y^* = F(m^*)$ and runs $\mathcal{A}(\mathsf{pk}, y^*)$. Once $\mathcal{A}$ returns m_0, m_1, it either aborts if one of them equals m^*, or sends them to its oracle otherwise. Upon receiving c, it forwards the reply to $\mathcal{A}$ and eventually returns the same bit $\mathcal{A}$ outputs upon halting. A detailed description of $\mathcal{B}$ is given in Figure 15.

$\mathcal{B}^{\mathcal{O}}(\mathsf{pk})$:

$m^* \xleftarrow{\$} M$
$y^* = F(m^*)$
Run $\mathcal{A}(\mathsf{pk}, y^*)$
$(m_0, m_1) \xleftarrow{\$} \mathcal{A}$
if $m_0 = m^* \vee m_1 = m^*$ **then**
    abort
$c \xleftarrow{\$} \mathcal{O}(m_0, m_1)$
Give c to $\mathcal{A}$
return $\mathcal{A}$'s output

Fig. 15. Adversary $\mathcal{B}$ for the IND-CPA of the original PKE from adversary $\mathcal{A}$ for the IND-CPA of the weakened PKE. $\mathcal{O}$ is the encryption oracle for the IND-CPA game provided to $\mathcal{B}$.

Define Abort as the event in which $\mathcal{B}$ aborts before making its oracle query, i.e., the event in which m_0 or m_1 is a preimage of y^*. Using the security of F we show it to occur with negligible probability. Let $\mathcal{C}$ be the following adversary attempting to invert F: on input y^* it generates $(\mathsf{pk}, \mathsf{sk})$ with $\mathsf{E.Gen}$, runs $\mathcal{A}(\mathsf{pk}, y^*)$ and, once it returns (m_0, m_1), checks whether $F(m_0) = y^*$ or $F(m_1) = y^*$. Clearly $\mathcal{C}$ simulates perfectly the view of $\mathcal{A}$ executed by $\mathcal{B}$ and it successfully inverts F if and only if $\mathcal{B}$ aborts. Thus $\Pr[\mathsf{Abort}] = \mathsf{Adv}(\mathcal{C})$ which is negligible because F is an injective OWF.

Finally, if $\neg\mathsf{Abort}$, $\mathcal{B}$ perfectly simulates the encryption oracle because $\mathsf{E.Enc}^*$ behaves as $\mathsf{E.Enc}$ on all messages but m^*. We thus conclude that $\mathsf{Adv}(\mathcal{A}) \leq \mathsf{Adv}(\mathcal{B}) + \Pr[\mathsf{Abort}] \leq \mathsf{negl}(\lambda)$. $\square$

C.2 Counterexample to Rejection Sampling

Before providing the proof of Claim 1, we recall the Markov lower-bound. Let X be a real random-variable with $0 \leq X \leq t$ and $\mu = \mathbb{E}[X]$. Then

$$\Pr[X \leq \alpha] \;\leq\; \frac{t-\mu}{t-\alpha}.$$

Proof of Claim 1. Without loss of generality, assume $\mathsf{RS.Enc}$ first computes ϑ ciphertexts, and later select the correct one if possible. Let C_0, C_1 the sets of those ϑ ciphertexts[14] $\mathsf{RS.Enc}$ computed by $\mathsf{RS.Enc}$ when encrypting $(m^*, 0)$ and $(m^*, 1)$. We will show that up to probability $1/4$, each set has size at least $\vartheta/2$ through a Markov argument. Indeed, as $|B| = 4\vartheta$, on expectation

$$\mathbb{E}[|C_\beta|] \;=\; 4\vartheta\left(1-\left(1-\frac{1}{4\vartheta}\right)^{\vartheta}\right) \;\geq\; \vartheta \cdot 4\left(1-\frac{1}{\sqrt[4]{e}}\right) \;\geq\; \vartheta \cdot \frac{7}{8}$$

where the first equality is taken summing the indicators $c \in C_\beta$ for $c \in B$, and the last can be verified numerically and is only used for notational convenience. Using Markov lower bound, as $0 \leq |C_\beta| \leq \vartheta$, we have that

$$\Pr[|C_\beta| \leq \vartheta/2] \;\leq\; \frac{\vartheta-(7/8)\vartheta}{\vartheta-(1/2)\vartheta} \;=\; \frac{1}{4}.$$

Up to probability $1/2$ we can then assume $|C_0| \geq \vartheta/2$ and $|C_1| \geq \vartheta/2$. Finally, under such condition, Fail_β only occurs if f_k assumes value $1-\beta$ for all elements in C_β. As this occurs with negligible probability for a truly random function, because $\vartheta/2 = \Omega(\lambda)$, it also occurs with negligible probability for f_k. We thus conclude that

$$\begin{aligned}
\Pr[\mathsf{Fail}] \;&\leq\; \Pr[\mathsf{Fail}_0] + \Pr[\mathsf{Fail}_1] \\
&\leq\; \Pr[\mathsf{Fail}_0 \,|\, |C_0| > \vartheta/2] + \Pr[|C_0| \leq \vartheta/2] \\
&\quad\; + \Pr[\mathsf{Fail}_1 \,|\, |C_1| > \vartheta/2] + \Pr[|C_1| \leq \vartheta/2] \\
&\leq\; 1/2 + \mathsf{negl}(\lambda).
\end{aligned}$$

□

C.3 IND-CPA of the Ideal PKE

Proof of Lemma 1. Perfect correctness immediately follows by the definition of ψ. Regarding IND-CPA, let $\mathcal{A}$ be a PPT adversary with oracle access to the four procedures in Figure 6. To fix notation let $\mathcal{A}(\mathsf{pk}) \to (m_0, m_1)$, and $c^* \xleftarrow{\$} \mathsf{E.Enc}(\mathsf{pk}, m_b)$ the challenge ciphertext sent, where b is the challenge bit and $\mathsf{pk} = \phi(\mathsf{sk})$. We assume $\mathcal{A}$ to perform at most $q = \mathsf{poly}(\lambda)$ oracle calls. Then we define three bad events. The first one BK captures $\mathcal{A}$ finding sk. The second one BM occurs when m_0 or m_1 is weak with respect to pk. The third one BC says that $\mathcal{A}$ find (pk, m, r) whose encryption yields c^*. Formally

[14] These sets may not be distinct.

– BK: $\mathcal{A}$ queries $\mathsf{E.Gen}$, $\mathsf{E.Dec}$ or $\mathsf{E.Find}$ on sk.
– BM: $m_\beta \in \{m_i^*(\mathsf{sk})\}_{i=1}^{\lambda}$ for some $\beta \in \{0,1\}$.
– BC: $\mathcal{A}$ queries $c^* \leftarrow \mathsf{E.Enc}(\mathsf{pk}, m; r)$.

We claim these to be negligible.

Claim 2. *Let* $\mathsf{Bad} = \mathsf{BK} \vee \mathsf{BM} \vee \mathsf{BC}$. *Then* $\Pr[\mathsf{Bad}] \leq \mathsf{negl}(\lambda)$.

Let v the view[15] of $\mathcal{A}$. Then we show that, for all v_0 satisfying $\neg\mathsf{Bad}$, $\mathcal{A}$ has almost no information on b, i.e., conditioning on $v = v_0$ then b is almost uniformly distributed from the point of view of $\mathcal{A}$. Toward this goal, let R_0 and R_1 be random coins not figuring in $\mathcal{A}$'s encryption queries respectively for m_0 and m_1 with public key pk. Further let us call $f_b(\cdot) = \mathsf{E.Enc}(\mathsf{pk}, m_b; \cdot)$. Then, from $\neg\mathsf{BK}$, c^* is uniformly distributed over $f_0(R_0) \cup f_1(R_1)$ conditioning on $v = v_0$, as it was never decrypted and never obtained through encryption queries. Moreover, as m_0, m_1 are not weak, $|f_\gamma(R_\beta)| = |R_\beta|$ (for $\gamma, \beta \in \{0,1\}$). Since $b = 0$ if and only if $c^* \in f_0(R_0)$ we have

$$\Pr[b = 0 \mid v = v_0] = \Pr[c^* \in f_0(R_0) \mid v = v_0] = \frac{|f_0(R_0)|}{|f_0(R_0) \cup f_1(R_1)|} = \frac{|R_0|}{|R_0| + |R_1|}$$

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

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

Note that the same bounds for $b = 1$ and that the second term of the sum is negligible for $\rho = \Omega(\lambda)$. We can thus conclude that, calling b' the final bit guessed by $\mathcal{A}$

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

□

Proof of Claim 2. Regarding BK, let $\mathsf{sk}_1, \ldots, \mathsf{sk}_q$ the secret keys queried. As $\phi : \mathsf{SK} \rightarrow \mathsf{PK}$ is a random bijection, we have that $\mathsf{sk} = \phi^{-1}(\mathsf{pk})$ is uniformly distributed among the keys not-yet-queried until correctly guessed. Hence

$$\begin{aligned}
\Pr[\exists j : \mathsf{sk}_j = \mathsf{sk}] &\leq \sum\nolimits_{j=1}^{q} \Pr[\mathsf{sk}_j = \mathsf{sk} \mid \mathsf{sk} \notin \{\mathsf{sk}_1, \ldots, \mathsf{sk}_{j-1}\}] \\
&\leq \sum\nolimits_{j=1}^{q} \frac{1}{|\mathsf{SK}| - (j-1)} \;\leq\; \frac{q}{|\mathsf{SK}| - q}
\end{aligned}$$

which is negligible as $|\mathsf{SK}| = \Omega(2^\lambda)$.
Next we study $\mathsf{BM} \wedge \neg\mathsf{BK}$. Let $m_1, \ldots, m_{q'}$ the messages involved in any query of

[15] i.e. the joint distribution of $\mathcal{A}$'s input, random coins and oracle replies.

$\mathcal{A}$. In order to include also the two challenge messages let $q = q' + 2$. As we condition on $\neg\mathsf{BK}$, $m_i^*(\mathsf{sk})$ is uniformly distributed among the non-yet queried messages (pessimistically assuming that each query involving a message immediately reveals whether it is weak or not). For ease of notation let $M^* = \{m_i^*(\mathsf{sk})\}_{i=1}^{\lambda}$. Then

$$\begin{aligned}\Pr\left[\exists j : m_j \in M^*,\, \neg\mathsf{BK}\right] &= \sum\nolimits_{j=1}^{q} \Pr\left[m_j \in M^*,\, \neg\mathsf{BK} \,|\, m_1, \ldots, m_{j-1} \notin M^*\right] \\ &= \sum\nolimits_{j=1}^{q} \frac{\lambda}{2^\mu - (j-1)} \;\leq\; \frac{q\lambda}{2^\mu - q}\end{aligned}$$

that is negligible as we assumed $\mu = \Omega(\lambda)$.
Finally we study $\mathsf{BC} \wedge \neg\mathsf{BK} \wedge \neg\mathsf{BM}$. In this case c^* is never decrypted and m_b is not a weak message (as neither m_0 or m_1 are). Thus, calling r^* the random coins used, we have that an encryption queries for (m, r) returns c^* if

$$\psi_\tau(\mathsf{pk}, m, r) = \psi_\tau(\mathsf{pk}, m_b, r^*) = \psi(\mathsf{pk}, m_b, r^*) \qquad \Longleftrightarrow \qquad m = m_b,\; r = r^*.$$

Finally, as r^* is uniformly random among the random tapes not yet queried due to the definition of ψ, we conclude that, calling $r_1, \ldots, r_q$ the randomness appearing in all $\mathcal{A}$'s queries

$$\begin{aligned}\Pr\left[\exists j : r_j = r^*,\, \neg\mathsf{BK}, \neg\mathsf{BM}\right] &= \sum\nolimits_{j=1}^{q} \Pr\left[r_j = r^*,\, \neg\mathsf{BK}, \neg\mathsf{BM} \,\middle|\, r^* \notin \{r_i\}_{i=1}^{j-1}\right] \\ &= \sum\nolimits_{j=1}^{q} \frac{1}{2^\rho - (j-1)} \;\leq\; \frac{q}{2^\rho - q}.\end{aligned}$$

This is negligible as $\rho = \Omega(\lambda)$.
Combining the three inequalities we get

$$\Pr\left[\mathsf{Bad}\right] \;\leq\; \frac{q}{|\mathsf{SK}| - q} + \frac{q\lambda}{2^\mu - q} + \frac{q}{2^\rho - q} \;=\; \mathsf{negl}(\lambda). \qquad \square$$

C.4 Impossibility of Stateless Black-Box AE

Proof of Theorem 1. We begin computing the probability that $\mathcal{A}$ returns 1 when executed in the real game. In this case $c_{0,i}$ and $c_{1,i}$ are 2ϑ ciphertexts computed with randomness $r_{0,i}, r_{1,i}$. Regarding the check in Line 4, two encryptions of the same messages collides only if their actual random coins (returned by τ, see Section 3.2) do. To simplify notation, let us call $\tau^*(\cdot) = \tau(\mathsf{pk}, m^*, \cdot)$. Then, we claim that a collision with respect to τ^* is likely.

Claim 3. *With the previous notation*

$$\Pr\left[\exists i, j \,:\, \tau^*(r_{0,i}) = \tau^*(r_{1,i})\right] \;\geq\; \frac{1}{2} - \frac{1}{2}\exp\left(-\frac{2\vartheta^2}{\nu}\right) - \mathsf{negl}(\lambda).$$

This concludes the first half of the proof as $\Pr[\mathcal{A}_{\vartheta,\nu} \to 1 \mid \mathsf{RealG}] =$

$$\begin{aligned} &= \Pr[\exists i,j \,:\, c_{0,i} = c_{0,j}] \\ &= \Pr[\exists i,j \,:\, \psi_\tau(\mathsf{pk}, m^*, r_{0,i}) = \psi_\tau(\mathsf{pk}, m^*, r_{1,j})] \\ &= \Pr[\exists i,j \,:\, \tau^*(r_{0,i}) = \tau^*(r_{1,j})] \;\geq\; (1-e^{-1})/2. \end{aligned}$$

Regarding the behavior of $\mathcal{A}$ in $\mathsf{AnamorphicG}$ we will prove it returns 1 with probability bounded by $o(\lambda^{-1})$. We do so first showing that the view of $\mathsf{AT.Enc}$ on input m^* is not statistically far from its view on a random message m. Then use ε-correctness on average to prove ciphertexts rarely collides. We recall that $q = \mathsf{poly}(\lambda)$ is the number of queries made by each algorithm of the black-box anamorphic triplet. For the first step we require the following claim:

Claim 4. *Let* View_b *and* View_b^* *be the joint views*[16] *of* $\mathsf{E.Gen}(\lambda) \overset{\$}{\to} (\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$ *and respectively of* $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, b)$ *and* $\mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m^*, b)$ *with* m *a random message. Then* $\Delta(\mathsf{View}_b, \mathsf{View}_b^*) \leq \frac{q^2}{2\nu} + \mathsf{negl}(\lambda)$.

Let $c'_{0,i}, c'_{1,j}$ be ciphertexts obtained encrypting a random message m instead of m^* during the execution of $\mathcal{A}$. The probability of $\mathcal{A}$ returning 1 can then be bounded by

$$\begin{aligned} \Pr[\mathcal{A}_{\vartheta,\nu} \to 1 \mid \mathsf{AnamorphicG}] &= \Pr[\exists i,j \,:\, c_{0,i} = c_{1,j}] \\ &\leq \Pr[\exists b,i \,:\, \mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c_{b,i}) \neq b] \\ &\leq \sum\nolimits_{b,i} \Pr[\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c_{b,i}) \neq b] \\ &\leq \sum\nolimits_{b,i} \left(\Pr\left[\mathsf{AT.Dec}(\mathsf{ask}, \mathsf{tk}, c'_{b,i}) \neq b\right] + \frac{q^2}{2\nu} + \mathsf{negl}(\lambda) \right) \\ &\leq \frac{\vartheta q^2}{\nu} + 2\vartheta\varepsilon + \mathsf{negl}(\lambda). \end{aligned}$$

The first inequality follows as any collision of the given type yields a ciphertext that decrypts incorrectly. The second is a union bound. The third is Claim 4 and the last uses ε-average correctness as mentioned.
Combining the two halves, and recalling $\nu = \lambda^2 q^4$, $\vartheta = \sqrt{\nu/2}$, a bound on the advantage of $\mathcal{A}$ can be derived as

$$\mathsf{Adv}(\mathcal{A}_{\vartheta,\nu}) \;\geq\; \frac{1-e^{-1}}{2} - \frac{1}{\lambda \cdot \sqrt{2}} - \mathsf{negl}(\lambda) \;=\; \Omega(1) - o(\lambda^{-1}). \qquad \square$$

Proof of Claim 3. First of all, to simplify notation, we call R_0 the set of $r_{0,i}$, R_1 the set of $r_{1,i}$ and R their union. We begin with a general result, that is, assuming all entries in R to be distinct, given a random function $f : R \to S$ and calling for simplicity $n = |R|$, $n_b = |R_b|$ and $\mathsf{Coll}(f, R_0, R_1)$ the event in which there exists $x_0 \in R_0$ and $x_1 \in R_1$ colliding w.r.t. f, then

$$\Pr[\mathsf{Coll}(f, R_0, R_1)] \;\geq\; \frac{2n_0 n_1}{n(n-1)} \cdot \Pr[|f(R)| < |R|].$$

[16] i.e. the joint distribution of inputs, random coins, and oracle replies.

To show this let F be the set of all functions from R to S, F^* the set of functions with a collision, and $\pi : R \to R$ a random permutation. Then

$$\begin{aligned}
\Pr[\mathsf{Coll}(f, R_0, R_1)] &= \Pr[\mathsf{Coll}(f \circ \pi, R_0, R_1)] \\
&= \sum\nolimits_{f_0 \in F} \Pr[\mathsf{Coll}(f_0 \circ \pi, R_0, R_1)] \Pr[f = f_0] \\
&= \sum\nolimits_{f_0 \in F^*} \Pr[\mathsf{Coll}(f_0 \circ \pi, R_0, R_1)] \Pr[f = f_0] \,.
\end{aligned}$$

The first equality follows as f and $f \circ \pi$ have the same distribution, while the last as when $f_0 \notin F^*$ then there is no collision at all. Next, given $f_0 \in F^*$, let $x, y \in R$ two points that collides. Then we observe that there are $2n_0n_1(n-2)!$ permutation mapping x in R_0 and y to R_1 or vice versa. As this condition would imply $\mathsf{Coll}(f_0 \circ \pi, R_0, R_1)$ we have

$$\begin{aligned}
&\geq \sum\nolimits_{f_0 \in F^*} \frac{2n_0n_1(n-2)!}{n!} \cdot \Pr[f = f_0] \\
&= \frac{2n_0n_1}{n(n-1)} \sum\nolimits_{f_0 \in F^*} \Pr[f = f_0] \;=\; \frac{2n_0n_1}{n(n-1)} \Pr[|f(R)| < |R|]
\end{aligned}$$

Where the last equality follows by our definition of F^*. This conclude the first part of the proof.

Returning now to our original problem, let Diff be the event that all $r_{b,i}$ are different, i.e. $|R| = 2\vartheta$. Note $\Pr[\neg\mathsf{Diff}] \leq \vartheta^2 \cdot 2^{-\rho}$, which is negligible. Note this does not occur with probability smaller than $\vartheta^2 2^{-\rho}$ that is negligible. Then, conditioning on Diff we can derive from the first part that

$$\begin{aligned}
\Pr[\exists i, j\ \tau^*(r_{0,i}) = \tau^*(r_{1,i}) \mid \mathsf{Diff}] &\geq \frac{2\vartheta^2}{2\vartheta(2\vartheta - 1)} \cdot \Pr[|\tau^*(R)| < 2\vartheta \mid \mathsf{Diff}] \\
&\geq \frac{\vartheta}{2\vartheta - 1} \cdot \left(1 - \exp\left(-\frac{(2\vartheta)^2}{2\nu}\right)\right) \\
&\geq \frac{1}{2} \cdot \left(1 - \exp\left(-\frac{2\vartheta^2}{\nu}\right)\right)
\end{aligned}$$

where the second inequality is the birthday paradox lower bound as τ^* has range of size ν. This completes the proof as we observed $\Pr[\neg\mathsf{Diff}] \leq \mathsf{negl}(\lambda)$. $\square$

Proof of Claim 4. Define $\mathsf{View}_b = (v_{\mathsf{gen}}, r, m, v_1, \ldots, v_q)$ and $\mathsf{View}_b^* = (v_{\mathsf{gen}}, r, m^*, v_1^*, \ldots, v_q^*)$ the two views, where v_{gen} is the view of $\mathsf{E.Gen}$, r is the random tape of $\mathsf{AT.Enc}$, and v_i, v_i^* are the oracle responses.

First we show m^* is observed (i.e. is involved in a decryption/encryption query) with negligible probability. Indeed, calling $m_1, \ldots, m_q$ the observed messages (at most one per query), as m^* is uniformly distributed since $\mathsf{AT.Gen}$ performs no query to $\mathsf{E.Find}$, we have that $\Pr[m^* \in \{m_1, \ldots, m_q\}] \leq q/2^\mu$.

Next, conditioning on $v_{\mathsf{gen}} = v_0$ such that m^* is not observed, we have that m is uniformly distributed by construction, whereas m^* is uniform over the set of non-observed messages. Thus

$$\Delta(m_{|v_{\mathsf{gen}}=v_0},\ m^*_{|v_{\mathsf{gen}}=v_0}) \;\leq\; \frac{q}{2^\mu} \quad\Rightarrow\quad \Delta((v_{\mathsf{gen}}, r, m), (v_{\mathsf{gen}}, r, m^*)) \;\leq\; \frac{2q}{2^\mu}$$

where the implication follows from the inductive hypothesis and Lemma 3. Next we show by induction on $h \in \{1, \ldots, q\}$ that the statistical distance between the given view until the h-th query of $\mathsf{AT.Enc}$ is

$$\Delta((v_{\mathsf{gen}}, r, m, v_1, \ldots, v_h),\ (v_{\mathsf{gen}}, r, m^*, v_1^*, \ldots, v_h^*)) \ \leq\ \frac{h^2}{2\nu} + \frac{(q+h)^2\lambda}{2^{\mu+1}} + \frac{2q}{2^\mu}$$

Let $\mathbf{v}, \mathbf{v}^*$ be the two vectors limited to the first $h-1$ queries. First of all we bound the probability that $\mathsf{AT.Enc}$ and $\mathsf{AT.Gen}$ observes a weak message, excluding the input message m. Calling $m_1, \ldots, m_{q+h-1}$ the observed messages, indeed, $m_i^* := m_i^*(\mathsf{sk})$, for $i \in \{1, \ldots, \lambda\}$, are uniformly distributed (until correctly guessed). Thus

$$\Pr\left[\exists i, j\ :\ m \neq m_i^*(\mathsf{sk}) = m_j\right] \ \leq\ \frac{(q+h)\lambda}{2^\mu}.$$

next, conditioning on $\mathbf{v} = \mathbf{v}_0 = \mathbf{v}^*$ for which the above does not happen, we study the statistical distance of v_h, v_h^*. According to the type of the h-th query, three cases have to be considered.

- $\mathsf{E.Gen}(\mathsf{sk}')$: The reply $\phi(\mathsf{sk}')$ is equally distributed in both views.
- $\mathsf{E.Dec}(\mathsf{sk}', c')$: If $\mathsf{sk}' \neq \mathsf{sk}$ the reply is the same in both cases. If c' was previously obtained as the encryption of m', the reply is consistent (i.e. it is m') in both views. Finally, if c' was not previously observed, then in both views the probability that c' is not decrypted to $\perp$ is smaller than $2^{\mu+\rho}/(2^\ell - (h+q))$. Thus in this case

  $$\Delta(v_{h\,|\,\mathbf{v}=\mathbf{v}_0},\ v^*_{h\,|\,\mathbf{v}^*=\mathbf{v}_0}) \ \leq\ \frac{2^{\mu+\rho}}{2^\ell - 2q} \ \leq\ \frac{h}{\nu}.$$

  The last inequality holds for sufficiently large λ as $\ell - (\mu + \rho) = \Omega(\lambda)$.
- $\mathsf{E.Enc}(\mathsf{pk}', m'; r')$: If the query was already performed the result is consistent. If $\mathsf{pk} \neq \mathsf{pk}'$ the response's distribution is the same. If $m' \neq m_0$ (where m_0 is the third entry of $\mathbf{v}_0$, defined above) and the query was not performed, let C be the set of observed ciphertexts. Then in both cases c is uniformly distributed over $\{0,1\}^\ell \setminus C$. Note this is also true as we assumed that weak messages (other than m_0) were not queried before.
  Finally, if the query is $\mathsf{E.Enc}(\mathsf{pk}, m_0; r')$, let C_0 be the ciphertext obtained so far as encryptions of m_0. Then in the first distribution c is uniform over $\{0,1\}^\ell \setminus C$. In the second one instead c collides with a previously observed encryption of m_0 with probability $1/\nu$ and is otherwise uniformly distributed over $\{0,1\}^\ell \setminus C$. More precisely

  $$\begin{aligned} c_0 \in C_0 \quad &\Rightarrow \quad \Pr\left[c = c_0 \,|\, \mathbf{v}^* = \mathbf{v}_0\right] \ =\ \frac{1}{v} \\ c_0 \in C \setminus C_0 \quad &\Rightarrow \quad \Pr\left[c = c_0 \,|\, \mathbf{v}^* = \mathbf{v}_0\right] \ =\ 0 \\ c_0 \in \{0,1\}^\ell \setminus C \quad &\Rightarrow \quad \Pr\left[c = c_0 \,|\, \mathbf{v}^* = \mathbf{v}_0\right] \ =\ \left(1 - \frac{|C_0|}{\nu}\right) \cdot \frac{1}{2^\ell - |C|}. \end{aligned}$$

  We thus conclude that in this case $\Delta(v_{h\,|\,\mathbf{v}=\mathbf{v}_0},\ v^*_{h\,|\,\mathbf{v}^*=\mathbf{v}_0}) \leq |C_0|/\nu \leq h/\nu$.

Combining this with the inductive hypothesis yields the thesis (by Lemma 3). Finally, this proves the inductive statement to hold for $h = q$ which is our thesis up to observing that the other two terms are negligible as $\mu = \Omega(\lambda)$. □

C.5 Rejection-Sampling is Semi-Adaptive AE

Proof of Theorem 2. We proceed through a sequence of hybrids starting from the anamorphic game. First we replace the PRF in RS with a truly random function, and later substitute each of those function invocation with the sampling of a fresh random value. Finally we conclude showing the last game's ciphertexts to follow the right distribution, i.e. that of freshly generated ones, information-theoretically.

H_0: The real Anamorphic Encryption game $\mathsf{AnamorphicG}$.

H_1: As in H_0 but the PRF is substituted by a truly random function f^*.

H_2: As in H_1 but instead of invoking f^*, sample a fresh random bit.

H_3: The real encryption game RealG.

$\mathsf{H}_0 \approx \mathsf{H}_1$ follows directly from the PRF security and the efficient simulatability of the PKE oracles. Indeed, if f is a PRF, then the advantage between the two games for an adversary $\mathcal{A}$ is $\mathsf{negl}(\lambda)$.

$\mathsf{H}_1 \approx \mathsf{H}_2$. Let $\mathcal{D}$ be a q queries distinguisher executed in G_{1+b} for a uniformly random bit $b \leftarrow^{\$} \{0,1\}$. To fix notation, let $(m_i, \widehat{m}_i)$ be the message involved in its q encryption queries, and $c_{i,j}$ for $j \in \{1, \ldots, \vartheta + 1\}$ the regular ciphertexts computed by the challenger to produce an Anamorphic Encryption of $(m_i, \widehat{m}_i)$ though rejection sampling. Then we define Coll the event that a collision occurs among those ciphertexts. If $\neg\mathsf{Coll}$, the random function f^* is always evaluated on distinct points, and thus computing $f^*(c_{i,j})$ is equivalent to sampling a random bit. Thus $\mathcal{D}$ has no advantage in this case and in particular $\mathsf{Adv}(\mathcal{D}) \leq \Pr[\mathsf{Coll}]$.

Next, we bound $\Pr[\mathsf{Coll}]$ using the PKE's security. Let $\mathcal{A}(\mathsf{pk})$ be the following IND-CPA adversary: initially it runs $\mathcal{D}(\mathsf{pk})$ and chooses a random pair of (distinct) indices $\alpha, \beta \in [q] \times [\vartheta + 1]$. Next, it simulates $\mathcal{D}$'s game. However when producing the α-th regular ciphertext it either encrypts m or 0 (according to the IND-CPA encryption oracle) with m the regular message requested by $\mathcal{D}$. Similarly, for the β-th ciphertext it either encrypts m or 1. Finally, it returns 1 if the α-th and β-th ciphertexts collided. A full description is presented in Figure 16.

Let b' be the IND-CPA's challenge bit, i.e. when $b' = 0$ the first message is encrypted, whereas the opposite occurs with $b' = 1$. Then it is immediate to see that when $b' = 0$, $\mathcal{A}$ perfectly simulates $\mathcal{D}$'s game until its last query. Indeed pk sampled from $\mathsf{E.Gen}$ matches the distribution of apk and all ciphertexts $c_{i,j}$ are computed as $\mathsf{E.Enc}(\mathsf{pk}, m_i)$. Finally, $\mathcal{D}$ has no information of α, β. Thus, setting $\widehat{q} = q(\vartheta + 1)$ the total number of encryption calls performed by $\mathcal{A}$, and χ a

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

```
Sample α, β ←$ [q] × [ϑ + 1] distinct couples, and b ←$ {0, 1}
Run D(pk)
when it queries (m_i, m̂_i):
  for j ∈ [ϑ + 1]: // Generate ciphertexts
    if (i, j) = α: c_{i,j} ←$ O(m_i, 0)
    elseif (i, j) = β: c_{i,j} ←$ O(m_i, 1)
    else : c_{i,j} ←$ E.Enc(pk, m_i)
  for j ∈ [ϑ]: // Rejection sampling
    if b = 0: b_{i,j} ← f*(c_{i,j})
    if b = 1: b_{i,j} ←$ {0, 1}
    if b_{i,j} = m̂_i: Reply with c_{i,j} and break
  // If no ciphertext was chosen through rejection sampling
  Reply with c_{i,ϑ+1}
// The execution of D is interrupted after the last query
return c_α == c_β
```

Fig. 16. Adversary $\mathcal{A}$ for IND-CPA from $\mathcal{D}$ distinguishing G_1 from G_2. $\mathcal{O}$ is the IND-CPA oracle encrypting either the first or the second message according to its challenge bit. f^* is a (lazily maintained) random function to $\{0,1\}$.

random variable denoting the number of ciphertexts couples colliding, then

$$\begin{aligned}
\Pr[\mathcal{A} \to 1 \mid b' = 0] &= \Pr[c_\alpha = c_\beta \mid b' = 0] \\
&= \sum_k \Pr[c_\alpha = c_\beta \mid b' = 0, \chi = k] \Pr[\chi = k] \\
&= \sum_{k \geq 1} k \cdot \binom{\widehat{q}}{2}^{-1} \cdot \Pr[\chi = k] \;\geq\; \binom{\widehat{q}}{2}^{-1} \cdot \sum_{k \geq 1} \Pr[\chi = k] \\
&= \binom{\widehat{q}}{2}^{-1} \cdot \Pr[\chi > 0] \;=\; \binom{\widehat{q}}{2}^{-1} \cdot \Pr[\mathsf{Coll}] .
\end{aligned}$$

The third equality follows as (α, β) is a uniformly distributed couple. The first inequality uses $k \geq 1$, while the last equality follows as $\chi > 0$ is the same event as Coll.

Conversely, when $b' = 1$, the two ciphertexts collides only if an encryption errors occurs. Indeed, as decryption is stateless and deterministic, when $c_\alpha = c_\beta$ either $\mathsf{E.Dec}(\mathsf{sk}, c_\alpha) \neq 0$ or $\mathsf{E.Dec}(\mathsf{sk}, c_\beta) \neq 1$. Using the scheme's correctness then

$$\begin{aligned}
\Pr[\mathcal{A} \to 1 \mid b' = 1] &\leq \Pr[\mathsf{E.Dec}(\mathsf{sk}, c_\alpha) \neq 0 \,\vee\, \mathsf{E.Dec}(\mathsf{sk}, c_\beta) \neq 1] \\
&\leq \Pr[\mathsf{E.Dec}(\mathsf{sk}, c_\alpha) \neq 0] + \Pr[\mathsf{E.Dec}(\mathsf{sk}, c_\beta) \neq 1] \;\leq\; \mathsf{negl}(\lambda).
\end{aligned}$$

Combining both part we finally get a bound which proves Coll only occurs with negligible probability, due the PKE's IND-CPA security.

$$\Pr[\mathsf{Coll}] \leq \binom{\widehat{q}}{2} \cdot \mathsf{Adv}(\mathcal{A}) + \mathsf{negl}(\lambda) \quad \Rightarrow \quad \mathsf{Adv}(\mathcal{D}) \leq \Pr[\mathsf{Coll}] \leq \mathsf{negl}(\lambda).$$

$\mathsf{G}_2 \overset{p}{=} \mathsf{G}_3$. We argue the two game to be equivalent, as rejection sampling in G_2 is performed on freshly sampled bits distributed independently from previously observed values, and upon failure a correctly generated ciphertext is returned. This is formally stated and proved in Lemma 4. □

C.6 Overcoming impossibility

Proof of Theorem 3. We proceed through a sequence of hybrids, relying first on the PRF security used in the rejection sampling construction (Figure 3), then we show an upper bound on the biased ciphertexts distribution.

H_0: The real Anamorphic Encryption game $\mathsf{AnamorphicG}$.
H_1: As in H_0 but the PRF is substituted by a truly random function f^*.
H_2: As in H_1 but instead of invoking f^*, sample a fresh random bit.
H_3: The real encryption game RealG.

$\mathsf{H}_0 \approx \mathsf{H}_1$ follows directly from the PRF security, and the efficient simulatability of the PKE oracles. Indeed, if the function used is a PRF, then the advantage in distinguishing the two games for an adversary $\mathcal{A}$ is $\mathsf{negl}(\lambda)$.

Proof of $\mathsf{H}_1 \approx \mathsf{H}_2$*.* The only way to distinguish the two games is distinguishing the distribution of the ciphertexts they produce. In both games (possibly) many ciphertexts are produced before choosing one of them. The only difference between the two games is that in the case of H_1 the choice of the ciphertext to return is biased from the output of the random function f^*, while in the case of H_2 the choice is biased from an uniformly sampled random bit.

Let $\mathsf{CollG1}$ be the event that in H_1 two encryption queries to $\mathsf{E.Enc}$ are answered with the same ciphertext at least one time, i.e., the probability that the encryption oracle returns two ciphertexts that collide on different messages. Given the fact that the PKE satisfies Definition 1, that $\mathsf{AT.Enc}$ tries ϑ times to find the right ciphertext, and that at most $q = \mathsf{poly}(\lambda)$ messages are queried, it holds that

$$\Pr[\mathsf{CollG1}] \leq \binom{q\vartheta}{2} \cdot 2^{-\mathrm{H}_\infty(\mathsf{E.Enc})} \leq q^2\vartheta^2 \cdot 2^{-\mathrm{H}_\infty(\mathsf{E.Enc})} \leq \mathsf{negl}(\lambda).$$

A similar bound holds for the event $\mathsf{CollG2}$, that is the same event as $\mathsf{CollG1}$ but defined regard to H_2. For the same argument above, it holds that

$$\Pr[\mathsf{CollG2}] \leq \binom{q\vartheta}{2} \cdot 2^{-\mathrm{H}_\infty(\mathsf{E.Enc})} \leq q^2\vartheta^2 \cdot 2^{-\mathrm{H}_\infty(\mathsf{E.Enc})} \leq \mathsf{negl}(\lambda).$$

Now, we can bound the advantage of an adversary distinguishing the two games as:

$$\begin{aligned}|\Pr[\mathsf{H}_1 = 1] - \Pr[\mathsf{H}_2 = 1]| = &|\Pr[\mathsf{H}_1 = 1 \mid \mathsf{CollG1}]\Pr[\mathsf{CollG1}] \\ &+ \Pr[\mathsf{H}_1 = 1 \mid \neg\mathsf{CollG1}]\Pr[\neg\mathsf{CollG1}] \\ &- \Pr[\mathsf{H}_2 = 1 \mid \mathsf{CollG2}]\Pr[\mathsf{CollG2}] \\ &- \Pr[\mathsf{H}_2 = 1 \mid \neg\mathsf{CollG2}]\Pr[\neg\mathsf{CollG2}]| \\ = &|\Pr[\mathsf{H}_1 = 1 \mid \mathsf{CollG1}]\Pr[\mathsf{CollG1}] \\ &- \Pr[\mathsf{H}_2 = 1 \mid \mathsf{CollG2}]\Pr[\mathsf{CollG2}] + \mathsf{negl}(\lambda)| \\ \leq &|\Pr[\mathsf{CollG1}] - \Pr[\mathsf{CollG2}] + \mathsf{negl}(\lambda)| = \mathsf{negl}(\lambda).\end{aligned}$$

where the second equality follows from the fact that, conditioning on not having collisions in both games, since in H_1 the value $f^*(c_i) = \widehat{m}_i$ is independet from c_i and the same happens for H_2 regarding the uniformly sampled bit, the two distributions of ciphertexts are exactly the same and $\Pr[\neg\mathsf{CollG1}] \approx \Pr[\neg\mathsf{CollG2}]$.

We can conclude that the two games are indistinguishable.

$\mathsf{H}_2 \overset{p}{=} \mathsf{H}_3$. We argue the two game to be equivalent, as rejection sampling in H_2 is performed on freshly sampled bits distributed independently from previously observed values, and upon failure a correctly generated ciphertext is returned. This is formally stated and proved in Lemma 4. □

C.7 Verifiable Obfuscation implications

First, we recall the definition of NIZK argument and VBB. Next, we prove the Lemma 2.

Definition 14 (NIZK argument [BFM88, BCC88]). *A Non Interactive Zero Knowledge (NIZK) argument for an* NP *relation* $\mathcal{R}$ *is a tuple of three algorithms* $(\mathsf{NIZK.S}, \mathsf{NIZK.P}, \mathsf{NIZK.V})$*, called prover and verifier, where*

- $\mathsf{NIZK.S}(\lambda) \xrightarrow{\$} \mathsf{crs}$ *on input the security parameter* λ *outputs a common reference string* crs*.*
- $\mathsf{NIZK.P}(\mathsf{crs}, x, w) \xrightarrow{\$} \pi$ *on input the common reference string* crs*, a statement* x *and a witness* w *outputs a proof* π *that* $(x, w) \in \mathcal{R}$*.*
- $\mathsf{NIZK.V}(\mathsf{crs}, x, \pi) \to b$ *on input the common reference string* crs*, a statement* x *and a proof* π *accept or reject the proof, i.e., output the bit* 1 *if it is a valid proof, else* 0*.*

and such that the following properties are satisfied

Perfect Completeness: *For all* $(x, w) \in \mathcal{R}$ *it holds that*

$$\Pr\left[\mathsf{NIZK.V}(\mathsf{crs}, x, \pi) \to 1 \,\middle|\, \pi \xleftarrow{\$} \mathsf{NIZK.P}(\mathsf{crs}, x, w)\right] = 1.$$

Computational Soundness: *For every x for which does not exists w such that $(x, w) \in \mathcal{R}$, and for every* PPT *adversaries $\mathcal{A}$, it holds that*

$$\Pr\left[\mathsf{NIZK.V}(\mathsf{crs}, x, \pi) \to 1 \,\middle|\, \pi \xleftarrow{\$} \mathcal{A}(x)\right] \leq \mathsf{negl}(\lambda).$$

Computational Zero Knowledge: *There exists a* PPT *simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$ such that, up to a negligible function* $\mathsf{negl}(\lambda)$*, for every $(x, w) \in \mathcal{R}$, for every* PPT *adversaries $\mathcal{A}$ it holds that*

$$\left|\Pr\left[\mathcal{A}(\mathsf{crs}_0, \pi_0) \to 1 \,\middle|\, \pi_0 \xleftarrow{\$} \mathsf{NIZK.P}(\mathsf{crs}_0, x, w)\right] - \right.$$
$$\left. - \Pr\left[\mathcal{A}(\mathsf{crs}_1, \pi_1) \to 1 \,\middle|\, \pi_1 \xleftarrow{\$} \mathcal{S}_2(\mathsf{crs}_1, x)\right]\right| \leq \mathsf{negl}(\lambda).$$

where $\mathsf{crs}_0 \xleftarrow{\$} \mathsf{NIZK.S}(\lambda)$ *and* $\mathsf{crs}_1 \xleftarrow{\$} \mathcal{S}_1(\lambda)$.

Definition 15 (VBB [BGI+12]). *A uniform* PPT *algorithm* O *is called a Virtual Black-Box Obfuscator (VBB) for a circuit class $\mathcal{C}_\lambda$ if the three following conditions are satisfied:*

– *For all $\lambda \in \mathbb{N}$, for all $C \in \mathcal{C}_\lambda$, for all inputs x, it holds that*

$$\Pr\left[C'(x) = C(x) : C' \xleftarrow{\$} \mathsf{O}(\lambda, C)\right] = 1.$$

– *There exists a polynomial p such that for all $C \in \mathcal{C}_\lambda$, it holds that*

$$|\mathsf{O}(C)| \leq p(|C|).$$

– *For any* PPT *adversaries $\mathcal{A}$, there exists a simulator $\mathcal{S}$ and a negligible ε such that for all $\lambda \in \mathbb{N}$ and for all circuits $C \in \mathcal{C}_\lambda$ then it holds that*

$$\Pr\left[\mathcal{A}(\mathsf{O}(C)) \to 1\right] - \Pr\left[\mathcal{S}^C(1^{|C|}) \to 1\right] \leq \varepsilon(|C|).$$

Proof of Lemma 2. We provide constructions for the two primitives separately.

NIZK. Let $\mathcal{R}$ be an NP relation relative to PKE oracles, and D a circuit relative to the same PKE oracles such that $(x, w) \in \mathcal{R}$ if and only if $D(x, w) = 1$. For a given x, let $C_{x,w}$ be a constant circuit that returns $D(x, w)$ on any input, and $P_x(C)$ the predicate that is true if $C = C_{x,w}$ for some w. Note that as w is plainly hard-coded in $C_{x,w}$, P is efficiently computable. We can then define a NIZK argument as follows:

– $\mathsf{NIZK.S}(\lambda)$: Return the empty string ϵ.
– $\mathsf{NIZK.P}(x, w)$: Return $\widetilde{C} \xleftarrow{\$} \mathsf{VO.Obf}(C_{x,w}, P_x)$.
– $\mathsf{NIZK.V}(x, \widetilde{C})$: Accept only if $\mathsf{VO.Vfy}(\widetilde{C}, P_x) \to 1$ and $\mathsf{VO.Eval}(\widetilde{C}, \perp) \to 1$.

Correctness follows as on input $(x, w) \in \mathcal{R}$, $C_{x,w}$ always returns 1. Perfect soundness hold as, given $\widetilde{C}$, if $\mathsf{VO.Vfy}(\widetilde{C}, P_x) \to 1$ then there exists w such that $\widetilde{C} = \mathsf{VO.Obf}(C_{x,w}, P_x)$. Moreover, $\mathsf{VO.Eval}(\widetilde{C}, \perp) \to 1$ means that $D(x, w) = 1$, and in particular $(x, w) \in \mathcal{R}$. Finally, to show computational zero-knowledge, we present an straight-line simulator $\mathcal{S}$ relative to the PKE interacting with a malicious verifier $\mathcal{V}^*$. $\mathcal{S}$ handles PKE queries forwarding them, and to VO ones by lazily maintaining a random length-preserving permutation ξ. In order to simulate a proof for x, it computes $\widetilde{C}^* \leftarrow \mathsf{VO.Obf}(C^*, P_x) = \xi(r, C^*, P_x)$ where r is a random λ-bit long string and C^* is the constant circuit always returning 1. Evaluations are carried out as prescribed by the oracles, while queries to $\mathsf{VO.Vfy}(\widetilde{C}^*, P_x)$ are answered with 1. The view $\mathcal{S}$ produces follows the same distribution observed with $\mathsf{NIZK.P}(x, w)$, as long as $\mathcal{V}^*$ never queries $\mathsf{VO.Obf}$ on an input that returns $\widetilde{C}^*$, the received proof. The latter case however occurs with probability at most $2^{-\lambda}$ for each query in both worlds. Calling q the total number of queries performed by $\mathcal{V}^*$ then, the statistical distance between the real and simulated view is smaller than $q \cdot 2^{-\lambda}$.

VBB. [17] This is simply realized obfuscating a program along with the predicate $\perp$ that is always false. Formally $\mathsf{O}^{\mathsf{VO}}(C) = \mathsf{VO.Obf}(C, \perp)$. To show this is a VBB we provide a simulator $\mathcal{S}$ relative to PKE oracles for a given adversary $\mathcal{A}$. As before, $\mathcal{S}$ will lazily maintain a length-preserving random permutation. Initially, given 1^ℓ with $\ell = |C|$, it sets $\widetilde{C} = \xi(r, 0^\ell, \perp)$ and executes $\mathcal{A}(\widetilde{C})$. When $\mathcal{A}$ queries the PKE oracles, $\mathcal{S}$ forwards them and their replies. When $\mathcal{A}$ queries to VO are replied honestly with the exception of $\mathsf{VO.Eval}(\widetilde{C}, x)$. In this case $\mathcal{S}$ queries $y = C(x)$ (recall $\mathcal{S}$ has oracle access to C) and returns y. Finally, $\mathcal{S}$ output the same bit as $\mathcal{A}$.

It is immediate to see that unless $\mathcal{A}$ obtains $\widetilde{C}$ from an oracle call, its view interacting with $\mathcal{S}$ is the same as when it interacts with the real VO oracles. As the first events occurs with probability $q \cdot 2^{\lambda}$ with q being the total number of queries, we have that

$$\left|\Pr\left[\mathcal{A}^{\mathsf{VO}}(\mathsf{O}^{\mathsf{VO}}(C)) \to 1\right] - \Pr\left[\mathcal{S}^{C}(1^{|C|}) \to 1\right]\right| \leq q \cdot 2^{-\lambda} = \mathsf{negl}(\lambda). \qquad \square$$

C.8 Compiling out Verifiable Obfuscation

Proof of Theorem 4. The only difference between the given triplet, and the one defined in Figure 10 lies in the inner verifiable obfuscation oracle. In particular the given scheme uses a truly random permutation ξ, whereas our compiler relies on a PRP with key k embedded in the double key.

In the following we only prove that our compiler preserves regular anamorphic security, as the case of Semi-Adaptive AE is analogous. Relative to any efficiently simulatable PKE oracle, we define two hybrid games: H_0, that is the anamorphic game with $(\mathsf{AT.Gen}^*, \mathsf{AT.Enc}^*, \mathsf{AT.Dec}^*)$, and H_1 that is the

[17] See [MMN16] for an in-depth discussion of VBB in idealized models.

anamorphic game with $(\mathsf{AT.Gen}, \mathsf{AT.Dec}, \mathsf{AT.Enc})$. Given a distinguisher $\mathcal{D}$ we describe $\mathcal{B}$ against the PRP security. At a high level, $\mathcal{B}$ executes $\mathcal{D}(\mathsf{apk}, \mathsf{ask})$ and $(\mathsf{AT.Gen}, \mathsf{AT.Enc}, \mathsf{AT.Dec})$ simulating the PKE oracles, which we assumed to be efficiently simulatable. To emulates the VO calls, it behaves as the ideal VO described in Figure 9, except that to evaluate ξ and ξ^{-1} it invokes the PRP oracles for f and f^{-1}. Note $\mathsf{apk}, \mathsf{ask}$ are generated via $\mathsf{AT.Gen}$ and can be computed as they do not depend on k (as opposed to dk^* in H_1).

It is immediate to observe that in the ideal world $\mathcal{B}$ perfectly emulates H_1 as the PRP oracles behave as a truly random length-preserving permutation f^*. Conversely, the PRP oracles gives access to f_k and f_k^{-1} meaning that $\mathcal{B}$ replies to VO queries as for VO_k described in Figure 10. Thus in this case it perfectly emulates H_0 and in particular $\mathsf{Adv}(\mathcal{D}) = \mathsf{Adv}(\mathcal{B}) = \mathsf{negl}(\lambda)$.

This concludes the proof as distinguishing the real game with the given PKE in Definition 3 from the anamorphic one, i.e. H_1, is computationally hard according to our hypothesis. □

C.9 First Construction from Obfuscation

Proof of Theorem 5. The proof is divided in two parts. First we show the basic anamorphism and next we prove Fully-Asymmetric security.

Basic Anamorphic Security. We proceed with a sequence of hybrids $\mathsf{H}_0, \ldots, \mathsf{H}_4$.

- H_0: The anamorphic game $\mathsf{AnamorphicG}$. Public parameters $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$ are generated through $\mathsf{AT.Gen}(\lambda)$. Encryption queries $(m, \widehat{m})$ are answered with a ciphertext $c \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$.
- H_1: As H_0 but replacing $G(s)$ with a random sampled $r \in \{0,1\}^\rho$.
- H_2: As H_1 but when executing $\mathsf{AT.Enc}$, replace the check in Line 3 with $f_k(c) = \widehat{m}$ where $c \leftarrow \mathsf{E.Enc}(\mathsf{apk}, m; r)$.
- H_3: As H_2 but $f_k(\cdot)$ is replaced with a truly random function f^*.
- H_4: As H_3 but encryption queries $(m, \widehat{m})$ are answered with $c \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{apk}, m)$.

Trivially, H_4 corresponds to the real game RealG as $\mathsf{apk}, \mathsf{ask}$ are sampled with $\mathsf{E.Gen}(\lambda)$. $\mathsf{H}_0 \approx \mathsf{H}_1$ follows from the PRG security. $\mathsf{H}_1 \approx \mathsf{H}_2$ follows from correctness of obfuscation. $\mathsf{H}_2 \approx \mathsf{H}_3$ as f_k is pseudorandom. Note in both experiments a distinguisher only observes $\mathsf{apk}, \mathsf{ask}$, both of which are generated independently from k, and evaluations of f_k, which are obtainable through oracle queries in the pseudorandomness game. Toward proving $\mathsf{H}_3 \approx \mathsf{H}_4$ let $c_1, \ldots, c_{q\vartheta}$ be the ciphertexts $\mathsf{AT.Enc}$ computes in H_2 to answer the q queries performed by a distinguisher. Then, as we assumed the PKE to satisfy Definition 1, the probability for a given pair of those ciphertexts to be equal is smaller than $2^{-\kappa}$. Thus, calling Coll the event $c_i = c_j$ for some $i \neq j$, a union bound yields $\Pr[\mathsf{Coll}] \leq q^2\vartheta^2 \cdot 2^{-\kappa}$. Conditioning on $\neg\mathsf{Coll}$, as all ciphertexts are different, the bits $f^*(c_1), \ldots, f^*(c_{q\vartheta})$ are uniformly and independently distributed. Thus $\mathsf{AT.Enc}$'s choice of the resulting ciphertext does not depend on those observed during its execution, meaning that its distribution is identical to the prescribed one.

Fully-Asymmetric Security. We proceed through a sequence of hybrids. To fix notation, we recall the game syntax. Initially the adversary $\mathcal{A}$ receives $(\mathsf{apk}, \mathsf{dk})$, where $\mathsf{dk} = \widetilde{C}$ in our case, outputs $(m_0, \widehat{m}_0, m_1, \widehat{m}_1)$ and receive c^* the Anamorphic Encryption of $(m_b, \widehat{m}_b)$ for a uniformly sampled challenge bit b.

H_0: The $\mathsf{FAsyAnam}$-IND-CPA game with challenge bit b.

H_1: As H_0 but c^* is computed as $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m_b, \widehat{m}_b)$, see Figure 17.

H_2: As H_1 but c^* is set to $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m^*, \widehat{m}_b)$ for a uniformly sampled m^*.

H_3: As H_2 but c^* is computed as $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m^*, b)$.

H_4: As H_3 but c^* is computed in the setup after $(\mathsf{apk}, \mathsf{ask}, k)$ are generated.

H_5: As H_4 but, calling $c_1, \ldots, c_{\vartheta+1}$ the intermediate ciphertexts computed by $\mathsf{AT.Enc}^*$, set $k^* \leftarrow \mathsf{PRF.Puncture}(k, c_1, \ldots, c_{\vartheta+1})$ and $\widetilde{C} \xleftarrow{\$} \mathsf{iO}(C^*_{\mathsf{apk},k^*})$.

H_6: As H_5 but c^* is computed as $\mathsf{E.Enc}(\mathsf{apk}, m^*)$.

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

1 : **for** $i \in \{1, \ldots, \vartheta\}$: Encrypt $c_i \xleftarrow{\$} \mathsf{E.Enc}(\mathsf{apk}, m)$

2 : **for** $i \in \{1, \ldots, \vartheta\}$: **if** $f_k(c_i) = \widehat{m}$: **return** c_i

3 : **return** $c_{\vartheta+1}$

Fig. 17. Alternative encryption used in the proof of Theorem 5.

Guessing b in H_6 is information-theoretically hard. $\mathsf{H}_3 = \mathsf{H}_4$ as only the order of operations is changed. We will show H_2 and H_3 to be equally hard, and the remaining hybrids to be indistinguishable.

$\mathsf{H}_0 \approx \mathsf{H}_1$. We reduce to the PRG security for $(\vartheta + 1)$ instances. Given a distinguisher $\mathcal{D}$ for the two games, let $\mathcal{B}$ be an adversary for the above problem. On input $r_1, \ldots, r_{\vartheta+1}$, it generates $\mathsf{apk}, \mathsf{ask}, k, \widetilde{C}$ as in H_0, get $(m_0, \widehat{m}_0, m_1, \widehat{m}_1)$ from $\mathcal{D}$ and computes $c_i \leftarrow \mathsf{E.Enc}(\mathsf{apk}, m_b; r_i)$ with b being a uniformly sampled random challenge. Then, it set c^* as the first ciphertext c_i such that $f_k(c_i) = \widehat{m}_b$, or to $c_{\vartheta+1}$ if not such ciphertext exists. Finally it sends c^* to $\mathcal{D}$ and eventually return the same bit returned by $\mathcal{D}$.

Clearly, if $r_i = G(s_i)$ for independently sampled s_i, then $\mathcal{B}$ perfectly simulates H_0, also thanks to iO's perfect correctness. Conversely, if r_i are uniformly random, $\mathcal{B}$ perfectly simulates H_1. Thus $\mathsf{Adv}(\mathcal{B}) = \mathsf{Adv}(\mathcal{D})$.

$\mathsf{H}_1 \approx \mathsf{H}_2$. Let $\mathcal{D}$ be a distinguisher for the two games. Then we define an adversary $\mathcal{B}$ breaking IND-CPA of the given scheme. On input pk it sets $\mathsf{apk} = \mathsf{pk}$ and generates $k, \widetilde{C}$ as in H_2. Once $\mathcal{D}(\mathsf{apk}, \widetilde{C}) \rightarrow (m_0, \widehat{m}_0, m_1, \widehat{m}_1)$, it samples a random bit b and computes, using its oracle, c_i as the encryption of either m_b or m^* for a uniformly sampled m^*. The challenge ciphertext c^* is then chosen among $c_1, \ldots, c_{\vartheta+1}$ as the first ciphertexts such that $f_k(c_i) = \widehat{m}_b$ or $c_{\vartheta+1}$ if none satisfy this condition. Finally, when $\mathcal{D}$ outputs a bit and halts, $\mathcal{B}$ returns the same bit.

It is immediate to see $\mathcal{B}$ perfectly emulates H_1 and H_2 when its oracle encrypts respectively the first or the second component of each query. Note this holds as in H_1 the ciphertexts c_i are computed using random coins that are uniformly sampled – as opposed as being generated through the PRG. Thus $\mathsf{Adv}(\mathcal{B}) = \mathsf{Adv}(\mathcal{D})$.

H_2 *is harder than* H_3. Given an adversary $\mathcal{A}$ for H_2, we define $\mathcal{B}$ guessing b in H_3. On input $(\mathsf{apk}, \mathsf{dk})$, it simply runs $\mathcal{A}(\mathsf{apk}, \mathsf{dk}) \xrightarrow{\$} (m_0, \widehat{m}_0, m_1, \widehat{m}_1)$. If $\widehat{m}_0 = \widehat{m}_1$, it aborts returning 0. Otherwise, it queries the encryption oracle with $(m_0, \widehat{m}_0, m_1, \widehat{m}_1)$ obtaining c^* and sends it to $\mathcal{A}$. Once $\mathcal{A}$ returns b', $\mathcal{B}$ returns $\widehat{m}_0 \oplus b'$.

Let Equal be the event $\mathcal{A}(\mathsf{apk}, \mathsf{dk})$ returns $\widehat{m}_0 = \widehat{m}_1$. Then conditioning on Equal, the advantage of $\mathcal{A}$ is 0 as it obtain no information on its challenge bit, which we call β. Hence, upper-bounding $\Pr[\neg\mathsf{Equal}] \leq 1$,

$$\begin{aligned} \mathsf{Adv}(\mathcal{A}) &= |\Pr[\mathcal{A} \to 1 \,|\, \beta = 0] - \Pr[\mathcal{A} \to 1 \,|\, \beta = 1]| \\ &\geq |\Pr[\mathcal{A} \to 1 \,|\, \beta = 0, \neg\mathsf{Equal}] - \Pr[\mathcal{A} \to 1 \,|\, \beta = 1, \neg\mathsf{Equal}]| \end{aligned}$$

Conversely, conditioning on $\neg\mathsf{Equal}$, we have $\widehat{m}_\beta = \widehat{m}_0 \oplus \beta = f_k(c^*) = b$. In particular $\mathcal{B}$ perfectly simulates the view of $\mathcal{A}$ given $\neg\mathsf{Equal}$ and challenge bit $\beta = \widehat{m}_0 \oplus b$. Thus

$$\begin{aligned} \mathsf{Adv}(\mathcal{B}) &= |\Pr[\mathcal{B} \to 1 \,|\, b = 0] - \Pr[\mathcal{B} \to 1 \,|\, b = 1]| \\ &= |\Pr[\mathcal{A} \to \widehat{m}_1 \,|\, \beta = \widehat{m}_0, \neg\mathsf{Equal}] - \Pr[\mathcal{B} \to \widehat{m}_1 \,|\, \beta = \widehat{m}_1, \neg\mathsf{Equal}]| \\ &= |\Pr[\mathcal{A} \to 1 \,|\, \beta = 0, \neg\mathsf{Equal}] - \Pr[\mathcal{B} \to 1 \,|\, \beta = 1, \neg\mathsf{Equal}]| \\ &\leq \mathsf{Adv}(\mathcal{A}). \end{aligned}$$

Where the third equality follows conditioning each term on $\widehat{m}_0 = 0$ and $\widehat{m}_1 = 1$, taking the negative event where necessary to always have $\mathcal{A} \to 1$ and rearranging.

$\mathsf{H}_4 \approx \mathsf{H}_5$. We begin by showing that, since each c_i is computed with real random coins, it is unlikely for them to be *reachable* by the circuit $C(m, s)$. More precisely we claim that

Claim 5. *Given* $\mathsf{apk}, \mathsf{ask} \xleftarrow{\$} \mathsf{AT.Gen}(\lambda)$*, a uniformly sampled message* m^**, and* $c \leftarrow \mathsf{E.Enc}(\mathsf{apk}, m^*; r)$ *with uniformly sampled coins* r*, then*

$$\Pr[\exists (m, s) : c = \mathsf{E.Enc}(\mathsf{apk}, m; G(s))] \;\leq\; \mathsf{negl}(\lambda).$$

Given the claim, let Bad_i the event that $\exists (m, s)$ such that $c_i = \mathsf{E.Enc}(\mathsf{apk}, m; G(s))$ and Bad the disjunction of $\mathsf{Bad}_1, \ldots, \mathsf{Bad}_{\vartheta+1}$. Then through a union bound we have that $\Pr[\mathsf{Bad}] \leq (\vartheta + 1)\mathsf{negl}(\lambda)$. Finally, due to puncturing correctness we have that $C_{\mathsf{apk},k}$ and C_{apk,k^*} agree on all inputs (m, s) such that $\mathsf{E.Enc}(\mathsf{apk}, m; G(s)) \notin$

$\{c_1, \ldots, c_{\vartheta+1}\}$. Conditioning on $\neg\mathsf{Bad}$ this is never the case. Security of the obfuscator can thus be invoked in this case. More specifically, calling $\mathcal{D}$ a distinguisher for the two games, simulating either H_4 or H_5 by obfuscating respectively $C_{\mathsf{apk},k}$ or C_{apk,k^*}, and computing correctly the other responses, and calling $\mathcal{B}$ an adversary against the obfuscation security, we can conclude that

$$\mathsf{Adv}(\mathcal{B}) \;\geq\; \mathsf{Adv}(\mathcal{D}) - 2\Pr[\mathsf{Bad}] \qquad \Rightarrow \qquad \mathsf{Adv}(\mathcal{D}) \leq \mathsf{negl}(\lambda).$$

Proof of Claim 5. At a high level, $c = \mathsf{E.Enc}(\mathsf{apk}, m; G(s))$ can happen for three reasons:

1. c is an incorrect ciphertext for m^*, i.e. $\mathsf{E.Dec}(\mathsf{ask}, c) \neq m^*$.
2. c is correct and correctly reachable, i.e. $c = \mathsf{E.Enc}(\mathsf{apk}, m^*; G(s))$.
3. c is correct but incorrectly reachable, i.e. $c = \mathsf{E.Enc}(\mathsf{apk}, m; G(s))$ for some $m \neq m^*$.

The first case occurs with negligible probability from ε-correctness. The second one too as there are at most $2^\sigma = 2^{\kappa/2} \geq 2^{\lambda/2}$ ciphertexts of the form $\mathsf{E.Enc}(\mathsf{apk}, m^*; G(s))$ for fixed apk and m^*, but $\kappa = \mathrm{H}_\infty(c) \geq \lambda$ as we assumed Definition 1 to hold. Regarding the third we use a Markov argument.

To fix notation, let $p(m_0, m_1)$, $S(m_0, m_1)$ and $B(m_0, m_1)$ be respectively the probability that an encryption (using G to generate the random coins) of m_0 yields a ciphertext decrypting to m_1, the set of seeds for which this happens and the set of *bad* ciphertexts obtained. Formally

$$\begin{aligned}
p(m_0, m_1) &= \Pr[\mathsf{E.Dec}(\mathsf{ask}, \mathsf{E.Enc}(\mathsf{apk}, m_0; G(s))) = m_1] \\
S(m_0, m_1) &= \{s_0 \in \{0,1\}^\sigma \;:\; \mathsf{E.Dec}(\mathsf{ask}, \mathsf{E.Enc}(\mathsf{apk}, m_0; G(s_0))) = m_1\} \\
B(m_0, m_1) &= \{\mathsf{E.Enc}(\mathsf{apk}, m_0; G(s_0)) \;:\; s_0 \in S(m_0, m_1)\}
\end{aligned}$$

Intuitively, this defines a weighted graph among messages, and our goal is to argue that an average vertex has *low* weighted in-degree. We define such weighted in-degrees as:

$$p^+(m_1) = \sum_{m_0 : m_0 \neq m_1} p(m_0, m_1)$$

$$S^+(m_1) = \bigcup_{m_0 : m_0 \neq m_1} S(m_0, m_1) \qquad B^+(m_1) = \bigcup_{m_0 : m_0 \neq m_1} B(m_0, m_1)$$

First of all we claim that $\mathsf{E.Enc}$ remains correct when using a PRG to sample its random coins on average. Formally that for a random message m and seed s

$$\Pr[\mathsf{E.Dec}(\mathsf{ask}, \mathsf{E.Enc}(\mathsf{apk}, m; G(s))) \neq m] \;\leq\; \varepsilon'(\lambda)$$

for a negligible ε'. This is proven studying an adversary for the PRG which generates $\mathsf{apk}, \mathsf{ask}$, samples a random message, and given r that is either $G(s)$ or random, checks the above condition to be true. Given this bound, we can study the expectation of $p^+(m^*)$:

$$\begin{aligned}
\varepsilon' &\geq \Pr\left[\mathsf{E.Dec}(\mathsf{ask}, \mathsf{E.Enc}(\mathsf{apk}, m; G(s))) \neq m\right] \\
&= \sum\nolimits_{m_0} \Pr\left[\mathsf{E.Dec}(\mathsf{ask}, \mathsf{E.Enc}(\mathsf{apk}, m_0; G(s))) \neq m_0\right] \cdot \frac{1}{|M|} \\
&= \frac{1}{|M|} \cdot \sum\nolimits_{m_0} \sum\nolimits_{m_1 : m_1 \neq m_0} p(m_0, m_1) \;=\; \frac{1}{|M|} \cdot \sum\nolimits_{m_1 \neq m_0} p(m_0, m_1) \\
&= \frac{1}{|M|} \cdot \sum\nolimits_{m_1} p^+(m_1) \;=\; \mathbb{E}\big[p^+(m^*)\big].
\end{aligned}$$

Let now T be the set of those $(\mathsf{apk}_0, \mathsf{ask}_0, m_0^*)$ such that $p^+(m_0^*) \leq 1$. Then Markov inequality implies that $\Pr\left[(\mathsf{apk}, \mathsf{ask}, m^*) \notin T\right] \leq \varepsilon'$. Conversely assuming $(\mathsf{apk}, \mathsf{ask}, m^*) \in T$, i.e. $p^+(m^*) \leq 1$, we give an upper bound on the number of "bad ciphertexts" $|B^+(m^*)|$. Indeed

$$\begin{aligned}
|B^+(m^*)| \;&\leq \sum_{m_0 : m_0 \neq m^*} |B^+(m_0, m^*)| \;\leq \sum_{m_0 : m_0 \neq m^*} |S^+(m_0, m^*)| \\
&\leq \sum_{m_0 : m_0 \neq m^*} 2^{\sigma} \cdot p(m_0, m^*) \;\leq\; 2^{\kappa/2} \cdot p^+(m^*) \;\leq\; 2^{\kappa/2}.
\end{aligned}$$

We are now ready to formally conclude our argument. For ease of notation, let $R(m^*)$ be the set of reachable ciphertexts from m^*, i.e. those c such that $c = \mathsf{E.Enc}(\mathsf{apk}, m^*; G(s))$. Moreover we set Bad_1 the event that $(\mathsf{apk}, \mathsf{ask}, m^*) \notin T$, where T was defined above, and Bad_2 the event that c is an incorrect encryption of m^* and Bad their logical disjunction. Then

$$\begin{aligned}
&\Pr\left[\exists (m, s) : c = \mathsf{E.Enc}(\mathsf{apk}, m; G(s))\right] \\
\leq\; &\Pr\left[\exists (m, s) : c = \mathsf{E.Enc}(\mathsf{apk}, m; G(s)) \wedge \neg\mathsf{Bad}\right] + \Pr\left[\mathsf{Bad}\right] \\
\leq\; &\Pr\left[\big(c \in B^+(m^*) \,\vee\, c \in R(m^*)\big) \wedge \neg\mathsf{Bad}\right] + \Pr\left[\mathsf{Bad}\right] \\
\leq\; &\Pr\left[c \in B^+(m^*) \wedge \neg\mathsf{Bad}\right] + \Pr\left[c \in R(m^*)\right] + \Pr\left[\mathsf{Bad}\right] \\
\leq\; &\frac{|B^+(m^*)|}{2^{\kappa}} + \frac{|R(m^*)|}{2^{\kappa}} + \varepsilon + \varepsilon' \;\leq\; \frac{2}{2^{\kappa/2}} + \varepsilon + \varepsilon'.
\end{aligned}$$

□

$\mathsf{H}_5 \approx \mathsf{H}_6$. We reduce any distinguisher $\mathcal{D}$ to $\mathcal{B}$ against the punctured pseudorandomness of the given PRF. Initially $\mathcal{B}$ generates $(\mathsf{apk}, \mathsf{ask})$ through $\mathsf{E.Gen}$, samples a random message m^*, a challenge bit b and computes $c_1, \ldots, c_{\vartheta+1}$ as encryptions of m^*. Then it queries a key k^* punctured over $c_1, \ldots, c_{\vartheta+1}$, and waits for the values $y_1, \ldots, y_{\vartheta+1}$. Next it sets c^* as the first ciphertexts c_i such that $y_i = b$, or to $c_{\vartheta+1}$ if no such ciphertext exists. It finally obfuscates $\widetilde{C} \leftarrow^{\$} \mathsf{iO}(C_{\mathsf{apk}, k^*})$ and runs $\mathcal{D}(\mathsf{apk}, \widetilde{C}, c^*)$, eventually returning the same bit as $\mathcal{D}$.

It is immediate to see that is $y_i = f_k(c_i)$ then $\mathcal{B}$ simulates H_5 perfectly. Conversely, call Coll the event in which there exists a collision among $c_1, \ldots, c_{\vartheta+1}$.

We have that conditioning on $\neg\mathsf{Coll}$, if the values y_i are uniformly random then the condition $b = y_i$ is independent from c_i. Thus the rejection sampling eventually return a ciphertext following the right distribution and in particular $\mathcal{B}$ perfectly simulates H_6. Because $\Pr[\mathsf{Coll}] \leq \vartheta^2 \cdot 2^{-\kappa}$, which follows as we assumed each ciphertexts to have min-entropy greater than λ, we can conclude that

$$\mathsf{Adv}(\mathcal{B}) \;\geq\; \mathsf{Adv}(\mathcal{D}) - 2\Pr[\mathsf{Coll}] \quad \Rightarrow \quad \mathsf{Adv}(\mathcal{D}) \;\leq\; \mathsf{Adv}(\mathcal{B}) - \frac{\vartheta^2}{2^\kappa}.$$

□

C.10 Second Construction from Obfuscation

Proof of Theorem 6. The proof is divided in two parts. First we show the basic anamorphism and next we prove Fully-Asymmetric security.

Basic Anamorphic Security. We proceed with a sequence of hybrids $\mathsf{H}_0, \ldots, \mathsf{H}_3$.

- H_0: The anamorphic game $\mathsf{AnamorphicG}$. Public parameters $(\mathsf{apk}, \mathsf{ask}, \mathsf{dk}, \mathsf{tk})$ are generated through $\mathsf{AT.Gen}(\lambda)$. Encryption queries $(m, \widehat{m})$ are answered with a ciphertext $c \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m, \widehat{m})$.
- H_1: As H_0 but when executing $\mathsf{AT.Enc}$, replace the check in Line 3 with $f_k(c) = \widehat{m}$ where $c \leftarrow \mathsf{E.Enc}(\mathsf{apk}, m; r)$.
- H_2: As H_1 but $f_k(\cdot)$ is replaced with a truly random function f^*.
- H_3: As H_2 but encryption queries $(m, \widehat{m})$ are answered with $c \leftarrow^{\$} \mathsf{E.Enc}(\mathsf{apk}, m)$.

Trivially, H_3 corresponds to the real game RealG as $\mathsf{apk}, \mathsf{ask}$ are sampled with $\mathsf{E.Gen}(\lambda)$. $\mathsf{H}_0 \approx \mathsf{H}_1$ follows from correctness of obfuscation. $\mathsf{H}_1 \approx \mathsf{H}_2$ as f_k is pseudorandom. Note in both experiments a distinguisher only observes $\mathsf{apk}, \mathsf{ask}$, both of which are generated independently from k, and evaluations of f_k, which are obtainable through oracle queries in the pseudorandomness game. Toward proving $\mathsf{H}_2 \approx \mathsf{H}_3$ let $c_1, \ldots, c_{q\vartheta}$ be the ciphertexts $\mathsf{AT.Enc}$ computes in H_2 to answer the q queries performed by a distinguisher. Then, as we assumed the PKE to satisfy Definition 1, the probability for a given pair of those ciphertexts to be equal is smaller than $2^{-\kappa}$. Thus, calling Coll the event $c_i = c_j$ for some $i \neq j$, a union bound yields $\Pr[\mathsf{Coll}] \leq q^2\vartheta^2 \cdot 2^{-\kappa}$. Conditioning on $\neg\mathsf{Coll}$, as all ciphertexts are different, the bits $f^*(c_1), \ldots, f^*(c_{q\vartheta})$ are uniformly and independently distributed. Thus $\mathsf{AT.Enc}$'s choice of the resulting ciphertext does not depend on those observed during its execution, meaning that its distribution is identical to the prescribed one.

Fully-Asymmetric Security. We recall the game syntax. The adversary $\mathcal{A}$, on input $(\mathsf{apk}, \mathsf{dk})$ generated via $\mathsf{AT.Gen}(\lambda)$, queries $(m_0, \widehat{m}_0)$, $(m_1, \widehat{m}_1)$. The challenger then replies with $c^* \leftarrow^{\$} \mathsf{AT.Enc}(\mathsf{apk}, \mathsf{dk}, m_b, \widehat{m}_b)$ for a randomly chosen challenge bit $b \in \{0, 1\}$. We prove the game to be hard through a sequence of hybrids. In the following we denote with m_0^*, m_1^* two distinct messages[18].

[18] We only require $m_0^* \neq m_1^*$, but they could potentially match the messages m_0, m_1 chosen by the adversary.

H_0: The $\mathsf{FAsyAnam}$-IND-CPA game with challenge bit b.

H_1: As H_0 but c^* is computed as $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m_b, \widehat{m}_b)$, see Figure 18.

H_2: As H_1 but c^* is computed as $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m_0^*, \widehat{m}_b)$.

H_3: As H_2 but c^* is computed as $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m_0^*, b)$.

H_4: As H_3 but c^* is computed during the setup after $(\mathsf{apk}, \mathsf{ask}, k)$ are generated.

H_5: As H_4 but, calling $\mathbf{c} = (c_1, \ldots, c_{\vartheta+1})$ the ciphertexts produced by $\mathsf{AT.Enc}^*$ to output c^*, then $\widetilde{C} \leftarrow \mathsf{iO}(C^*_{\mathsf{apk},k,\mathbf{c}})$ where C^* is described in Figure 18.

H_6: As H_5, but c^* is computed as $\mathsf{AT.Enc}^*(\mathsf{apk}, k, m_1^*, b)$.

H_7: As H_6, but during the setup compute $k^* \leftarrow \mathsf{PRF.Puncture}(k, c_1, \ldots, c_{\vartheta+1})$ and obfuscate $\widetilde{C} \xleftarrow{\$} \mathsf{iO}(C^*_{\mathsf{apk},k^*,\mathbf{c}})$.

H_8: As H_7, but c^* is computed as $\mathsf{E.Enc}(\mathsf{apk}, m_1^*)$.

$\mathsf{AT.Enc}^*(\mathsf{apk}, k, m, \widehat{m})$	$C^*_{\mathsf{pk},k,\mathbf{c}}(m, r)$:
for $i \in \{1, \ldots, \vartheta\}$:	Parse $\mathbf{c} = (c_1, \ldots, c_\vartheta, c_{\vartheta+1})$
Encrypt $c_i \xleftarrow{\$} \mathsf{E.Enc}(\mathsf{apk}, m)$	Encrypt $c \leftarrow \mathsf{E.Enc}(\mathsf{pk}, m; r)$
for $i \in \{1, \ldots, \vartheta\}$:	**if** $c = c_i$ for some i and $m = m_1^*$:
if $f_k(c_i) = \widehat{m}$: **return** c_i	**return** 0
return $c_{\vartheta+1}$	**else** : **return** $f_k(c)$

Fig. 18. Alternative encryption (left) and circuit (right) used in the proof of Theorem 6.

Guessing b in H_8 is information-theoretically hard. Moreover $\mathsf{H}_0 \approx \mathsf{H}_1$ due to the obfuscator's correctness and $\mathsf{H}_3 = \mathsf{H}_4$ as only the order of operations is changed[19]. To conclude we prove the remaining hybrids to be indistinguishable, with the exception of $(\mathsf{H}_2, \mathsf{H}_3)$, where we show that guessing b is equally hard.

$\mathsf{H}_1 \approx \mathsf{H}_2$. Any distinguisher $\mathcal{D}$ can be reduced to $\mathcal{B}$ breaking the IND-CPA security of the underlying scheme in $\vartheta+1$ encryption queries. It initially generates $k, \widetilde{C}$ honestly and runs $\mathcal{D}$. When $\mathcal{D}$ returns $(m_0, \widehat{m}_0, m_1, \widehat{m}_1)$, it uses its own encryption oracle to produce $\vartheta+1$ ciphertexts either encrypting m_b (for a random b chosen by $\mathcal{B}$) or m_0^*. A full description is given in Figure 19.

It is immediate to see $\mathcal{B}$ perfectly simulates H_1 and H_2 respectively when its challenger encrypts the first or the second message in each queried couple. Thus $\mathsf{Adv}(\mathcal{D}) = \mathsf{Adv}(\mathcal{B})$, which is negligible.

H_2 *is harder than* H_3. The proof is identical to the one presented in the proof of Theorem 5.

[19] Note this is possible as c^* does not depend on $\mathsf{dk} = \widetilde{C}$, nor on the challenge messages.

$\mathcal{B}(\mathsf{pk})$:

Sample a PRF key k, $\widetilde{C} \xleftarrow{\$} \mathsf{iO}(C_{\mathsf{pk},k})$ and run $\mathcal{D}(\mathsf{pk}, \widetilde{C}) \to (m_0, \widehat{m}_0, m_1, \widehat{m}_1)$
Sample a random bit $b \xleftarrow{\$} \{0,1\}$
for ϑ times:
  Query (m_b, m_0^*) to the challenger and wait for c
  if $f_k(c) = \widehat{m}_b$: Set $c^* \leftarrow c$ and **break**
if c^* was not defined in the previous loop:
  Query (m_b, m_0^*) to the challenger and set c^* to the response.
Reply c^* to $\mathcal{D}$
when $\mathcal{D}$ returns b': **return** b'

Fig. 19. Reduction $\mathcal{B}$ of a distinguisher $\mathcal{D}$ for H_1, H_2 to IND-CPA.

$\mathsf{H}_4 \approx \mathsf{H}_5$. We reduce to the obfuscator security. Indeed for any (m, r) the circuits $C_{\mathsf{apk},k}$ and $C^*_{\mathsf{apk},k,\mathbf{c}}$ evaluate to $f_k(c)$ with $c = \mathsf{E.Enc}(\mathsf{apk}, m; r)$ unless $c \in \{c_1, \ldots, c_{\vartheta+1}\}$ and $m = m_1^*$. However, each c_i is the encryption of $m_0^* \neq m_1^*$. Thus, from perfect correctness, the above condition is impossible and the two circuits are functionally equivalent.

$\mathsf{H}_5 \approx \mathsf{H}_6$. We again reduce any distinguisher $\mathcal{D}$ to $\mathcal{B}$ breaking IND-CPA for the underlying PKE. The strategy is analogous to that for $\mathsf{H}_1 \approx \mathsf{H}_2$: in this case $\mathcal{B}$ initially generates the PRF key k and queries $\vartheta + 1$ ciphertexts c_i that are either the encryption of m_0^* or m_1^*. It then chooses the first c_i such that $f_k(c_i) = b$ for a randomly chosen bit b, and obfuscate $\widetilde{C} = \mathsf{iO}(C^*_{\mathsf{apk},k,\mathbf{c}})$ with $\mathbf{c} = (c_1, \ldots, c_{\vartheta+1})$. As $\mathcal{B}$ perfectly simulates respectively $\mathsf{H}_5, \mathsf{H}_6$ according to its challenge bit, we conclude $\mathsf{Adv}(\mathcal{D}) = \mathsf{Adv}(\mathcal{B})$.

$\mathsf{H}_6 \approx \mathsf{H}_7$. Again we reduce to the obfuscator security. Indeed, from Definition 13 (specifically, the first point) the two circuits are identical on (m, r) such that $\mathsf{E.Enc}(\mathsf{apk}, m; r) \notin \{c_1, \ldots, c_{\vartheta+1}\}$. Conversely, when $\mathsf{E.Enc}(\mathsf{apk}, m; r)$ lies in the above set, from perfect correctness of the given PKE, this means $m = m_1^*$ as each c_i is an encryption of m_1^* and in particular both circuits return 0.

$\mathsf{H}_7 \approx \mathsf{H}_8$. We reduce a distinguisher $\mathcal{D}$ to an adversary $\mathcal{B}$ for the pseudorandomness of the punctured PRF. Initially it generates a random challenge bit $b \in \{0,1\}$, keys $\mathsf{apk}, \mathsf{ask}$, and $\vartheta + 1$ ciphertexts $c_1, \ldots, c_{\vartheta+1}$ as $\mathsf{E.Enc}(\mathsf{apk}, m_1^*)$ (each with fresh random coins). Then it queries a key punctured in those ciphertexts. Upon receiving k^* and the values $y_1, \ldots, y_{\vartheta+1}$ from the challenger, it computes c^* as the first c_i such that $y_i = b$ or $c_{\vartheta+1}$ if the $y_1 = \ldots = y_\vartheta \neq b$. Finally, it obfuscates $\widetilde{C} = \mathsf{iO}(C^*_{\mathsf{apk},k^*,\mathbf{c}})$, runs $\mathcal{D}(\mathsf{apk}, \widetilde{C}, c^*)$ and eventually returns $\mathcal{D}$'s output. It is immediate to see that if $y_i = f_k(c_i)$ then $\mathcal{B}$ perfectly simulates H_7. Conversely, in the ideal experiment $y_1, \ldots, y_{\vartheta+1}$ are uniformly and independent bits assuming no collisions among the ciphertexts. In this case performing rejection sampling on the condition $b = y_i$ does not alter the distribution of c^* as both b and y_i are independent from c_i. Thus c^* is distributed as a correct encryption of m_1^* and in particular, $\mathcal{B}$ perfectly simulates H_8. Finally, calling Coll the event in which any two ciphertexts collide, as we assume Definition 1 to apply to the given PKE, $\Pr[\mathsf{Coll}] \leq \vartheta^2 2^{-\kappa}$. Calling β the challenge bit for $\mathcal{B}$ (i.e. when $\beta = 1$ then $y_i = f_k(c_i^*)$), we conclude that $\mathsf{Adv}(\mathcal{B}) =$

$$\begin{aligned}
&= |\Pr[\mathcal{B} \to 1 \,|\, \beta = 1] - \Pr[\mathcal{B} \to 1 \,|\, \beta = 0]| \\
&\geq \Pr[\neg\mathsf{Coll}] \, |\Pr[\mathcal{B} \to 1 \,|\, \beta = 1, \neg\mathsf{Coll}] - \Pr[\mathcal{B} \to 1 \,|\, \beta = 0, \neg\mathsf{Coll}]| - \Pr[\mathsf{Coll}] \\
&= \Pr[\neg\mathsf{Coll}] \, |\Pr[\mathcal{D} \to 1 \,|\, \mathsf{H}_7, \neg\mathsf{Coll}] - \Pr[\mathcal{D} \to 1 \,|\, \mathsf{H}_8, \neg\mathsf{Coll}]| - \Pr[\mathsf{Coll}] \\
&= |\Pr[\mathcal{D} \to 1, \neg\mathsf{Coll} \,|\, \mathsf{H}_7] - \Pr[\mathcal{D} \to 1, \neg\mathsf{Coll} \,|\, \mathsf{H}_8]| - \Pr[\mathsf{Coll}] \\
&\geq |\Pr[\mathcal{D} \to 1 \,|\, \mathsf{H}_7] - \Pr[\mathcal{D} \to 1 \,|\, \mathsf{H}_8]| - 3\Pr[\mathsf{Coll}] \\
&\geq \mathsf{Adv}(\mathcal{D}) - 3\vartheta^2/2^\kappa
\end{aligned}$$

where the second to last step follows adding and subtracting $\Pr[\mathcal{D} \to 1 \,|\, \mathsf{H}_7]$, using inverse triangular inequality[20] and observing that the remaining terms are smaller than $\Pr[\mathsf{Coll}]$ (which is the same in H_7 and H_8). $\square$

[20] $|x + y| \geq |x| - |y|$ for all reals.