Anamorphic Proofs of Knowledge from Prover Randomness

Zhi Lu[1], Songfeng Lu[2], and Jiaheng Zhang[1]

[1] National University of Singapore, Singapore
zhilu@nus.edu.sg, jhzhang@nus.edu.sg
[2] Huazhong University of Science and Technology, Wuhan, China
lusongfeng@hust.edu.cn

Abstract. We introduce anamorphic proofs of knowledge. Each proof publicly certifies its normal statement and certifies a separately registered knowledge claim to a designated checker. Our construction uses a commitment that the normal prover already makes to its randomness. The prover masks a witness for the registered statement and supplies the result as those random coins. The public statement, proof format and length, and normal verifier are unchanged. An external observer given both witnesses and the actual normal prover tape cannot distinguish the two modes. We identify two properties of the normal proof system that support the construction. Extracting an opening of the initial randomness commitment gives anamorphic knowledge soundness. Simulating the complete proof at a prescribed commitment gives zero knowledge under the actual checking key. Non-encodability also prevents the checker from forging a proof at an unused label for an independently registered target, even after obtaining honest proofs at other labels. We instantiate the construction in random-base Compressed Σ and aggregated Bulletproofs in the random-oracle model. For three-move protocols, we prove an exclusion for oracle-free local checks when the normal witness permits challenge transfer and finding two accepting responses to the same challenge and first message is hard. The exclusion covers Schnorr and uncompressed random-base representation protocols.

Keywords: Anamorphism · Proofs of Knowledge · Compressed Σ · Bulletproofs · Fiat–Shamir

1 Introduction

Anamorphic encryption, introduced by Persiano, Phan, and Yung [37], studies covert communication through established encryption schemes against a dictator who can obtain the normal decryption key or dictate the message to be encrypted. A ciphertext retains its normal meaning while carrying an anamorphic message recoverable by a designated receiver. Anamorphic signatures extend this idea to settings in which authentication is permitted but encryption is prohibited, carrying an anamorphic message through a valid normal signature [34]. Extending this idea to proofs of knowledge requires a prescribed public proof

to support an additional knowledge claim, with knowledge extraction and zero knowledge for that claim. Covert message transmission alone does not directly provide these guarantees.

Consider a participant submitting a range proof for a committed value. A dictator specifies the proof system and can demand the normal witness and the actual prover tape, then reproduce the whole proof. Before the audited session, the participant privately registers a second commitment with a designated checker. During the session the checker sees the same range proof as the dictator. Can the checker use that proof to verify a knowledge claim about the registered commitment while learning neither opening?

We introduce **anamorphic proofs of knowledge** (APoKs). A single proof publicly certifies its normal statement and also certifies knowledge of a witness for a separately registered anamorphic statement to a designated checker. Private auxiliary setup gives the prover an embedding key and the checker a checking key. The normal public parameters, statement, proof format, and verification algorithm remain unchanged. An external observer holding neither auxiliary key cannot distinguish the two modes even when given both witnesses and the actual random tape supplied to the normal prover.

Our construction uses a commitment that the normal prover already makes to its own randomness. We call this commitment the **carrier**. Masking the anamorphic witness lets the prover set the carrier to the value expected by the checker using the normal proving algorithm. Two properties connect this sampling rule to a knowledge claim. An extractor must recover an opening of the carrier in the prover's initial execution. A simulator must reproduce the complete proof at the carrier fixed by the checking key and registered statement. We establish these properties for Compressed Σ [4,3] and Bulletproofs [13].

1.1 Our Results

Section 2 formalizes four security requirements for APoK.

- **External covertness.** The normal and anamorphic modes remain indistinguishable to an external observer even when both witnesses and the normal prover's actual random tape are disclosed.
- **Strong anamorphic zero knowledge.** Under the actual checking key, the checker's view across adaptive valid requests can be simulated without either witness or the embedding key.
- **Anamorphic knowledge soundness.** Even given both auxiliary keys, a successful prover admits extraction of a witness for the anamorphic statement in its initial execution.
- **Non-encodability after honest proofs.** A checker given both witnesses and honest proofs at other labels cannot produce an accepting anamorphic proof at an unused label together with a valid witness for its chosen target, except with negligible probability.

Non-encodability also protects a target registered independently of the checking key. For such a target, producing an accepting proof at an unused label already constitutes a forgery, even if the checker supplies no witness. Corollary 1 gives a lossless reduction from this experiment to non-encodability.

1. **Generic construction.** For anamorphic preimage relations under a public homomorphism, we give an APoK construction that changes only the normal prover's randomness sampling (Theorem 1). Two properties of the normal proof system are required, namely carrier-opening extraction and conditional simulation of the full proof at a prescribed carrier. Under the stated hardness assumptions the construction satisfies all four security requirements. The non-encodability reduction guesses one of the labels fixed at auxiliary setup.
2. **Two instantiations.** We instantiate APoK with random-base Compressed Σ and Bulletproofs in the random-oracle model (Theorems 2 and 3). The main technical step in each case is to simulate the complete proof and random-oracle view at a prescribed randomness commitment. For Compressed Σ, a surjectivity condition gives equally sized fibers after the first compression. For Bulletproofs, a bijection with the conditioned prover tape lets the simulator sample the inner-product vectors and blinders with S fixed. Carrier-opening extraction uses fewer outer branches than the full range-witness tree analyzed here (Table 1) and extends to range aggregation (Corollary 2).
3. **A structural boundary.** In three-move protocols with witness-assisted challenge transfer, a checker can transfer a simulated response to the real challenge. For an oracle-free check of the full transcript, Theorem 4 bounds the attack in terms of the prover's query budget and the difficulty of finding two accepting responses for the same first message and challenge. The bound excludes Schnorr and uncompressed random-base representation protocols in our model. Appendix D.3 gives a four-coordinate Compressed Σ instance certifying discrete-log knowledge, allowing a direct comparison with Schnorr.

Both instantiations change only the normal prover's randomness sampling and add no proof bytes. Independent uniform masks give perfect external covertness and statistical strong anamorphic zero knowledge. Masks derived from a pseudorandom function (PRF) give computational guarantees.

The auxiliary setup allocates a polynomial number T of labels, each used for at most one proof. The checking key stores one group element per label, while the embedding key stores independent masks or a PRF seed. Beyond normal verification, anamorphic verification adds one group multiplication and one equality test. Appendix D.2 gives proof sizes and auxiliary overheads.

1.2 Technical Overview

We certify the additional knowledge claim through the normal prover's existing randomness commitment. Let $C = \Phi(r)$, where r is uniform in an additive group U and $\Phi : U \to \mathcal{G}$ is a public homomorphism into a multiplicative group $\mathcal{G}$. For

one-time use, the auxiliary setup gives a uniform mask η to the prover and gives $M = \Phi(\eta)$ to the checker. For a registered anamorphic statement $\widehat{C} = \Phi(\widehat{u})$, anamorphic verification combines normal verification with the check $C = \widehat{C} \cdot M$.

Rejection sampling, a feasibility example for anamorphic encryption in [37], would repeatedly generate normal proofs until their carrier meets this equation. A normal proof succeeds with probability $1/|\Phi(U)|$, requiring $|\Phi(U)| = q$ candidates in expectation in our instantiations. Lemma 1 formalizes the sampling bound. Our prover sets $r = \widehat{u} + \eta$ directly. At a fresh label, the independent uniform mask preserves the normal tape distribution even for inputs chosen adaptively from prior observer views.

Knowledge certification comes from extracting a carrier opening through the complete proof's verification constraints. The resulting extractor serves two reductions. Let C_0 be the carrier in the prover's initial execution, suppose that execution passes anamorphic verification, and let $\Phi(r^\star) = C_0$.

- The **anamorphic-knowledge reduction** knows the embedding mask η. Subtracting it gives an anamorphic witness because $\Phi(r^\star - \eta) = C_0 M^{-1} = \widehat{C}$.
- The **non-encodability reduction** uses the anamorphic witness $\widehat{u}$ supplied by the attacking checker. Subtracting it gives an opening of the mask image because $\Phi(r^\star - \widehat{u}) = C_0 \widehat{C}^{-1} = M$.

For multi-round Fiat–Shamir proofs, we use the $(\Gamma_1, \ldots, \Gamma_\mu)$-special-sound extraction framework of [5] and bind extraction to the carrier and claim of the initial execution. Section 3 develops both reductions.

Witness privacy poses a separate problem because the actual checking key fixes M and the registered anamorphic statement fixes $\widehat{C}$, so the simulator must reproduce a complete proof at the prescribed carrier $C = \widehat{C} \cdot M$. Normal zero knowledge does not impose this constraint. We simulate the complete proof and random-oracle view conditioned on $\Phi(r) = C$, with r uniform over the corresponding fiber, without either witness or the secret mask.

The two instantiations realize this conditional distribution differently. For Compressed Σ, a surjectivity condition gives equally sized fibers for the state after the first compression. The simulator can sample that state directly and then continue with the normal prover. The same accepting tree used to recover the normal witness also opens the first message A_0 (Section 4). For Bulletproofs, a bijection between the conditioned prover tape and the simulated inner-product vectors and blinding values yields the complete proof distribution with S fixed. We isolate from the original range-extraction proof [13] the step needed to recover an opening of S. Fixing the first two challenges and using two polynomial-evaluation branches yields a carrier-opening tree with fewer outer branches than the full range-witness tree analyzed in Section 5.

A homomorphic first message is insufficient for APoK. In three-move protocols such as Schnorr [40,23], the normal witness can transfer an accepting response to another challenge while leaving the first message unchanged. We turn a strong anamorphic zero-knowledge simulator into a checker that produces a proof passing anamorphic verification. The checker uses the normal witness to transfer

the simulator's response to the real challenge. For unique accepting responses the transferred response is the only one the check can see. For the general case, a forking argument shows that a check distinguishing the transferred response from the honest one yields two accepting responses for the same first message and challenge. Section 6 formalizes this transformation in the local-checking model and bounds its success probability, even when the honest anamorphic prover adaptively searches over its first message.

1.3 Related Work

Following Persiano, Phan, and Yung [37], anamorphic encryption has been studied with post-deployment auxiliary setup [7], homomorphic and asymmetric constructions [16], public-key anamorphism [38], fully asymmetric LWE constructions [27], and multi-message embedding [6]. The demand to disclose actual randomness appears in PPY's sender-side motivation and is incorporated into later security definitions [42]. Catalano et al. [17,18] study limitations of black-box constructions and generic rejection sampling, and positive results under weaker notions or assumptions. Choi et al. [21] compare definitions and analyze dictator–sender collusion. APoK concerns an extractable knowledge claim with witness privacy against the checker. Deniable encryption allows a ciphertext to be explained as encrypting another message with different randomness [15].

Receiver unforgeability in anamorphic signatures parallels our separation of checking from generation. Kutyłowski et al. [34] use pseudorandom symmetric ciphertexts as signature randomness. Jaeger and Stracovsky [32] analyze double-key holders, with the revised analysis covering attacks by an active dictator. Deo and Libert [26] retain receiver unforgeability beyond lengths supported by randomness replacement. APoK requires non-encodability at an unused label even when the checker knows both witnesses and obtains honest proofs elsewhere. Fully asymmetric anamorphic encryption separates the keys in the other direction, allowing embedding without decoding.

Dao and Grubbs [24] establish simulation extractability for Spartan and Bulletproofs in the random-oracle model, including aggregated range proofs. Abe et al. [2] study simulation given advance knowledge of a critical challenge, with applications to partial-knowledge proofs and trapdoor commitments. Our simulator receives a prescribed randomness commitment, while our extractor recovers its opening to support a separately registered anamorphic statement. Section 3 develops these two interfaces.

Designated-verifier proofs limit transferability by allowing the verifier to simulate proofs [33]. Chaidos and Couteau [19] construct designated-verifier non-interactive zero-knowledge proofs of knowledge for algebraic languages, sound under arbitrary verification feedback. APoK retains normal public verification for the normal statement while requiring non-encodability against the checker. Its zero-knowledge simulator uses programmability of the random oracle, unavailable to the checker. Witness authentication NIZKs [29] allow a candidate witness or identifier to determine whether a proof was produced using that wit-

ness, with privacy expressed through entropic zero knowledge. APoK checking concerns a separately registered knowledge claim.

Subliminal channels and kleptography use randomness for covert communication [41,45]. Biryukov, Feher, and Vitto [12] study channels in Zcash zk-SNARKs and Pedersen commitments, while anamorphic commitments consider coercion and parameter subversion [44]. Algorithm-substitution attacks use subverted encryption outputs to recover plaintexts [10]. Rewindable Bulletproofs in the Grin fork of secp256k1-zkp derive proof blinders from a nonce and allow its holder to recover committed values and embedded messages from single-commitment proofs [43]. APoK also uses prover randomness, but adds knowledge extraction, witness privacy, and non-encodability against the checker. Sampling bounds for black-box steganography [25] relate to our analysis of screening normal proof candidates.

Subliminal-free zero knowledge [14], steganography-free zero knowledge [1], and anamorphism-resistant encryption [28] study how to suppress covert functionality. Quach, Tyner, and Wichs [39] obtain computationally unique NIZK proofs for statements with computationally unique witnesses once the prover setup is fixed. In the protocols studied here, the normal prover's freedom to sample fresh coins supports an additional knowledge claim.

2 Anamorphic Proofs of Knowledge

2.1 Syntax

We work in the random-oracle model (ROM) [11,30] with security parameter λ. Unless stated otherwise, algorithms are probabilistic polynomial time (PPT), and $\mathrm{negl}(\lambda)$ denotes a negligible function. Knowledge soundness is required against PPT provers, so the proof systems here include arguments of knowledge.

We use normal and anamorphic to distinguish both functionalities, following PPY's mode terminology [37]. Fix a normal non-interactive proof of knowledge [8] $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$, with public parameters $pp \leftarrow \mathsf{Setup}(1^\lambda)$ and normal relation $\mathcal{R} = \mathcal{R}_{pp}$. The additional relation $\widehat{\mathcal{R}} = \widehat{\mathcal{R}}_{pp}$ is the anamorphic relation. Both are polynomial-time decidable. The publicly verified normal statement x has normal witness w with $(x, w) \in \mathcal{R}$. The registered anamorphic statement $\widehat{x}$ has anamorphic witness $\widehat{w}$ with $(\widehat{x}, \widehat{w}) \in \widehat{\mathcal{R}}$. These names identify the statements' roles. Security experiments specify which statements and witnesses are disclosed.

The normal prover $\mathsf{P}^{\mathbf{H}}(pp, x, w; \omega)$ outputs a proof π. In the normal mode, its tape is sampled from $\Omega_\Pi(pp, x, w)$. The normal verifier $\mathsf{V}^{\mathbf{H}}(pp, x, \pi)$ is public and returns 1 for acceptance and 0 for rejection. Superscript $\mathbf{H}$ denotes access to the independent random oracles $\mathbf{H} = (H_1, \ldots, H_\nu)$, with query domains $\mathcal{M}_j$ and response spaces $\mathcal{Y}_j$ for $H_j : \mathcal{M}_j \to \mathcal{Y}_j$. Simulators and extractors receive no setup trapdoor. Parameter hardness is measured over the output of Setup, and we omit pp from auxiliary interfaces when clear.

For $I = (x, w, \widehat{x}, \widehat{w})$, define $\mathsf{Valid}(I) = [(x, w) \in \mathcal{R} \wedge (\widehat{x}, \widehat{w}) \in \widehat{\mathcal{R}}]$, where brackets denote the indicator of the enclosed condition. Before auxiliary setup, fix distinct labels $\mathcal{C} = (ctx_1, \ldots, ctx_T)$ independently of $\mathbf{H}$, for polynomial $T = T(\lambda)$. Index $i \in [T] = \{1, \ldots, T\}$ identifies ctx_i.

Definition 1 (Anamorphic mode). *An anamorphic mode for Π and $\widehat{\mathcal{R}}$ consists of three algorithms for auxiliary key generation, anamorphic random-tape generation, and anamorphic verification.*

- $\mathsf{AKGen}(pp, \mathcal{C})$ *generates an embedding key ek for the prover and a checking key ck for the designated checker. It has no access to* $\mathbf{H}$.
- $\mathsf{ACoins}^{\mathbf{H}}(ek, i, I)$ *takes the embedding key, an index, and a valid input, and returns a random tape ω in the support of $\Omega_{\Pi}(pp, x, w)$.*
- $\mathsf{AVerify}^{\mathbf{H}}(ck, i, x, \widehat{x}, \pi)$ *takes the checking key, index, both statements, and proof and returns an acceptance bit. It is stateless and deterministic given* $\mathbf{H}$.

2.2 Modes of Operation and Correctness

Auxiliary setup is an honest execution of AKGen with private delivery of ek to the prover and ck to the designated checker.

In the normal mode, sample $\omega \leftarrow \Omega_{\Pi}(pp, x, w)$ and return $\mathsf{P}^{\mathbf{H}}(pp, x, w; \omega)$. In the anamorphic mode, the derived prover $\mathsf{AProve}^{\mathbf{H}}(ek, i, I)$ samples $\omega \leftarrow \mathsf{ACoins}^{\mathbf{H}}(ek, i, I)$ and returns $\mathsf{P}^{\mathbf{H}}(pp, x, w; \omega)$. Thus both modes execute P and retain V for normal verification.

The mode preserves parameter generation, the normal statement, messages, proof format and length, Fiat–Shamir inputs, challenge spaces, and normal verification. Auxiliary keys, labels, and anamorphic statements use an auxiliary channel and are not added to the proof or its hash inputs. For honestly generated auxiliary keys, every valid input I and fresh index i must satisfy

$$\mathsf{AVerify}^{\mathbf{H}}\big(ck, i, x, \widehat{x}, \mathsf{AProve}^{\mathbf{H}}(ek, i, I)\big) = 1$$

with probability 1. For every candidate proof, we also require

$$\mathsf{V}^{\mathbf{H}}(pp, x, \pi) = 0 \implies \mathsf{AVerify}^{\mathbf{H}}(ck, i, x, \widehat{x}, \pi) = 0.$$

An index is fresh until used for an honest proof. The prover and checker associate $(i, \widehat{x})$ at registration, before the audited session. A long-term credential may use several distinct indices. The application enforces this association and one-time use. Anamorphic verification is stateless and accepts a replay of the same valid proof with the same inputs.

2.3 Security

Experiment probabilities average over $pp \leftarrow \mathsf{Setup}(1^{\lambda})$, the random oracles, auxiliary setup, and algorithm randomness, including when pp is written as an input. All parties have query access to $\mathbf{H}$. Permissions are specified below.

Each proof-request experiment initializes an empty set Used of consumed indices. A request (i, I) with $i \notin [T] \setminus \mathsf{Used}$ or $\mathsf{Valid}(I) = 0$ returns $\perp$ without changing the set. Otherwise, it adds i to Used and computes the prescribed response. This rule applies to real and simulated responses. The honest anamorphic proving interface under this rule is denoted by $\mathsf{AProve}^{\mathbf{H}}_{\mathsf{Used}}$. Malformed adversarial outputs or out-of-range indices score 0.

External Covertness. The distinguisher $\mathcal{D}$ receives both witnesses and the actual prover tape ω, but neither auxiliary key. This tape excludes auxiliary setup randomness and internal randomness of ACoins. The experiment allows adaptive valid-input selection and compares normal ($b = 0$) and anamorphic ($b = 1$) tapes.

$\underline{\mathsf{Exp}^{b}_{\mathcal{D}}(pp)}$

1: Send pp to $\mathcal{D}$.
2: $(ek, ck) \leftarrow \mathsf{AKGen}(pp, \mathcal{C})$. Give neither auxiliary key nor the auxiliary setup randomness to $\mathcal{D}$.
3: On each adaptive request (i, I) from $\mathcal{D}$, apply the common request rule. If $b = 0$, sample $\omega \leftarrow \Omega_{\Pi}(pp, x, w)$. If $b = 1$, sample $\omega \leftarrow \mathsf{ACoins}^{\mathbf{H}}(ek, i, I)$. Set $\pi \leftarrow \mathsf{P}^{\mathbf{H}}(pp, x, w; \omega)$ and return (π, ω).
4: Return $\mathcal{D}$'s output bit.

Write $\mathsf{Exp}^{\mathrm{normal}} = \mathsf{Exp}^{0}$ and $\mathsf{Exp}^{\mathrm{ana}} = \mathsf{Exp}^{1}$ for the normal and anamorphic experiments, respectively. If their probabilities of outputting 1 are $p^{\mathrm{normal}}_{\mathcal{D}}(\lambda)$ and $p^{\mathrm{ana}}_{\mathcal{D}}(\lambda)$, define $\mathrm{Adv}^{\mathrm{cov}}(\mathcal{D}) = |p^{\mathrm{normal}}_{\mathcal{D}} - p^{\mathrm{ana}}_{\mathcal{D}}|$.

Definition 2 (External covertness). *External covertness holds if every PPT distinguisher $\mathcal{D}$ has negligible advantage $\mathrm{Adv}^{\mathrm{cov}}(\mathcal{D})$. It is perfect if the complete observer views in both experiments have identical distributions.*

Strong Anamorphic Zero Knowledge. The environment $\mathcal{E}$ holds the actual checking key, adaptively chooses valid requests, and may retain both witnesses. It receives proofs but no tapes. The stateful simulator receives only the inputs displayed in the following experiment, excluding both witnesses, ek, and auxiliary setup randomness.

The simulator has programming access to $\mathbf{H}$. It may test whether $m \in \mathcal{M}_j$ is defined in H_j, without assigning a value. Only at an undefined input may it set $H_j(m) \leftarrow y$ for $y \in \mathcal{Y}_j$. All parties use the same updated tables.

$\underline{\mathsf{Exp}^{b}_{\mathcal{E},\mathsf{Sim}}(pp)}$

1: $(ek, ck) \leftarrow \mathsf{AKGen}(pp, \mathcal{C})$. Send (pp, ck) to $\mathcal{E}$ and initialize Sim with empty state.
2: On each adaptive request (i, I) from $\mathcal{E}$, apply the common request rule. In the real experiment $b = 0$, return $\mathsf{AProve}^{\mathbf{H}}(ek, i, I)$. In the simulated experiment $b = 1$, return $\mathsf{Sim}^{\mathbf{H}}(pp, ck, i, x, \widehat{x})$.
3: Return $\mathcal{E}$'s output bit.

Let $\mathsf{View}_{\mathrm{real}}$ and $\mathsf{View}_{\mathrm{sim}}$ be the complete environment views, including inputs, randomness, proof responses, and oracle queries and answers. Define

$$\mathrm{Adv}^{\mathrm{azk}}(\mathcal{E}, \mathsf{Sim}) := \left|\Pr[\mathsf{Exp}^{0}_{\mathcal{E},\mathsf{Sim}}(pp) = 1] - \Pr[\mathsf{Exp}^{1}_{\mathcal{E},\mathsf{Sim}}(pp) = 1]\right|.$$

Definition 3 (Strong anamorphic zero knowledge). *The mode has strong anamorphic zero knowledge if there exists a stateful PPT simulator with the programming access above such that* $\mathrm{Adv}^{\mathrm{azk}}(\mathcal{E}, \mathsf{Sim})$ *is negligible for every PPT environment* $\mathcal{E}$*.*

The mode has δ-statistical strong anamorphic zero knowledge if a stateful PPT simulator with the same interface and permissions achieves $\Delta(\mathsf{View}_{\mathrm{real}}, \mathsf{View}_{\mathrm{sim}}) \leq \delta(\lambda, N, Q)$ for every environment making at most N proof requests and Q direct oracle queries, including computationally unbounded environments. Here Δ is total variation distance.

Anamorphic Knowledge Soundness. A malicious prover $\mathcal{A}$ receives both auxiliary keys and chooses both statements and the proof. Extraction targets the anamorphic statement in its initial execution, even if later replays declare another one.

The extractor $\mathcal{X}$ has fixed-randomness replay access [5, Remark 2 and Section 5.2]. For a bound $\ell_{\mathcal{A}}$ on the prover's random bits, the experiment samples $\rho \xleftarrow{\$} \{0,1\}^{\ell_{\mathcal{A}}}$ and records the initial output o_0 and ordered query-and-answer trace τ_0. Besides this record, the extractor receives the restart oracle $\mathcal{A}'_\rho(\mathbf{F}) := \mathcal{A}^{\mathbf{F}}(pp, ek, ck; \rho)$. Each replay starts from the initial state with the same inputs and unknown tape ρ. The extractor chooses oracle procedures $\mathbf{F} = (F_1, \ldots, F_\nu)$ with the interfaces of $\mathbf{H}$, answers consistently within each replay, and may abort at any query. It neither learns ρ nor modifies the real oracles $\mathbf{H}$. The extraction experiment is as follows.

$\underline{\mathsf{Exp}^{\mathrm{Ext}}_{\mathcal{A},\mathcal{X}}(pp)}$

1: $(ek, ck) \leftarrow \mathsf{AKGen}(pp, \mathcal{C})$ and $\rho \xleftarrow{\$} \{0,1\}^{\ell_{\mathcal{A}}}$.
2: Run $o_0 = (i_0, x_0, \widehat{x}_0, \pi_0) \leftarrow \mathcal{A}^{\mathbf{H}}(pp, ek, ck; \rho)$, recording the trace τ_0.
3: If $\mathsf{AVerify}^{\mathbf{H}}(ck, i_0, x_0, \widehat{x}_0, \pi_0) = 0$, return 0.
4: $\widehat{w}^* \leftarrow \mathcal{X}^{\mathbf{H}, \mathcal{A}'_\rho}(pp, ek, ck, o_0, \tau_0)$.
5: Return $[(\widehat{x}_0, \widehat{w}^*) \in \widehat{\mathcal{R}}]$.

Let $\varepsilon_{\mathrm{ana}}(\mathcal{A})$ be the probability of passing step 3. The bound Q includes queries in the initial execution and its anamorphic verification. The extractor's expected running time is polynomial in λ and the adversary's running-time and query bounds, counting all replays and implemented oracle procedures.

Definition 4 (Anamorphic knowledge soundness). *The mode has anamorphic knowledge error* $\kappa : \mathbb{N}^2 \to [0,1]$ *if there exists a black-box extractor* $\mathcal{X}$ *with the replay access and running-time bound above such that, for every PPT adversary* $\mathcal{A}$*,* $\Pr[\mathsf{Exp}^{\mathrm{Ext}}_{\mathcal{A},\mathcal{X}}(pp) = 1] \geq \varepsilon_{\mathrm{ana}}(\mathcal{A}) - \kappa(\lambda, Q) - \mathrm{negl}(\lambda)$.

Non-Encodability after Honest Proofs. A checker may know both witnesses and obtain honestly generated anamorphic proofs at other indices. The following experiment tests generation at an unused index. The final output must supply an anamorphic witness $\widehat{w}_{\text{out}}$ for $\widehat{x}^*$, but need not supply a normal witness for x^*.

$\underline{\mathsf{Exp}^{\text{NonEnc}}_{\mathcal{A}}(pp)}$

1: $(ek, ck) \leftarrow \mathsf{AKGen}(pp, \mathcal{C})$. Send (pp, ck) to $\mathcal{A}$.
2: On each adaptive request (i, I) from $\mathcal{A}$, return $\mathsf{AProve}^{\mathbf{H}}_{\mathsf{Used}}(ek, i, I)$.
3: $\mathcal{A}$ outputs $(i, x^*, \widehat{x}^*, \widehat{w}_{\text{out}}, \pi^*)$. If $i \notin [T] \setminus \mathsf{Used}$ or $(\widehat{x}^*, \widehat{w}_{\text{out}}) \notin \widehat{\mathcal{R}}$, return 0. Otherwise, return $\mathsf{AVerify}^{\mathbf{H}}(ck, i, x^*, \widehat{x}^*, \pi^*)$.

Define $\text{Adv}^{\text{NonEnc-HP}}(\mathcal{A}) = \Pr[\mathsf{Exp}^{\text{NonEnc}}_{\mathcal{A}}(pp) = 1]$, where HP denotes honest-proof access. Write $\text{Adv}^{\text{NonEnc}}(\mathcal{A})$ for the same probability with the honest-proof query stage removed.

Definition 5 (Non-encodability after honest proofs). *The mode is non-encodable after honest proofs if* $\text{Adv}^{\text{NonEnc-HP}}(\mathcal{A})$ *is negligible for every PPT adversary* $\mathcal{A}$*. It is non-encodable if the same holds after removing the honest-proof query stage.*

Definition 6 (Secure APoK). *An anamorphic mode satisfying the syntax of Definition 1, correctness, and protocol preservation is a secure APoK if it satisfies external covertness, strong anamorphic zero knowledge, anamorphic knowledge soundness with error* $\kappa(\lambda, Q)$ *negligible for every polynomial* Q*, and non-encodability after honest proofs.*

A registered target. In this experiment an oracle-free PPT sampler $\mathcal{T}(pp)$ draws a valid pair $(\widehat{x}^\dagger, \widehat{w}^\dagger)$ using coins independent of auxiliary setup. The checker receives $(pp, ck, \widehat{x}^\dagger)$. It may make ordinary honest-proof requests or request $(i, x, w, \dagger)$, which the experiment completes with $(\widehat{x}^\dagger, \widehat{w}^\dagger)$ under the same request rule and shared Used set. Its output (i, x^*, π^*) wins if $i \in [T] \setminus \mathsf{Used}$ and $\mathsf{AVerify}^{\mathbf{H}}(ck, i, x^*, \widehat{x}^\dagger, \pi^*) = 1$. Write $\text{Adv}^{\text{TU-HP}}(\mathcal{A}, \mathcal{T})$ for its success probability.

Corollary 1 (Unforgeability for an independently registered target). *For every checker* $\mathcal{A}$ *and sampler* $\mathcal{T}$ *in this experiment, a non-encodability adversary* $\mathcal{B}$ *attains* $\text{Adv}^{\text{NonEnc-HP}}(\mathcal{B}) = \text{Adv}^{\text{TU-HP}}(\mathcal{A}, \mathcal{T})$.

Proof. On receiving (pp, ck), $\mathcal{B}$ runs $\mathcal{T}(pp)$, retains $\widehat{w}^\dagger$, and gives $\mathcal{A}$ its prescribed input. It completes each $\dagger$-request with the sampled pair and forwards all requests to its honest-proof oracle. On output (i, x^*, π^*), it returns $(i, x^*, \widehat{x}^\dagger, \widehat{w}^\dagger, \pi^*)$. Independence gives the same joint input distribution, and the two experiments use identical request and acceptance rules. □

The experiments distinguish a tape observer from a checker. The observer receives the prover's tape and both witnesses. The checker receives the checking key and proof. Only the zero-knowledge simulator may program the shared random oracles.

3 A Construction from Homomorphic Commitments

3.1 The Carrier and the Construction

Fix the normal proof system Π of Section 2, with perfect completeness. Its tape decomposes as $\omega = (r, \omega_{\mathrm{rest}})$, where r is uniform in a finite additive abelian group U and $\omega_{\mathrm{rest}} \leftarrow \Omega_{\mathrm{rest}}(pp, x, w)$ is independent. The public parameters specify a multiplicative group $\mathcal{G}$ and a homomorphism $\Phi : U \to \mathcal{G}$, with efficient evaluation, group operations, membership tests, and uniform sampling from U.

A deterministic polynomial-time parser $\mathsf{Car}(pp, x, \pi)$ returns a designated commitment, the carrier, from a well-formed proof, and $\perp$ otherwise. Every honest execution satisfies $\mathsf{Car}\big(pp, x, \mathsf{P}^{\mathbf{H}}(pp, x, w; (r, \omega_{\mathrm{rest}}))\big) = \Phi(r)$. This structure is fixed by the normal protocol before auxiliary setup. Write $\Phi(U) = \{\Phi(u) : u \in U\}$. The anamorphic statement $\widehat{C}$ and witness $\widehat{u}$ have relation $\widehat{\mathcal{R}}_{\Phi} = \left\{ (\widehat{C}, \widehat{u}) \in \mathcal{G} \times U \;\middle|\; \Phi(\widehat{u}) = \widehat{C} \right\}$.

Use the fixed labels $\mathcal{C} = (ctx_1, \ldots, ctx_T)$ of Section 2, each supporting at most one proof. The mask variant is selected before auxiliary setup. For the seed variant, $F : \mathcal{K} \times \mathsf{Ctx} \to U$ is a PRF independent of $\mathbf{H}$, with seed space $\mathcal{K}$ and label space Ctx. Let ϵ_F bound the advantage of replacing its mask table on these labels by independent uniform elements of U, including sampling-conversion loss. Appendix D.1 implements this derivation from a standard PRF [31] for $U = \mathbb{F}_q^D$.

The two modes differ only in the sampling of r:

Normal prover	Anamorphic prover
1. $r \xleftarrow{\$} U$	1. $r \leftarrow \widehat{u} + \eta_i$
2. $\omega_{\mathrm{rest}} \leftarrow \Omega_{\mathrm{rest}}(pp, x, w)$	2. $\omega_{\mathrm{rest}} \leftarrow \Omega_{\mathrm{rest}}(pp, x, w)$
3. Run $\mathsf{P}^{\mathbf{H}}(pp, x, w; (r, \omega_{\mathrm{rest}}))$	3. Run $\mathsf{P}^{\mathbf{H}}(pp, x, w; (r, \omega_{\mathrm{rest}}))$
$C = \Phi(r)$	$C = \Phi(\widehat{u} + \eta_i) = \widehat{C} M_i$

Construction 1 (APoK from a homomorphic commitment).
Common inputs are pp and $\mathcal{C}$. A valid proving input is $I = (x, w, \widehat{C}, \widehat{u})$.

1. $\mathsf{AKGen}(pp, \mathcal{C})$. Sample independent $\eta_i \xleftarrow{\$} U$ and set $ek = (\eta_i)_{i=1}^{T}$, or sample $k \xleftarrow{\$} \mathcal{K}$ and set $ek = k$, $\eta_i = F_k(ctx_i)$. Return ek and $ck = (M_i)_{i=1}^{T}$, where $M_i = \Phi(\eta_i)$.
2. $\mathsf{ACoins}(ek, i, I)$. Recover η_i from ek, sample $\omega_{\mathrm{rest}} \leftarrow \Omega_{\mathrm{rest}}(pp, x, w)$ independently, and return $(\widehat{u} + \eta_i, \omega_{\mathrm{rest}})$. This algorithm makes no oracle queries. The derived prover $\mathsf{AProve}^{\mathbf{H}}$ runs $\mathsf{P}^{\mathbf{H}}$ on this tape.
3. $\mathsf{AVerify}^{\mathbf{H}}(ck, i, x, \widehat{C}, \pi)$. Reject invalid indices, malformed inputs, or $\widehat{C} \notin \mathcal{G}$. If $\mathsf{V}^{\mathbf{H}}(pp, x, \pi) = 0$, return 0. Otherwise parse $C = \mathsf{Car}(pp, x, \pi)$ and return $[C = \widehat{C} M_i]$, rejecting if parsing fails.

Homomorphism gives $C = \Phi(\widehat{u} + \eta_i) = \widehat{C} M_i$ for every honest proof. The index and target use the session metadata of Section 2, with no additional proof fields or hash inputs. The checking key stores T group elements in both variants.

Beyond normal verification and carrier parsing, anamorphic checking costs one group multiplication and one equality test. Reusing a mask in this construction would force the carrier blocks of the two tapes to satisfy $r - r' = \widehat{u} - \widehat{u}'$. An observer with the tapes and witnesses can test it. Independently sampled normal blocks satisfy it with probability $1/|U|$. Appendix D.1 gives the embedding-key entropy requirement for perfect joint tape preservation.

The timing of target selection matters. A checker can choose an unused i, sample $r \xleftarrow{\$} U$, register $\widehat{C} = \Phi(r)M_i^{-1}$, then generate an honest normal proof using r and a known normal witness. The checker passes anamorphic verification with probability 1. Producing a witness $\widehat{u}$ for this target would also open M_i as $r - \widehat{u}$. Thus self-chosen targets require the witness-output condition in Definition 5, whereas Corollary 1 applies to independently registered targets.

3.2 Two Properties of the Normal Proof System

Conditional Carrier Simulation. For $C \in \Phi(U)$, write $\Phi^{-1}(C) = \{r \in U : \Phi(r) = C\}$. Given pp, an efficient sampler produces auxiliary information z independently of $\mathbf{H}$. An environment receives (pp, z), retains its witnesses, and makes at most N adaptive valid requests (x, w, C) and Q direct oracle queries. Validity requires $(x, w) \in \mathcal{R}$ and $C \in \Phi(U)$. Compare these two services:

1. $\mathcal{O}_{\mathrm{real}}(x, w, C)$. Sample $r \xleftarrow{\$} \Phi^{-1}(C)$ and independently $\omega_{\mathrm{rest}} \leftarrow \Omega_{\mathrm{rest}}(pp, x, w)$. Return $\mathsf{P}^{\mathbf{H}}(pp, x, w; (r, \omega_{\mathrm{rest}}))$.
2. $\mathcal{O}_{\mathrm{sim}}(x, w, C)$. Return $\mathsf{Sim}^{\mathbf{H}}_{\mathrm{car}}(pp, x, C)$.

The real service specifies a reference distribution and need not be efficient. The simulator is stateful and PPT, has the programming access of Section 2, and receives neither w nor z.

Definition 7 (Conditional carrier simulation). *We say that Π admits conditional carrier simulation with error $\Delta^{\mathrm{car}}(\lambda, N, Q)$ if there exists such a simulator for which the distinguishing advantage between $\mathcal{O}_{\mathrm{real}}$ and $\mathcal{O}_{\mathrm{sim}}$ is at most $\Delta^{\mathrm{car}}(\lambda, N, Q)$ for every admissible PPT environment and auxiliary-input sampler.*

Simulation is statistical if the same bound holds in total variation between complete views for all environments within these query bounds, including unbounded ones. Write Δ^{car}_N when λ and Q are clear.

Extracting the initial carrier. The predicate $\mathsf{Bad}(pp, v)$ checks in polynomial time whether v is a parameter relation. Set $\mathsf{Bad} \equiv 0$ when this branch is unused. An adversary $\mathcal{A}^{\mathbf{H}}(pp, z)$ receives auxiliary input z sampled independently of $\mathbf{H}$ given pp and returns (x, π, d). The polynomial-length declaration d is recorded but need not appear in the proof.

Run $\mathcal{A}$ once with the real oracles, recording (x_0, π_0, d_0) and its query trace. If $\mathsf{V}^{\mathbf{H}}(pp, x_0, \pi_0) = 0$, extraction fails. Otherwise set $C_0 = \mathsf{Car}(pp, x_0, \pi_0)$ and give the extractor the initial record and replay access of Section 2. Success is an opening $r^* \in U$ with $\Phi(r^*) = C_0$, or a relation v with $\mathsf{Bad}(pp, v) = 1$.

Let $\varepsilon_{\mathrm{norm}}$ denote the initial normal-acceptance probability, and let Q bound all queries in one execution, including verification. Fix $\kappa(\lambda, Q) \in [0,1]$ and $0 < \beta \leq 1$.

Definition 8 (Extraction of the initial carrier). *We say that Π admits extraction of the initial carrier if there exists a black-box extractor that succeeds with probability at least $(\varepsilon_{\mathrm{norm}} - \kappa(\lambda, Q))/\beta$ for every adversary above, in expected time polynomial in λ and the adversary's running-time and query bounds.*

Probability is over the complete experiment, and running time includes all replays and oracle implementations. We use the extractor of Attema, Fehr, Klooß, and Resch (AFKR) [5, Theorem 5 and Remark 13]. For an extraction structure $(\Gamma_1, \ldots, \Gamma_\mu)$, its search parameter is $\mathcal{T} = \prod_{i=1}^{\mu}(t_{\Gamma_i} + 1)$, where t_{Γ_i} is the maximum number of successive useful challenges needed to reach Γ_i [5, Definition 8]. For a round with threshold k_i, $t_{\Gamma_i} = k_i$; for a round requiring four challenges with pairwise distinct squares, $t_{\Gamma_i} = 4$. For an interactive extraction structure with error $\kappa_0 < 1$, this gives

$$\kappa(Q) = \min\{1, (Q+1)\kappa_0\}, \quad \beta = 1 - \kappa_0, \quad L_Q = \frac{\mathcal{T} + (\mathcal{T}-1)Q}{1-\kappa_0}. \tag{1}$$

Here L_Q bounds the expected number of wrapped-prover invocations. Each instance establishes its extraction structure and efficient useful-challenge sampling, as required by the AFKR result. Its early-choice property fixes the tree's first message to the initial invocation [5, Definition 13 and Lemma 10]. In both instantiations, this message contains the normal statement and carrier, so extraction opens the initial carrier. Appendix B.3 gives the replay and query adapter.

We say that Φ has joint hardness for Bad if no PPT adversary given pp and $M = \Phi(\eta)$ for uniform $\eta \xleftarrow{\$} U$ outputs a preimage of M or a valid parameter relation with non-negligible probability. Denote its success by $\mathrm{Adv}^{\mathsf{J}_{\Phi,\mathsf{Bad}}}(\mathcal{B})$, and write $\mathrm{Adv}^{\mathsf{Bad}}(\mathcal{B})$ for finding one from pp alone. Joint hardness implies the latter task's hardness and reduces to random-image preimage hardness when $\mathsf{Bad} \equiv 0$. ROM reductions lazily sample their oracles, counted in their running time.

3.3 Security

Theorem 1 (APoK from a homomorphic commitment). *Under this section's carrier structure and perfect-completeness assumption, if Π satisfies Definitions 7 and 8, Φ satisfies joint hardness, and κ and Δ^{car} are negligible for polynomially bounded interactions, then Construction 1 with independent masks is a secure APoK with anamorphic knowledge error κ. The seed variant has corresponding guarantees under PRF security.*

Security is for $\widehat{\mathcal{R}}_\Phi$ in the sense of Definition 6. Independent masks give perfect external covertness and anamorphic zero-knowledge error Δ_N^{car}, statistical when the carrier simulation is statistical. The seed variant gives computational bounds ϵ_F and $2\epsilon_F + \Delta_N^{\mathrm{car}}$, respectively, with the same anamorphic knowledge error κ.

The non-encodability reduction reserves one of T indices. Corollary 1 transfers its bound to the registered-target experiment. Appendix B.4 specifies the reserved-index reduction.

Proof. The construction invokes $\mathsf{P}^{\mathbf{H}}$ on a tape in the support of Ω_{Π}. Homomorphism gives $\Phi(\widehat{u} + \eta_i) = \widehat{C}M_i$ for every valid input and mask, so perfect completeness gives APoK correctness.

For external covertness, condition on the complete prior view and the adaptively chosen valid request (i, I) at a fresh index i. With independent masks, η_i remains independent and uniform in U under this conditioning. Hence $r = \widehat{u}+\eta_i$ is uniform in U for any adaptively chosen $\widehat{u}$. The remaining prover coins have the same distribution as in the normal mode, so (π, ω) has the same conditional distribution in both modes. Induction over requests gives identical joint views, hence perfect external covertness. For the seed variant, replacing the adaptively accessed PRF values by independent uniform masks costs at most ϵ_F.

For anamorphic zero knowledge, condition on the actual checking table, the retained witnesses, and the prior view. At an unused index i, the mask has been exposed only through M_i, so η_i remains uniform in $\Phi^{-1}(M_i)$. Addition of $\widehat{u}$ bijectively maps this fiber to $\Phi^{-1}(\widehat{C}M_i)$. The real proof thus has precisely the reference distribution in Definition 7. Its simulator returns $\mathsf{Sim}^{\mathbf{H}}_{\mathrm{car}}(pp, x, \widehat{C}M_i)$ and retains its state. The actual checking table is admissible auxiliary information because setup is independent of $\mathbf{H}$. This gives joint-view error Δ^{car}_N for all adaptive requests. Replacing the PRF table in the real and simulated experiments separately adds at most $2\epsilon_F$.

For anamorphic knowledge soundness, give the carrier extractor auxiliary input $z = (ek, ck)$ and the initial declaration $d_0 = (i_0, \widehat{C}_0)$. Both variants use the actual embedding key, fixed across replays. Let N_0, A_0 be the initial normal and anamorphic acceptance events, and E successful carrier extraction. Since $E, A_0 \subseteq N_0$, the carrier guarantee gives

$$\Pr[E \cap A_0] \geq \Pr[E] - \Pr[N_0 \setminus A_0] \geq \frac{\varepsilon_{\mathrm{ana}} - \kappa(Q)}{\beta}. \tag{2}$$

For the last inequality, substitute $\Pr[E] \geq (\varepsilon_{\mathrm{norm}} - \kappa(Q))/\beta$ and use $\beta \leq 1$. On an opening r^* of C_0, return $\widehat{u}^* = r^* - \eta_{i_0}$. Initial anamorphic acceptance yields $\Phi(\widehat{u}^*) = C_0 M_{i_0}^{-1} = \widehat{C}_0$. On a parameter relation, a reduction that generates the same auxiliary keys and runs the same extraction returns that relation instead. By joint hardness, the relation branch has negligible probability.

For non-encodability, first work with independent masks and denote attack success by $\varepsilon_{\mathrm{rnd}}$. Given a random-image challenge M, reserve $J \xleftarrow{\$} [T]$, set $M_J = M$, and generate every other mask honestly. A wrapper W runs the attacker on the resulting checking key. It answers valid fresh requests away from J with the known masks by running $\mathsf{P}^{\mathbf{H}}$ on the corresponding anamorphic tape, aborts on a valid request at J, and outputs the attacker's proof with declaration $(\widehat{C}_0, \widehat{u}_0)$ only when the final attack succeeds at J.

Couple W's initial execution with the independent-mask non-encodability experiment, using the same checking table, random oracles, and random tapes

for the attacker and honest prover. Every successful trajectory targets an unused label, so it survives exactly when J equals that label. Hence $\varepsilon_{\mathrm{norm}}(W) = \varepsilon_{\mathrm{rnd}}/T$. The reserved index, auxiliary keys, and random tape stay fixed during replay, while each run starts with an empty consumed set. Carrier extraction on W returns a parameter relation or an opening r^* of its initial carrier. Using the initial declared witness gives $\Phi(r^* - \widehat{u}_0) = C_0\widehat{C}_0^{-1} = M$. Thus the opening used to recover an anamorphic witness opens the challenge mask image. Let Q_{all} count wrapper queries, including honest-proof responses and final verification. The joint-hardness reduction succeeds with probability at least $(\varepsilon_{\mathrm{rnd}}/T - \kappa(Q_{\mathrm{all}}))/\beta$. For the seed variant, one PRF replacement loses at most ϵ_F.

Both reductions run in expected polynomial time. Polynomial truncation converts non-negligible success into a strict PPT attack on joint hardness. Thus the relation branch is negligible, (2) gives anamorphic knowledge error κ, and the reserved-index reduction gives non-encodability for polynomial T. □

Concrete Security Bounds. Let $\mathcal{B}_{\mathrm{ext}}$ and $\mathcal{B}_{\mathrm{ne}}$ be the relation and reserved-index reductions in the proof. Let R_{ext} and R_{ne} bound their expected total running times, including initial execution, verification, extraction, and oracle implementation. For an integer cutoff $b \geq 1$, write $\mathcal{B}_{\mathrm{ext},\leq b}$ and $\mathcal{B}_{\mathrm{ne},\leq b}$ for their truncated versions, and let $K_{\leq b}$ denote successful anamorphic knowledge extraction by that time. Applying Markov's inequality to the opening, relation, and timeout branches of (2) gives the bound

$$\Pr[K_{\leq b}] \geq \frac{\varepsilon_{\mathrm{ana}} - \kappa(Q)}{\beta} - \mathrm{Adv}^{\mathsf{Bad}}(\mathcal{B}_{\mathrm{ext},\leq b}) - \frac{R_{\mathrm{ext}}}{b}. \tag{3}$$

Applying the same timeout bound to the reserved-index reduction and including the PRF replacement gives the non-encodability bound

$$\mathrm{Adv}^{\text{NonEnc-HP}}(\mathcal{A}) \leq \epsilon_F + T\left[\kappa(Q_{\mathrm{all}}) + \beta\left(\mathrm{Adv}^{\mathsf{J}_{\Phi,\mathsf{Bad}}}(\mathcal{B}_{\mathrm{ne},\leq b}) + \frac{R_{\mathrm{ne}}}{b}\right)\right]. \tag{4}$$

For independent masks, set $\epsilon_F = 0$. The wrapper generates honest proofs with known masks, so (4) incurs no simulation error. The same bound applies to $\mathrm{Adv}^{\text{TU-HP}}(\mathcal{A}, \mathcal{T})$ by using the adversary of Corollary 1 in the reserved-index reduction. Its running time includes target sampling.

Conditional carrier simulation also gives normal zero knowledge. On input (pp, x), sample $r \xleftarrow{\$} U$, set $C = \Phi(r)$, discard r, and invoke $\mathsf{Sim}^{\mathbf{H}}_{\mathrm{car}}(pp, x, C)$. Averaging over the carrier distribution recovers the normal proof distribution up to the corresponding simulation error.

To compare direct embedding with selecting among normal proofs, let a PPT procedure with only pp, the actual checking key ck, and normal random-oracle access generate a valid state $Z = (i, x, w, \widehat{C}, \widehat{u})$, where i is unused. Let $\mathcal{A}_{\mathrm{norm}}$ execute this procedure, generate one normal proof using a fresh tape from $\Omega_{\Pi}(pp, x, w)$, and output it together with $i, x, \widehat{C}, \widehat{u}$.

Consider an encoder that generates candidates $\pi_1, \ldots, \pi_B$ and may hold additional sender secrets. For each fixed j, require that the joint distribution of π_j

with $pp, ck, Z, \mathbf{H}$ equal that of the normal proof generated by $\mathcal{A}_{\text{norm}}$ with the corresponding data. The candidates need not be independent.

Lemma 1 (Normal-candidate sampling). *Under the preceding conditions, an encoder that outputs only failure or an unmodified candidate has anamorphic-acceptance probability* $\delta_{\text{enc}} \leq B \ \text{Adv}^{\text{NonEnc}}(\mathcal{A}_{\text{norm}})$.

Proof. If the encoder succeeds, at least one candidate passes anamorphic verification. Apply the union bound to these B acceptance events. By the joint-distribution condition, each candidate's acceptance probability equals the success probability of $\mathcal{A}_{\text{norm}}$. □

For polynomial B, non-encodability rules out a correct encoder in this sampling model. For fixed valid $\widehat{C}$ and M_i, the carrier of a normal proof is uniform in $\Phi(U)$ and equals $\widehat{C}M_i$ with probability $1/|\Phi(U)|$. Independent sampling thus requires $|\Phi(U)|$ candidates in expectation, which is q in both instantiations.

Our prover instead sets $r = \widehat{u} + \eta_i$, satisfying the carrier check directly. Conditioned on the checking data, its distribution differs from the normal-candidate distribution and thus lies outside the sampling model of Lemma 1.

4 Anamorphic Compressed Σ

The first Compressed Σ message $A_0 = \Psi(r)$ is exactly the carrier of Construction 1. We show that a prescribed A_0 admits conditional simulation after one compression, while the normal accepting tree opens both the normal statement and A_0 without extra challenge branches.

4.1 The Carrier and the Construction

We use Compressed Σ of Attema and Cramer [4] in the formulation of Attema's Protocol 3 and its Fiat–Shamir transform for adaptively chosen statements [3, Protocol 3 and Definition 2.44]. Let $(\mathcal{U}, +)$ and $(\mathcal{H}, \cdot)$ be finite abelian groups of prime exponent $q \geq 3$, with efficient operations, uniform sampling, and membership tests. Scalars lie in $\mathbb{F}_q$. Set $n = 2^d \geq 4$ and $m = n/2$, so at least one compression precedes the final two-coordinate response. Public parameters $pp = (\mathcal{U}, \mathcal{H}, q, n, \Psi)$ specify an efficient homomorphism $\Psi : \mathcal{U}^n \to \mathcal{H}$.

The normal statement and witness are (P, w), and the anamorphic statement and witness are $(\widehat{C}, \widehat{u})$, with

$$\mathcal{R}_{pp} = \{(P, w) \in \mathcal{H} \times \mathcal{U}^n : \Psi(w) = P\}, \quad \widehat{\mathcal{R}}_{\Psi} = \{(\widehat{C}, \widehat{u}) \in \mathcal{H} \times \mathcal{U}^n : \Psi(\widehat{u}) = \widehat{C}\}.$$

The normal tape is a uniform $r \in \mathcal{U}^n$, with no remaining randomness. For Construction 1, set $U = \mathcal{U}^n$ and $\Phi = \Psi$, and let the carrier parser return A_0 from a well-formed proof and $\perp$ otherwise. The independent random oracles $\mathbf{H} = (H_0, \ldots, H_{d-1})$ have output space $\mathbb{F}_q$. Challenges are computed from canonical message prefixes:

$$c_1 = H_0(pp, P, A_0), \quad c_{j+1} = H_j(pp, P, A_0, (A_t, B_t)_{t=1}^{j}), \quad 1 \leq j < d. \tag{5}$$

At level j, let v_j, F_j, and Θ_j denote the response vector, homomorphism, and verification target. Subscripts L, R denote equal halves, with $F_{j,L}(t) = F_j(t, 0)$ and $F_{j,R}(t) = F_j(0, t)$.

Public input (pp, P). Valid request $I = (P, w, \widehat{C}, \widehat{u})$ at index i.

$\underline{\mathsf{AKGen}(pp, \mathcal{C})}$

Run AKGen of Construction 1 with $U = \mathcal{U}^n$ and $\Phi = \Psi$, obtaining ek and $ck = (M_i)_{i=1}^T$, where $M_i = \Psi(\eta_i)$.

$\underline{\mathsf{AProve}^{\mathbf{H}}(ek, i, I)}$

1: Recover η_i from ek and set $r \leftarrow \widehat{u} + \eta_i$. **[normal mode: $r \xleftarrow{\$} \mathcal{U}^n$]**
2: Compute $A_0 \leftarrow \Psi(r)$ and c_1 using (5). Set $F_1 \leftarrow \Psi$ and $v_1 \leftarrow r + c_1 w$.
3: For $j = 1, \ldots, d-1$, split $v_j = (v_{j,L}, v_{j,R})$ and compute

$$A_j \leftarrow F_{j,R}(v_{j,L}), B_j \leftarrow F_{j,L}(v_{j,R}), c_{j+1} \leftarrow H_j(pp, P, A_0, (A_t, B_t)_{t=1}^j),$$

$$v_{j+1} \leftarrow v_{j,L} + c_{j+1} v_{j,R}, F_{j+1}(t) \leftarrow F_{j,L}(t)^{c_{j+1}} F_{j,R}(t).$$

4: Return $\pi = (A_0, (A_j, B_j)_{j=1}^{d-1}, v_d)$.

$\underline{\mathsf{V}_{\mathrm{CS}}^{\mathbf{H}}(pp, P, \pi)}$

Check canonical encoding and membership $P, A_0, A_j, B_j \in \mathcal{H}$ and $v_d \in \mathcal{U}^2$, rejecting on failure. Recompute the challenges using (5) and set $F_1 \leftarrow \Psi$, $\Theta_1 \leftarrow A_0 P^{c_1}$. For $j = 1, \ldots, d-1$, update

$$\Theta_{j+1} \leftarrow A_j \Theta_j^{c_{j+1}} B_j^{c_{j+1}^2}, \quad F_{j+1}(t) \leftarrow F_{j,L}(t)^{c_{j+1}} F_{j,R}(t).$$

Return $[F_d(v_d) = \Theta_d]$.

$\underline{\mathsf{AVerify}^{\mathbf{H}}(ck, i, P, \widehat{C}, \pi)}$

Return 0 if i is invalid, $\widehat{C} \notin \mathcal{H}$, or $\mathsf{V}_{\mathrm{CS}}^{\mathbf{H}}(pp, P, \pi) = 0$. Otherwise return $[A_0 = \widehat{C} M_i]$.

Fig. 1. Anamorphic Compressed Σ.

4.2 Conditional Simulation through One Compression

Fix A_0 and set $\Theta = A_0 P^{c_1}$. The first compression maps $v = (v_L, v_R)$ with $\Psi(v) = \Theta$ to $z = v_L + e v_R$, where $e = c_2$. Its messages are $A = \Psi_R(v_L)$ and $B = \Psi_L(v_R)$, and $\Psi_e(t) = \Psi_L(t)^e \Psi_R(t)$, so $\Psi_e(z) = A \Theta^e B^{e^2}$.

Lemma 2 (Equal fibers). *If $J : \mathcal{U}^m \to \mathcal{H}^2$, $J(t) = (\Psi_L(t), \Psi_R(t))$, is surjective, then for fixed Θ, e and uniform $v \in \Psi^{-1}(\Theta)$, the pair (z, B) is uniform in $\mathcal{U}^m \times \mathcal{H}$, and*

$$A = \Psi_e(z)\Theta^{-e}B^{-e^2}. \tag{6}$$

Proof. Fix z, B, write $t = v_R$ and $v_L = z - et$. Consistent vectors are exactly the solutions of $J(t) = \big(B, \Theta\Psi_L(z)^{-1}B^e\big)$. Surjectivity gives $|\ker J|$ preimages for every (z, B), hence a uniform pair. Substituting into $A = \Psi_R(v_L)$ gives (6). □

The surjectivity assumption cannot be dropped from Lemma 2. If $\Psi_L = \Psi_R$, honest messages satisfy $AB = \Theta$, whereas independently sampled A, B need not.

The simulator $\mathsf{SimCar}^{\mathbf{H}}_{\mathrm{CS}}(pp, P, A_0)$ proceeds as follows:

1. Query $c_1 \leftarrow H_0(pp, P, A_0)$ and set $\Theta \leftarrow A_0 P^{c_1}$.
2. Sample independently $e \xleftarrow{\$} \mathbb{F}_q$, $z \xleftarrow{\$} \mathcal{U}^m$, and $B \xleftarrow{\$} \mathcal{H}$, and compute A by (6).
3. Set $\sigma = (pp, P, A_0, (A, B))$. If $H_1(\sigma)$ is defined, return $\perp$. Otherwise program $H_1(\sigma) \leftarrow e$.
4. Set $(A_1, B_1) = (A, B)$, $F_2 = \Psi_e$, and $v_2 = z$. Continue Figure 1 from compression $j = 2$ and return the proof. If $d = 2$, use z as the final response.

Lemma 3 (Simulation at a prescribed carrier). *Let δ_{pp} be the probability that J is not surjective. Then $\mathsf{SimCar}^{\mathbf{H}}_{\mathrm{CS}}$ satisfies Definition 7. For N adaptive requests and at most Q_H direct environment queries,*

$$\Delta^{\mathrm{car}}_N \le \delta_{pp} + \frac{NQ_H + N(N-1)/2}{|\mathcal{H}|^2}. \tag{7}$$

Proof. In the real reference distribution, r is uniform in $\Psi^{-1}(A_0)$. Translation by $c_1 w$ makes v uniform in $\Psi^{-1}(\Theta)$, so Lemma 2 gives the simulated compressed-state distribution.

Surjectivity of J makes every Ψ_e surjective. For fixed e, B, uniform z makes A uniform in $\mathcal{H}$, so (A, B) is uniform and independent of e. Before request k, at most $Q_H + k - 1$ inputs of H_1 are defined. The programming collision probability is therefore at most $(Q_H + k - 1)/|\mathcal{H}|^2$. At a fresh point, the programmed challenge has the real distribution and the remaining normal proving steps give the same view. Summing over the N requests gives (7). Parameter-rank failure contributes δ_{pp} once. □

4.3 Two Openings from One Accepting Tree

Compressed Σ has two branches at the first challenge and three at each compression level [3, Theorems 3.2–3.3]. The same accepting tree opens both P and A_0.

Lemma 4 (Two-opening extraction). *For any public homomorphism Ψ, an accepting tree with two distinct first challenges and three distinct challenges at each compression node yields $w^*, r^* \in \mathcal{U}^n$ such that $\Psi(w^*) = P$ and $\Psi(r^*) = A_0$, in time polynomial in the tree size.*

Proof. At a compression node with map $F : \mathcal{U}^{2\ell} \to \mathcal{H}$, target Θ, and messages (A, B), three child subtrees give $F(e_i z_i, z_i) = A\Theta^{e_i} B^{e_i^2}, i = 1, 2, 3$. For distinct e_i, choose λ_i satisfying $\sum_i \lambda_i = 0, \sum_i \lambda_i e_i = 1, \sum_i \lambda_i e_i^2 = 0$. Then $\sum_i \lambda_i e_i z_i$, $\sum_i \lambda_i z_i$ is a preimage of Θ under F. Recursing to the first challenge yields, for distinct c, c', $\Psi(v) = A_0 P^c, \Psi(v') = A_0 P^{c'}$. Set

$$w^* = (c' - c)^{-1}(v' - v), \quad r^* = v - cw^*.$$

Then $\Psi(w^*) = P$ and $\Psi(r^*) = A_0$. □

The extractor continues from the initial execution of Section 2. Its accepting tree preserves the prefix (pp, P, A_0) and therefore opens the initial carrier. Appendix B.3 gives the recursive invariant [5, Sections 5.2 and 6.4].

The interactive knowledge error and AFKR search parameter are

$$\kappa_d = 1 - (1 - 1/q)(1 - 2/q)^{d-1} \le (2d - 1)/q, \beta_d = 1 - \kappa_d, \mathcal{T}_d = 3 \cdot 4^{d-1} = 3n^2/4.$$

At each node, sample uniformly from challenge values not yet collected. Define $\kappa_{\mathrm{CS}}(Q)$ and $L_{d,Q}$ by substituting $\kappa_0 = \kappa_d$ and $\mathcal{T} = \mathcal{T}_d$ into (1). By the AFKR bound [5, Theorem 5 and Remark 13], this establishes Definition 8 with $\kappa = \kappa_{\mathrm{CS}}$, $\beta = \beta_d$, and $\mathsf{Bad} \equiv 0$.

4.4 Random Bases and Security

Set $\mathcal{U} = \mathbb{F}_q$ and $\mathcal{H} = \mathbb{G} = \langle g \rangle$ of prime order q. Setup samples independent uniform $g_i \in \mathbb{G} \setminus \{1\}$ and defines $\Psi(t) = \prod_{i=1}^n g_i^{t_i}$. Let $\boldsymbol{\alpha}, \boldsymbol{\gamma} \in (\mathbb{F}_q^*)^m$ be the exponent rows of the two halves. They are linearly independent exactly when J is surjective. For fixed $\boldsymbol{\alpha}$, exactly $q-1$ of the $(q-1)^m$ possible $\boldsymbol{\gamma}$ are proportional, so $\delta_{pp} = (q-1)^{1-m}$. Lemma 3 therefore gives

$$\delta_{\mathrm{CS}}(N, Q_H) = (q-1)^{1-n/2} + \frac{NQ_H + N(N-1)/2}{q^2}.$$

Preimage hardness reduces to DL. Given $X = g^a$ for uniform $a \in \mathbb{F}_q$, sample independent $\alpha_i \xleftarrow{\$} \mathbb{F}_q^*$, set $g_i = g^{\alpha_i}$ and $M = X$. Any t satisfying $\Psi(t) = M$ yields $a = \sum_i \alpha_i t_i \bmod q$. The rank condition enters only the privacy bound.

Assume $n = 2^d \ge 4$, the label count and interaction bounds are polynomial in λ, and $1/q$ is negligible.

Theorem 2 (Compressed Σ APoK). *Under independent uniform nonidentity bases and the DL assumption, the independent-mask construction of Figure 1 is a secure APoK for $\widehat{\mathcal{R}}_\Psi$ with perfect external covertness, δ_{CS}-statistical strong anamorphic zero knowledge, and anamorphic knowledge error κ_{CS}.*

Proof. Lemma 3 gives conditional carrier simulation, and Lemma 4 gives initial-carrier extraction. The DL reduction above gives preimage hardness. Honest executions satisfy $F_j(v_j) = \Theta_j$ at every level and $A_0 = \widehat{C} M_i$, giving perfect correctness. Since $\mathcal{T}_d = 3n^2/4$ and β_d^{-1} are polynomial, Theorem 1 applies. □

For bases derived from public data, hardness and rank require separate arguments. Assume it is hard, given the bases and their generation data, to find a nonzero v with $\Psi(v) = 1$. Sample $\eta \xleftarrow{\$} \mathbb{F}_q^n$ and give $M = \Psi(\eta)$ to a preimage finder. Any opening $t \neq \eta$ yields the relation $t - \eta$, and $\Pr[t = \eta] \leq q^{1-n}$ since η is uniform on its fiber. For privacy, use $\delta_{pp} = \Pr[\operatorname{rank} J < 2]$ for the actual setup in (7). Independent hash-to-group outputs uniform on $\mathbb{G} \setminus \{1\}$ at distinct fixed inputs give the rank bound of Theorem 2 when the setup oracle is independent of the Fiat–Shamir oracles.

The seed variant inherits Theorem 1 with $\Delta_N^{\mathrm{car}} = \delta_{\mathrm{CS}}$. Both variants use the same carrier-opening extractor. For at most Q_H direct random-oracle queries, set $Q_{\mathrm{ext}} = Q_H + d$. Anamorphic knowledge extraction succeeds with probability at least $\max\left\{0, \frac{\varepsilon_{\mathrm{ana}} - \kappa_{\mathrm{CS}}(Q_{\mathrm{ext}})}{\beta_d}\right\}$. For at most N honest-proof requests, set $Q_{\mathrm{all}} = Q_H + (N+1)d$. Equation (4) and Corollary 1 bound non-encodability and target unforgeability with $\kappa = \kappa_{\mathrm{CS}}$, $\beta = \beta_d$, and the DL advantage in place of joint hardness. For independent masks, $\epsilon_F = 0$. The expected number of wrapped invocations is at most $L_{d,Q_{\mathrm{all}}}$.

5 Anamorphic Bulletproofs

The first-round Bulletproofs message S commits to the prover's blinding randomness and is the carrier of Construction 1. We simulate the full proof at a prescribed S. For extraction, we isolate the step of the original range argument that opens S at two challenges [13, Appendix C].

5.1 The Carrier and the Construction

We use the Bulletproofs range protocol and inner-product argument (IPA) [13]. The normal proof system is the Fiat–Shamir protocol specified in Appendix A, where each challenge hashes the public input and preceding transcript.

The generator $\mathsf{Setup}_{\mathrm{BP}}(1^\lambda)$ outputs pp, specifying a group $\mathbb{G}$ of prime order $q > 7$, a range bit length $s = 2^L$ with $2^s < q$, and ordered nonidentity bases

$$\mathbf{B} = (g_0, h_0, g_1, \ldots, g_s, h_1, \ldots, h_s, u_{\mathsf{IP}}).$$

Write $\mathbf{g} = (g_j)_{j=1}^s$, $\mathbf{h} = (h_j)_{j=1}^s$, and $\mathbf{g}^{\mathbf{a}} = \prod_j g_j^{a_j}$. The symbol $\mathbf{1}$ denotes the all-ones vector, and $\langle \cdot, \cdot \rangle$ denotes the inner product. The operation $\circ$ and vector inversion act coordinatewise. The normal statement is a Pedersen commitment V [36] with witness (v, γ) and relation

$$\mathcal{R}_{\mathrm{BP}} = \{(V, (v, \gamma)) : 0 \leq v < 2^s,\ \gamma \in \mathbb{F}_q,\ V = g_0^v h_0^\gamma\}.$$

The normal tape contains carrier randomness $(\rho, \mathbf{s}_L, \mathbf{s}_R) \in U = \mathbb{F}_q^{2s+1}$ and blinding scalars $(\alpha, \tau_1, \tau_2) \in \mathbb{F}_q^3$. Define $\Phi(\rho, \mathbf{s}_L, \mathbf{s}_R) = h_0^\rho \mathbf{g}^{\mathbf{s}_L} \mathbf{h}^{\mathbf{s}_R}$ and $\widehat{\mathcal{R}}_\Phi = \{(\widehat{S}, \widehat{u}) \in \mathbb{G} \times U : \Phi(\widehat{u}) = \widehat{S}\}$. The carrier parser returns the proof's S field.

The outer proof contains A, S, T_1, T_2 and responses $\tau_x, \mu, \widehat{t}$. Independent random oracles produce y, z, x, the initial IPA challenge ξ_0, and recursive challenges $\xi_1, \ldots, \xi_L$, all uniformly in $\mathbb{F}_q^*$. In this section the normal statement is V, and x denotes the third challenge, following the notation of [13]; the generic statement x of Section 2 does not appear. These challenge domains are fixed in the normal protocol and used in both modes. Write

$$\mathbf{y} = (1, y, \ldots, y^{s-1}), \quad \mathbf{p} = (1, 2, \ldots, 2^{s-1}), \quad \mathbf{h}' = (h_j^{y^{-(j-1)}})_{j=1}^{s}.$$

Set $\delta(y,z) = (z - z^2)\langle \mathbf{1}, \mathbf{y} \rangle - z^3 \langle \mathbf{1}, \mathbf{p} \rangle$. Normal verification checks

$$g_0^{\widehat{t}} h_0^{\tau_x} = V^{z^2} g_0^{\delta(y,z)} T_1^{x} T_2^{x^2}, \tag{8}$$

and runs the IPA on $P_{\mathrm{IP}} = A S^x \mathbf{g}^{-z\mathbf{1}} (\mathbf{h}')^{z\mathbf{y} + z^2 \mathbf{p}} h_0^{-\mu}$, proving knowledge of $(\mathbf{l}, \mathbf{r})$ such that $P_{\mathrm{IP}} = \mathbf{g}^{\mathbf{l}} (\mathbf{h}')^{\mathbf{r}}$ and $\langle \mathbf{l}, \mathbf{r} \rangle = \widehat{t}$. The initial IPA challenge sets $u = u_{\mathsf{IP}}^{\xi_0}$. Appendix A gives the complete prover and verifier.

Construction 2 (Anamorphic Bulletproofs).
A valid request at index i is $I = (V, (v, \gamma), \widehat{S}, \widehat{u})$.

1. $\mathsf{AKGen}(pp, \mathcal{C})$. Run AKGen of Construction 1 with $U = \mathbb{F}_q^{2s+1}$ and Φ above, obtaining ek and $ck = (M_i)_{i=1}^{T}$ with $M_i = \Phi(\eta_i)$.
2. $\mathsf{ACoins}(ek, i, I)$. Recover η_i from ek, sample $\alpha, \tau_1, \tau_2 \xleftarrow{\$} \mathbb{F}_q$, and return $\omega = (\widehat{u} + \eta_i, (\alpha, \tau_1, \tau_2))$. **[normal mode:** $(\rho, \mathbf{s}_L, \mathbf{s}_R) \xleftarrow{\$} U$**]**
3. $\mathsf{AProve}^{\mathbf{H}}(ek, i, I)$. Sample $\omega \leftarrow \mathsf{ACoins}(ek, i, I)$ and run $\mathsf{P}_{\mathrm{BP}}^{\mathbf{H}}(pp, V, (v, \gamma); \omega)$. Return its proof π, whose carrier satisfies $S = \widehat{S} M_i$.
4. $\mathsf{AVerify}^{\mathbf{H}}(ck, i, V, \widehat{S}, \pi)$. Reject if i is invalid, $\widehat{S} \notin \mathbb{G}$, or $\mathsf{V}_{\mathrm{BP}}^{\mathbf{H}}(pp, V, \pi) = 0$. Otherwise return $[S = \widehat{S} M_i]$.

5.2 Conditional Simulation at a Prescribed Carrier

The original range simulator samples A and solves for S [13, Appendix C]. Here S is prescribed. For every $(\mathbf{s}_L, \mathbf{s}_R)$ there is a unique consistent ρ, since h_0 has order q. Thus, conditioned on S, the pair $(\mathbf{s}_L, \mathbf{s}_R)$ remains uniform. The simulator samples the IPA vectors directly and determines the remaining commitments from the normal verification equations.

Let $I_y = ((pp, V), A, S)$. Appending y gives I_z, and appending (z, T_1, T_2) to this prefix gives I_x. On input (pp, V, S) the simulator $\mathsf{SimCar}_{\mathrm{BP}}^{\mathbf{H}}$ takes the following steps.

1. Sample $y, z, x \xleftarrow{\$} \mathbb{F}_q^*$, $\mathbf{l}, \mathbf{r} \xleftarrow{\$} \mathbb{F}_q^s$, $\mu, \tau_x \xleftarrow{\$} \mathbb{F}_q, T_2 \xleftarrow{\$} \mathbb{G}$. Compute $\widehat{t} = \langle \mathbf{l}, \mathbf{r} \rangle$ and

$$A = h_0^{\mu} \mathbf{g}^{\mathbf{l} + z\mathbf{1}} (\mathbf{h}')^{\mathbf{r} - z\mathbf{y} - z^2 \mathbf{p}} S^{-x}, T_1 = \left(g_0^{\widehat{t}} h_0^{\tau_x} V^{-z^2} g_0^{-\delta(y,z)} T_2^{-x^2} \right)^{x^{-1}}. \tag{9}$$

2. If I_y, I_z, or I_x is already defined in its oracle, return $\perp$. Otherwise program $H_y(I_y) = y$, $H_z(I_z) = z$, and $H_x(I_x) = x$.

3. Run the IPA with witness $(\mathbf{l}, \mathbf{r})$ under the remaining random oracles and return the complete proof.

The sampled values satisfy (8) and make $P_{\mathrm{IP}} = \mathbf{g}^{\mathbf{l}}(\mathbf{h}')^{\mathbf{r}}$.

Lemma 5 (Bulletproofs simulation conditioned on the carrier). *For every setup output, valid normal statement V, and $S \in \mathbb{G}$, the simulator satisfies Definition 7 statistically. If Q_y, Q_z, Q_x bound the initial sizes of the outer oracle tables, its single-proof error is at most $\Delta_{\mathrm{BP},1} = \frac{Q_y}{q} + \frac{Q_z}{q(q-1)} + \frac{Q_x}{q^2}$.*

Proof. Fix a valid witness (v, γ), let $\mathbf{a}_L$ be the bit vector of v, and set $\mathbf{a}_R = \mathbf{a}_L - \mathbf{1}$. For fixed nonzero y, z, x, normal tapes conditioned on S are in bijection with $(\mathbf{l}, \mathbf{r}, \mu, T_2, \tau_x)$. The inverse begins with

$$\mathbf{s}_L = x^{-1}(\mathbf{l} - \mathbf{a}_L + z\mathbf{1}), \mathbf{s}_R = x^{-1}\big(\mathbf{y}^{-1} \circ (\mathbf{r} - z^2\mathbf{p}) - \mathbf{a}_R - z\mathbf{1}\big). \tag{10}$$

The fixed S determines ρ, then $\alpha = \mu - x\rho$. With $t_2 = \langle \mathbf{s}_L, \mathbf{y} \circ \mathbf{s}_R \rangle$, the equation $T_2 = g_0^{t_2} h_0^{\tau_2}$ determines τ_2, and $\tau_1 = x^{-1}(\tau_x - z^2\gamma - x^2\tau_2)$. These assignments give (9). Both spaces have q^{2s+3} elements, so uniform simulation variables correspond to uniform conditional tapes.

Condition on the prior view. Uniform μ makes A uniform, giving collision bounds Q_y/q and $Q_z/[q(q-1)]$ for I_y and I_z. Uniform τ_x, T_2 make (T_1, T_2) uniform, giving Q_x/q^2 for I_x. At fresh points, the programmed challenges have the real distribution. Running the IPA on the corresponding vectors and the same remaining oracles gives identical subsequent views. A union bound gives the claimed error. □

Let Q_j^{env} bound the environment's direct queries to H_j for $j \in \{y, z, x\}$. For N adaptive requests,

$$\Delta_{\mathrm{BP}}(N) \le N\left(\frac{Q_y^{\mathrm{env}}}{q} + \frac{Q_z^{\mathrm{env}}}{q(q-1)} + \frac{Q_x^{\mathrm{env}}}{q^2}\right) + \binom{N}{2}\left(\frac{1}{q} + \frac{1}{q(q-1)} + \frac{1}{q^2}\right). \tag{11}$$

Each preceding proof adds at most one entry to each outer oracle table.

5.3 Carrier Opening from the IPA Tree

Following the DLR formulation of Bulletproofs [13, Definition 1], define $\mathsf{Bad}_{\mathrm{DLR}}(pp, \mathbf{v}) = [\mathbf{v} \in \mathbb{F}_q^{2s+3} \setminus \{\mathbf{0}\} \wedge \mathbf{B}^{\mathbf{v}} = 1]$. Write $\mathrm{Adv}^{\mathsf{DLR}}(\mathcal{B})$ for the probability that $\mathcal{B}(pp)$ outputs such a relation for $pp \leftarrow \mathsf{Setup}_{\mathrm{BP}}(1^\lambda)$.

Proposition 1 (IPA extraction interface). *For an IPA input $(P_{\mathrm{IP}}, \widehat{t}, \mathbf{g}, \mathbf{h}')$, an accepting tree with two distinct initial challenges and four challenges with pairwise distinct squares at each recursive node yields vectors $(\mathbf{a}, \mathbf{b})$ satisfying $P_{\mathrm{IP}} = \mathbf{g}^{\mathbf{a}}(\mathbf{h}')^{\mathbf{b}}$ and $\langle \mathbf{a}, \mathbf{b} \rangle = \widehat{t}$, or a nonzero DLR relation on $\mathbf{B}$.*

Proof. At a recursive node, write P for the current target, $(\mathbf{g}, \mathbf{h}, u)$ for the current bases, and $X = \xi^2$. Three distinct values of X recover representations of L, R, and P, since the matrix with rows $(X, 1, X^{-1})$ is invertible. A fourth child either disagrees, giving a nonzero relation, or is consistent with these representations.

Let $\mathbf{a}^{\mathsf{L}}, \mathbf{a}^{P}, \mathbf{a}^{\mathsf{R}}$ be the recovered coefficient vectors on the current $\mathbf{g}$ bases. For a child response $\mathbf{a}'$ at $X = \xi^2$,

$$\xi^{-1}\mathbf{a}' = \mathbf{a}_L^{\mathsf{L}} X + \mathbf{a}_L^{P} + \mathbf{a}_L^{\mathsf{R}} X^{-1}, \quad \xi\mathbf{a}' = \mathbf{a}_R^{\mathsf{L}} X + \mathbf{a}_R^{P} + \mathbf{a}_R^{\mathsf{R}} X^{-1}.$$

Eliminating $\mathbf{a}'$ gives

$$\mathbf{a}_L^{\mathsf{L}} X^3 + (\mathbf{a}_L^{P} - \mathbf{a}_R^{\mathsf{L}}) X^2 + (\mathbf{a}_L^{\mathsf{R}} - \mathbf{a}_R^{P}) X - \mathbf{a}_R^{\mathsf{R}} = \mathbf{0}. \tag{12}$$

Four pairwise distinct squared challenges force all coefficients to vanish. The same argument on the $\mathbf{h}$ coordinates gives the folding identities, while the scalar coordinate gives $c^P = \langle \mathbf{a}^P, \mathbf{b}^P \rangle$. More explicitly, the recovered child vectors satisfy

$$\mathbf{a}' = \xi \mathbf{a}_L^P + \xi^{-1} \mathbf{a}_R^P, \qquad \mathbf{b}' = \xi^{-1} \mathbf{b}_L^P + \xi \mathbf{b}_R^P.$$

Substituting these identities into $\langle \mathbf{a}', \mathbf{b}' \rangle = c^{\mathsf{L}} X + c^P + c^{\mathsf{R}} X^{-1}$ and multiplying by X gives

$$\big(\langle \mathbf{a}_L^P, \mathbf{b}_R^P \rangle - c^{\mathsf{L}}\big) X^2 + \big(\langle \mathbf{a}^P, \mathbf{b}^P \rangle - c^P\big) X + \langle \mathbf{a}_R^P, \mathbf{b}_L^P \rangle - c^{\mathsf{R}} = 0.$$

This quadratic vanishes at all four values of X, so its coefficients vanish and the parent vectors satisfy the required inner-product relation. Recursing yields either a relation or vectors $(\mathbf{a}, \mathbf{b})$ with $P_{\mathrm{IP}} = \mathbf{g}^{\mathbf{a}} (\mathbf{h}')^{\mathbf{b}} u_{\mathsf{IP}}^{\xi_0(\langle \mathbf{a}, \mathbf{b} \rangle - \widehat{t})}$. For two distinct initial challenges, different recovered vectors give a nonzero relation. If they agree, nonidentity u_{IP} forces $\langle \mathbf{a}, \mathbf{b} \rangle = \widehat{t}$. A relation $(\mathbf{v}_g, \mathbf{v}_h, v_u)$ on folded bases lifts one level to

$$(\xi^{-1}\mathbf{v}_g, \xi \mathbf{v}_g, \xi \mathbf{v}_h, \xi^{-1}\mathbf{v}_h, v_u)$$

on the parent bases. This map is injective for $\xi \neq 0$. Repeated lifting, followed by the nonzero scalings defining $\mathbf{h}'$ and $u = u_{\mathsf{IP}}^{\xi_0}$, preserves nonzero relations. Padding zeros on g_0, h_0 gives a relation on the original public bases. Appendix B.1 gives the full coordinate argument and relation lifting. □

To open S, fix y, z and apply the two-branch step from the range extractor [13, Appendix C]. If the outer range protocol sends $(\mathbf{l}, \mathbf{r})$ explicitly, the same subtraction below applies directly. The IPA compresses these vectors and supplies them to the extractor through Proposition 1.

Lemma 6 (Bulletproofs carrier-opening extraction). *An accepting tree with one branch at each of y, z, two distinct branches at each of x, ξ_0, and four challenges with pairwise distinct squares at every recursive IPA node yields an opening of S or a nonzero DLR relation on $\mathbf{B}$.*

Proof. Apply Proposition 1 to the two x branches. It returns a DLR relation or vectors $(\mathbf{l}_i, \mathbf{r}_i)$ satisfying

$$AS^{x_i}\mathbf{g}^{-z\mathbf{1}}(\mathbf{h}')^{z\mathbf{y}+z^2\mathbf{p}} = h_0^{\mu_i}\mathbf{g}^{\mathbf{l}_i}(\mathbf{h}')^{\mathbf{r}_i}.$$

In the latter case, divide the equations and let $a = (x_1 - x_2)^{-1}$. Set

$$\rho^* = a(\mu_1 - \mu_2), \quad \mathbf{s}_L^* = a(\mathbf{l}_1 - \mathbf{l}_2), \quad \mathbf{s}_R^* = a\mathbf{y}^{-1} \circ (\mathbf{r}_1 - \mathbf{r}_2).$$

Since $(\mathbf{h}')^{\mathbf{v}} = \mathbf{h}^{\mathbf{y}^{-1}\circ\mathbf{v}}$, $r^* = (\rho^*, \mathbf{s}_L^*, \mathbf{s}_R^*)$ satisfies $\Phi(r^*) = S$. □

Full range extraction uses three x branches and varies y, z [24, Lemma 6.3]. After interpolation in x, consistency reduces to

$$\langle \mathbf{a}_L \circ \mathbf{a}_R, \mathbf{y}\rangle + z\langle \mathbf{a}_L - \mathbf{a}_R - \mathbf{1}, \mathbf{y}\rangle + z^2(\langle \mathbf{a}_L, \mathbf{p}\rangle - v) = 0.$$

Three distinct z values separate its coefficients, and s distinct y values recover the coordinatewise range constraints. Appendix B.2 gives the full extraction.

Table 1. Extraction in the same normal Bulletproofs protocol.

Parameter	Carrier opening	Full range witness
y branches	1	s
z branches	1	3
x branches	2	3
Initial IPA branches	2	2
Branches per IPA node	4	4
AFKR search parameter	$36 \cdot 5^L$	$48(s+1)5^L$
Interactive knowledge-error bound	$(2+6L)/(q-1)$	$(s+4+6L)/(q-1)$

The two trees in Table 1 have $4 \cdot 4^L$ and $18s \cdot 4^L$ leaves, respectively, giving a ratio of $2/(9s)$. Their AFKR search parameters have ratio $3/[4(s+1)]$ because each round contributes $t_{\Gamma_i} + 1$ to the product. At each recursive IPA node, three distinct squared challenge values exclude at most six nonzero challenges, contributing $6/(q-1)$. The x and initial IPA rounds each contribute $1/(q-1)$. Hence

$$\kappa_{\mathrm{BP}} = 1 - \left(1 - \frac{1}{q-1}\right)^2 \left(1 - \frac{6}{q-1}\right)^L \le \frac{2+6L}{q-1}. \tag{13}$$

Set $\beta_{\mathrm{BP}} = 1 - \kappa_{\mathrm{BP}}$ and $\mathcal{T}_{\mathrm{BP}} = 36 \cdot 5^L$. Define $\kappa_{\mathrm{BP}}^{\mathrm{FS}}(Q)$ and $L_{\mathrm{BP},Q}$ by substituting $\kappa_0 = \kappa_{\mathrm{BP}}$ and $\mathcal{T} = \mathcal{T}_{\mathrm{BP}}$ into (1). The AFKR bound [5, Theorem 5 and Remark 13] gives Definition 8 with $\beta = \beta_{\mathrm{BP}}$ and $\mathsf{Bad} = \mathsf{Bad}_{\mathrm{DLR}}$. Appendix B.3 gives the Fiat–Shamir query adapter and useful-challenge sampling.

5.4 Security and Aggregation

Assume s, the label count T, and the interaction bounds are polynomial in λ, and $1/q$ is negligible.

Theorem 3 (Bulletproofs APoK). *Under the DLR assumption for the fixed normal protocol and parameter distribution above, Construction 2 with independent masks is a secure APoK for $\widehat{\mathcal{R}}_{\Phi}$, with perfect external covertness, Δ_{BP}-statistical strong anamorphic zero knowledge, and anamorphic knowledge error $\kappa_{\mathrm{BP}}^{\mathrm{FS}}$.*

Proof. Honest tapes satisfy the normal verification equations and $S = \widehat{S}M_i$, giving perfect correctness. Lemma 5 gives conditional carrier simulation, and Lemma 6 gives initial-carrier extraction.

For joint hardness, sample $\eta \xleftarrow{\$} U$, give $M = \Phi(\eta)$ to the attacker, and retain η. Return any DLR relation produced by the attacker. If it instead returns a different opening t of M, then $t-\eta$, padded with zeros on g_0 and u_{IP}, is a nonzero relation on the public bases. Conditioned on the attacker's view, η is uniform among the q^{2s} openings of M, so coincident openings occur with probability at most q^{-2s}. Thus $\mathrm{Adv}^{\mathsf{J}_{\Phi},\mathsf{Bad}_{\mathrm{DLR}}}(\mathcal{B}_{\mathrm{J}}) \le \mathrm{Adv}^{\mathsf{DLR}}(\mathcal{B}_{\mathrm{DLR}}) + q^{-2s}$. Since $5^L = s^{\log_2 5}$ is polynomial, Theorem 1 applies. □

The certified witness. The extractor returns a representation on $(h_0, \mathbf{g}, \mathbf{h})$. Suppose an issuer samples $k \xleftarrow{\$} \mathbb{F}_q$, registers $\widehat{S} = h_0^k$, and uses $(k, 0, \ldots, 0)$ as its opening. A reduction sampling the same k can turn any different extracted opening t into the nonzero DLR relation $t - (k, 0, \ldots, 0)$. Thus, except with negligible probability under DLR, a successful extraction for this issuance procedure recovers k. Independence of the registered target from ck alone gives no such conclusion. A registrant choosing $\widehat{S} = g_1^a$ for uniform a obtains the same uniform public-key distribution and knows $(0, a, 0, \ldots, 0)$. Appendix D.3 gives a Compressed Σ instance where every opening yields the scalar for every registered target.

The seed variant inherits the PRF guarantees of Theorem 1 with $\Delta_N^{\mathrm{car}} = \Delta_{\mathrm{BP}}(N)$. For at most Q_H direct random-oracle queries, set $Q_{\mathrm{ext}} = Q_H + \nu$, $Q_{\mathrm{all}} = Q_H + (N+1)\nu$, and $\nu = L + 4$. Equations (3) and (4) and Corollary 1 give the corresponding bounds with $\kappa = \kappa_{\mathrm{BP}}^{\mathrm{FS}}$, $\beta = \beta_{\mathrm{BP}}$, and joint-hardness advantage bounded by the DLR advantage plus q^{-2s}. The corresponding $L_{\mathrm{BP},Q}$ bounds the expected number of wrapped invocations.

The aggregate protocol retains the same carrier. For commitments $V_j = g_0^{v_j} h_0^{\gamma_j}$ to m values $0 \le v_j < 2^n$, let $D = mn = 2^L$ and define

$$\mathbf{c}_z = (z^2\mathbf{p}_n, \ldots, z^{m+1}\mathbf{p}_n), \quad \delta_{\mathrm{agg}}(y,z) = (z - z^2)\langle \mathbf{1}_D, \mathbf{y}_D \rangle - \sum_{j=1}^{m} z^{j+2}\langle \mathbf{1}_n, \mathbf{p}_n \rangle,$$

where $\mathbf{y}_D = (1, y, \ldots, y^{D-1})$ and $\mathbf{p}_n = (1, 2, \ldots, 2^{n-1})$. The simulator replaces $z^2\mathbf{p}$ with $\mathbf{c}_z$, V^{z^2} with $\prod_{j=1}^{m} V_j^{z^{j+1}}$, and δ with δ_{agg}. These public terms are fixed across the two x branches and cancel in carrier extraction.

Assume polynomial m, n, $D = mn = 2^L$, $2^n < q$, and $\max\{D, m+2\} \le q-1$.

Corollary 2 (APoK for range aggregation). *Theorem 3 applies to the normal aggregate protocol, replacing s with D in the carrier dimension and extraction parameters. The conditional simulation bound is unchanged, and the coincident-opening term becomes q^{-2D}.*

Proof. For prescribed S, apply Lemma 5 in dimension D with the aggregate substitutions above. Fix a valid normal witness and nonzero y, z, x. The inverse map in (10) uses $\mathbf{c}_z$ in place of $z^2\mathbf{p}$, and the aggregate scalar response gives $\tau_1 = x^{-1}(\tau_x - x^2\tau_2 - \sum_{j=1}^{m} z^{j+1}\gamma_j)$. The remaining inverse assignments give a bijection between two spaces of q^{2D+3} elements. Uniform samples reproduce the conditional tape distribution, and the same blinders give the programming bound (11).

For carrier extraction, the aggregate offset is fixed across the two x branches. Lemma 6 therefore yields an opening of S or a nonzero DLR relation in dimension D. The carrier fiber has q^{2D} elements, so the coincident-opening term in the joint-hardness reduction is q^{-2D}. Theorem 1 gives the claimed APoK guarantees. □

Each aggregate proof consumes one label, certifies knowledge of a $(2D+1)$-coordinate anamorphic witness, and contains $2\log_2 D + 4$ group elements and five scalars. Appendix C gives the complete aggregate equations and normal range-witness extraction.

6 A Structural Boundary for Three-Move Protocols

A checker that knows the normal witness may be able to transfer a simulated response to the real Fiat–Shamir challenge. For the three-move protocols below, this converts a strong anamorphic zero-knowledge simulator into a checking-side encoder. For a first-message check, transfer preserves acceptance. For a full-transcript check, the bound also depends on the number of oracle queries and the difficulty of finding two accepting responses to the same challenge at the same first message.

Write (a, c, z) for the first message, challenge, and complete canonical response. Normal verification uses a deterministic predicate $V_{\mathrm{int}}(x, a, c, z)$ and the Fiat–Shamir rule $c = H(\mathsf{st}, a)$, with uniform challenges in a finite set Ch. The fixed context st contains the normal protocol fields, its encoding is unambiguous in a, and normal verification makes no further random-oracle queries. Let Q bound the honest anamorphic prover's oracle queries, including any made while selecting its tape.

Consider one fresh index, omitted from the notation. Setup and initial state are independent of H. Given (pp, ck), a PPT input sampler chooses valid pairs (x, w) and $(\widehat{x}, \widehat{w})$ without querying H and retains both witnesses. A local anamorphic check is

$$V_{\mathrm{int}}(x, a, c, z) \wedge \mathsf{Chk}(ck, x, \widehat{x}, a, c, z),$$

where Chk is a deterministic PPT predicate using only its displayed inputs and making no oracle queries. Let ε be the honest anamorphic acceptance probability, averaged over setup, inputs, and proving. Fix the simulator guaranteed by Definition 3. For each PPT test below, let δ_{zk} denote its absolute real-versus-simulated gap under the actual checking key. By strong anamorphic zero knowledge, every such δ_{zk} is negligible.

Definition 9 (Witness-assisted challenge transfer). *Let w be any valid witness for x and (a, c, z) any accepting transcript. A protocol admits transfer if an oracle-free PPT algorithm Tr produces a response $z' = \mathsf{Tr}(x, w, a, c, z, c')$ with $V_{\mathrm{int}}(x, a, c', z') = 1$ for every $c' \in \mathsf{Ch}$.*

Transfer preserves a without requiring the normal prover's coins.

Proposition 2 (First-message anamorphic checks). *If the protocol admits transfer and Chk depends only on $(ck, x, \widehat{x}, a)$, some PPT checking-side attacker has non-encodability success at least $\varepsilon - \delta_{\mathrm{zk}}$.*

Proof. The strong anamorphic zero-knowledge test accepts exactly when the proof passes normal and anamorphic verification, so its simulated acceptance is at least $\varepsilon - \delta_{\mathrm{zk}}$. The attacker samples the same inputs and runs the simulator with the actual ck under an independent local oracle using a private table. On local acceptance, it queries $c' = H(\mathsf{st}, a)$ and outputs the normal proof encoding of $(a, c', \mathsf{Tr}(x, w, a, c, z, c'))$, with its anamorphic witness. Transfer preserves normal acceptance, while the first-message predicate is unchanged. The attack uses no honest proofs and at most one real-oracle query. □

For a full-transcript check, the transferred response and the response accepted in a second execution may differ. A fork at the selected first message lets us bound the effect of this difference.

Definition 10 (Unique accepting response). *A protocol has unique accepting responses if $|\{z : V_{\mathrm{int}}(x, a, c, z) = 1\}| \leq 1$ for every valid x, first message a, and challenge c.*

Definition 11 (Two-response advantage). *For a PPT algorithm $\mathcal{F}^H(pp)$ that receives $pp \leftarrow \mathsf{Setup}(1^\lambda)$ and outputs (x, w, a, c, z_1, z_2), let $\mathrm{Adv}^{\mathrm{2resp}}(\mathcal{F})$ be the probability that $(x, w) \in \mathcal{R}$, $z_1 \neq z_2$, and $V_{\mathrm{int}}(x, a, c, z_1) = V_{\mathrm{int}}(x, a, c, z_2) = 1$.*

Unique accepting responses give $\mathrm{Adv}^{\mathrm{2resp}}(\mathcal{F}) = 0$ for every $\mathcal{F}$. For the representation protocol of Section 6.1, two accepting responses give a nonzero relation among the public bases.

Lemma 7 (Transfer acceptance after forking). *Suppose the protocol admits transfer and the honest anamorphic prover makes at most Q random-oracle queries. Let Acc be honest anamorphic acceptance with output transcript $(a_{\mathrm{out}}, c_{\mathrm{out}}, z_{\mathrm{out}})$, let $c' \xleftarrow{\$} \mathsf{Ch}$ be sampled afterwards, and let*

$$z' = \mathsf{Tr}(x, w, a_{\mathrm{out}}, c_{\mathrm{out}}, z_{\mathrm{out}}, c').$$

Then there is a PPT algorithm $\mathcal{F}$ with

$$\Pr\big[\mathsf{Acc} \wedge \mathsf{Chk}(ck, x, \widehat{x}, a_{\mathrm{out}}, c', z') = 1\big] \geq \frac{\varepsilon^2}{Q+1} - \mathrm{Adv}^{\mathrm{2resp}}(\mathcal{F}).$$

Proof. Use the replay experiment of Bellare and Neven [9, Lemma 1], allowing the forked challenge to equal the original. Put $K = Q + 1$. Use independent uniform answers $c_1, \ldots, c_K$ for first queries at distinct (st, a_i), appending the verification query if needed and leaving unused answers unread. Let R contain pp, the auxiliary setup, sampler and prover coins, and independent coins for lazy sampling of queries outside this context. The execution is determined by $(R, c_1, \ldots, c_K)$. Define $I = 0$ on rejection and, on acceptance, let I index the first query at $(\mathsf{st}, a_{\mathrm{out}})$.

Given pp, $\mathcal{F}$ samples the rest of the execution. If $I = j > 0$, it replays with the same $R, c_1, \ldots, c_{j-1}$ and a fresh independent suffix $c'_j, \ldots, c'_K$. Put $p_j = \Pr[I = j \mid R, c_1, \ldots, c_{j-1}]$. Conditional independence of the two suffixes and Cauchy–Schwarz give

$$\Pr[\mathsf{Fork}] = \sum_{j=1}^{K} \mathbb{E}[p_j^2] \geq \frac{1}{K} \left(\sum_{j=1}^{K} \mathbb{E}[p_j] \right)^2 = \frac{\varepsilon^2}{K},$$

where Fork means that both runs select the same nonzero index j. Their common prefix fixes the same query at index j. On Fork the second run thus outputs an accepting $(a_{\mathrm{out}}, c'_j, z'')$ that passes Chk.

Transfer the first response to c'_j, obtaining z'. If Fork occurs and $z' \neq z''$, $\mathcal{F}$ returns $(x, w, a_{\mathrm{out}}, c'_j, z', z'')$. Otherwise, on Fork the transferred response passes Chk. Since c'_j is uniform and independent of the first run, this gives the claimed bound for the independent challenge c'. □

Theorem 4 (Transfer exclusion). *Under this section's interfaces, suppose the protocol admits transfer. For every honest anamorphic prover making at most Q queries, there exist a PPT checking-side attacker $\mathcal{A}$ and a PPT algorithm $\mathcal{F}$ with*

$$\mathrm{Adv}^{\mathrm{NonEnc}}(\mathcal{A}) \geq \frac{\varepsilon^2}{Q+1} - \mathrm{Adv}^{\mathrm{2resp}}(\mathcal{F}) - \delta_{\mathrm{zk}}.$$

If accepting responses are unique, then $\mathrm{Adv}^{\mathrm{NonEnc}}(\mathcal{A}) \geq \varepsilon^2/(Q+1) - \delta_{\mathrm{zk}}$. The attacker makes no honest-proof queries and at most one real random-oracle query.

Proof. Consider the strong anamorphic zero-knowledge test that retains both witnesses and rejects unless its proof passes normal and anamorphic verification. On acceptance, sample $c' \xleftarrow{\$} \mathsf{Ch}$, compute $z' = \mathsf{Tr}(x, w, a, c, z, c')$, and output $\mathsf{Chk}(ck, x, \widehat{x}, a, c', z')$ without changing the oracle. Lemma 7 bounds its real success. Unique responses give $\mathrm{Adv}^{\mathrm{2resp}}(\mathcal{F}) = 0$. Simulated success is at least real success minus δ_{zk}.

The attacker runs the simulator under an independent local oracle and, after local acceptance, transfers the simulated response to the single real challenge

$H(\mathsf{st}, a)$. Setup and input sampling are independent of the real oracle, and no real query has yet been made, so this challenge is uniform and independent of the simulator computation. The attack therefore has the simulated test's success probability and runs in strict polynomial time. □

For polynomial Q, strong anamorphic zero knowledge and a negligible two-response advantage thus make non-encodability incompatible with non-negligible honest anamorphic acceptance.

6.1 Examples and the Scope of the Exclusion

Schnorr verifies $g^z = aY^c$ for $Y = g^w$ in a prime-order group. Its response $z \in \mathbb{F}_q$ is unique, and $z' = z + (c' - c)w$ transfers it to c' [40,23]. Chaum–Pedersen [20] has the same scalar transfer and uniqueness when its bases have order q. More generally, an injective homomorphism proof $\Phi(z) = aX^c$ satisfies both conditions when its response domain is closed under the corresponding transfer operation [35].

For a surjective representation map $\Phi : \mathbb{F}_q^n \to \mathbb{G}$, where $|\mathbb{G}| = q$ and $n > 1$, each accepting-response fiber has q^{n-1} elements. Transfer still holds by adding $(c' - c)w$, but uniqueness fails. Two accepting responses $z_1 \neq z_2$ at the same (x, a, c) satisfy $\Phi(z_1 - z_2) = 1$. Returning the nonzero vector $z_1 - z_2$ therefore solves the discrete logarithm relation problem whenever $\mathcal{F}$ succeeds. Under DLR, Theorem 4 excludes every oracle-free local check for the uncompressed representation protocol, including checks of the full transcript.

The OR composition of Cramer, Damgård, and Schoenmakers [22] preserves transfer by changing the witnessed branch's challenge share and response while keeping the other branch fixed [23, Section 4]. Proposition 2 therefore excludes first-message checks. In the Schnorr OR example of Appendix D.4, a party choosing both witnesses can produce two responses to the same challenge. Thus a negligible two-response advantage cannot establish a general exclusion for full-transcript OR checks.

For $n = 2$, random-base Compressed Σ sends $v_1 = r + c_1 w$ and is the uncompressed representation protocol above. Its local checks are excluded under DLR, which follows from DL for independent random bases. For $n = 2^d \geq 4$, the normal proof sends a folded response v_d after further challenges. Bulletproofs likewise sends a folded IPA transcript for its range statement. Appendix D.3 gives a four-coordinate Compressed Σ host whose normal and anamorphic claims are both knowledge of a discrete logarithm. Schnorr proves the same normal knowledge claim in three moves and falls under Theorem 4.

The theorem requires an oracle-free Chk. A check that queries the oracle may obtain different answers in the local simulation and in the real execution. Whether such checks can circumvent the exclusion remains open.

7 Conclusion

A proof can certify a second knowledge claim even when an observer can reproduce the normal proving execution from its witness and coins. In our construc-

tions the normal verifier already checks a commitment to prover randomness. Extracting an opening of that commitment in the initial execution certifies the anamorphic claim, while simulation at the prescribed commitment protects both witnesses from the checker.

For a fixed discrete-logarithm claim, the choice of normal protocol matters. The four-coordinate Compressed Σ protocol in Appendix D.3 admits an anamorphic mode. For Schnorr, the checker can forge under any oracle-free local check. A broader characterization of the proof systems that admit such a mode remains open. Another question is whether the checking key can be made compact while retaining security across many fresh labels. Here the checker stores one group element per label, even when the prover derives its masks from a short seed.

References

1. Abdolmaleki, B., Fleischhacker, N., Goyal, V., Jain, A., Malavolta, G.: Steganography-free zero-knowledge. In: TCC 2022, Part I. LNCS, vol. 13747, pp. 143–172. Springer (2022). https://doi.org/10.1007/978-3-031-22318-1_6
2. Abe, M., Balbás, D., Bui, D., Ohkubo, M., Shang, Z., Takahashi, A., Tibouchi, M.: Critical rounds in multi-round proofs: Proof of partial knowledge and trapdoor commitments. Cryptology ePrint Archive, Paper 2024/1766 (2026), https://eprint.iacr.org/2024/1766, revised version of EUROCRYPT 2026, revision dated 23 February 2026
3. Attema, T.: Compressed Σ-Protocol Theory. Ph.D. thesis, Leiden University (2023), https://ir.cwi.nl/pub/33032/33032D.pdf
4. Attema, T., Cramer, R.: Compressed Σ-protocol theory and practical application to plug & play secure algorithmics. In: CRYPTO 2020, Part III. LNCS, vol. 12172, pp. 513–543. Springer (2020). https://doi.org/10.1007/978-3-030-56877-1_18, https://eprint.iacr.org/2020/152, full version available as Cryptology ePrint Archive, Paper 2020/152
5. Attema, T., Fehr, S., Klooß, M., Resch, N.: The Fiat–Shamir transformation of $(\Gamma_1, \ldots, \Gamma_\mu)$-special-sound interactive proofs. Journal of Cryptology **39**, 19 (2026). https://doi.org/10.1007/s00145-025-09569-y
6. Banerjee, S., Pal, T., Rupp, A., Slamanig, D.: Simple asymmetric anamorphic encryption and signature using multi-message extensions. In: CRYPTO 2026, Part X. LNCS, vol. 16809, pp. 432–464. Springer (2026). https://doi.org/10.1007/978-3-032-35428-0_14
7. Banfi, F., Gegier, K., Hirt, M., Maurer, U., Rito, G.: Anamorphic encryption, revisited. In: EUROCRYPT 2024, Part II. LNCS, vol. 14652, pp. 3–32. Springer (2024). https://doi.org/10.1007/978-3-031-58723-8_1
8. Bellare, M., Goldreich, O.: On defining proofs of knowledge. In: CRYPTO '92. LNCS, vol. 740, pp. 390–420. Springer (1993). https://doi.org/10.1007/3-540-48071-4_28
9. Bellare, M., Neven, G.: Multi-signatures in the plain public-key model and a general forking lemma. In: ACM CCS 2006. pp. 390–399. ACM (2006). https://doi.org/10.1145/1180405.1180453
10. Bellare, M., Paterson, K.G., Rogaway, P.: Security of symmetric encryption against mass surveillance. In: CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 1–19. Springer (2014). https://doi.org/10.1007/978-3-662-44371-2_1
11. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security (CCS 1993). pp. 62–73. ACM (1993). https://doi.org/10.1145/168588.168596
12. Biryukov, A., Feher, D., Vitto, G.: Privacy aspects and subliminal channels in Zcash. In: ACM SIGSAC Conference on Computer and Communications Security (CCS 2019). pp. 1813–1830. ACM (2019). https://doi.org/10.1145/3319535.3345663
13. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. In: 2018 IEEE Symposium on Security and Privacy. pp. 315–334. IEEE (2018). https://doi.org/10.1109/SP.2018.00020, https://web.stanford.edu/~buenz/pubs/bulletproofs.pdf, full version

14. Burmester, M., Desmedt, Y.G., Itoh, T., Sakurai, K., Shizuya, H.: Divertible and subliminal-free zero-knowledge proofs for languages. Journal of Cryptology **12**(3), 197–223 (1999). https://doi.org/10.1007/s001459900053
15. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable encryption. In: CRYPTO 1997. LNCS, vol. 1294, pp. 90–104. Springer (1997). https://doi.org/10.1007/BFb0052229
16. Catalano, D., Giunta, E., Migliaro, F.: Anamorphic encryption: New constructions and homomorphic realizations. In: EUROCRYPT 2024, Part II. LNCS, vol. 14652, pp. 33–62. Springer (2024). https://doi.org/10.1007/978-3-031-58723-8_2
17. Catalano, D., Giunta, E., Migliaro, F.: Limits of black-box anamorphic encryption. In: CRYPTO 2024, Part II. LNCS, vol. 14921, pp. 352–383. Springer (2024). https://doi.org/10.1007/978-3-031-68379-4_11
18. Catalano, D., Giunta, E., Migliaro, F.: Generic anamorphic encryption, revisited: New limitations and constructions. In: EUROCRYPT 2025, Part II. LNCS, vol. 15602, pp. 275–303. Springer (2025). https://doi.org/10.1007/978-3-031-91124-8_10
19. Chaidos, P., Couteau, G.: Efficient designated-verifier non-interactive zero-knowledge proofs of knowledge. In: EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 193–221. Springer (2018). https://doi.org/10.1007/978-3-319-78372-7_7
20. Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: CRYPTO '92. LNCS, vol. 740, pp. 89–105. Springer (1993). https://doi.org/10.1007/3-540-48071-4_7
21. Choi, W., Collins, D., Liu, X., Stracovsky, R., Zikas, V.: A unified treatment of anamorphic encryption. In: CRYPTO 2026, Part X. LNCS, vol. 16809, pp. 465–496. Springer (2026). https://doi.org/10.1007/978-3-032-35428-0_15
22. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: CRYPTO '94. LNCS, vol. 839, pp. 174–187. Springer (1994). https://doi.org/10.1007/3-540-48658-5_19
23. Damgård, I.: On Σ-protocols. CPT 2010, version 2 (2010), https://www.cs.au.dk/~ivan/Sigma.pdf
24. Dao, Q., Grubbs, P.: Spartan and Bulletproofs are simulation-extractable (for free!). In: EUROCRYPT 2023, Part II. LNCS, vol. 14005, pp. 531–562. Springer (2023). https://doi.org/10.1007/978-3-031-30617-4_18
25. Dedić, N., Itkis, G., Reyzin, L., Russell, S.: Upper and lower bounds on black-box steganography. Journal of Cryptology **22**(3), 365–394 (2009). https://doi.org/10.1007/s00145-008-9020-3
26. Deo, A., Libert, B.: Anamorphic signatures with dictator and recipient unforgeability for long messages. In: ASIACRYPT 2025, Part VI. LNCS, vol. 16250, pp. 370–401. Springer (2025). https://doi.org/10.1007/978-981-95-5119-4_12
27. Deo, A., Libert, B.: Fully asymmetric anamorphic homomorphic encryption from LWE. In: EUROCRYPT 2026, Part V. LNCS, vol. 16545, pp. 609–639. Springer (2026). https://doi.org/10.1007/978-3-032-25330-9_21
28. Dodis, Y., Goldin, E.: Anamorphic-resistant encryption: Or why the encryption debate is still alive. In: CRYPTO 2025, Part III. LNCS, vol. 16002, pp. 440–471. Springer (2025). https://doi.org/10.1007/978-3-032-01881-6_14
29. Feng, H., Tang, Q.: Witness authenticating NIZKs and applications. In: CRYPTO 2021, Part I. LNCS, vol. 12828, pp. 3–33. Springer (2021). https://doi.org/10.1007/978-3-030-84259-8_1

30. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer (1987). https://doi.org/10.1007/3-540-47721-7_12
31. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. Journal of the ACM **33**(4), 792–807 (1986). https://doi.org/10.1145/6490.6503
32. Jaeger, J., Stracovsky, R.: Dictators? Friends? Forgers. Breaking and fixing unforgeability definitions for anamorphic signature schemes. In: ASIACRYPT 2024, Part II. LNCS, vol. 15485, pp. 105–137. Springer (2024). https://doi.org/10.1007/978-981-96-0888-1_4, https://eprint.iacr.org/2025/2123, revised version: Cryptology ePrint Archive, Paper 2025/2123 (2025)
33. Jakobsson, M., Sako, K., Impagliazzo, R.: Designated verifier proofs and their applications. In: EUROCRYPT 1996. LNCS, vol. 1070, pp. 143–154. Springer (1996). https://doi.org/10.1007/3-540-68339-9_13
34. Kutyłowski, M., Persiano, G., Phan, D.H., Yung, M., Zawada, M.: Anamorphic signatures: Secrecy from a dictator who only permits authentication! In: CRYPTO 2023, Part II. LNCS, vol. 14082, pp. 759–790. Springer (2023). https://doi.org/10.1007/978-3-031-38545-2_25
35. Maurer, U.: Zero-knowledge proofs of knowledge for group homomorphisms. Designs, Codes and Cryptography **77**, 663–676 (2015), https://crypto.ethz.ch/publications/Maurer15.html
36. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer (1992). https://doi.org/10.1007/3-540-46766-1_9
37. Persiano, G., Phan, D.H., Yung, M.: Anamorphic encryption: Private communication against a dictator. In: EUROCRYPT 2022, Part II. LNCS, vol. 13276, pp. 34–63. Springer (2022). https://doi.org/10.1007/978-3-031-07085-3_2, https://eprint.iacr.org/2022/639
38. Persiano, G., Phan, D.H., Yung, M.: Public-key anamorphism in (CCA-secure) public-key encryption and beyond. In: CRYPTO 2024, Part II. LNCS, vol. 14921, pp. 422–455. Springer (2024). https://doi.org/10.1007/978-3-031-68379-4_13
39. Quach, W., Tyner, L., Wichs, D.: Unique NIZKs and steganography detection. In: EUROCRYPT 2025, Part IV. LNCS, vol. 15604, pp. 34–63. Springer (2025). https://doi.org/10.1007/978-3-031-91134-7_2
40. Schnorr, C.P.: Efficient identification and signatures for smart cards. In: CRYPTO '89. LNCS, vol. 435, pp. 239–252. Springer (1990). https://doi.org/10.1007/0-387-34805-0_22
41. Simmons, G.J.: The prisoners' problem and the subliminal channel. In: CRYPTO 1983. pp. 51–67. Plenum Press (1984). https://doi.org/10.1007/978-1-4684-4730-9_5
42. Tanishita, Y., Matsuda, T., Matsuura, K.: Revisiting security definitions of sender-anamorphic encryption. In: PKC 2026, Part IV. LNCS, vol. 16554, pp. 130–161. Springer (2026). https://doi.org/10.1007/978-3-032-26740-5_5
43. The secp256k1-zkp contributors: secp256k1-zkp: Bulletproofs module (Grin fork). Source code (2023), https://github.com/mimblewimble/secp256k1-zkp/blob/b247e1ec8ed62b9abf123dc83189d253e17d488d/include/secp256k1_bulletproofs.h, revision b247e1ec8ed6
44. Wang, W., Zheng, Y., Xu, P., Wang, W., Chen, R., Yang, Y., Yung, M.: Anamorphic commitment: Robust privacy under coercion and parameter tampering. Science China Information Sciences **69**(5), 152107 (2026). https://doi.org/10.1007/s11432-025-4751-y

45. Young, A., Yung, M.: The dark side of "black-box" cryptography or: Should we trust Capstone? In: CRYPTO 1996. LNCS, vol. 1109, pp. 89–103. Springer (1996). https://doi.org/10.1007/3-540-68697-5_8

Supplementary Material

A Complete Anamorphic Bulletproofs

We expand the fixed normal Fiat–Shamir protocol used in Construction 2, including its recursive IPA [13, Protocols 1–3 and Section 4]. Parameters and challenge domains are those of Section 5. In the random-base instantiation, group generation is followed by independent uniform sampling of the listed bases in $\mathbb{G} \setminus \{1\}$. A fixed hash-derived base distribution is covered under DLR for its output together with the public generation data.

Use the auxiliary algorithms of Construction 2 and the request rule of Section 2. After consuming a valid fresh index, obtain $\omega \leftarrow \mathsf{ACoins}(ek, i, I)$ and run the range prover below. In the normal mode, sample the same tape uniformly from $\mathbb{F}_q^{2s+1} \times \mathbb{F}_q^3$.

A.1 Range Proving

The range prover uses $\mathsf{st} = (pp, V)$ and the notation $\mathbf{y}, \mathbf{p}, \mathbf{h}'$ of Section 5. Its IPA subroutine is given in Appendix A.2.

$\underline{\mathsf{P}^{\mathbf{H}}_{\mathrm{BP}}(pp, V, (v, \gamma); \omega)}$
Input tape $\omega = ((\rho, \mathbf{s}_L, \mathbf{s}_R), (\alpha, \tau_1, \tau_2))$.

1: Set $\mathbf{a}_L = (\lfloor v/2^{j-1} \rfloor \bmod 2)_{j=1}^{s}$ and $\mathbf{a}_R = \mathbf{a}_L - \mathbf{1}$. Compute $A = h_0^{\alpha}\mathbf{g}^{\mathbf{a}_L}\mathbf{h}^{\mathbf{a}_R}$ and $S = h_0^{\rho}\mathbf{g}^{\mathbf{s}_L}\mathbf{h}^{\mathbf{s}_R}$, then query $y = H_y(\mathsf{st}, A, S)$ and $z = H_z(\mathsf{st}, A, S, y)$.

2: Compute $\mathbf{y}, \mathbf{p}, \mathbf{h}'$ and set

$$\mathbf{l}_0 = \mathbf{a}_L - z\mathbf{1}, \qquad \mathbf{r}_1 = \mathbf{y} \circ \mathbf{s}_R, \qquad \mathbf{r}_0 = \mathbf{y} \circ (\mathbf{a}_R + z\mathbf{1}) + z^2\mathbf{p}.$$

Set $t_1 = \langle \mathbf{l}_0, \mathbf{r}_1 \rangle + \langle \mathbf{s}_L, \mathbf{r}_0 \rangle$, $t_2 = \langle \mathbf{s}_L, \mathbf{r}_1 \rangle$, and $T_j = g_0^{t_j} h_0^{\tau_j}$ for $j = 1, 2$. Query $x = H_x(\mathsf{st}, A, S, y, z, T_1, T_2)$.

3: Compute

$$\begin{aligned} &\mathbf{l}(x) = \mathbf{a}_L - z\mathbf{1} + x\mathbf{s}_L, \quad \mathbf{r}(x) = \mathbf{y} \circ (\mathbf{a}_R + z\mathbf{1} + x\mathbf{s}_R) + z^2\mathbf{p}, \\ &\widehat{t} = \langle \mathbf{l}(x), \mathbf{r}(x) \rangle, \qquad \mu = \alpha + x\rho, \qquad \tau_x = z^2\gamma + x\tau_1 + x^2\tau_2. \end{aligned} \tag{14}$$

4: Set $\mathsf{tr}_0 = (\mathsf{st}, A, S, y, z, T_1, T_2, x, \tau_x, \mu, \widehat{t})$.
Run $\mathsf{IPA.Prove}^{\mathbf{H}}(pp, \mathsf{tr}_0, \mathbf{g}, \mathbf{h}', \mathbf{l}(x), \mathbf{r}(x))$
Obtain $\pi_{\mathrm{IP}} = ((\mathsf{L}_j, \mathsf{R}_j)_{j=1}^{L}, a_{\mathrm{IP}}, b_{\mathrm{IP}})$.
Return $\pi = (A, S, T_1, T_2, \tau_x, \mu, \widehat{t}, (\mathsf{L}_j, \mathsf{R}_j)_{j=1}^{L}, a_{\mathrm{IP}}, b_{\mathrm{IP}})$.

A.2 The Normal Inner-Product Prover

Each recursive level halves the vector dimension. Subscripts L, R refer to the current halves, and t_j records the outer transcript and the first j IPA rounds.

$\underline{\mathsf{IPA.Prove}^{\mathbf{H}}(pp, \mathsf{tr}_0, \mathbf{g}, \mathbf{h}', \mathbf{l}, \mathbf{r})}$

1: Query $\xi_0 = H_{\mathrm{IP},0}(\mathsf{tr}_0)$. Set $u = u_{\mathsf{IP}}^{\xi_0}$, $\mathsf{t}_0 = (\mathsf{tr}_0, \xi_0)$, and $(\mathbf{a}, \mathbf{b}, \mathbf{G}, \mathbf{K}) = (\mathbf{l}, \mathbf{r}, \mathbf{g}, \mathbf{h}')$.

2: For $j = 1, \ldots, L$, compute

$$\mathsf{L}_j = \mathbf{G}_R^{\mathbf{a}_L} \mathbf{K}_L^{\mathbf{b}_R} u^{\langle \mathbf{a}_L, \mathbf{b}_R \rangle}, \qquad \mathsf{R}_j = \mathbf{G}_L^{\mathbf{a}_R} \mathbf{K}_R^{\mathbf{b}_L} u^{\langle \mathbf{a}_R, \mathbf{b}_L \rangle},$$

and query $\xi_j = H_{\mathrm{IP},j}(\mathsf{t}_{j-1}, \mathsf{L}_j, \mathsf{R}_j)$. Update

$$\mathbf{a}' = \xi_j \mathbf{a}_L + \xi_j^{-1} \mathbf{a}_R, \quad \mathbf{b}' = \xi_j^{-1} \mathbf{b}_L + \xi_j \mathbf{b}_R,$$
$$\mathbf{G}' = \mathbf{G}_L^{\xi_j^{-1}} \circ \mathbf{G}_R^{\xi_j}, \qquad \mathbf{K}' = \mathbf{K}_L^{\xi_j} \circ \mathbf{K}_R^{\xi_j^{-1}}.$$

Replace the vectors by their primed values and set $\mathsf{t}_j = (\mathsf{t}_{j-1}, \mathsf{L}_j, \mathsf{R}_j, \xi_j)$.

3: Return $\pi_{\mathrm{IP}} = ((\mathsf{L}_j, \mathsf{R}_j)_{j=1}^{L}, a_1, b_1)$, where $a_{\mathrm{IP}} = a_1$ and $b_{\mathrm{IP}} = b_1$.

A.3 Normal and Anamorphic Verification

The same normal verifier serves both modes. The designated checker runs $\mathsf{AVerify}$ from Construction 2 and tests the registered relation $S = \widehat{S}M_i$.

$\underline{\mathsf{V}_{\mathrm{BP}}^{\mathbf{H}}(pp, V, \pi)}$

1: Parse $\pi = (A, S, T_1, T_2, \tau_x, \mu, \widehat{t}, (\mathsf{L}_j, \mathsf{R}_j)_{j=1}^{L}, a_{\mathrm{IP}}, b_{\mathrm{IP}})$. Reject noncanonical encodings, incorrect lengths, or failed group and field membership checks on V and π.

2: Set $\mathsf{st} = (pp, V)$ and recompute $y = H_y(\mathsf{st}, A, S)$, $z = H_z(\mathsf{st}, A, S, y)$, and $x = H_x(\mathsf{st}, A, S, y, z, T_1, T_2)$. Compute $\mathbf{y}, \mathbf{p}, \mathbf{h}'$ and $\delta(y, z)$ as in Section 5. Reject unless

$$g_0^{\widehat{t}} h_0^{\tau_x} = V^{z^2} g_0^{\delta(y,z)} T_1^x T_2^{x^2}.$$

3: Compute

$$P_{\mathrm{IP}} = A S^x \mathbf{g}^{-z\mathbf{1}} (\mathbf{h}')^{z\mathbf{y} + z^2 \mathbf{p}} h_0^{-\mu}.$$

Reconstruct tr_0 as in the range prover. Set $\xi_0 = H_{\mathrm{IP},0}(\mathsf{tr}_0)$, $u = u_{\mathsf{IP}}^{\xi_0}$, $\mathsf{t}_0 = (\mathsf{tr}_0, \xi_0)$, $P = P_{\mathrm{IP}} u^{\widehat{t}}$, $\mathbf{G} = \mathbf{g}$, and $\mathbf{K} = \mathbf{h}'$.

4: For $j = 1, \ldots, L$, query $\xi_j = H_{\mathrm{IP},j}(\mathsf{t}_{j-1}, \mathsf{L}_j, \mathsf{R}_j)$ and update

$$P \leftarrow \mathsf{L}_j^{\xi_j^2} P \mathsf{R}_j^{\xi_j^{-2}},$$
$$\mathbf{G}' = \mathbf{G}_L^{\xi_j^{-1}} \circ \mathbf{G}_R^{\xi_j}, \qquad \mathbf{K}' = \mathbf{K}_L^{\xi_j} \circ \mathbf{K}_R^{\xi_j^{-1}}.$$

Protocol (continued)
Replace the bases by their primed values and set $\mathsf{t}_j = (\mathsf{t}_{j-1}, \mathsf{L}_j, \mathsf{R}_j, \xi_j)$. 5: Return $[P = G_1^{a_{\text{IP}}} K_1^{b_{\text{IP}}} u^{a_{\text{IP}} b_{\text{IP}}}]$.

B Supplementary Extraction and Reduction Proofs

B.1 Proof of the IPA Extraction Interface

We give the full coordinate argument for Proposition 1 and make relation lifting to the original public bases explicit.

Proof of Proposition 1. Use the recursive verification equation of Appendix A.3. At a recursive node, write the current bases as $(\mathbf{g}, \mathbf{h}, u)$ and the target as P. Initially $\mathbf{h} = \mathbf{h}'$, $u = u_{\mathsf{IP}}^{\xi_0}$, and $P = P_{\text{IP}} u^{\widehat{t}}$.

Suppose four child subtrees return accepting vector pairs $(\mathbf{a}'_i, \mathbf{b}'_i)$; return immediately if any child instead produces a parameter relation. Let $X_i = \xi_i^2$, with the four X_i distinct. Expanding the folded bases gives the representation

$$\mathbf{v}_i = (\xi_i^{-1}\mathbf{a}'_i,\ \xi_i \mathbf{a}'_i,\ \xi_i \mathbf{b}'_i,\ \xi_i^{-1}\mathbf{b}'_i,\ \langle \mathbf{a}'_i, \mathbf{b}'_i \rangle)$$

of $\mathsf{L}^{X_i} P \mathsf{R}^{X_i^{-1}}$ on $(\mathbf{g}_L, \mathbf{g}_R, \mathbf{h}_L, \mathbf{h}_R, u)$.

The matrix with rows $(X_i, 1, X_i^{-1})$ for $i = 1, 2, 3$ is invertible, since multiplying each row by X_i gives a quadratic Vandermonde matrix. Hence it recovers representations $Z = \mathbf{g}^{\mathbf{a}^Z} \mathbf{h}^{\mathbf{b}^Z} u^{c^Z}, Z \in \{\mathsf{L}, P, \mathsf{R}\}$. Compare every child with the corresponding linear combination of these three representations. Any disagreement gives a nonzero relation on the current bases.

Otherwise all four branches agree coordinatewise. On the $\mathbf{g}$ coordinates,

$$\xi^{-1}\mathbf{a}' = \mathbf{a}_L^{\mathsf{L}} X + \mathbf{a}_L^P + \mathbf{a}_L^{\mathsf{R}} X^{-1}, \quad \xi \mathbf{a}' = \mathbf{a}_R^{\mathsf{L}} X + \mathbf{a}_R^P + \mathbf{a}_R^{\mathsf{R}} X^{-1}.$$

Eliminating $\mathbf{a}'$ gives (12). Its four distinct roots imply $\mathbf{a}_L^{\mathsf{L}} = \mathbf{0}, \mathbf{a}_R^{\mathsf{L}} = \mathbf{a}_L^P, \mathbf{a}_L^{\mathsf{R}} = \mathbf{a}_R^P, \mathbf{a}_R^{\mathsf{R}} = \mathbf{0}$. The $\mathbf{h}$ coordinates satisfy

$$\xi \mathbf{b}' = \mathbf{b}_L^{\mathsf{L}} X + \mathbf{b}_L^P + \mathbf{b}_L^{\mathsf{R}} X^{-1}, \quad \xi^{-1}\mathbf{b}' = \mathbf{b}_R^{\mathsf{L}} X + \mathbf{b}_R^P + \mathbf{b}_R^{\mathsf{R}} X^{-1}.$$

Eliminating $\mathbf{b}'$ yields $\mathbf{b}_R^{\mathsf{L}} X^3 + (\mathbf{b}_R^P - \mathbf{b}_L^{\mathsf{L}}) X^2 + (\mathbf{b}_R^{\mathsf{R}} - \mathbf{b}_L^P) X - \mathbf{b}_L^{\mathsf{R}} = \mathbf{0}$.

Thus $\mathbf{b}_R^{\mathsf{L}} = \mathbf{0}, \mathbf{b}_L^{\mathsf{L}} = \mathbf{b}_R^P, \mathbf{b}_R^{\mathsf{R}} = \mathbf{b}_L^P, \mathbf{b}_L^{\mathsf{R}} = \mathbf{0}$. Every child witness is therefore the normal fold of $(\mathbf{a}^P, \mathbf{b}^P)$:

$$\mathbf{a}' = \xi \mathbf{a}_L^P + \xi^{-1} \mathbf{a}_R^P, \qquad \mathbf{b}' = \xi^{-1} \mathbf{b}_L^P + \xi \mathbf{b}_R^P.$$

The scalar coordinate gives $\langle \mathbf{a}', \mathbf{b}' \rangle = c^{\mathsf{L}} X + c^P + c^{\mathsf{R}} X^{-1}$. Substituting the folded vectors and multiplying by X gives a quadratic vanishing at the four X_i, whose coefficients imply $c^{\mathsf{L}} = \langle \mathbf{a}_L^P, \mathbf{b}_R^P \rangle, c^{\mathsf{R}} = \langle \mathbf{a}_R^P, \mathbf{b}_L^P \rangle, c^P = \langle \mathbf{a}^P, \mathbf{b}^P \rangle$. Hence $(\mathbf{a}^P, \mathbf{b}^P)$ is a valid parent witness. Recursing reaches the root.

For an initial challenge ξ_0, the recovered root witness satisfies

$$P_{\mathrm{IP}} = \mathbf{g}^{\mathbf{a}}(\mathbf{h}')^{\mathbf{b}} u_{\mathsf{IP}}^{\xi_0(\langle \mathbf{a},\mathbf{b}\rangle - \widehat{t})}.$$

Two distinct initial challenges give two such representations. If the recovered vector pairs differ, dividing the representations gives a nonzero relation. If they agree, nonidentity u_{IP} and $\xi_0 \neq \xi_0'$ force $\langle \mathbf{a}, \mathbf{b}\rangle = \widehat{t}$.

Finally, a relation (v_g, v_h, v_u) on folded bases lifts one level to

$$(\xi^{-1} v_g, \xi v_g, \xi v_h, \xi^{-1} v_h, v_u)$$

on the parent bases. This map is injective because $\xi \neq 0$ and the coordinate blocks are disjoint. Repeated lifting preserves nonzero relations. The transformation $\mathbf{h} \mapsto \mathbf{h}'$ and the rescaling $u = u_{\mathsf{IP}}^{\xi_0}$ are also nonzero coordinate scalings. Padding zeros on g_0, h_0 gives a nonzero relation on the original ordered bases $\mathbf{B}$. □

B.2 Normal Bulletproofs Range Extraction

Lemma 8 (The full range tree). *An accepting tree with s distinct y challenges, three distinct z challenges at each y, three distinct x challenges at each (y, z), two distinct initial IPA challenges, and four challenges with pairwise distinct squares at every IPA node yields a valid range witness for V or a nonzero DLR relation.*

Proof. Apply Proposition 1 at every outer branch. Return any relation it produces. Otherwise each outer leaf gives $(\mathbf{l}_i, \mathbf{r}_i)$ satisfying both the representation equation and $\widehat{t}_i = \langle \mathbf{l}_i, \mathbf{r}_i\rangle$.

Fix (y, z) and two distinct x_1, x_2. Lemma 6 extracts an opening $(\rho, \mathbf{s}_L, \mathbf{s}_R)$ of S. The first branch then gives an opening of A:

$$\alpha = \mu_1 - x_1\rho, \quad \mathbf{a}_L = \mathbf{l}_1 + z\mathbf{1} - x_1\mathbf{s}_L, \quad \mathbf{a}_R = \mathbf{y}^{-1} \circ (\mathbf{r}_1 - z^2\mathbf{p}) - z\mathbf{1} - x_1\mathbf{s}_R.$$

If different branch pairs produce different formal coefficient vectors for A or S, their difference gives a nonzero relation on $(h_0, \mathbf{g}, \mathbf{h})$ and hence on $\mathbf{B}$. Otherwise retain the common openings.

At every remaining branch, compare its representation of AS^x with that obtained from the common openings. Any difference gives a nonzero DLR relation. Hence, absent a parameter relation, all recovered vectors satisfy (14) with common $\mathbf{a}_L, \mathbf{a}_R, \mathbf{s}_L, \mathbf{s}_R, \alpha, \rho$.

For fixed (y, z), the three scalar equations are

$$g_0^{\widehat{t}_i - \delta(y,z)} h_0^{\tau_{x_i}} = V^{z^2} T_1^{x_i} T_2^{x_i^2}, \qquad i = 1, 2, 3.$$

Invert the matrix with rows $(1, x_i, x_i^2)$. Applying the inverse matrix to the known exponent pairs opens V^{z^2}, T_1, T_2. Since $z \neq 0$, the first opens V as (v, γ). Different recovered openings of V give a nonzero relation on (g_0, h_0); otherwise retain the common (v, γ) across all (y, z). For each fixed (y, z),

$$\widehat{t}_i = \delta(y, z) + z^2 v + t_1 x_i + t_2 x_i^2.$$

The recovered quadratic and the inner-product polynomial determined by the common vector openings agree at three distinct points, so their constant coefficients agree. Cancelling the public correction gives

$$\langle \mathbf{a}_L \circ \mathbf{a}_R, \mathbf{y} \rangle + z\langle \mathbf{a}_L - \mathbf{a}_R - \mathbf{1}, \mathbf{y} \rangle + z^2(\langle \mathbf{a}_L, \mathbf{p} \rangle - v) = 0.$$

For every selected y, three distinct z values force all coefficients of this quadratic to vanish. The first two coefficients, as polynomials in y, have degree at most $s-1$ and vanish at s distinct points. Hence $\mathbf{a}_L \circ \mathbf{a}_R = \mathbf{0}$ and $\mathbf{a}_R = \mathbf{a}_L - \mathbf{1}$, while $v = \langle \mathbf{a}_L, \mathbf{p} \rangle$. Thus every coordinate of $\mathbf{a}_L$ is in $\{0, 1\}$. Since $2^s < q$, its integer value lies in $[0, 2^s)$ without wraparound. Together with $V = g_0^v h_0^\gamma$, this gives the required witness. □

For the AFKR bound (1), let $p = q - 1$. The full range tree has

$$\kappa_{\text{full}} = 1 - \left(1 - \frac{s-1}{p}\right)\left(1 - \frac{2}{p}\right)^2\left(1 - \frac{1}{p}\right)\left(1 - \frac{6}{p}\right)^L \le \frac{s + 4 + 6L}{q-1},$$

and $\mathcal{T}_{\text{full}} = 48(s+1)5^L$. These are the full-range values in Table 1 and give normal knowledge soundness for the fixed Fiat–Shamir protocol under DLR.

B.3 Initial Execution and the Fiat–Shamir Adapter

We use AFKR's recursive extractor $\mathcal{F}_\mu^{\mathcal{A}}$ for μ challenge rounds [5, Sections 5.2 and 6.4]. Fix the prover's private tape and its oracle-independent auxiliary input. Treat oracle-dependent preprocessing and final verification as part of $\mathcal{A}$. Its first invocation uses uniform random oracles; rejection of that invocation causes extraction failure, and every returned tree retains its first prover message. In the adaptive transform, this message also contains the normal statement.

These claims follow by induction on μ. For one round, AFKR first runs the prover under uniform H_1, rejects if that run is rejected, and otherwise searches for accepting responses at the recorded message a_1. For $\mu > 1$, fix H_1 and combine the first two messages in $\mathcal{A}[H_1]$, whose verification fixes the first challenge to $H_1(a_1)$. The recursion is $\mathcal{B}^{\mathcal{A}}(H_1) = \mathcal{F}_{\mu-1}^{\mathcal{A}[H_1]}$, $\mathcal{F}_\mu^{\mathcal{A}} = \mathcal{F}_1^{\mathcal{B}^{\mathcal{A}}}$. The outer extractor starts with uniform H_1. By induction, the inner extractor's first invocation uses independent uniform remaining oracles, rejects if that invocation rejects, and preserves its initial pair (a_1, a_2). The outer search retains a_1.

With the access of Section 2, initialize this recursion from the recorded first output and query trace. Extend tables by lazy sampling and use the restart interface on the same private tape for subsequent calls. This resumes the recursion after its initial invocation without changing its probability or cost. For Compressed Σ, the retained prefix is (pp, P, A_0).

For Bulletproofs, number the oracles in transcript order. The first three are H_y, H_z, and H_x. Put $H_{4+j} = H_{\text{IP},j}$ for $0 \le j \le L$. The normalized messages are

$$\begin{aligned} &\widetilde{a}_1 = (\mathsf{st}, A, S), \quad \widetilde{a}_2 = y, \quad \widetilde{a}_3 = (z, T_1, T_2), \\ &\widetilde{a}_4 = (x, \tau_x, \mu, \widehat{t}), \quad \widetilde{a}_{4+j} = (\xi_{j-1}, \mathsf{L}_j, \mathsf{R}_j), \quad 1 \le j \le L. \end{aligned} \tag{15}$$

The outer challenges y, z, x precede the IPA challenges $\xi_0, \ldots, \xi_L$. The prover's final response is $(a_{\mathrm{IP}}, b_{\mathrm{IP}})$. The virtual verifier checks echoed challenges and applies the normal scalar and IPA checks. Including the statement in $\widetilde{a}_1$ gives AFKR's adaptive form [5, Remark 13]. For each well-typed virtual prefix, reconstruct the corresponding normal-oracle input by

$$\begin{aligned} E_1(\widetilde{a}_1) &= (\mathsf{st}, A, S), & E_2(\widetilde{a}_{\le 2}) &= (\mathsf{st}, A, S, y), \\ E_3(\widetilde{a}_{\le 3}) &= (\mathsf{st}, A, S, y, z, T_1, T_2), & E_4(\widetilde{a}_{\le 4}) &= \mathsf{tr}_0, \\ E_{4+j}(\widetilde{a}_{\le 4+j}) &= (\mathsf{t}_{j-1}, \mathsf{L}_j, \mathsf{R}_j), & & 1 \le j \le L. \end{aligned}$$

The transcripts $\mathsf{tr}_0, \mathsf{t}_j$ are defined in Appendix A. Canonical encoding makes each E_k a polynomial-time bijection onto the corresponding well-typed oracle inputs, with inverse obtained by separating the fields in (15).

Define $\widetilde{H}_k(\widetilde{a}_{\le k}) = H_k(E_k(\widetilde{a}_{\le k}))$. A query adapter replaces $H_k(I)$ by $\widetilde{H}_k(E_k^{-1}(I))$ on this image. Queries outside the image retain their independent values and count toward the budget. The adapter completes normal verification, obtains missing challenges, and returns virtual messages with consistent echoes. Thus it preserves acceptance and maps each encoded query to one query. Final verification adds at most $L + 4$ queries.

During replay, supply the restart interface for the normal protocol with $\widetilde{F}_k(I) = \widetilde{F}_k(E_k^{-1}(I))$ on encoded inputs, using the real oracle outside the image. Translate the recorded initial trace by the same maps. A virtual accepting tree projects to a normal tree with identical branch sets. The initial-execution argument fixes its prefix to the recorded (pp, V, A, S), so Lemma 6 opens that execution's S or returns a nonzero DLR relation.

At an IPA node with collected challenges C, sample uniformly from

$$\{\xi \in \mathbb{F}_q^* : \xi^2 \notin \{c^2 : c \in C\}\}.$$

Before four distinct squared challenge values are collected, rejection sampling excludes at most six elements and needs at most $(q-1)/(q-7)$ trials in expectation. The remaining carrier rounds require one or two distinct challenges. These samplers meet AFKR Theorem 5 [5]. Applying (1) to (13) gives $\kappa_{\mathrm{BP}}^{\mathrm{FS}}(Q)$ and $L_{\mathrm{BP},Q}$.

B.4 Concrete Bounds and the Reserved-Index Reduction

The proof of Theorem 1 gives two reductions. Here we specify the checking-side wrapper and truncation. Let N_0, A_0, E denote initial normal acceptance, initial anamorphic acceptance, and successful carrier-opening extraction. The reduction $\mathcal{B}_{\mathrm{ext}}$ generates the actual auxiliary keys and returns any parameter relation obtained on $E \cap A_0$.

Let R_{ext} bound its expected total running time, including initial execution, verification, extraction, and oracle implementation. For an integer cutoff $b \ge 1$, let $K_{\le b}$ denote successful anamorphic knowledge extraction by that time and $\mathcal{B}_{\mathrm{ext},\le b}$ the truncated relation reduction. Partitioning (2) into an opening, a

relation, and timeout gives the concrete bound (3) in the main text. The timeout term follows from Markov's inequality.

For non-encodability, work after the PRF replacement and let $\varepsilon_{\mathrm{rnd}}$ denote attack success with independent masks. The following reduction receives the joint-hardness challenge M.

$\underline{\mathcal{B}_{\mathrm{ne}}(pp, M)}$

1: Sample $J \xleftarrow{\$} [T]$, set $M_J = M$, and generate independent $\eta_i \xleftarrow{\$} U$ and $M_i = \Phi(\eta_i)$ for $i \neq J$. Set $ck = (M_i)_{i=1}^{T}$.
2: Apply the initial-carrier extractor to the following fixed wrapper $W^{\mathbf{H}}$, including the attacker's and honest prover's randomness in its replay tape.
Each run initializes $\mathsf{Used} = \varnothing$ and executes the attacker on (pp, ck). On a request (i, I), return $\perp$ without changing Used if the request is invalid or the index is unavailable. For a valid fresh request at J, abort.
Otherwise consume i, sample ω_{rest} as in the normal protocol, and return the normal prover's output on $(\widehat{u} + \eta_i, \omega_{\mathrm{rest}})$.
For final output $(j, x, \widehat{C}, \widehat{u}, \pi)$, return $\perp_{\mathrm{normal}}$ unless $j = J$, $j \notin \mathsf{Used}$, $\Phi(\widehat{u}) = \widehat{C}$, and anamorphic verification accepts. Otherwise return (x, π, d) with $d = (\widehat{C}, \widehat{u})$.
3: If extraction returns a valid parameter relation, output it. If it returns an opening r^* of the recorded carrier C_0, return $r^* - \widehat{u}_0$, using the initial declaration $d_0 = (\widehat{C}_0, \widehat{u}_0)$. Otherwise fail.

Here $\perp_{\mathrm{normal}}$ is a rejecting normal-protocol output. All requests and outputs undergo the syntax and membership checks of Section 2.

The checking table has the independent-mask distribution. In the coupling of Theorem 1, a successful trajectory at index j is preserved exactly when $J = j$. Invalid requests at J leave the coupling unchanged because validity is checked before the reserved-index test. Every replay retains the index, checking table, known masks, and random tape, and reinitializes Used. The declaration fixes $(\widehat{C}_0, \widehat{u}_0)$ to the initial run, giving $\Phi(r^* - \widehat{u}_0) = M$ on an opening. Let Q_{all} bound the number of queries in one complete wrapped execution in any replay, including honest-proof responses and final verification. The untruncated reduction succeeds with probability at least $(\varepsilon_{\mathrm{rnd}}/T - \kappa(Q_{\mathrm{all}}))/\beta$. If R_{ne} bounds its expected total running time, cutting off the whole reduction at b loses at most R_{ne}/b. Rearranging and including the PRF replacement gives the concrete bound (4) in the main text. Honest-proof responses use known masks and the normal prover, with no simulation error.

Both reductions have expected polynomial total running time. If an untruncated reduction succeeds with inverse-polynomial probability on infinitely many security parameters, a sufficiently large polynomial cutoff makes the timeout probability less than half its success. This yields a strict PPT attack on joint

hardness. Hence the untruncated relation probability is negligible, giving the claimed extraction and non-encodability bounds.

C Aggregation

We use the normal unpadded aggregate protocol and the notation $D = mn = 2^L$, $\mathbf{c}_z, \delta_{\mathrm{agg}}$ of Section 5, with $2^n < q$ and $\max\{D, m+2\} \le q-1$. The normal input is $(V_1, \dots, V_m)$, where $V_j = g_0^{v_j} h_0^{\gamma_j}$ and $0 \le v_j < 2^n$. Concatenate the bit vectors into $\mathbf{a}_L$ and set $\mathbf{a}_R = \mathbf{a}_L - \mathbf{1}_D$. The normal protocol [13, Section 4.3] uses

$$\mathbf{l}(X) = \mathbf{a}_L - z\mathbf{1}_D + X\mathbf{s}_L, \mathbf{r}(X) = \mathbf{y}_D \circ (\mathbf{a}_R + z\mathbf{1}_D + X\mathbf{s}_R) + \mathbf{c}_z,$$

and $\tau_x = x\tau_1 + x^2\tau_2 + \sum_{j=1}^{m} z^{j+1}\gamma_j$. The inner product has constant term $\sum_j z^{j+1}v_j + \delta_{\mathrm{agg}}(y,z)$, yielding the scalar verification equation

$$g_0^{\widehat{t}} h_0^{\tau_x} = (\prod_{j=1}^{m} V_j^{z^{j+1}}) g_0^{\delta_{\mathrm{agg}}(y,z)} T_1^x T_2^{x^2}.$$

All Fiat–Shamir prefixes contain the entire normal commitment vector. The carrier and $\mu = \alpha + x\rho$ retain their forms.

Normal range-witness extraction. Use D values of y, $m+2$ values of z, and three values of x [13, Appendix C]. The argument of Lemma 8 gives common openings of A, S or a nonzero DLR relation. Interpolating the scalar checks in x opens $\prod_j V_j^{z^{j+1}}$. At m distinct nonzero z values, the matrix $(z_i^{j+1})_{i,j=1}^{m}$ is invertible: factoring z_i^2 from each row leaves a Vandermonde matrix. This opens every V_j.

Compare all other recovered representations with these openings and return a nonzero DLR relation if they differ. The remaining constant-term identity has degree at most $m+1$ in z. Its $m+2$ roots give the per-value equations and, using D roots in y, the coordinatewise bit constraints. Since $2^n < q$, each value lies in the required integer range. □

The full aggregate tree has search parameter $12(D+1)(m+3)5^L$ and knowledge error at most $(D+m+3+6L)/(q-1)$. The carrier tree retains $36 \cdot 5^L$ and $(2+6L)/(q-1)$. One aggregate proof consumes one label, with keys bound to its total dimension and public-base list.

D Parameters, Costs, and an OR Example

D.1 PRF-Derived Masks

For mask dimension D and group order q, use a standard PRF [31] $F_0 : \mathcal{K} \times \{0,1\}^* \to \{0,1\}^\ell$, where $\ell = \lambda + \lceil \log_2 q \rceil$, and define

$$F_k(ctx) = (\mathsf{int}(F_{0,k}(\mathsf{tag}_F \| ctx \| \langle j \rangle_b)) \bmod q)_{j=1}^{D}, \quad b = \lceil \log_2(D+1) \rceil.$$

The function int reads a bit string as a nonnegative integer. Labels have fixed length, and $\langle j \rangle_b$ denotes the b-bit encoding of coordinate index j. The tag tag_F separates mask derivation from the random-oracle domains. Setup and embedding use the same derivation.

Reducing a uniform ℓ-bit integer modulo q has statistical distance at most $2^{-\lambda}$ from uniform in $\mathbb{F}_q$. Hence one may take

$$\epsilon_F(\lambda, T) = \epsilon_{F_0}(\lambda, TD) + TD\, 2^{-\lambda},$$

where ϵ_{F_0} bounds the corresponding PRF advantage, including random-oracle access. The TD inputs are the distinct label-coordinate pairs.

For deterministic derivation at fixed public parameters, perfect preservation of the joint normal-mode tapes requires $H(ek) \geq TD \log_2 q$. Indeed, setting all anamorphic openings to zero makes the prover's randomness blocks equal the mask table, which must be uniform in $\mathbb{F}_q^{TD}$. A deterministic map cannot increase entropy. This lower bound concerns the embedding key; the compact PRF seed gives only a computational tape guarantee.

D.2 Auxiliary-Key and Proof Costs

Let b_K, b_F, b_G be the encoding lengths in bytes of a PRF seed, field element, and group element. For the random-base instances, the mask dimension is $D = n$ in Compressed Σ and $D = 2s + 1$ in Bulletproofs. The costs are

Item	Size in bytes
Independent-mask embedding key	TDb_F
PRF embedding key	b_K
Checking key, either variant	Tb_G
Compressed Σ proof, $n = 2^d$	$(2d-1)b_G + 2b_F$
Bulletproofs proof, $s = 2^L$	$(2L+4)b_G + 5b_F$

Proof sizes exclude statements, parameters, and auxiliary data. The prover can track used indices with a T-bit bitmap. PRF setup evaluates TD inputs and computes T mask images. Embedding evaluates D inputs and adds D scalars before invoking the normal prover, which already computes the carrier. Beyond normal verification and carrier parsing, $C = \widehat{C}M_i$ costs one table lookup, one group multiplication, and one equality test.

For a 64-bit Bulletproof, $D = 129$. With illustrative encoding lengths $b_K = b_F = b_G = 32$ and $T = 2^{16}$, the proof is 672 bytes, the seed is 32 bytes, and the checking table is 2 MiB. The reserved-index factor in the reduction is $T = 2^{16}$, separate from its oracle and extraction losses.

D.3 A Compressed Σ Instance Certifying a Discrete Logarithm

With the following bases, both Compressed Σ and Schnorr prove knowledge of a discrete logarithm.

Let $n = 4$, so $d = 2$ and $m = 2$, with $\mathcal{U} = \mathbb{F}_q$, $\mathcal{H} = \mathbb{G} = \langle g \rangle$ of prime order $q > 5$, and

$$\Psi(t) = g^{t_1+2t_2+3t_3+5t_4}, \qquad \text{that is, bases } (g, g^2, g^3, g^5).$$

The exponent rows of the two halves are $(1, 2)$ and $(3, 5)$, with determinant -1. Hence J is surjective, $\delta_{pp} = 0$ in Lemma 3, and Lemma 4 opens the carrier. A discrete logarithm k gives a representation $(k, 0, 0, 0)$, and the displayed linear form maps every representation back to the discrete logarithm. Thus random-image preimage hardness is equivalent to DL. Theorem 1 gives the same guarantees with $\delta_{\mathrm{CS}}(N, Q_H) = (NQ_H + N(N-1)/2)/q^2$ and the same κ_{CS}.

Corollary 3 (Certifying a discrete logarithm). *In this instance, for a registered public key $Y = g^k$, the anamorphic prover embeds $\widehat{u} = (k, 0, 0, 0)$, and every representation t returned by the anamorphic knowledge extractor satisfies $k = t_1 + 2t_2 + 3t_3 + 5t_4 \bmod q$. The hidden claim is knowledge of $\log_g Y$ without any computational uniqueness argument, and the normal claim is knowledge of $\log_g P$.*

Both protocols therefore prove knowledge of $\log_g P$. Schnorr has transfer and a unique response, so Theorem 4 excludes an APoK under its oracle-free local check. Here an additional challenge folds the four-coordinate response before its two remaining coordinates are sent. This yields the conditional simulator of Lemma 3.

D.4 Challenge Transfer in OR Composition

In standard OR composition [22], transfer fixes the first messages and all other branches. With XOR splitting, set the witnessed branch's share to $c'_j = c_j \oplus c \oplus c'$ and update its response [23, Section 4].

For two Schnorr branches $Y_j = g_j^{w_j}$ with t-bit challenges and $1 < 2^t < q$, fix $a_j = g_j^{r_j}$ and global challenge c. For every t-bit share b, the complete response

$$c_0 = b, \quad c_1 = b \oplus c, \quad z_0 = r_0 + c_0 w_0 \pmod q, \quad z_1 = r_1 + c_1 w_1 \pmod q$$

accepts at the same (a_0, a_1, c). Distinct b give distinct canonical responses, even if one share is omitted and reconstructed. An algorithm can choose both branch witnesses, provide either as its OR witness, and use two choices of b to win the two-response game with probability 1. Thus Theorem 4 cannot exclude all full-transcript checks for this protocol by a uniform bound on two-response advantage. One branch witness suffices for transfer, so Proposition 2 still excludes first-message checks.