SCIENCE CHINA
Information Sciences

• RESEARCH PAPER •

May 2026, Vol. 69, Iss. 5, 152107:1–152107:16
https://doi.org/10.1007/s11432-025-4751-y

Anamorphic commitment: robust privacy under coercion and parameter tampering

Weiqi WANG[1], Yubo ZHENG[1], Peng XU[1*], Wei WANG[2], Rongmao CHEN[3], Yifan YANG[1] & Moti YUNG[4,5]

[1]*Hubei Key Laboratory of Distributed System Security, School of Cyber Science and Engineering, Huazhong University of Science and Technology, Wuhan 430074, China*
[2]*Cyber-Physical-Social Systems Laboratory, School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan 430074, China*
[3]*College of Computer Science and Technology, National University of Defense Technology, Changsha 410073, China*
[4]*Google LLC, New York NY 10011, USA*
[5]*Department of Computer Science, Columbia University, New York NY 10027, USA*

Received 25 July 2025/Revised 15 October 2025/Accepted 8 January 2026/Published online 9 April 2026

Abstract A commitment protocol enables a sender to fix a message without revealing or altering it until later, which is crucial for various cryptographic protocols and applications. However, traditional commitments risk coercion, where an entity, named dictator, can force premature opening or involuntary commitments. This coercive entity was initially explored in encryption by Persiano et al. at EUROCRYPT'22. This work investigates how to circumvent the dictator in commitment, considering the dictator's enhanced ability to influence the sender by tampering with public parameters. We note that from inception to today, commitments were defined and employed under a passive adversary, and this work, therefore, conceptually demonstrates how the "anamorphic perspective" can enrich the relevant adversarial settings of basic primitives. Specifically, we formalize dictator's capabilities for undermining commitments—ranging from weak to strong models—and establish reductions among them. To counter coercion, we propose "anamorphic commitment", a novel primitive that allows the sender to embed a covert message within an innocent-looking commitment to alert to coercion or transmit secret information. We define the syntax, security models, and robustness of anamorphic commitment, and present two generic constructions with distinct performance advantages that can integrate seamlessly with the standard commitments.

Keywords commitment protocol, anamorphic cryptography, coercion resistance, parameter tampering, subliminal channel

Citation Wang W Q, Zheng Y B, Xu P, et al. Anamorphic commitment: robust privacy under coercion and parameter tampering. Sci China Inf Sci, 2026, 69(5): 152107, https://doi.org/10.1007/s11432-025-4751-y

1 Introduction

A commitment protocol is a cryptographic primitive that allows a sender to commit to a selected message without immediate disclosure, ensuring the message remains unaltered and concealed until the sender reveals it. Commitment plays a foundational role in various cryptographic protocols and applications; even where encryption is prohibited, commitments can still be employed to implement essential processes such as auctions [1] and ZK proofs [2], essentially ensuring integrity, honesty, and preventing premature disclosure for fairness and independence reasons. To date, commitment protocols have focused on resisting against passive attackers (i.e., the committer is assumed to be tamper-free and its state secure to the environment). However, coercive entities, named dictators as discussed by Persiano et al. [3], introduce the risk of manipulating the messages associated with encryption, potentially undermining their fairness and reliability. In fact, existing commitment protocols also fail to secure against such dictator's attacks. Note that prior work typically assumes the sender is tamper-proof. This work relaxes this assumption and argues that the anamorphic perspective, together with a strong and coercive adversary model, naturally motivates the study of commitments under active attacks on the sender.

Specifically, this work models a dictator in a commitment that can (i) force the sender to reveal the committed message in advance (breaking the receiver-privacy assumption) and (ii) coerce the sender into committing to a message chosen by the dictator (breaking the sender-freedom assumption). Since the dictator does not directly join the commitment process, the receiver cannot detect the dictator's covert malicious actions, which severely

* Corresponding author (email: xupeng@mail.hust.edu.cn)

undermines the security and fairness of commitment protocols. Furthermore, a commitment protocol usually has some public parameters (including binding factors) generated by the receiver. The dictator can also corrupt the receiver to choose particular public parameters. This corruption allows the dictator to test the sender's honesty. Specifically, this work models. Therefore, it raises a natural question: *How to resist the dictator in the field of commitment protocol?*

Given the covert interference of very powerful dictators, this work focuses on utilizing the idea of anamorphic cryptography to augment standard commitment protocols to covertly resist covertly acting dictator adversaries (i.e. "fighting fire with fire"). Specifically, the contributions are as follows.

Anamorphic commitment definitions. We define the concept of an anamorphic commitment protocol. This new cryptographic primitive allows a sender to embed a covert message in a commitment. The covert message can inform the expected receiver that a dictator is coercing the sender. Before the commitment phase, the sender and the receiver share a double key. At the opening phase of a commitment, following the model for anamorphic encryption, only the receiver with the double key can extract the covert message from the commitment. Otherwise, anyone without the double key, including the dictators, cannot obtain any information about the covert message or distinguish whether the commitment is regular or anamorphic. Such indistinguishability is necessary. Otherwise, the dictators can find that a covert communication is occurring and force the sender to hand over the double key.

We detail the processes of an anamorphic commitment. Unlike standard commitments such as [4–11], an anamorphic commitment protocol uses a three-algorithm structure named anamorphic triplet, including double-key generation, anamorphic commitment, and anamorphic decryption (de-commitment). Here, the sender and the receiver run double-key generation to share a double key `dk`. Suppose the sender wants to convey a covert message to the receiver through a commitment. In that case, the sender runs the anamorphic commitment algorithm with the double key `dk`, an expected-to-be-committed message `msg`, and the covert message `amsg`, and produces a blind factor `x` and an anamorphic commitment `ac`. Anyone with `x` can verify the validity of the commitment `ac`. Upon executing anamorphic decryption with `dk` and `x`, the receiver discloses the covert message `amsg`.

Various dictator models. Corresponding to the anamorphic properties required under varying dictators, this work categorizes the dictator's capabilities and defines three models of increasing strength. The weakest, the Transcriber, breaks the receiver-privacy assumption by compelling the sender to reveal the commitment at creation time. The Dominator violates the sender-freedom assumption by coercing the sender into committing to a dictator-chosen value. The strongest, the *Corrupter*, combines the previous two powers and additionally tampers with public parameters of the commitment. Section 3 proves reductions among these models and shows the necessity of a double key under the *Corrupter* model.

This work also discusses the semantic security of a covert message, the semantic security of a committed message, and the robustness of an anamorphic commitment. The semantic security of a covert message means that an anamorphic commitment does not reveal any information about the covert message to anyone without the double key, similar to the indistinguishability under chosen plaintext attack (IND-CPA) security of the ciphertext. The semantic security of a committed message means that the committed message of an anamorphic commitment is safe under an ordinary adversary without the double key. Robustness ensures that when receiving a non-anamorphic commitment, the anamorphic decryption algorithm will deterministically output $\perp$, indicating that the commitment (intentionally) contains no covert message. Section 3 proves that anti-dictator security naturally implies the semantic security of a covert message, while anti-dictator security does not imply the semantic security of a committed message.

Generic anamorphic constructions. This work aims to construct generic, robust anamorphic commitment transforms that can augment any standard commitment with covert-channel capability. In standard commitments, the sender's only controllable element under a dictator who manipulates messages or public parameters is the random blind factor `x`. Because the sender must reveal `x` at opening and the commitment must remain binding and hiding, the sender cannot simply substitute an incorrect `x` or equivocate without detection; this makes embedding a covert channel under strong dictator models technically challenging. Our solution is to embed the covert message directly into `x` while decoupling the generation process of `x` from receiver-chosen public parameters (thus foiling a *Corrupter* who tampers with those parameters). We additionally introduce an explicit "robustness token" as part of the embedded payload; the receiver uses the token to distinguish anamorphic from regular commitments. Based on these ideas, we present two constructions that realize anti-*Corrupter* anamorphism while preserving the semantic security of the committed message and robustness.

The first construction replaces the random blind factor with a specially constructed one. The new blind factor has two parts: one embeds a covert message `amsg`, and another serves as a robustness token. We introduce a synchronization mechanism and use the pseudo-random function to generate a synchronized pseudo-random mask. Then this mask will split into two parts: one serves as a one-time key to protect `amsg`, and the other is the robustness

token. The second construction regards the robustness token as a part of the double key. The sender encrypts `amsg` and the robustness token by symmetric encryption of the pseudo-random-ciphertext feature and uses the generated ciphertext as the blind factor to generate an anamorphic commitment as a regular commitment does. We discuss in detail the varying performances and advantages of these two constructions and the application scenarios for which they are suitable.

2 Related work

The notion of anamorphic commitment shares similarities with several other cryptographic primitives. However, anamorphic commitment diverges significantly from them in several vital aspects.

Anamorphic cryptography. The concept and preliminary idea of anamorphic cryptography is *Covert Cryptographic Primitives and Protocols* proposed by Phan and Yung in [12]. Ideally, protocol Q is hidden within another protocol P, appearing to the adversary as participants simply carrying out protocol P regularly. PPY22 [3] then formalized the novel notion of anamorphic encryption and introduced sender- and receiver-anamorphic encryption schemes, sparking subsequent extensions.

For receiver-anamorphism, KPPY23$_1$ [13] distinguishes single/multiple-receiver models via double key functionality, while BGHM24 [14] discusses the robustness of anamorphic encryption. These two studies refine the security model of anamorphic encryption from the perspectives of semantic security for covert messages and preventing misinterpretation of covert messages, respectively. CGM24 [15] shows how anamorphic encryptions support homomorphism. PPY24 [16] presents the new notion of "public-key anamorphic encryption". Public-key anamorphic encryption enables parties to conduct covert transmissions without any prior secret exchange, significantly enhancing usability. However, the proposed scheme suffers from excessively long ciphertexts and lacks universality. Refs. [17,18] discussed the upper bound on channel capacity and limitations for black-box anamorphic encryption. For sender-anamorphism, Ref. [19] developed a robust sender-anamorphic encryption construction, but this construction requires multiple encryptions to complete one covert message transmission. To mitigate the potential misuse of anamorphic techniques by malicious entities, Refs. [20, 21] introduced anamorphic-resistant encryption, an encryption that cannot be transformed into high-efficiency anamorphic encryption.

In addition, WHL24 [22] proposes anamorphic authentication key exchange to share double key covertly. KPPY23$_2$ [23] extends anamorphic cryptography from encryption to signatures, exploring how to use signatures to transmit covert messages when dictators forbid encryption entirely, and further JS24 [24] refines the security model of anamorphic signature and achieves unforgeability. Note that this motivation for anamorphic signature is similar to our case, where there is a need to embed covert messages in other non-banned primitives.

Compared to the studies mentioned above, this work extends the concept of anamorphic cryptography to commitment protocols. This work considers not only dictators who can break the sender-freedom and receiver-privacy assumptions, but also considers another type of dictator that can corrupt the receiver to tamper with the public parameters of a commitment protocol.

Subliminal channel. Subliminal channels leverage redundancy or specific attributes of cryptographic communication to enable covert communication through public channels that remain undetectable to unintended recipients. Simmons' pioneering work [25] used prisoners' problem to introduce the concept of subliminal channel for the first time and later proposed a series of methods to embed covert messages within publicly transmitted messages. Implementations span digital signatures [26, 27], visual cryptography [28], and blockchain systems [29]. To mitigate the risk that an adversary may exploit the subliminal channel to launch subversion attacks, Mironov et al. [30] introduced the notion of cryptographic reverse firewall (CRF). Building on this paradigm, Chen et al. [31] further constructed a subversion-resilient instantiation of the SM9 IBE scheme.

While subliminal channels focus on high-bandwidth covert communication under standard adversarial assumptions, this work constructs generic anamorphic commitment schemes targeting dictatorial environments with stronger adversarial models, beyond standard assumptions, and more harmful.

Trapdoor and mercurial commitment. Brassard et al. [32] introduced trapdoor commitments, in which a sender holding a secret trapdoor can alter the committed message at opening—a capability subsequently harnessed by Catalano and Visconti [33] and Hanaoka and Schuldt [34] to build zero-knowledge proofs, non-malleable commitments and signatures, and by Baghery [35] to thwart parameter subversion. In parallel, Chase et al. [36] proposed mercurial commitments with two modes: a binding "hard" mode indistinguishable from standard commitments, and a "soft" mode allowing equivocation while preserving computational indistinguishability. Later studies [37–40] extended this primitive to new variants and vector commitments. Together, these two lines of work illustrate how adding "trapdoors" or "soft-opening" modes can enrich commitment schemes for advanced cryptographic protocols.

Table 1 Comparison across anamorphic cryptography schemes and related primitives.

	Anamorphic cryptography								Subliminal channel	Trapdoor/mercurial commitment
	PPY22 [3]	KPPY23$_1$ [13]	BGHM24 [14]	CGM24 [15]	PPY24 [16]	WHL24 [22]	KPPY23$_2$ [23]	JS24 [24]		
Primitive	Anamorphic encryption					Anamorphic AKE	Anamorphic signature		Signature	Commitment
Covert transmission	✓	✓	✓	✓	✓	Only random key	✓	✓	✓	✗
Generality	✗	✓	✓	✗	✗	✓	✓	✓	✗	✗
Anti-dictator	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
Additional feature	–	–	Robustness	Homomorphism	No double key	No double key	–	Unforgeability	–	Ambiguity

However, it is noteworthy that both trapdoor and mercurial commitments remain defenseless against dictator adversaries. For trapdoor commitments, if a dictator inspects or even chooses the message before opening (or corrupts the honest third party that generated the trapdoor), the sender loses any ability to modify it. Likewise, mercurial commitments inherit this weakness in their hard mode, and although soft mode lets the sender equivocate to mislead a dictator, such soft-opened messages carry no validity for the honest receiver—leaving the dictator's goals intact.

To provide a more straightforward overview of the conceptual relationships among anamorphic cryptography, subliminal channels, and trapdoor/mercurial commitments, this work summarizes their core properties and functional differences in Table 1.

3 Preliminaries

This section first introduces the notions that will be used in the subsequent sections. Given a security parameter 1^λ, $\mathsf{poly}(\lambda)$ denotes a polynomial function over λ and $\mathsf{negl}(\lambda)$ denotes a negligible function over λ. Let $s \leftarrow_\$ \mathcal{S}$ denote sampling an element s from an unempty set $\mathcal{S}$ randomly. $y \leftarrow_\$ \mathsf{alg}_1(x)$ denotes that alg_1 is a probabilistic algorithm, which takes x as input and output y. If $\mathsf{alg}_2(x)$ is a deterministic algorithm, $y := \mathsf{alg}_2(x)$ denote its running. Lowercase letters (such m and n) denote bit lengths.

3.1 Standard commitment

A cryptographic standard commitment can be understood as a box with a lock. In the commitment phase, the sender places a piece of paper with a message into the box, locks it with a key, and hands it over to the receiver. The commitment must simultaneously ensure binding, preventing the sender from altering the message, and hiding, preventing the receiver from learning the message before opening.

This work generalizes standard commitments by introducing a parameter generation algorithm producing public parameters $\mathtt{pp}$ (possibly empty, as in [37]), which may include binding factors in interactive protocols. Non-interactive and interactive variants share core structures: the commitment algorithm implicitly uses a random blind factor $\mathtt{x}$, which we explicitly formalize as an input to enable embedding covert messages in anamorphic extensions; for opening, since standard commitments allow message recovery from randomness [41], we define the open algorithm to take a commitment as input and output the corresponding message and $\mathtt{x}$, leveraging the sender's knowledge of both. This framework systematically unifies protocol variants. By explicitly parameterizing $\mathtt{x}$, we establish a foundation for extending commitments with anamorphic capabilities without altering their core functionality.

Definition 1. A standard commitment $\mathcal{SC}$ consists of four PPT algorithms $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$, such that the following.

(1) The generation algorithm Gen takes as input a security parameter 1^λ and returns public parameters $\mathtt{pp} \leftarrow_\$ \mathsf{Gen}(1^\lambda)$.

(2) The commitment algorithm Com takes as input public parameters $\mathtt{pp}$, a message $\mathtt{msg}$, and a randomly-chosen blind factor $\mathtt{x}$, and returns a commitment $\mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$.

(3) The open algorithm Open takes as input a commitment $\mathtt{c}$ and finds and returns the corresponding message $\mathtt{msg}$ and blind factor $\mathtt{x}$ locally, namely $(\mathtt{msg}, \mathtt{x}) := \mathsf{Open}(\mathtt{c})$.

(4) The verify algorithm Ver takes as input a commitment $\mathtt{c}$, a message $\mathtt{msg}$ and its blind factor $\mathtt{x}$, and public commitment parameters $\mathtt{pp}$, and outputs $1 := \mathsf{Ver}(\mathtt{c}, \mathtt{msg}, \mathtt{x}, \mathtt{pp})$ if $\mathtt{c}$ is a valid commitment, otherwise outputs 0.

A standard commitment must satisfy the following two security properties.

(1) Hiding property. A commitment does not reveal any information about the committed message. Formally, for any PPT adversary $\mathcal{A}$, we have

$$\left| \Pr\left[b = b' \,\middle|\, \begin{array}{l} \mathtt{pp} \leftarrow_\$ \mathcal{SC}.\mathsf{Gen}(1^\lambda), (\mathtt{msg}_0, \mathtt{msg}_1) \leftarrow_\$ \mathcal{A}(1^\lambda) \\ b \leftarrow_\$ \{0,1\}, \mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}_b, \mathtt{x}), b' \leftarrow_\$ \mathcal{A}(\mathtt{c}) \end{array} \right] - \frac{1}{2} \right| \leqslant \mathsf{negl}(\lambda).$$

(2) Binding property. The sender who generates a commitment cannot open it into two different messages. Formally for any PPT adversary $\mathcal{A}$, we have

$$\Pr\left[\begin{array}{l|l} \mathcal{SC}.\mathsf{Ver}(\mathtt{c},\mathtt{msg}_0,\mathtt{x}_0,\mathtt{pp})=1\wedge & \mathtt{pp}\leftarrow_{\$}\mathsf{Gen}(1^\lambda) \\ \mathcal{SC}.\mathsf{Ver}(\mathtt{c},\mathtt{msg}_1,\mathtt{x}_1,\mathtt{pp})=1\wedge & (\mathtt{msg}_0,\mathtt{msg}_1,\mathtt{x}_0,\mathtt{x}_1)\leftarrow_{\$}\mathcal{A}(1^\lambda) \\ \mathtt{msg}_0\neq\mathtt{msg}_1 & \mathtt{c}:=\mathcal{SC}.\mathsf{Com}(\mathtt{pp},\mathtt{msg}_0,\mathtt{x}_0) \end{array}\right]\leqslant\mathsf{negl}(\lambda).$$

3.2 Pseudo-random function

This work also employs a pseudo-random function (PRF) to protect the covert message and to achieve the robustness of the anamorphic commitment. Its definition is as follows.

Definition 2. An efficient function $\mathcal{F}:\mathcal{K}\times\mathcal{X}\to\mathcal{Y}$ with key space $\mathcal{K}$ is a PRF if for $k\leftarrow_{\$}\mathcal{K}$, any PPT adversary $\mathcal{A}$ has $|\Pr[\mathcal{A}^{\mathcal{F}(k,\cdot)}(1^\lambda)=1]-\Pr[\mathcal{A}^{f(\cdot)}(1^\lambda)=1]|\leqslant\mathsf{negl}(\lambda)$, where $f:\mathcal{X}\to\mathcal{Y}$ is a real random function, $\mathcal{X}$ and $\mathcal{Y}$ denote the input space and output space of f, respectively.

4 Anamorphic commitment and its security definitions

An anamorphic commitment consists of a standard commitment $\mathcal{SC}$ and an associated anamorphic triplet $\mathcal{AMC}=(\mathsf{aKG},\mathsf{aCom},\mathsf{aDec})$. Hence, two parties (a sender and a receiver) can employ $\mathcal{SC}$ of the anamorphic commitment to achieve a standard commitment. When suffering coercion from a dictator, the two parties can use $\mathcal{AMC}$ to transmit a covert message via an "innocent" commitment. In this case, the two parties initially run algorithm $\mathcal{AMC}.\mathsf{aKG}$ to share a double key $\mathtt{dk}$. When the dictator forces the sender to commit an unwilling message, the sender runs the algorithm $\mathcal{AMC}.\mathsf{aCom}$ to generate a commitment $\mathtt{ac}$ which embeds a covert message $\mathtt{amsg}$. When the sender decides to open $\mathtt{ac}$, he can reveal $\mathtt{msg}$ and the corresponding blind factor $\mathtt{x}$ to the receiver. Then, the receiver can extract the covert message $\mathtt{amsg}$ from $\mathtt{x}$. To avoid arousing the dictator's suspicion, $\mathtt{ac}$ and its corresponding $\mathtt{x}$ must be indistinguishable from the regular commitment and blind factor both of $\mathtt{msg}$, respectively.

4.1 Anamorphic triplet for any standard commitment

Definition 3. An anamorphic triplet for any standard commitment consists of three PPT algorithms $\mathcal{AMC}=(\mathsf{aKG},\mathsf{aCom},\mathsf{aDec})$ such that the following.

(1) The double-key generation algorithm aKG takes a security parameter 1^λ and returns a double key $\mathtt{dk}\leftarrow_{\$}\mathsf{aKG}(1^\lambda)$.

(2) The anamorphic commit algorithm aCom takes as input the double key $\mathtt{dk}$, a message $\mathtt{msg}$, a covert message $\mathtt{amsg}$ and public parameters $\mathtt{pp}$, and outputs an anamorphic commitment $\mathtt{ac}$ and a blind factor $\mathtt{x}$ used to open the anamorphic commitment, that is $(\mathtt{ac},\mathtt{x})\leftarrow_{\$}\mathsf{aCom}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})$.

(3) The anamorphic decrypt algorithm aDec takes as input the double key $\mathtt{dk}$, the anamorphic commitment $\mathtt{ac}$, and the blind factor $\mathtt{x}$, and outputs the covert message $\mathtt{amsg}:=\mathsf{aDec}(\mathtt{dk},\mathtt{ac},\mathtt{x})$.

In addition, anamorphic commitments need to satisfy correctness. Specifically, for any message $\mathtt{msg}$, covert message $\mathtt{amsg}$, $\mathtt{dk}\leftarrow_{\$}\mathcal{AMC}.\mathsf{aKG}(1^\lambda)$, and $(\mathtt{ac},\mathtt{x})\leftarrow_{\$}\mathcal{AMC}.\mathsf{aCom}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})$, we have $\Pr[\mathcal{AMC}.\mathsf{aDec}(\mathtt{dk},\mathtt{ac},\mathtt{x})\neq\mathtt{amsg}]\leqslant\mathsf{negl}(\lambda)$.

4.2 Dictator models and their reductions

This section discusses the dictators' capabilities to coerce a standard commitment from weak to strong. We categorize the dictators into three types: the first one is to break the receiver-privacy assumption; the second one is to break the sender-freedom assumption; the final one is the strongest one, which encompasses the previous two types and can corrupt the receiver to tamper with public parameters. Finally, we prove the reductions among these dictators.

***Transcriber*: the first type.** *Transcriber* can coerce the sender to reveal the committed message $\mathtt{msg}$ and the corresponding blind factor $\mathtt{x}$ after the sending has generated the commitment $\mathtt{c}$. Then, *Transcriber* can verify the validity of $\mathtt{c}$ by running $\mathcal{SC}.\mathsf{Ver}$ before the sender opens it.

Transcriber models continuous surveillance, embodying basic coercion mechanisms faced in regulated industries or authoritarian states where monitored channels exist. This entity (e.g., governmental authorities) enforces compliance through transparent commitment registration while respecting formal non-intervention boundaries. Although abstaining from direct message manipulation, its capacity to mandate disclosure of commitment details inherently

$\underline{\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}(\lambda)}$
1. $\mathtt{dk}\leftarrow_{\$}\mathcal{AMC}.\mathsf{aKG}(1^{\lambda})$
2. $\mathtt{pp}\leftarrow_{\$}\mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
3. $\mathtt{msg}\leftarrow_{\$}\mathcal{M}$
4. $b\leftarrow_{\$}\mathcal{T}^{\mathsf{Oa}(\mathtt{dk},\mathtt{msg},\cdot,\mathtt{pp})}(\mathtt{msg},\mathtt{pp})$
5. return b

$\underline{\mathsf{Oa}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})}$
1. $(\mathtt{ac},\mathtt{x})\leftarrow_{\$}\mathcal{AMC}.\mathsf{aCom}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})$
2. return $(\mathtt{ac},\mathtt{x})$

$\underline{\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}(\lambda)}$
1. $\mathtt{pp}\leftarrow_{\$}\mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
2. $\mathtt{msg}\leftarrow_{\$}\mathcal{M}$
3. $b\leftarrow_{\$}\mathcal{T}^{\mathsf{Oc}(\mathtt{msg},\cdot,\mathtt{pp})}(\mathtt{msg},\mathtt{pp})$
4. return b

$\underline{\mathsf{Oc}(\mathtt{msg},\mathtt{amsg},\mathtt{pp})}$
1. $\mathtt{c}:=\mathcal{SC}.\mathsf{Com}(\mathtt{pp},\mathtt{msg},\mathtt{x})$
2. return $(\mathtt{c},\mathtt{x})$

Figure 1 Games $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}(\lambda)$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}(\lambda)$, where $\mathcal{M}$ denotes message space.

$\underline{\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}(\lambda)}$
1. $\mathtt{dk}\leftarrow_{\$}\mathcal{AMC}.\mathsf{aKG}(1^{\lambda})$
2. $\mathtt{pp}\leftarrow_{\$}\mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
3. $b\leftarrow_{\$}\mathcal{D}^{\mathsf{Oa}(\mathtt{dk},\cdot,\cdot,\mathtt{pp})}(\mathtt{pp})$
4. return b

$\underline{\mathsf{Oa}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})}$
1. $(\mathtt{ac},\mathtt{x})\leftarrow_{\$}\mathcal{AMC}.\mathsf{aCom}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})$
2. return $(\mathtt{ac},\mathtt{x})$

$\underline{\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}(\lambda)}$
1. $\mathtt{pp}\leftarrow_{\$}\mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
2. $b\leftarrow_{\$}\mathcal{D}^{\mathsf{Oc}(\cdot,\cdot,\mathtt{pp})}(\mathtt{pp})$
3. return b

$\underline{\mathsf{Oc}(\mathtt{msg},\mathtt{amsg},\mathtt{pp})}$
1. $\mathtt{c}:=\mathcal{SC}.\mathsf{Com}(\mathtt{pp},\mathtt{msg},\mathtt{x})$
2. return $(\mathtt{c},\mathtt{x})$

Figure 2 Games $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}(\lambda)$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}(\lambda)$.

undermines the hiding property. *Transcriber* illustrates that even without direct interference, the presence of an entity that can compel disclosure of commitment details fundamentally undermines the privacy of the commitment schemes.

Anamorphic commitment allows the sender to counter persistent monitoring or the threat posed by *Transcriber*. In this approach, a commitment contains an 'innocent' message, while a covert message is delivered through an anamorphic channel. This strategy helps mitigate the threat posed by *Transcriber*. We formalize the security model under *Transcriber* in Definition 4 through the following two games involving *Transcriber* $\mathcal{T}$ as illustrated in Figure 1. Note that in the following content, Oa and Oc are two oracles that dictators can query for different issues.

Definition 4. A standard commitment protocol $\mathcal{SC}=(\mathsf{Gen},\mathsf{Com},\mathsf{Open},\mathsf{Ver})$ can be transformed to an anti-*Transcriber* anamorphic commitment if there exists an anamorphic triplet $\mathcal{AMC}=(\mathsf{aKG},\mathsf{aCom},\mathsf{aDec})$ and for any PPT Transcriber $\mathcal{T}$, we have $|\Pr[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}(\lambda)=1]-\Pr[\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}(\lambda)=1]|\leqslant\mathsf{negl}(\lambda)$.

***Dominator*: the second type.** The second type is named as *Dominator*. He represents the conventional capabilities attributed to a dictator, consistent with prior formulations. *Dominator* can restrict the sender's freedom by eliminating the sender's autonomy in a commitment. For example, *Dominator* can coerce the sender into committing to the message chosen by *Dominator* rather than allowing the sender to select the message freely. Moreover, *Dominator* can also run the algorithm $\mathcal{SC}.\mathsf{Ver}$ to verify if the sender follows the *Dominator*'s orders.

Anamorphic commitment allows the sender to generate a commitment to an unwilling message, namely following the *Dominator*'s order and avoiding arousing *Dominator*'s suspicion. Meanwhile, the commitment contains an anamorphic channel, invisible to *Dominator*, to transfer an expected message to the receiver. The formal security model under *Dominator* is presented in Figure 2 and Definition 5.

Definition 5. A standard commitment protocol $\mathcal{SC}=(\mathsf{Gen},\mathsf{Com},\mathsf{Open},\mathsf{Ver})$ can be transformed to an anti-*Dominator* anamorphic commitment if there exists an anamorphic triplet $\mathcal{AMC}=(\mathsf{aKG},\mathsf{aCom},\mathsf{aDec})$ and for any PPT *Dominator* $\mathcal{D}$, we have $|\Pr[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}(\lambda)=1]-\Pr[\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}(\lambda)=1]|\leqslant\mathsf{negl}(\lambda)$.

According to games $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}$ and $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}$, both *Dominator* and *Transcriber* can know the committed message and aim to distinguish an anamorphic commitment from a regular one. In contrast, *Dominator* is stronger than *Transcriber* since *Dominator* can adaptively choose a committed message, even an unwilling one of the sender. Hence, intuitively, we have that an anamorphic commitment of the anti-*Dominator* security is also anti-*Transcriber*. The formal result is shown in Theorem 1. Due to space constraints, we postpone the proof of this theorem to Appendix A.

Theorem 1. If a standard commitment $\mathcal{SC}=(\mathsf{Gen},\mathsf{Com},\mathsf{Open},\mathsf{Ver})$ with an anamorphic triplet $\mathcal{AMC}=(\mathsf{aKG},\mathsf{aCom},\mathsf{aDec})$ is an anti-*Dominator* anamorphic commitment, this commitment is also an anti-*Transcriber* anamorphic commitment.

***Corrupter*: the third type.** The third type, which we refer to as *Corrupter*, possesses capabilities beyond both *Transcriber* and *Dominator*. In addition to coercing the sender, *Corrupter* can manipulate the receiver by tampering with the public parameters. The existence of a *Corrupter* is a realistic and well-motivated assumption. In practice, beyond confirming that the sender has committed honestly, a dictator is also concerned about whether the sender does some additional actions, such as signaling that the sender is under coercion. *Corrupter* captures

$\underline{\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)}$	$\underline{\mathsf{Oa}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})}$	$\underline{\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}(\lambda)}$	$\underline{\mathsf{Oc}(\mathtt{msg},\mathtt{amsg},\mathtt{pp})}$
1. $\mathtt{dk}\leftarrow_{\$}\mathcal{AMC}.\mathsf{aKG}(1^{\lambda})$	1. $(\mathtt{ac},\mathtt{x})\leftarrow_{\$}$ $\mathcal{AMC}.\mathsf{aCom}(\mathtt{dk},\mathtt{msg},\mathtt{amsg},\mathtt{pp})$	1. Do nothing	1. $\mathtt{c}:=\mathcal{SC}.\mathsf{Com}(\mathtt{pp},\mathtt{msg},\mathtt{x})$
2. return $\mathcal{CO}^{\mathsf{Oa}(\mathtt{dk},\cdot,\cdot,\cdot)}()$	2. return $(\mathtt{ac},\mathtt{x})$	2. return $\mathcal{CO}^{\mathsf{Oc}(\cdot,\cdot,\cdot)}()$	2. return $(\mathtt{c},\mathtt{x})$

Figure 3 Games $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}(\lambda)$.

such a suspicious dictator. By tampering with the public parameters, *Corrupter* can either impersonate the receiver to test the sender's honesty or directly take over the receiver to force the sender to commit in a completely controlled environment.

Although *Corrupter* has the coercive capability more potent than that both of *Transcriber* and *Dominator*, anamorphic commitment still allows the sender to resist *Corrupter*. The sender can "honestly" construct an anamorphic commitment, indistinguishable from the regular one, to pass the *Corrupter*'s stricter verification. Meanwhile, the sender can transmit a covert message through an anamorphic channel to counter *Corrupter*. The security model under *Corrupter* is formalized in Definition 6 through the following two games depicted in Figure 3.

Definition 6. A standard commitment $\mathcal{SC}$ is an anti-*Corrupter* anamorphic commitment if there exists an anamorphic triplet $\mathcal{AMC}$ and for any PPT Corrupter $\mathcal{CO}$, we have $|\Pr[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)=1]-\Pr[\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}(\lambda)=1]|\leqslant \mathsf{negl}(\lambda)$.

Compared with game $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}$, game $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}$ shows that *Corrupter* can not only adaptively choose both the committed message $\mathtt{msg}$ and the covert message $\mathtt{amsg}$, but also can adaptively select the public parameter $\mathtt{pp}$. Hence, *Corrupter* is more potent than *Dominator*. Intuitively, we have the following theorem. Since the proof of Theorem 2 is analogous to the proof of Theorem 1, we omit the proof process.

Theorem 2. If a standard commitment protocol $\mathcal{SC}=(\mathsf{Gen},\mathsf{Com},\mathsf{Open},\mathsf{Ver})$ with an anamorphic triplet $\mathcal{AMC}=(\mathsf{aKG},\mathsf{aCom},\mathsf{aDec})$ is an anti-*Corrupter* anamorphic commitment, it is also an anti-*Dominator* anamorphic commitment.

Double key is needed or not. In the anamorphic commitment and dictator models mentioned above, the security of the covert message hinges on the secrecy of the double key. This dependence naturally raises a question: *Is double key essential for an anamorphic commitment?*

The necessity of the double key depends on the entity responsible for generating the public parameters. Because anamorphic commitment relies on standard commitment methods, it presents four cases: *no public parameter*, *sender-generated*, *third-party generated*, and *receiver-generated*. In the first three scenarios, dictators—even the weakest, such as the Transcriber—enjoy the same receiver-level capabilities without a double key, accessing the covert message. In this case, dictators can access the covert message as the receiver does, and it is impossible to construct a secure anamorphic commitment without the double key.

In the fourth case, since the most vigorous dictator, like *Corrupter*, can manipulate the receiver by tampering with the public parameters, *Corrupter* has the same capability as the receiver to access the anamorphic message if no double key. So, we say that the double key is also necessary to guarantee the security of anamorphic commitment under the coercion of *Corrupter*. However, for the two weaker models, like *Transcriber* and *Dominator*, if the sender and the receiver do not share a double key, the receiver's public parameters must enable the sender to construct an asymmetric (or public-key) anamorphic channel, which is similar to the public-key anamorphic model mentioned in [16].

As this work is the first to consider dictator models in commitment protocols, we focus on defending against the most potent dictator *Corrupter*. Therefore, in this work, the anamorphic commitment model must contain the double key. It is also an interesting future work to construct anamorphic commitment without the double key in *Transcriber* and *Dominator* models.

4.3 Semantic securities of messages

An anamorphic commitment contains two messages: a committed message $\mathtt{msg}$ and a covert message $\mathtt{amsg}$. The semantic securities of both $\mathtt{msg}$ and $\mathtt{amsg}$ are vital. If an anamorphic commitment fails to keep the semantic security of $\mathtt{amsg}$, it implies that dictators can discover the covert communication between the sender and the receiver by exploiting the semantic leakage of $\mathtt{amsg}$ and distinguish an anamorphic commitment from a standard commitment. Meanwhile, since an anamorphic commitment can also be a standard commitment in use, the anamorphic commitment must keep the semantic security of $\mathtt{msg}$ to the ordinary adversary as the standard commitment does.

Anti-dictator security implies semantic security of $\mathtt{amsg}$. This reduction is straightforward to establish. For any dictator of the above three dictator models, he can distinguish an anamorphic commitment from a standard

$\underline{\mathsf{IndcpaG}^{\beta}_{\mathcal{AMC},\mathcal{A}}(\lambda)}$
1. $\mathtt{dk} \leftarrow_{\$} \mathcal{AMC}.\mathsf{aKG}(1^{\lambda})$, $\mathtt{pp} \leftarrow_{\$} \mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
2. $(\mathtt{msg}_0, \mathtt{msg}_1, \mathtt{amsg}) \leftarrow_{\$} \mathcal{A}^{\mathsf{Ocpa}(\mathtt{dk},\cdot,\cdot,\mathtt{pp})}(\mathtt{pp})$
3. $(\mathtt{ac}, \mathtt{x}) \leftarrow_{\$} \mathcal{AMC}.\mathsf{aCom}(\mathtt{dk}, \mathtt{msg}_{\beta}, \mathtt{amsg}, \mathtt{pp})$
4. return $\mathcal{A}(\mathtt{ac}, \mathtt{x})$

$\underline{\mathsf{Ocpa}(\mathtt{dk}, \mathtt{msg}, \mathtt{amsg}, \mathtt{pp})}$
1. $(\mathtt{ac}, \mathtt{x}) \leftarrow_{\$} \mathcal{AMC}.\mathsf{aCom}(\mathtt{dk}, \mathtt{msg}, \mathtt{amsg}, \mathtt{pp})$
2. return $(\mathtt{ac}, \mathtt{x})$

Figure 4 Game $\mathsf{IndcpaG}_{\mathcal{AMC},\mathcal{A}}$.

$\underline{\mathsf{RobG}^{0}_{\mathcal{AMC},\mathcal{A}}(\lambda)}$
1. $\mathtt{dk} \leftarrow_{\$} \mathcal{AMC}.\mathsf{aKG}(1^{\lambda})$
2. $\mathtt{pp} \leftarrow_{\$} \mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
3. $b \leftarrow_{\$} \mathcal{A}^{O^0(\mathtt{dk},\cdot,\mathtt{pp})}(\mathtt{pp})$
4. return b

$\underline{O^0(\mathtt{dk}, \mathtt{msg}, \mathtt{pp})}$
1. $\mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$
2. $\mathtt{amsg}' := \mathcal{AMC}.\mathsf{aDec}(\mathtt{dk}, \mathtt{c}, \mathtt{x})$
3. return $\mathtt{amsg}'$

$\underline{\mathsf{RobG}^{1}_{\mathcal{AMC},\mathcal{A}}(\lambda)}$
1. $\mathtt{pp} \leftarrow_{\$} \mathcal{SC}.\mathsf{Gen}(1^{\lambda})$
2. $b \leftarrow_{\$} \mathcal{A}^{O^1(\cdot,\mathtt{pp})}(\mathtt{pp})$
3. return b

$\underline{O^1(\mathtt{msg}, \mathtt{pp})}$
1. return $\perp$

Figure 5 Games $\mathsf{RobG}^{0}_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^{1}_{\mathcal{AMC},\mathcal{A}}(\lambda)$. Note that O^0 and O^1 are oracles that can be queried by $\mathcal{A}$ with any msg.

commitment if existing the semantic leakage of amsg since the dictator can adaptively choose amsg and the standard commitment does not contain any content of amsg.

Anti-dictator security vs. semantic security of msg. All dictator models, including *Transcriber*, *Dominator*, and *Corrupter* model, are distinct with the semantic security of msg. The dictator models define an adversary as one who knows all committed messages, including the challenge message of the challenge commitment. On the contrary, the adversary in semantic security cannot see the challenge commitment's message. Hence, on the one hand, dictator models trivially allow the dictators to know all committed messages. In other words, the dictator models ignore the semantic security msg. Hence, the anti-dictator security does not imply the semantic security of msg.

As an illustrative example, consider an insecure commitment that directly includes the committed message msg, and then transform this commitment to an anamorphic one. Since the anamorphic commitment must be indistinguishable from the insecure standard commitment, the anamorphic commitment also explicitly contains msg. Hence, the anamorphic commitment fails to meet the semantic security of msg even if it is anti-dictator secure.

On the other hand, we cannot combine any dictator model and the semantic security model since their adversaries have distinct capabilities. Hence, we must independently define the semantic security of msg for anamorphic commitment. Let us review the hiding property mentioned in Subsection 3.1. The semantic security of msg follows the hiding property definition of a standard commitment. But, it additionally allows the adversary $\mathcal{A}$ to adaptively select the issued covert message amsg and replace all instances of $\mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{msg}, \mathtt{pp}, \mathtt{x})$ with $(\mathtt{ac}, \mathtt{x}) \leftarrow_{\$} \mathcal{AMC}.\mathsf{aCom}(\mathtt{dk}, \mathtt{msg}, \mathtt{amsg}, \mathtt{pp})$. We use the game IndcpaG presented in Figure 4 to capture the semantic security of msg and give the following definition.

Definition 7. A standard commitment protocol $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with an anamorphic triplet $\mathcal{AMC} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is semantically secure if for any PPT adversary $\mathcal{A}$, $|\Pr[\mathsf{IndcpaG}^{0}_{\mathcal{AMC},\mathcal{A}} = 1] - \Pr[\mathsf{IndcpaG}^{1}_{\mathcal{AMC},\mathcal{A}} = 1]| \leqslant \mathsf{negl}(\lambda)$ holds.

4.4 Robustness

The robustness is complementary to the correctness of an anamorphic commitment. The correctness guarantees that the receiver can correctly decrypt a covert message if the sender has embedded the message in an anamorphic commitment. The robustness guarantees that if the sender does not embed a covert message (namely, the sender deploys the anamorphic commitment as a regular one), the receiver should not extract any covert message. In practice, the robustness ensures that the receiver does not misread a regular commitment. We define the robustness of an anamorphic commitment as follows.

Definition 8. A standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with an anamorphic triplet $\mathcal{AMC} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is robust if $\left|\Pr\left[\mathsf{RobG}^{0}_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right] - \Pr\left[\mathsf{RobG}^{1}_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right]\right| \leqslant \mathsf{negl}(\lambda)$ holds for any PPT adversary $\mathcal{A}$, where games $\mathsf{RobG}^{0}_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^{1}_{\mathcal{AMC},\mathcal{A}}(\lambda)$ are defined in Figure 5.

$\underline{\mathsf{aKG}(1^\lambda)}$
1. $k \leftarrow_\$ \mathcal{K}$, $\mathtt{ctr}:=0$
2. $\mathtt{dk}:=(k, \mathtt{ctr})$
3. **return** $\mathtt{dk}$

$\underline{\mathsf{aCom}(\mathtt{pp}, \mathtt{dk}, \mathtt{msg}, \mathtt{amsg})}$
1. **parse** $\mathtt{dk}:=(k, \mathtt{ctr})$
2. $\mathtt{t}_1 \| \mathtt{t}_2 := \mathcal{F}(k, \mathtt{ctr})$
3. $\mathtt{x}:=(\mathtt{amsg} \| \mathtt{t}_2) \oplus \mathtt{t}_1$
4. $\mathtt{ac}:=\mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$
5. **return** $(\mathtt{ac}, \mathtt{x}; \mathtt{ctr}+1)$

$\underline{\mathsf{aDec}(\mathtt{dk}, \mathtt{x})}$
1. **parse** $\mathtt{dk}:=(k, \mathtt{ctr})$
2. $\mathtt{t}_1 \| \mathtt{t}_2 := \mathcal{F}(k, \mathtt{ctr})$
3. $\mathtt{tmp}:=\mathtt{x} \oplus \mathtt{t}_1$
4. **parse** $\mathtt{tmp}:=\mathtt{amsg} \| \mathtt{t}_2'$
5. **if** $\mathtt{t}_2' \neq \mathtt{t}_2$ **then** **return** $\perp$
6. **return** $(\mathtt{amsg}; \mathtt{ctr}+1)$

Figure 6 $\mathcal{AMC}_\mathcal{P}$ construction.

5 Generic constructions of anamorphic commitment

This section designs two generic anamorphic commitment constructions named $\mathcal{AMC}_\mathcal{P}$ and $\mathcal{AMC}_\mathcal{A}$. Both are anti-*Corrupter*, semantically secure, and robust. $\mathcal{AMC}_\mathcal{P}$ utilizes a pre-computation strategy and a synchronization mechanism to achieve high performance. In contrast, $\mathcal{AMC}_\mathcal{A}$ adopts an asynchronous approach. They are suitable for various application scenarios, and we will discuss them in Subsection 5.3.

$\mathcal{AMC}_\mathcal{P}$ and $\mathcal{AMC}_\mathcal{A}$ have some common design ideas. Under the *Corrupter* model, the only part of the commitment that the sender can control is the random blind factor $\mathtt{x}$. Therefore, the sender can embed the covert message $\mathtt{amsg}$ in $\mathtt{x}$ to construct an anamorphic commitment. To achieve robustness, we introduce an auxiliary token or additional verification process in $\mathcal{AMC}.\mathsf{aCom}$ and a corresponding explicit robustness detection mechanism in $\mathcal{AMC}.\mathsf{aDec}$. Moreover, to construct an anti-*Corrupter* anamorphic commitment, the generation process of $\mathtt{x}$ should be independent of the public parameters $\mathtt{pp}$ chosen by *Corrupter*.

5.1 $\mathcal{AMC}_\mathcal{P}$: an anti-*Corrupter* construction with pre-computation

This section presents $\mathcal{AMC}_\mathcal{P}$, a pre-computable anamorphic commitment leveraging synchronized counters. Both participants initialize a shared counter $\mathtt{ctr}$ (part of the double key $\mathtt{dk}$) to the same value (e.g., 0, timestamp or some other specific value). During commitment generation, $\mathtt{ctr}$ feeds into a pseudo-random function (PRF), producing a pseudo-random value split into two components: one encrypts the covert message $\mathtt{amsg}$, while the other acts as an explicit robustness token. After successful covert message retrieval, sender and receiver synchronously increment $\mathtt{ctr}$ by 1, maintaining state alignment.

Given a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$, a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, Figure 6 shows the whole construction of $\mathcal{AMC}_\mathcal{P} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$. In $\mathcal{AMC}_\mathcal{P}$, the sender and the receiver share a double key dk, where dk consists of a key k of the pseudo-random function $\mathcal{F}$ and a synchronized counter $\mathtt{ctr}$. When embedding a covert message $\mathtt{amsg}$ into an anamorphic commitment, the sender first computes $\mathtt{t}_1 \| \mathtt{t}_2 := \mathcal{F}(k, \mathtt{ctr})$. Next, the sender serves $\mathtt{t}_2$ as the robustness-check part and combines $\mathtt{amsg}$ and $\mathtt{t}_2$. The sender then uses $\mathtt{t}_1$ as a one-time key to encrypt $\mathtt{amsg}$ by performing an XOR operation on $\mathtt{amsg} \| \mathtt{t}_2$ and $\mathtt{t}_1$ which serves as the random factor $\mathtt{x}$ of the commitment. To extract $\mathtt{amsg}$ from $\mathtt{x}$, the receiver first parses $\mathtt{x}$ to obtain the robustness check part $\mathtt{t}_2'$ and then generates $\mathtt{t}_1$ and $\mathtt{t}_2$ in the same manner as the sender; then, if $\mathtt{t}_2' = \mathtt{t}_2$, it means that the sender is processing an anamorphic commitment with the receiver; finally, the receiver uses $\mathtt{t}_1$ to extract $\mathtt{amsg}$.

The correctness of $\mathcal{AMC}_\mathcal{P}$. The correctness property directly follows from the synchronization mechanism. When the sender and receiver are synchronized, the sender and the receiver have the same $\mathtt{t}_1$ and $\mathtt{t}_2$. $\mathcal{AMC}_\mathcal{P}.\mathsf{aDec}$ can successfully extract $\mathtt{amsg}$ by phasing $\mathcal{F}(\mathtt{k}, \mathtt{ctr}) := \mathtt{t}_1 \| \mathtt{t}_2$ and computing $\mathtt{amsg} \| \mathtt{t}_2 = \mathtt{x} \oplus \mathtt{t}_1 = ((\mathtt{amsg} \| \mathtt{t}_2) \oplus \mathtt{t}_1) \oplus \mathtt{t}_1$.

The semantic security analysis of $\mathcal{AMC}_\mathcal{P}$. We propose Theorem 3 to claim that $\mathcal{AMC}_\mathcal{P}$ can transform a standard commitment into a semantic secure anamorphic commitment. Note that the standard commitment has the hiding property shown in Definition 1 by default.

Theorem 3. Given a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with $\mathcal{AMC}_\mathcal{P} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is semantically secure, where $\mathcal{SC}$ meets the hiding property.

The correctness of Theorem 3 is straightforward. First, by Theorem 4, $\mathcal{AMC}_\mathcal{P}$ is an anti-*Corrupter* anamorphic commitment, meaning that *Corrupter* cannot distinguish $\mathcal{AMC}_\mathcal{P}$ from the standard commitment $\mathcal{SC}$. Additionally, according to games $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}$ and $\mathsf{IndcpaG}_{\mathcal{AMC},\mathcal{A}}$, *Corrupter* is strictly stronger than the regular adversary $\mathcal{A}$ in a commitment protocol. We have that $\mathcal{A}$ is also unable to distinguish $\mathcal{AMC}_\mathcal{P}$ from $\mathcal{SC}$. Hence, according to the hiding property of $\mathcal{SC}$, $\mathcal{AMC}_\mathcal{P}$ must be semantically secure in the view of $\mathcal{A}$.

$\underline{\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda)}$
1. $\mathtt{sk} \leftarrow_{\$} \mathcal{PSE}.\mathsf{KG}(1^\lambda)$
2. $b \leftarrow_{\$} \mathcal{A}^{\mathsf{Oe}(\mathtt{sk},\cdot)}()$
3. return b

$\underline{\mathsf{Oe}(\mathtt{sk},\mathtt{msg})}$
1. $\mathtt{ct} \leftarrow_{\$} \mathcal{PSE}.\mathsf{Enc}(\mathtt{sk},\mathtt{msg})$
2. return $\mathtt{ct}$

$\underline{\mathsf{RandG}_{\mathcal{R},\mathcal{A}}(\lambda)}$
1. $b \leftarrow_{\$} \mathcal{A}^{\mathsf{Or}(\cdot)}()$
2. return b

$\underline{\mathsf{Or}(\mathtt{msg})}$
1. $c \leftarrow_{\$} \mathcal{C}$
2. return c

Figure 7 Games $\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda)$ and $\mathsf{Random}_{\mathcal{R},\mathcal{A}}(\lambda)$.

$\underline{\mathsf{KG}(1^\lambda)}$
1. $k \leftarrow_{\$} \mathcal{K}$
2. $\mathtt{sk} := k$
3. **return** $\mathtt{sk}$

$\underline{\mathsf{Enc}(\mathtt{sk},\mathtt{msg})}$
1. $r \leftarrow_{\$} \{0,1\}^n$
2. $\mathtt{ct} := r \| (\mathtt{msg} \oplus \mathcal{F}(\mathtt{sk}, r))$
3. **return** $\mathtt{ct}$

$\underline{\mathsf{Dec}(\mathtt{sk},\mathtt{ct})}$
1. **parse** $\mathtt{ct} := r \| t$
2. $\mathtt{msg} := t \oplus \mathcal{F}(\mathtt{sk}, r)$
3. **return** $\mathtt{msg}$

Figure 8 $\mathcal{PSE}$ introduced in [13].

The security and rubstness analysis of $\mathcal{AMC}_{\mathcal{P}}$. We give the following two theorems to show the anti-*Corrupter* security and robustness of $\mathcal{AMC}_{\mathcal{P}}$, and due to space constraints, we postpone the proof procedure to Appendix B.

Theorem 4. Given a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_{\mathcal{P}} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is an anti-*Corrupter* anamorphic one.

Theorem 5. Given a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with $\mathcal{AMC}_{\mathcal{P}} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is robust.

Pre-computation strategy. According to $\mathcal{AMC}_{\mathcal{P}}$ shown in Figure 6, the pseudo-random function $\mathcal{F}$ takes only the double key as input. The computation of $\mathcal{F}$ is independent of the commitment message $\mathtt{msg}$ and the covert message $\mathtt{amsg}$. Thus, the offline sender and receiver can pre-compute and store the values of $\mathtt{t}_1 \| \mathtt{t}_2 := \mathcal{F}(k, \mathtt{ctr})$ in advance. When generating the blind factor $\mathtt{x}$ for a given $\mathtt{amsg}$, the online sender only implements some lightweight operations, such as an XOR operation, a concatenation operation, and two lookup operations. Thus, the online efficiency of $\mathcal{AMC}_{\mathcal{P}}.\mathsf{aCom}$ is close to the standard commitment.

5.2 $\mathcal{AMC}_{\mathcal{A}}$: an asynchronous anti-*Corrupter* construction

$\mathcal{AMC}_{\mathcal{A}}$ is an asynchronous anamorphic commitment construction that leverages pseudo-random symmetric encryption ($\mathcal{PSE}$) to eliminate the synchronization process used in $\mathcal{AMC}_{\mathcal{P}}$. We start by revisiting $\mathcal{PSE}$. $\mathcal{PSE}$ ensures that, for any message chosen by a PPT adversary without the private key, the ciphertext is indistinguishable from a random value. The existence of $\mathcal{PSE}$ relies on the assumption of the existence of a one-way function, one of the most fundamental cryptographic assumptions and applies to common algorithms like AES in CTR mode. The formal definition of $\mathcal{PSE}$ is as follows.

Definition 9. An IND-CPA symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ with ciphertext space $\mathcal{C}$ is a pseudo-random symmetric encryption if for any PPT adversary $\mathcal{A}$, we have $|\Pr[\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda) = 1] - \Pr[\mathsf{RandG}_{\mathcal{R},\mathcal{A}}(\lambda) = 1]| \leqslant \mathsf{negl}(\lambda)$, where games $\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda)$ and $\mathsf{RandG}_{\mathcal{R},\mathcal{A}}(\lambda)$ are in Figure 7.

The underlying $\mathcal{PSE}$ must have the feature of "uniform decryption" to achieve the robustness of $\mathcal{AMC}_{\mathcal{A}}$. This feature means that when decrypting a randomly chosen ciphertext (not the one generated by encryption) from the ciphertext space, the decryption result is uniform in the message space. We define the uniform decryption feature of $\mathcal{PSE}$ as follows.

Theorem 6. An IND-CPA pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\mathcal{M}$ and ciphertext space $\mathcal{C}$ has uniform decryption if for any PPT adversary $\mathcal{A}$, $|\Pr[\mathcal{A}(c, \mathtt{msg}) = 1] - \Pr[\mathcal{A}(c, \mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, c)) = 1]| \leqslant \mathsf{negl}(\lambda)$ holds, where $c \leftarrow_{\$} \mathcal{C}$, $\mathtt{sk} \leftarrow_{\$} \mathcal{PSE}.\mathsf{KG}(1^\lambda)$, and $\mathtt{msg} \leftarrow_{\$} \mathcal{M}$.

We take the $\mathcal{PSE}$ scheme introduced in [13] as an example and demonstrate its uniform decryption feature. Hence, the underlying $\mathcal{PSE}$ of $\mathcal{AMC}_{\mathcal{A}}$ is efficient. Given a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, the $\mathcal{PSE}$ scheme is shown in Figure 8. It is easy to find that if $\mathtt{ct}$ is randomly chosen in the ciphertext space rather than a regular ciphertext, namely r and t are random, $\mathtt{msg} := t \oplus \mathcal{F}(\mathtt{sk}, r)$ is also indistinguishable from a pseudo-random string. This result means that the $\mathcal{PSE}$ scheme meets the uniform decryption property.

Compared with $\mathcal{AMC}_{\mathcal{P}}$, $\mathcal{AMC}_{\mathcal{A}}$ introduces a novel mechanism to embed a covert message $\mathtt{amsg}$ into a $\mathcal{PSE}$'s ciphertext to replace the blind factor of $\mathcal{SC}$. Initially, the sender and the receiver share a secret key $\mathtt{sk}$ of $\mathcal{PSE}$. During a commitment process, the sender encrypts a covert message $\mathtt{amsg}$ using $\mathcal{PSE}.\mathsf{Enc}$ and takes the resulting

$\underline{\mathsf{aKG}(1^\lambda)}$
1. $k \leftarrow_\$ \{0,1\}^m$
2. $\mathtt{sk} \leftarrow_\$ \mathcal{PSE}.\mathsf{KG}(1^\lambda)$
3. $\mathtt{dk} := (k, \mathtt{sk})$
4. **return** $\mathtt{dk}$

$\underline{\mathsf{aCom}(\mathtt{pp}, \mathtt{dk}, \mathtt{msg}, \mathtt{amsg})}$
1. **parse** $\mathtt{dk} := (k, \mathtt{sk})$
2. $\mathtt{x} \leftarrow_\$ \mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{amsg} \| k)$
3. $\mathtt{ac} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$
4. **return** $(\mathtt{ac}, \mathtt{x})$

$\underline{\mathsf{aDec}(\mathtt{dk}, \mathtt{x})}$
1. **parse** $\mathtt{dk} := (k, \mathtt{sk})$
2. $M := \mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, \mathtt{x})$
3. **parse** $M := \mathtt{amsg} \| k'$
4. **if** $k' \neq k$ **then**
   return $\perp$
5. **return** $\mathtt{amsg}$

Figure 9 $\mathcal{AMC}_\mathcal{A}$ construction.

pseudo-random ciphertext as a blind factor $\mathtt{x}$ to generate an anamorphic commitment $\mathtt{ac}$. In the opening phase, upon receiving $\mathtt{x}$, the receiver decrypts $\mathtt{x}$ using $\mathtt{sk}$ to obtain $\mathtt{amsg}$. To ensure robustness, $\mathcal{AMC}_\mathcal{A}$ adopts the same generic robust-transformation principle as in $\mathcal{AMC}_\mathcal{P}$. In $\mathcal{AMC}_\mathcal{A}$, the sender and receiver share a robustness token k as a part of the double key. The sender uses $\mathcal{PSE}$ to encrypt $\mathtt{amsg} \| k$. During the anamorphic decryption phase, the receiver first decrypts $\mathtt{x}$. Then, he determines there is a covert message in the commitment if the token k appears; otherwise, he outputs $\perp$.

Given a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$, a pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ with uniform decryption, the construction of anamorphic triplet $\mathcal{AMC}_\mathcal{A} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is shown in Figure 9. Let m denote the binary length of a robustness token.

The correctness of $\mathcal{AMC}_\mathcal{A}$. If a blind factor $\mathtt{x}$ is generated by $\mathcal{AMC}_\mathcal{A}.\mathsf{aCom}$, $\mathtt{x}$ is a $\mathcal{PSE}$ ciphertext. $\mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, \mathtt{x})$ will correctly output $\mathtt{amsg} \| k$ with an overwhelming probability.

The semantic security analysis of $\mathcal{AMC}_\mathcal{A}$. Theorem 7 claims that $\mathcal{AMC}_\mathcal{A}$ can also transform a standard commitment into a semantic secure anamorphic commitment.

Theorem 7. Given a pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_\mathcal{A} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is semantically secure, where $\mathcal{SC}$ meets the hiding property.

The reason is similar to that of Theorem 3. To begin, $\mathcal{AMC}_\mathcal{A}$ also meets anti-*Corrupter* anamorphism; it naturally follows that a regular adversary $\mathcal{A}$ in commitment also cannot distinguish between $\mathcal{AMC}_\mathcal{A}$ and $\mathcal{SC}$. Hence, $\mathcal{AMC}_\mathcal{A}$ inherits $\mathcal{SC}$'s hiding property and achieves the semantic security.

The security and robustness analysis of $\mathcal{AMC}_\mathcal{A}$. We propose Theorems 8 and 9 to claim that $\mathcal{AMC}_\mathcal{A}$ can transform a standard commitment into an anti-*Corrupter* anamorphic one while achieving robustness. Due to space constraints, we postpone the proof procedure to Appendix C.

Theorem 8. Given a pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_\mathcal{A} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is an anti-*Corrupter* anamorphic one.

Theorem 9. Given a pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_\mathcal{A} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is robust.

5.3 Comparisons between $\mathcal{AMC}_\mathcal{P}$ and $\mathcal{AMC}_\mathcal{A}$

Both $\mathcal{AMC}_\mathcal{P}$ and $\mathcal{AMC}_\mathcal{A}$ provide anti-*Corrupter* anamorphism, semantic security, and robustness, ensuring resilience against even the strongest dictator and preventing the receiver from misinterpreting covert messages. $\mathcal{AMC}_\mathcal{P}$ uses pre-computation to reduce online costs, achieving performance nearly equivalent to the underlying $\mathcal{SC}.\mathsf{Com}$. Although $\mathcal{AMC}_\mathcal{A}$ incurs slightly higher online cost, it remains practical by requiring only one additional $\mathcal{PSE}.\mathsf{Enc}$ operation, as illustrated in Figure 8.

Anamorphic bandwidth analysis. Anamorphic bandwidth refers to the binary length of the covert message that an anamorphic commitment can transmit. Evidently, a larger covert-message length corresponds to higher covert-communication performance. In short, both $\mathcal{AMC}_\mathcal{P}$ and $\mathcal{AMC}_\mathcal{A}$ have a high anamorphic bandwidth. They can transmit a long covert message, even if their underlying standard commitment is a bit commitment.

Let r represent the binary length of the blind factor of the underlying $\mathcal{SC}.\mathsf{Com}$, and l indicate the binary length of the $\mathcal{PSE}$ message space. Note that in the preceding discussion, m denotes the bit length of the robustness token. Similarly, n represents the bit length of the randomness space of $\mathcal{PSE}.\mathsf{Enc}$. It is easy to find that the anamorphic bandwidths of $\mathcal{AMC}_\mathcal{P}$ and $\mathcal{AMC}_\mathcal{A}$ are $\mathcal{O}(r-m)$ and $\mathcal{O}(l-m)$, respectively, where we have $l = r - n$ since $\mathcal{PSE}$ is an IND-CPA secure symmetric encryption. In theory, both r and l can be polynomial-level. Hence, both $\mathcal{AMC}_\mathcal{P}$

Table 2 Theoretical comparison across four dimensions of previous anamorphic studies and our work.

	Ref. [3]		Ref. [13]	Ref. [14]		Our work	
	PPY22a	PPY22b	KPPY23	Σ_2	Σ_3	$\mathcal{AMC}_{\mathcal{P}}$	$\mathcal{AMC}_{\mathcal{A}}$
Generality	✓	✗	✓	✓	✓	✓	✓
Robustness	✗	✗	✗	✓	✓	✓	✓
Sync./Async.	Async.	Async.	Async.	Sync.	Async.	Sync.	Async.
`amsg`'s length	$O(\log n)$	$O(n)$	$O(n)$	$O(\log n)$	$O(\log n)$	$O(n)$	$O(n)$

and $\mathcal{AMC}_{\mathcal{A}}$ can achieve polynomial-level anamorphic bandwidth, which is the optimum level in the field of current anamorphic cryptography.

In practice, the user can set m and n independently (e.g., $m = n = 128$) to guarantee that the probability of the receiver misunderstanding a regular commitment is negligible and keep the IND-CPA security of $\mathcal{PSE}$. Standard commitments typically use an r value much larger than m and n—for example, Pedersen commitments usually sets $r = 256$ at least to guarantee hiding. Hence, we can rely on the specific application requirement about the anamorphic bandwidth to initialize a standard commitment having a big enough r as the underlying primitive to construct a robust anamorphic commitment.

Synchronization vs. asynchronization. The key difference between $\mathcal{AMC}_{\mathcal{P}}$ and $\mathcal{AMC}_{\mathcal{A}}$ lies in state synchronization. In $\mathcal{AMC}_{\mathcal{P}}$, a synchronous design using a shared counter and precomputed pseudo-random values ensures that commitment instances are generated and verified sequentially, ideal for scenarios requiring continuous state updates and optimized performance. In contrast, $\mathcal{AMC}_{\mathcal{A}}$ employs an asynchronous approach in which the sender and the receiver operate independently. This makes $\mathcal{AMC}_{\mathcal{A}}$ adaptable to environments that handle multiple concurrently independent commitments, providing enhanced flexibility and scalability. These two kinds of scenarios are popular in practice.

$\mathcal{AMC}_{\mathcal{P}}$'s synchronous architecture benefits settings where sequential integrity is crucial. In secure multi-round contract negotiations, parties must commit to terms sequentially. With a synchronized counter, $\mathcal{AMC}_{\mathcal{P}}$ ensures commitments occurring in order, thereby preventing replay attacks and preserving negotiation integrity. Offline pre-computation and low online overhead make $\mathcal{AMC}_{\mathcal{P}}$ ideal for cryptographic applications that require strict sequencing and fast responses. This rigorous sequence enhances the overall security and reliability of cryptographic operations in practice.

On the other hand, $\mathcal{AMC}_{\mathcal{A}}$'s asynchronous design suits scenarios requiring independent, concurrent commitments. For example, in decentralized systems, nodes may commit to parameters or proposals at arbitrary times without global synchronization. Its asynchronous nature allows commitments to be generated and verified in isolation, simplifying state management and enabling parallel processing. This flexibility is especially valuable in scalable systems that must handle multiple independent commitment processes efficiently and concurrently.

6 Analysis and experimental results

This section presents both theoretical and experimental analyses to demonstrate the advantages of our two proposed anamorphic constructions. From a theoretical standpoint, our work is the first to introduce anamorphic commitments. Hence, we have to compare our schemes with the most closely related primitive anamorphic encryption for clarifying their functional distinctions. Specifically, we evaluate $\mathcal{AMC}_{\mathcal{P}}$ and $\mathcal{AMC}_{\mathcal{A}}$ against representative anamorphic encryption schemes, including the rejection-sampling and Naor-Yung variants from [3] (we name these two schemes as PPY22a and PPY22b, respectively), the generic paradigm of [13] (we name this scheme as KPPY23), and the Σ_2/Σ_3 schemes from [14]. The comparison spans four dimensions, such as generality (applicability across underlying schemes), robustness (ability to avoid false positives), synchrony/asynchrony, and covert-message length. Table 2 shows that our $\mathcal{AMC}_{\mathcal{P}}$ and $\mathcal{AMC}_{\mathcal{A}}$ constructions provide generic and robust transformations that are applicable to a broad class of standard commitments while achieving the highest transmission efficiency compared with existing designs.

Next, we proceed with the experimental analysis. The above sections have proven that the dictator cannot distinguish $\mathcal{SC}$ and $\mathcal{AMC}$ based on commitments and blind factors. This section experimentally demonstrates that the dictator also cannot distinguish $\mathcal{SC}$ and $\mathcal{AMC}$ based on their runtimes.

To compare the performance of standard and anamorphic commitments, we choose Naor bit commitment and the widely used Pedersen commitment based on the multiplicative group of a finite field [5] as the benchmarks and extend each commitment to three generic anamorphic commitment schemes, namely the synchronized construction $\mathcal{AMC}_{\mathcal{P}}$ and pre-compute variant p-$\mathcal{AMC}_{\mathcal{P}}$ and the non-synchronized construction $\mathcal{AMC}_{\mathcal{A}}$, as the method described

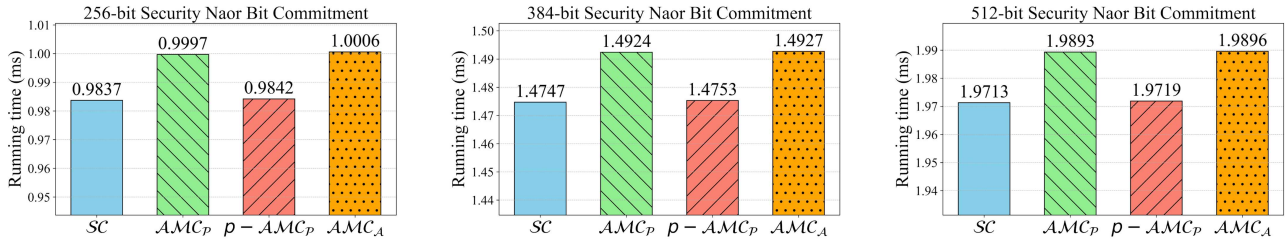

Figure 10 (Color online) Performance comparisons between standard and anamorphic Naor bit commitments.

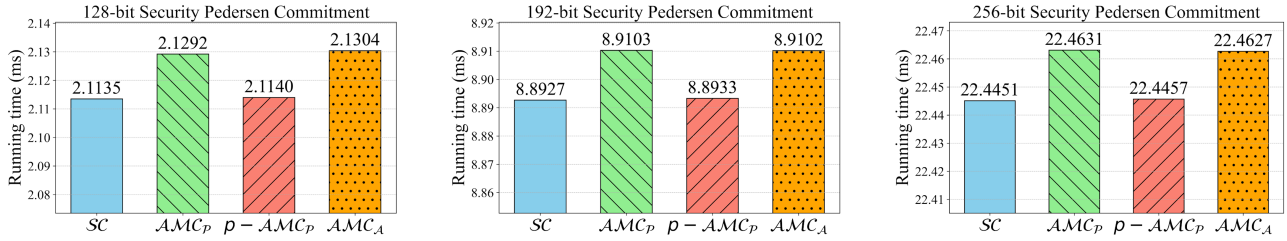

Figure 11 (Color online) Performance comparisons between standard and anamorphic Pedersen commitments.

in Section 5. Both $\mathcal{AMC}_{\mathcal{P}}$ and p-$\mathcal{AMC}_{\mathcal{P}}$ realize their pseudo-random functions as $\mathcal{F}(k,m) = \text{SHAKE128}(k||m)$. $\mathcal{AMC}_{\mathcal{A}}$ uses AES in CTR mode to realize the pseudo-random symmetric encryption. We implement the above four schemes with the Crypto library[1)] in Python and run them on an Intel Core i5-12500H (2.5 GHz, 16 GB RAM) under Windows 10.

To illustrate that the dictator cannot distinguish whether the sender is using an anamorphic or standard commitment based on the running efficiency of the commitment algorithm, we run $\mathcal{SC}.\mathsf{Com}$, $\mathcal{AMC}_{\mathcal{P}}.\mathsf{aCom}$, $p{-}\mathcal{AMC}_{\mathcal{P}}.\mathsf{aCom}$, and $\mathcal{AMC}_{\mathcal{A}}.\mathsf{aCom}$ 10000 times for 128, 256, and 384 bits `amsg` (for Naor bit commitment, under 256-, 384-, and 512-bit security; for Pedersen commitment, under 128-, 192-, and 256-bit security) and record their average execution times. Figures 10 and 11 show the average execution times of Naor bit commitment and Pedersen commitment and their anamorphic variant, respectively. These results show that the runtimes of the three anamorphic variants are very close to that of the standard commitment, growing with approximately 0.03%–1.72% time cost of Naor bit commitment and 0.02%–0.74% time cost of Pedersen commitment. The experiment also demonstrates that our anamorphic schemes are very efficient.

7 Conclusion and future work

This work demonstrates that, through anamorphic cryptography, one can reveal previously ignored potential adversarial scenarios. Technically, it considers the coercive dictator for the first time in commitment protocols. To address this dictator, we model anamorphic commitment and various levels of dictator strength, from weak to strong. We construct two anamorphic commitment protocols that allow the sender to alert coercion in secret or deliver covert messages, even under the most aggressive dictator. These anamorphic protocols are not exotic but are generically allowed within standard commitment protocols (originally designed without anamorphism in mind). They can also effectively counter even stronger dictators who can corrupt the receiver. Our contribution implies that, despite the dictator's increases in surveillance and coercion, covert messaging can still be maintained (to regain privacy) unless all standard commitments are banned, which is impractical (as they imply that basic protocols like auctions, elections, or zero-knowledge proofs, and any procedure which requires "independence of inputs" are eliminated).

Our work also strongly supports the meta-conjecture proposed in [3] that "*For any standard scheme, there is a technical demonstration of the futility of the dictator's demands.*" Noting the widespread use of commitment protocols as building blocks in cryptography, we believe that, in addition to resisting dictators in the commitment, one can construct more anamorphic cryptosystems based on our anamorphic commitment protocol.

This work is the first to study anamorphic commitment. Several interesting future studies deserve exploration. For example, it is still unclear how to precisely characterize the power and limitations of double-key-free anamorphic commitment, and an important question is whether public-key or keyless variants can resist stronger dictators capable of tampering with parameters. In addition, as the standard commitment has done in the past, anamorphic

1) https://github.com/Legrandin/pycryptodome.

commitment should also be a promising primitive for promoting the development of other cryptographic schemes, such as more efficient ZK proofs and coercion-resistant voting systems. Hence, exploring the applications of anamorphic commitment is also a valuable direction.

Acknowledgements This work was supported by National Natural Science Foundation of China (Grant Nos. 62272186, 62372201) and Science and Technology Innovation Program of Hunan Province (Grant No. 2025RC1040).

References

1 Feng T, Li F H, Ma J F, et al. A new approach for UC security concurrent deniable authentication. Sci China Ser F-Inf Sci, 2008, 51: 352–367
2 Yan Z B, Deng Y. A novel approach to public-coin concurrent zero-knowledge and applications on resettable security. Sci China Inf Sci, 2019, 62: 032110
3 Persiano G, Phan D H, Yung M. Anamorphic encryption: private communication against a dictator. In: Advances in Cryptology—EUROCRYPT 2022. Berlin: Springer, 2022. 34–63
4 Fujisaki E, Okamoto T. Statistical zero knowledge protocols to prove modular polynomial relations. In: Advances in Cryptology—CRYPTO'97. Berlin: Springer, 1997. 16–30
5 Pedersen T P. Non-interactive and information-theoretic secure verifiable secret sharing. In: Advances in Cryptology—CRYPTO'91. Berlin: Springer, 1991. 129–140
6 Naor M. Bit commitment using pseudorandomness. J Cryptology, 1991, 4: 151–158
7 Damgård I, Fujisaki E. A statistically-hiding integer commitment scheme based on groups with hidden order. In: Advances in Cryptology—ASIACRYPT 2002. Berlin: Springer, 2002. 125–142
8 Fischlin M, Fischlin R. Efficient non-malleable commitment schemes. J Cryptol, 2011, 24: 203–244
9 Damgård I. Practical and provably secure release of a secret and exchange of signatures. In: Advances in Cryptology—EUROCRYPT'93. Berlin: Springer, 1993. 200–217
10 Attema T, Lyubashevsky V, Seiler G. Practical product proofs for lattice commitments. In: Advances in Cryptology—CRYPTO 2020. Berlin: Springer, 2020. 470–499
11 Damgård I. Commitment schemes and zero-knowledge protocols. In: Lectures on Data Security. Berlin: Springer, 1998. 63–86
12 Hieu P D, Yung M. Privacy in advanced cryptographic protocols: prototypical examples. J Comput Sci Cybern, 2021, 37: 429–451
13 Kutylowski M, Persiano G, Phan D H, et al. The self-anti-censorship nature of encryption: on the prevalence of anamorphic cryptography. PoPETs, 2023, 2023: 170–183
14 Banfi F, Gegier K, Hirt M, et al. Anamorphic encryption, revisited. In: Advances in Cryptology—EUROCRYPT 2024. Berlin: Springer, 2024. 3–32
15 Catalano D, Giunta E, Migliaro F. Anamorphic encryption: new constructions and homomorphic realizations. In: Advances in Cryptology—EUROCRYPT 2024. Berlin: Springer, 2024. 33–62
16 Persiano G, Phan D H, Yung M. Public-key anamorphism in (CCA-secure) public-key encryption and beyond. In: Advances in Cryptology—CRYPTO 2024. Berlin: Springer, 2024. 422–455
17 Catalano D, Giunta E, Migliaro F. Generic anamorphic encryption, revisited: new limitations and constructions. Cryptology ePrint Archive, 2024. https://eprint.iacr.org/2024/1119
18 Catalano D, Giunta E, Migliaro F. Limits of black-box anamorphic encryption. In: Advances in Cryptology—CRYPTO 2024. Berlin: Springer, 2024. 352–383
19 Wang Y, Chen R M, Huang X Y, et al. Sender-anamorphic encryption reformulated: achieving robust and generic constructions. In: Advances in Cryptology—ASIACRYPT 2023. Berlin: Springer, 2023. 135–167
20 Dodis Y, Goldin E. Anamorphic-resistant encryption; or why the encryption debate is still alive. In: Advances in Cryptology—CRYPTO 2025. Berlin: Springer, 2025. 440–471
21 Carnemolla D, Catalano D. Anamorphic resistant encryption: the good, the bad and the ugly. In: Advances in Cryptology—CRYPTO 2025. Berlin: Springer, 2025. 472–503
22 Wang W, Han S, Liu S L. Anamorphic authenticated key exchange: double key distribution under surveillance. In: Advances in Cryptology—ASIACRYPT 2024. Berlin: Springer, 2024. 168–200
23 Kutylowski M, Persiano G, Phan D H, et al. Anamorphic signatures: secrecy from a dictator who only permits authentication! In: Advances in Cryptology—CRYPTO 2023. Berlin: Springer, 2023. 759–790
24 Jaeger J, Stracovsky R. Dictators? Friends? Forgers.—Breaking and fixing unforgeability definitions for anamorphic signature schemes. In: Advances in Cryptology—ASIACRYPT 2024. Berlin: Springer, 2024. 105–137
25 Simmons G J. The prisoners' problem and the subliminal channel. In: Advances in Cryptology—CRYPTO'83. New York: Plenum Press, 1983. 51–67
26 Simmons G J. Subliminal communication is easy using the DSA. In: Advances in Cryptology—EUROCRYPT'93. Berlin: Springer, 1993. 218–232
27 Simmons G J. A secure subliminal channel (?). In: Advances in Cryptology— CRYPTO '85. Berlin: Springer, 1985. 33–41
28 Koptyra K, Ogiela M R. Subliminal channels in visual cryptography. Cryptography, 2022, 6: 46
29 Biryukov A, Feher D, Vitto G. Privacy aspects and subliminal channels in Zcash. In: Proceedings of ACM SIGSAC Conference on Computer and Communications Security, London, 2019. 1795–1811
30 Mironov I, Stephens-Davidowitz N. Cryptographic reverse firewalls. In: Proceedings of the 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT 2015), 2015. 657–686
31 Chen R M, Chen J R, Huang X Y, et al. RCCA-SM9: securing SM9 on corrupted machines. Sci China Inf Sci, 2024, 67: 212103
32 Brassard G, Chaum D, Crépeau C. Minimum disclosure proofs of knowledge. J Comput Syst Sci, 1988, 37: 156–189
33 Catalano D, Visconti I. Hybrid trapdoor commitments and their applications. In: Automata, Languages and Programming, Lisbon, Portugal. Berlin: Springer, 2005. 298–310
34 Hanaoka G, Schuldt J C N. Signatures from trapdoor commitments with strong openings. In: Proceedings of International Symposium on Information Theory and its Applications, Monterey, 2016. 81–85
35 Baghery K. Subversion-resistant commitment schemes: definitions and constructions. In: Proceedings of International Workshop on Security and Trust Management, Guildford, 2020. 106–122
36 Chase M, Healy A, Lysyanskaya A, et al. Mercurial commitments with applications to zero-knowledge sets. In: Advances in Cryptology—EUROCRYPT 2005. Berlin: Springer, 2005. 422–439
37 Catalano D, Dodis Y, Visconti I. Mercurial commitments: minimal assumptions and efficient constructions. In: Theory of Cryptography. Berlin: Springer, 2006. 120–144
38 Chen X, Susilo W, Zhang F, et al. Identity-based trapdoor mercurial commitments and applications. Theor Comput Sci, 2011, 412: 5498–5512
39 Libert B, Yung M. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In: Theory of Cryptography. Berlin: Springer, 2010. 499–517
40 Wang H, Yiu S M, Zhao Y, et al. Updatable, aggregatable, succinct mercurial vector commitment from lattice. IACR Cryptol ePrint Arch, 2024. https://eprint.iacr.org/2024/027

41 Khurana D, Waters B. On the CCA compatibility of public-key infrastructure. In: Public-Key Cryptography. Berlin: Springer, 2021. 235–260

Appendix A Relationship between dictator models

Theorem A1. If a standard commitment protocol $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with an anamorphic triplet $\mathcal{AMC} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is an anti-*Dominator* anam-orphic commitment, it is also an anti-*Transcriber* anamorphic commitment.

Proof. Let $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ be a commitment protocol with an anamorphic triplet $\mathcal{AMC} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$, an anti-*Transcriber* anamorphic commitment. We assume that there exists a *Transcriber* $\mathcal{T}$ that can distinguish between $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}$ with non-negligible probability $1/\mathsf{poly}(\lambda)$ and then we can construct a *Dominator* $\mathcal{D}$ who can distinguish $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}$ with the same probability.

According to Definition 5, $\mathcal{D}$ can access its oracle O, where either $O = \mathsf{Oa}$ that returns an anmorphic commitment embedding a covert message, or $O = \mathsf{Oc}$ that returns a regular commitment. We can construct $\mathcal{D}$ as follows: (1) $\mathcal{D}$ receives public parameters $\mathtt{pp}$ and runs $\mathcal{T}$ with input $\mathtt{pp}$, where $\mathcal{T}$ can issue $\mathtt{amsg}$ to his own oracle; (2) for each issue from $\mathcal{T}$, $\mathcal{D}$ simulates $\mathcal{T}$'s oracle by choosing a message $\mathtt{msg}$ and feeding $(\mathtt{msg}, \mathtt{amsg})$ to O; (3) after receiving $(\mathtt{c}, \mathtt{x})$ from O, $\mathcal{D}$ answers to $\mathcal{T}$ with $(\mathtt{c}, \mathtt{x})$; (4) when $\mathcal{T}$ returns his result $b \in \{0, 1\}$, $\mathcal{D}$ returns the same b as the result.

We now analyze $\mathcal{T}$'s view relative to his own oracle and outputs received from $\mathcal{D}$. First, the public parameters $\mathtt{pp}$ are generated by $\mathcal{D}$'s oracle using the same $\mathcal{SC}.\mathsf{Gen}$, just like in the games $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}$.

Next, in response to the game played by $\mathcal{D}$, let us analyze the situation in detail. If $\mathcal{D}$ plays $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}$, then the commitment $\mathcal{T}$ receives is an anamorphic commitment embedded with his chosen covert message $\mathtt{amsg}$ and $\mathcal{T}$'s view is the same as $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}$. Therefore, $\Pr\left[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}(\lambda) = 1\right] = \Pr\left[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{T}}(\lambda) = 1\right]$ holds. Then if $\mathcal{D}$ plays $\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}$, the commitment $\mathcal{T}$ receives is a regular commitment independent of his choice of anamorphic message $\mathtt{amsg}$. Thus $\mathcal{T}$'s view is the same as $\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}$, we have $\Pr\left[\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}(\lambda) = 1\right] = \Pr\left[\mathsf{RealG}_{\mathcal{SC},\mathcal{T}}(\lambda) = 1\right]$.

To sum up, $\mathcal{D}$ simulates the view of $\mathcal{T}$. Since $\mathcal{T}$ can distinguish these two games, $\mathcal{D}$ can also distinguish his own games with the same probability, which is $\left|\Pr\left[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{D}}(\lambda) = 1\right] - \Pr\left[\mathsf{RealG}_{\mathcal{SC},\mathcal{D}}(\lambda) = 1\right]\right| = 1/\mathsf{poly}(\lambda)$. In such a case, $\mathcal{D}$ can distinguish between an anti-*Dominator* anamorphic commitment and a regular one, which contradicts Definition 5.

Appendix B Anamorphism and robustness of $\mathcal{AMC}_{\mathcal{P}}$

Theorem B1. Given a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_{\mathcal{P}} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is an anti-*Corrupter* anamorphic one.

Proof. In order to complete the proof, we first introduce an intermediate game $\mathsf{ImG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$. It is the same as $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ except $\mathtt{t}_1\|\mathtt{t}_2$ is generated by a true random functions f rather than by the pseudo-random function $\mathcal{F}$.

Lemma B1. Given a PRF $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, for any PPT adversary $\mathcal{CO}$, $|\Pr[\mathsf{ImG}_{\mathcal{AMC},\mathcal{CO}}(\lambda) = 1] - \Pr[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda) = 1]| \leqslant \mathsf{negl}(\lambda)$.

Proof. We suppose that there exists a PPT *Corrupter* $\mathcal{CO}$ that can distinguish between $\mathsf{ImG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ and $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ with a non-negligible probability. We can then construct a PPT adversary $\mathcal{A}$ to distinguish between a pseudo-random function and a real random function. Without loss of generality, we set this non-negligible probability to be $1/\mathsf{poly}(\lambda)$, which implies $|\Pr[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda) = 1] - \Pr[\mathsf{ImG}_{\mathcal{SC},\mathcal{CO}}(\lambda) = 1]| = 1/\mathsf{poly}(\lambda)$.

Suppose that $\mathcal{A}$ can access an oracle O which outputs either a pseudo-random value or a real random number. We can construct $\mathcal{A}$ contradicts the assumption of pseudo-random function like the following: (1) $\mathcal{A}$ initializes $\mathtt{ctr}$, where $\mathcal{CO}$ can issue $(\mathtt{msg}, \mathtt{amsg}, \mathtt{pp})$ to his oracle; (2) for each issue $(\mathtt{msg}, \mathtt{amsg}, \mathtt{pp})$ from $\mathcal{CO}$, $\mathcal{A}$ simulates the response of $\mathcal{CO}$'s oracle by feeding $\mathtt{ctr}$ to O and updates $\mathtt{ctr}$; (3) after receiving y from O, $\mathcal{A}$ parses $y := \mathtt{t}_1\|\mathtt{t}_2$, computes $\mathtt{x} := (\mathtt{amsg}\|\mathtt{t}_2) \oplus \mathtt{t}_1$ and generates $\mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$; then $\mathcal{A}$ uses $(\mathtt{c}, \mathtt{x})$ to answer $\mathcal{CO}'s$ issues; (4) after $\mathcal{CO}$ outputs his answer b, $\mathcal{A}$ uses b as his own judgment of O.

Without loss of generality, we can assume that $\mathtt{ctr}$ is never repeating. It is obviously that no matter in $\mathsf{ImG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ or $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$, $\mathtt{c}$ is generated by Com of regular commitment, $\mathtt{msg}$ and $\mathtt{pp}$ are both chosen by $\mathcal{CO}$, just as $\mathcal{A}$ does. Then when O returns a random number, $\mathcal{A}$ perfectly simulates $\mathsf{ImG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ for $\mathcal{CO}$. For k is a random key of the pseudo-random function $\mathcal{F}$, when O returns a pseudo-random value, $\mathcal{A}$ simulates $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ for $\mathcal{CO}$ well. In this case, $\left|\Pr\left[\mathcal{A}^{\mathcal{F}(k,\cdot)}(1^\lambda) = 1\right] - \Pr\left[\mathcal{A}^{f(\cdot)}(1^\lambda) = 1\right]\right| = 1/\mathsf{poly}(\lambda)$, which contradicts the assumption of the pseudo-random function. Hence this lemma holds.

We now turn to $\mathsf{ImG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)$ and $\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}(\lambda)}$. Since the outputs of f are uniform random strings, it is obvious that these two games share the same distribution. Thus, $\left|\Pr\left[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda) = 1\right] - \Pr\left[\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}(\lambda) = 1\right]\right| \leqslant \mathsf{negl}(\lambda)$.

Theorem B2. Given a pseudo-random function $\mathcal{F} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with $\mathcal{AMC}_{\mathcal{P}} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is robust.

Proof. In order to complete the proof, we first introduce an intermediate game $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$ for reduction. This game is similar to $\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)$, except that $\mathtt{t}_1\|\mathtt{t}_2$ is generated by a true random function f. Without loss of generality, we set the binary length of $\mathtt{t}_2$ to m. Since $\mathtt{ctr}$ does not repeat, according to Lemma B1, there is no PPT adversary $\mathcal{A}$ that can distinguish between $\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$, that is $\left|\Pr\left[\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right] - \Pr\left[\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right]\right| \leqslant \mathsf{negl}(\lambda)$.

Now we consider $\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$. For $\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda)$ will always output $\perp$, regardless of inputs, $\mathcal{A}$ can distinguish these two games if and only if the output of $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$ is not $\perp$. This occurs only when $\mathtt{x} \oplus \mathtt{t}_1 = \mathtt{amsg}\|\mathtt{t}'_2$ and $\mathtt{t}'_2 = \mathtt{t}_2$. The probability of this occurring is $1/2^m$. For $\mathcal{A}$ can make at most q queries, the advantage of winning the game $\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$ is $\left|\Pr\left[\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right] - \Pr\left[\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right]\right| = q/2^m$, a negligible probability. In this case, we have $\left|\Pr\left[\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right] - \Pr\left[\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda) = 1\right]\right| \leqslant \mathsf{negl}(\lambda)$.

Appendix C Anamorphism and robustness of $\mathcal{AMC}_{\mathcal{A}}$

Theorem C1. Given a pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_{\mathcal{A}} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is an anti-*Corrupter* anamorphic one.

Proof. To complete the proof, we suppose that there exists a corrupt $\mathcal{CO}$ that can break anti-*Corrupter* anamorphic of $\mathcal{AMC}_{\mathcal{A}}$. And we set the advantage of $\mathcal{CO}$ as a non-negligible probability, that is, $\left|\Pr\left[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)=1\right]-\Pr\left[\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}(\lambda)=1\right]\right| = 1/\mathsf{poly}(\lambda)$.

According to Definition 6, *Corrupter* $\mathcal{CO}$ can access his oracle $O = \mathsf{Oa}$ that returns an anamorphic commitment embedding a covert message $\mathtt{amsg}$, or $O = \mathsf{Oc}$ that returns a regular one. We then construct an adversary $\mathcal{A}$ that can break the pseudo-randomness assumption of $\mathcal{PSE}$ as follows: (1) $\mathcal{A}$ generates $k \leftarrow_{\$} \mathcal{K}$, where $\mathcal{CO}$ can issue $(\mathtt{msg}, \mathtt{amsg}, \mathtt{pp})$ to his oracle; (2) for each issue $(\mathtt{msg}, \mathtt{amsg}, \mathtt{pp})$ from $\mathcal{CO}$, $\mathcal{A}$ simulates the response of $\mathcal{CO}$'s oracle by joining $\mathtt{amsg}$ and k and feeding $\mathtt{amsg}\|k$ to $\mathcal{A}$'s oracle, and after receiving $\mathtt{x}$ from $\mathcal{A}$'s oracle, $\mathcal{A}$ generates $\mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$ and uses $(\mathtt{c}, \mathtt{x})$ to answer $\mathcal{CO}$'s issues; (3) $\mathcal{A}$ returns $\mathcal{CO}$'s output b as his own judgement.

We now analyze the view of the $\mathcal{CO}$ when $\mathcal{A}$ accesses with different types of the oracle, Oe or Or. If $\mathcal{A}$ plays $\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}$, then $\mathcal{A}$ will receive a pseudo-random ciphertext of $\mathtt{amsg}\|k$. Moreover, $\mathtt{pp}$ and $\mathtt{msg}$ are the same as $\mathcal{CO}$ own choice, $\mathtt{c} = \mathsf{Com}(\mathtt{pp}, \mathtt{msg}, \mathtt{x})$ the same in $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}$. In this case, $\mathcal{CO}$'s view that $\mathcal{A}$ provided is the same as $\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}$. Therefore, we have $\Pr\left[\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda)=1\right] = \Pr\left[\mathsf{AnaG}_{\mathcal{AMC},\mathcal{CO}}(\lambda)=1\right]$; and if $\mathcal{A}$ plays $\mathsf{RandG}_{\mathcal{R},\mathcal{A}}$, $\mathcal{A}$ will receive a random string. $\mathtt{c} := \mathcal{SC}.\mathsf{Com}(\mathtt{pp}\ , \mathtt{msg}, \mathtt{x})$ generated by $\mathcal{A}$ is a regular commitment, identical to the output of $\mathcal{CO}$'s oracle Oc. In this case, $\mathcal{A}$'s reply is a complete simulation of $\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}$ for $\mathcal{CO}$, hence $\Pr\left[\mathsf{RandG}_{\mathcal{R},\mathcal{A}}(\lambda)=1\right] = \Pr\left[\mathsf{RealG}_{\mathcal{SC},\mathcal{CO}}(\lambda)=1\right]$.

In this case, the advantage of $\mathcal{A}$ in distinguishing between $\mathsf{RandG}_{\mathcal{R},\mathcal{A}}$ and $\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}$ is $\left|\Pr\left[\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda)=1\right]-\Pr\left[\mathsf{RandG}_{\mathcal{R},\mathcal{A}}(\lambda)=1\right]\right| = 1/\mathsf{poly}(\lambda)$, the same as the probability that $\mathcal{CO}$ breaks the corruptor-anamorphic $\mathcal{AMC}_{\mathcal{A}}$, which contradicts the pseudo-randomness assumption of $\mathcal{PSE}$. Thus, this theorem holds.

Theorem C2. Given a pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$, a standard commitment $\mathcal{SC} = (\mathsf{Gen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Ver})$ with anamorphic triplet $\mathcal{AMC}_{\mathcal{A}} = (\mathsf{aKG}, \mathsf{aCom}, \mathsf{aDec})$ is robust.

Proof. To prove this theorem, we first show that for a $\mathcal{PSE}$ with uniform decryption, the following lemma holds.

Lemma C1. If an IND-CPA pseudo-random symmetric encryption $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ with message space $\mathcal{M}$ and ciphertext space $\mathcal{C}$ has uniform decryption property, for any PPT adversary $\mathcal{A}$, it holds that $|\Pr[\mathcal{A}(c, \mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, c)) = 1] - \Pr[\mathcal{A}(\mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{msg}), \mathtt{msg}) = 1]| \leqslant \mathsf{negl}(\lambda)$, where $c \leftarrow_{\$} \mathcal{C}, \mathtt{sk} \leftarrow_{\$} \mathcal{PSE}.\mathsf{KG}(1^{\lambda}), \mathtt{msg} \leftarrow_{\$} \mathcal{M}$.

Proof. Let $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ a pseudo-random symmetric encryption has uniform decryption. To complete the proof, for any PPT adversary $\mathcal{A}$, we first introduce an intermediate state $\mathcal{A}(c, \mathtt{msg})$ and assume that there exists an adversary $\mathcal{A}$ that $|\Pr[\mathcal{A}(c, \mathtt{msg}) = 1] - \Pr[\mathcal{A}(\mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{msg}), \mathtt{msg}) = 1]| = 1/\mathsf{poly}(\lambda)$, and there must exist an adversary $\mathcal{B}$ that can distinguish PRandG and RandG in the same probability.

According to Definition 9, $\mathcal{B}$ can access its oracle O, where either $O = \mathsf{Oe}$ that returns a ciphertext, or $O = \mathsf{Or}$ that returns a random string. Then, another adversary $\mathcal{B}$ can leverages $\mathcal{A}$ to distinguish PRandG and RandG as follows: (1) $\mathcal{B}$ simulates $\mathcal{A}$'s receiving by randomly choosing $\mathtt{msg} \leftarrow_{\$} \mathcal{M}$ and feeding $\mathtt{msg}$ to the oracle O; (2) after receiving $\mathtt{x}$ from O, $\mathcal{B}$ sends $(\mathtt{x}, \mathtt{msg})$ to $\mathcal{A}$; (3) when $\mathcal{A}$ returns his result $b \in \{0, 1\}$, $\mathcal{B}$ returns the same b as the result.

We next consider the view of $\mathcal{A}$. If $O = \mathsf{Oe}$, $\mathtt{x} := \mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{msg})$, the distribution of $(\mathtt{x}, \mathtt{msg})$ is the same as that of $(\mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{msg}), \mathtt{msg})$, and if $O = \mathsf{Or}$ and $\mathtt{x}$ is a random string, the distribution of $(\mathtt{x}, \mathtt{msg})$ is the same as that of $(c, \mathtt{msg})$. Thus $\mathcal{B}$ perfectly simulates $\mathcal{A}$'s view. As a result, $\left|\Pr\left[\mathsf{PRandG}_{\mathcal{PSE},\mathcal{A}}(\lambda)=1\right]-\Pr\left[\mathsf{RandG}_{\mathcal{R},\mathcal{A}}(\lambda)=1\right]\right| = 1/\mathsf{poly}(\lambda)$.

In this way, $\mathcal{B}$ successfully distinguishes PRandG and RandG in a non-negligible probability by leveraging $\mathcal{A}$. However, this contradicts Definition 9. Thus, we have $|\Pr[\mathcal{A}(c, \mathtt{msg}) = 1] - \Pr[\mathcal{A}(\mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{msg}), \mathtt{msg}) = 1]| \leqslant \mathsf{negl}(\lambda)$. And according to Theorem 6, we have $|\Pr[\mathcal{A}(c, \mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, c)) = 1] - \Pr[\mathcal{A}(c, \mathtt{msg}) = 1]| \leqslant \mathsf{negl}(\lambda)$. By the triangle inequality, we have $|\Pr[\mathcal{A}(c, \mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, c)) = 1] - \Pr[\mathcal{A}(\mathcal{PSE}.\mathsf{Enc}(\mathtt{sk}, \mathtt{msg}), \mathtt{msg}) = 1]| \leqslant \mathsf{negl}(\lambda)$. Therefore, this lemma holds.

It is similar to the proof of $\mathcal{AMC}_{\mathcal{P}}$'s robustness that we also introduce an intermediate game $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$. We set $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$ the same as $\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)$ except that the random blind factor $\mathtt{x}$ is replaced with $\mathtt{x} \leftarrow_{\$} \mathcal{PSE}.\mathsf{Enc}(M, \mathtt{sk})$ where $M \leftarrow_{\$} \mathcal{M}$. For the robust adversary $\mathcal{A}$, if he plays $\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)$, his oracle O^0 will return a random value x on the ciphertext space and its decryption result $\mathcal{PSE}.\mathsf{Dec}(\mathtt{sk}, \mathtt{x})$, and if he plays $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$, his oracle will return the ciphertext of M, and M itself. In this case, $\mathcal{A}$'s perspective is the same as the adversary's in Lemma C1. Thus $\mathcal{A}$ is unable to distinguish between $\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$. Thus, we have the following lemma.

Lemma C2. Given a $\mathcal{PSE} = (\mathsf{KG}, \mathsf{Enc}, \mathsf{Dec})$ with uniform decryption, for any PPT robust game adversary $\mathcal{A}$, it holds that $\left|\Pr\left[\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)=1\right]-\Pr\left[\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)=1\right]\right| \leqslant \mathsf{negl}(\lambda)$.

We now turn to the advantage of $\mathcal{A}$ in distinguishing between $\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda)$ and $\mathsf{RobG}^2_{\mathcal{AMC},\mathcal{A}}(\lambda)$. Since $\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda)$ outputs $\perp$ for any input, the only case that these two games output differently is when and only when $k = k'$ that $\mathtt{amsg}\|k' = M$. Without loss of generality, we set the binary length of the robustness token k to be m and the binary length of M to be l. This would mean that for a single query, the advantage of $\mathcal{A}$ in distinguishing these two games is $2^{(l-m)}/2^l = 1/2^m$. Since $\mathcal{A}$ queries up to q times, the probability is $q/2^m$, a negligible probability. In summary, we have $\left|\Pr\left[\mathsf{RobG}^0_{\mathcal{AMC},\mathcal{A}}(\lambda)=1\right]-\Pr\left[\mathsf{RobG}^1_{\mathcal{AMC},\mathcal{A}}(\lambda)=1\right]\right| \leqslant \mathsf{negl}(\lambda)$. Hence, according to Definition 8, this theorem holds.