Non-Interactive CCA-Secure Threshold Cryptosystems with Adaptive Security:
New Framework and Constructions

Benoît Libert[1] [*] and Moti Yung[2]

[1]Université catholique de Louvain, ICTEAM Institute (Belgium)
[2] Google Inc. and Columbia University (USA)

Abstract. In threshold cryptography, private keys are divided into n shares, each one of which is given to a different server in order to avoid single points of failure. In the case of threshold public-key encryption, at least $t \leq n$ servers need to contribute to the decryption process. A threshold primitive is said *robust* if no coalition of t malicious servers can prevent remaining honest servers from successfully completing private key operations. So far, most practical non-interactive threshold cryptosystems, where no interactive conversation is required among decryption servers, were only proved secure against static corruptions. In the adaptive corruption scenario (where the adversary can corrupt servers at any time, based on its complete view), all existing robust threshold encryption schemes that also resist chosen-ciphertext attacks (CCA) till recently require interaction in the decryption phase. A specific method (in composite order groups) for getting rid of interaction was recently suggested, leaving the question of more generic frameworks and constructions with better security and better flexibility (*i.e.*, compatibility with distributed key generation).
This paper describes a general construction of adaptively secure robust non-interactive threshold cryptosystems with chosen-ciphertext security. We define the notion of *all-but-one perfectly sound* threshold hash proof systems that can be seen as (threshold) hash proof systems with publicly verifiable and simulation-sound proofs. We show that this notion generically implies threshold cryptosystems combining the aforementioned properties. Then, we provide efficient instantiations under well-studied assumptions in bilinear groups (e.g., in such groups of prime order). These instantiations have a tighter security proof and are indeed compatible with distributed key generation protocols.

Keywords. Threshold cryptography, adaptive corruptions, public-key encryption, chosen-ciphertext security, non-interactivity, robustness.

1   Introduction

Threshold cryptography [22, 23, 12] avoids single points of failure by splitting keys into $n > 1$ shares which are held by servers in such a way that at least t out of n servers should contribute to private key operations. In (t, n)-threshold cryptosystems, an adversary breaking into up to $t - 1$ servers should not jeopardize the security of the system.

Chosen-ciphertext security [45] (or IND-CCA for short) is widely recognized as the standard security notion for public-key encryption. Securely distributing the decryption procedure of CCA-secure public key schemes has proved to be a challenging task. As discussed in, e.g., [49, 25], the difficulty is that decryption servers should return their partial decryption results, called "decryption shares", before knowing whether the incoming ciphertext is valid or not and partial decryptions of ill-formed ciphertexts may leak useful information to the adversary.

The first solution to this problem was put forth by Shoup and Gennaro [49] and it requires the

---

random oracle model [5], notably to render valid ciphertexts publicly recognizable. In the standard model, Canetti and Goldwasser [15] gave a threshold variant of the Cramer-Shoup encryption scheme [16]. Unfortunately, their scheme requires interaction among decryption servers to obtain robustness (*i.e.*, ensure that no coalition of $t - 1$ malicious servers can prevent uncorrupted servers from successfully decrypting) as well as to render invalid ciphertexts harmless. The approach of [15] consists in randomizing the decryption process in such a way that partial decryptions of invalid ciphertexts are uniformly random and thus meaningless to the adversary. To avoid the need to jointly generate randomizers at each decryption, shareholders can alternatively store a large number (*i.e.*, proportional to the expected number of decryptions) of pre-shared secrets, which does not scale well. Cramer, Damgård and Ishai suggested [20] a method to generate randomizers without interaction but it is only efficient for a small number of servers.

Other threshold variants of Cramer-Shoup were suggested [1, 40] and Abe notably showed [1] how to achieve optimal resilience (namely, guarantee robustness as long as the adversary corrupts a minority of $t < n/2$ servers) in the Canetti-Goldwasser system [15]. In the last decade, generic constructions of CCA-secure threshold cryptosystems with static security were put forth [24, 52].

NON-INTERACTIVE SCHEMES. As an application of the Canetti-Halevi-Katz (CHK) paradigm [18], Boneh, Boyen and Halevi [8] came up with the first fully *non-interactive* robust CCA-secure threshold cryptosystem with a security proof in the standard model: in their scheme, decryption servers can generate their decryption shares *without* any communication with other servers. Their scheme takes advantage of bilinear maps to publicly check the validity of ciphertexts, which considerably simplifies the task of proving security in the threshold setting. In addition, the validity of decryption shares can be verified in the same way, which provides robustness. Similar applications of the CHK methodology to threshold cryptography were studied in [13, 36].

Recently, Wee [52] defined a framework allowing to construct non-interactive threshold signatures and (chosen-ciphertext secure) threshold cryptosystems in a static corruption model. He left as an open problem the extension of his framework in the scenario of adaptive corruptions.

ADAPTIVE CORRUPTIONS. Most threshold systems (including [49, 15, 24, 25, 8]) have been analyzed in a static corruption model, where the adversary chooses which servers it wants to corrupt *before* the scheme is set up. Unfortunately, adaptive adversaries – who can choose whom to corrupt at any time, as a function of their entire view of the protocol execution – are known (see, e.g., [19]) to be strictly stronger. As discussed in [15], properly dealing with adaptive corruptions often comes at some substantial expense like a lower resilience. For example, the Canetti-Goldwasser system can be proved robust and adaptively secure when the threshold t is sufficiently small (typically, when $t = O(n^{1/2})$) but supporting an optimal number of faulty servers is clearly preferable.

Assuming reliable erasures, Canetti *et al.* [14] devised adaptively secure protocols for the distributed generation of discrete-logarithm-based keys and DSA signatures. Their techniques were re-used later on [3] in proactive [44] RSA signatures. In 1999, Frankel, MacKenzie and Yung [26, 27] independently showed different methods to achieve adaptive security in the erasure-enabled setting.

Subsequently, Jarecki and Lysyanskaya [34] eliminated the need for erasures and gave an adaptively secure variant of the Canetti-Goldwasser threshold cryptosystem [15] which appeals to interactive zero-knowledge proofs but is designed to remain secure in concurrent environments. Unfortunately, their scheme requires a fair amount of interaction among decryption servers. Abe and Fehr [2] showed how to dispense with zero-knowledge proofs in the Jarecki-Lysyanskaya construction so as to prove it secure in (a variant of) the universal composability framework but without completely eliminating interaction from the decryption procedure. As in most threshold variants of

Cramer-Shoup, hedging against invalid decryption queries requires an interactive (though off-line) randomness generation phase for each ciphertext, unless many pre-shared secrets are stored.

Recently, the authors of this paper showed [39] an adaptively secure variant of the Boneh-Boyen-Halevi construction [8] using groups of composite order and the dual system encryption approach [50, 38] that was initially applied to identity-based encryption [48, 10]. The scheme of [39] is based on a very specific use of the Lewko-Waters techniques [38], which limits its applicability to composite order groups and makes it hard to combine with existing adaptively secure distributed key generation techniques. Also, the concrete security of this initial scheme is not optimal as its security reduction is related to the number of decryption queries made by the adversary. To solve these problems, we need a new approach and different methods to analyze the security of schemes.

OUR CONTRIBUTION. Motivated by an open question raised by Wee [52] and the limitations of [39], we define a general framework for constructing robust, adaptively secure and fully non-interactive threshold cryptosystems with chosen-ciphertext security. Our goal is to have simple and practical client/server protocols, as advocated in [49][Section 2.5], and even avoid the off-line interactive randomness generation stage which is usually needed in threshold versions of Cramer-Shoup.

To this end, we also appeal to hash proof systems (HPS) [17] and take advantage of the property that, in security reductions using the techniques of [16, 17], the simulator knows the private keys, which is convenient to answer adaptive corruption queries. Indeed, when the reduction has to reveal the internal state of dynamically-corrupted servers, it is not bound to a particular set of available shares since it knows them all. At the same time, we depart from [15] in that the validity of ciphertexts is made publicly verifiable – which eliminates the need to randomize the decryption operation – using non-interactive proofs satisfying some form of simulation-soundness [46]: in the security reduction, the simulator should be able to generate a proof for a possibly false statement but the adversary should be unable to do it on its own, even after having seen a fake proof.

To this end, we define the notion of *all-but-one perfectly sound threshold hash proof systems* that can be seen as (threshold) hash proof systems [17] with *publicly* verifiable proofs (as opposed to designed-verifier proofs used in traditional HPS [17]). More precisely, each proof is associated with a tag, in the same way as ciphertexts are associated with tags in [41, 36]. Real public parameters are indistinguishable from alternative parameters that are generated in an *all-but-one* mode, which is only used in the security analysis. In the latter mode, non-interactive proofs are perfectly sound on all tags, except for a single specific tag where some trapdoor makes it possible to simulate proofs for false statements. While our primitive bears similarities with Wee's extractable hash proof systems [51, 52] (where hash proof systems are also associated with tags), it is different in that no extractability property is required and proofs are always used as proofs of membership.

Using all-but-one perfectly sound threshold hash proof systems, we generically construct adaptively secure robust non-interactive threshold cryptosystems with optimal resilience. An additional benefit of this approach is to provide a better concrete security as the security proof requires a constant number of game transitions whereas, in [39], the number of games is proportional to the number of decryption queries.

Then, we show three concrete instantiations using number theoretic assumptions in bilinear groups. The first one uses groups whose order is a product of two primes (whereas three primes are needed in [39]). Our second and third schemes rely on the Groth-Sahai proof systems [31] in their instantiations based on the Decision Linear [9] and symmetric eXternal Diffie-Hellman assumptions [47]. The latter two constructions operate over bilinear groups of prime order, which allows for a significantly better efficiency than composite order groups (as discussed in [28]) and makes them

much easier to combine with known adaptively secure discrete-log-based distributed key generation protocols. For example, in the erasure-free setting, the protocols of [34, 2] can be used so as to eliminate the need for a trusted dealer at the same time as the reliance on reliable erasures.

2 Background and Definitions

2.1 Definitions for Threshold Public Key Encryption

A non-interactive (t, n)-threshold encryption scheme is a set of algorithms with these specifications.

Setup(λ, t, n)**:** given a security parameter λ and integers $t, n \in \mathsf{poly}(\lambda)$ (with $1 \leq t \leq n$) denoting the number of decryption servers n and the decryption threshold t, this algorithm outputs $(PK, \mathbf{VK}, \mathbf{SK})$, where PK is the public key, $\mathbf{SK} = (SK_1, \ldots, SK_n)$ is a vector of private-key shares and $\mathbf{VK} = (VK_1, \ldots, VK_n)$ is a vector of verification keys. Decryption server i is given the private key share (i, SK_i). For each $i \in \{1, \ldots, n\}$, the verification key VK_i will be used to check the validity of decryption shares generated using SK_i.

Encrypt(PK, M)**:** is a randomized algorithm that, given a public key PK and a plaintext M, outputs a ciphertext C.

Ciphertext-Verify(PK, C)**:** takes as input a public key PK and a ciphertext C. It outputs 1 if C is deemed valid w.r.t. PK and 0 otherwise.

Share-Decrypt(PK, i, SK_i, C)**:** on input of a public key PK, a ciphertext C and a private-key share (i, SK_i), this (possibly randomized) algorithm outputs a special symbol $(i, \perp)$ if **Ciphertext-Verify**$(PK, C) = 0$. Otherwise, it outputs a decryption share $\mu_i = (i, \hat{\mu}_i)$.

Share-Verify(PK, VK_i, C, μ_i)**:** takes in PK, the verification key VK_i, a ciphertext C and a purported decryption share $\mu_i = (i, \hat{\mu}_i)$. It outputs either 1 or 0. In the former case, μ_i is said to be a *valid* decryption share. We adopt the convention that $(i, \perp)$ is an invalid decryption share.

Combine$(PK, \mathbf{VK}, C, \{\mu_i\}_{i \in S})$**:** given PK, $\mathbf{VK}$, C and a subset $S \subset \{1, \ldots, n\}$ of size $t = |S|$ with decryption shares $\{\mu_i\}_{i \in S}$, this algorithm outputs either a plaintext M or $\perp$ if the set contains invalid decryption shares.

CHOSEN-CIPHERTEXT SECURITY. We use a game-based definition of chosen-ciphertext security which is akin to the one of [49, 8] with the difference that the adversary can adaptively decide which parties it wants to corrupt.

Definition 1. *A non-interactive (t, n)-Threshold Public Key Encryption scheme is secure against chosen-ciphertext attacks (or IND-CCA2 secure) and adaptive corruptions if no PPT adversary has non-negligible advantage in this game:*

1. *The challenger runs **Setup**(λ, t, n) to obtain a public key PK, a vector of private key shares $\mathbf{SK} = (SK_1, \ldots, SK_n)$ and verification keys $\mathbf{VK} = (VK_1, \ldots, VK_n)$. It gives PK and $\mathbf{VK}$ to the adversary $\mathcal{A}$ and keeps $\mathbf{SK}$ to itself.*

2. *The adversary $\mathcal{A}$ adaptively makes the following kinds of queries:*

   - *Corruption query: $\mathcal{A}$ chooses $i \in \{1, \ldots, n\}$ and obtains SK_i. No more than $t - 1$ private key shares can be obtained by $\mathcal{A}$ in the whole game.*
   - *Decryption query: $\mathcal{A}$ chooses an index $i \in \{1, \ldots, n\}$ and a ciphertext C. The challenger replies with $\mu_i = **Share-Decrypt**(PK, i, SK_i, C)$.*

3. *The adversary $\mathcal{A}$ chooses two equal-length messages M_0, M_1 and obtains $C^\star = **Encrypt**(PK, M_\beta)$ for some random bit $\beta \xleftarrow{R} \{0, 1\}$.*

4. $\mathcal{A}$ *makes further queries as in step 2 but is not allowed to make decryption queries on* $C^\star$.

5. $\mathcal{A}$ *outputs a bit* β' *and is deemed successful if* $\beta' = \beta$. *As usual,* $\mathcal{A}$'s *advantage is measured as the distance* $\mathbf{Adv}(\mathcal{A}) = |\Pr[\beta' = \beta] - \frac{1}{2}|$.

CONSISTENCY. A (t, n)-Threshold Encryption scheme provides decryption consistency if no PPT adversary has non-negligible advantage in a three-stage game where stages 1 and 2 are identical to those of Definition 1 with the difference that the adversary $\mathcal{A}$ is allowed to obtain *all* private key shares (alternatively, $\mathcal{A}$ can directly obtain $\mathbf{SK}$ at the beginning of the game). In stage 3, $\mathcal{A}$ outputs a ciphertext C and two t-sets of decryption shares $\Gamma = \{\mu_1, \ldots, \mu_t\}$ and $\Gamma' = \{\mu'_1, \ldots, \mu'_t\}$. The adversary $\mathcal{A}$ is declared successful if

1. **Ciphertext-Verify**$(PK, C) = 1$.
2. Γ and Γ' only consist of valid decryption shares.
3. **Combine**$(PK, \mathbf{VK}, C, \Gamma) \neq$ **Combine**$(PK, \mathbf{VK}, C, \Gamma')$.

We note that condition 1 prevents an adversary from trivially winning by outputting an invalid ciphertext, for which distinct sets of key shares may give different results. This definition of consistency is identical to the one of [49, 8] with the difference that $\mathcal{A}$ can adaptively corrupt servers.

2.2 Hardness Assumptions in Composite Order Groups

In one occasion, we appeal to groups $(\mathbb{G}, \mathbb{G}_T)$ of order $N = p_1 p_2$, where p_1 and p_2 are primes, with a bilinear map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ (*i.e.*, for which $e(g^a, h^b) = e(g, h)^{ab}$ for any $g, h \in \mathbb{G}$ and $a, b \in \mathbb{Z}_N$). In the notations hereafter, for each $i \in \{1, 2\}$, $\mathbb{G}_{p_i}$ stands for the subgroup of order p_i in $\mathbb{G}$.

Definition 2 ([11]). *In a group* $\mathbb{G}$ *of composite order* N, *the* **Subgroup Decision** *(SD) problem is given* $(g \in \mathbb{G}_{p_1}, h \in \mathbb{G})$ *and* η, *to decide whether* $\eta \in \mathbb{G}_{p_1}$ *or* $\eta \in_R \mathbb{G}$. *The* **Subgroup Decision** *assumption states that, for any PPT distinguisher* $\mathcal{D}$, *the SD problem is infeasible.*

2.3 Assumptions in Prime Order Groups

We also use bilinear maps $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$ over groups of prime order p. We will work in symmetric pairing configurations, where $\mathbb{G} = \hat{\mathbb{G}}$, and sometimes in asymmetric configurations, where $\mathbb{G} \neq \hat{\mathbb{G}}$.

In the symmetric setting $(\mathbb{G}, \mathbb{G}_T)$, we rely on the following assumption.

Definition 3 ([9]). *In a group* $\mathbb{G}$ *of prime order* p, *the* **Decision Linear Problem** *(DLIN) is to distinguish the distributions* $(g, g^a, g^b, g^{ac}, g^{bd}, g^{c+d})$ *and* $(g, g^a, g^b, g^{ac}, g^{bd}, g^z)$, *with* $a, b, c, d, z \xleftarrow{R} \mathbb{Z}_p$. *The* **Decision Linear Assumption** *is the intractability of DLIN for any PPT distinguisher* $\mathcal{D}$.

The problem amounts to deciding if vectors $\vec{g_1} = (g^a, 1, g)$, $\vec{g_2} = (1, g^b, g)$ and $\vec{g_3} = (g^{ac}, g^{bd}, g^\delta)$ are linearly dependent (*i.e.*, if $\delta = c + d$) or not.

In *asymmetric* bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$, we assume the hardness of the Decision Diffie-Hellman (DDH) problem in $\mathbb{G}$ and $\hat{\mathbb{G}}$. This implies the unavailability of efficiently computable isomorphisms between $\hat{\mathbb{G}}$ and $\mathbb{G}$. This assumption is called **Symmetric eXternal Diffie-Hellman** (SXDH) assumption. Given vectors $\vec{u}_1 = (g, h)$, $\vec{u}_2 = (g^a, h^c)$ in $\mathbb{G}^2$ or $\hat{\mathbb{G}}^2$, the SXDH assumption asserts the infeasibility of deciding whether $\vec{u}_1$ and $\vec{u}_2$ are linearly dependent (*i.e.*, whether $a = c \bmod p$).

3 All-But-One Perfectly Sound Threshold Hash Proof Systems

Let $\mathcal{C}$, $\mathcal{K}$ and $\mathcal{K}'$ be sets and let $\mathcal{V} \subset \mathcal{C}$ be a subset. Let also $\mathcal{R}$ be a space where random coins can be chosen. We mandate that $\mathcal{V}$, $\mathcal{K}$, $\mathcal{K}'$ and $\mathcal{R}$ be of exponential size in λ, where $\lambda \in \mathbb{N}$ is a security parameter. In addition, $\mathcal{C}$, $\mathcal{V}$ and $\mathcal{C} \backslash \mathcal{V}$ should be efficiently samplable and we also require the set $\mathcal{K}$ to form a group for some binary operation, which is denoted by $\odot$ hereafter.

An *all-but-one perfectly sound threshold hash proof system* for the sets $(\mathcal{C}, \mathcal{V}, \mathcal{K}, \mathcal{K}', \mathcal{R})$ is a tuple (SetupSound, SetupABO, Sample, Prove, SimProve, Verify, PubEval, SharePrivEval, ShareEvalVerify, Combine) of efficient algorithms with the following specifications.

SetupSound(λ, t, n)**:** given a security parameter $\lambda \in \mathbb{N}$ and integers $t, n \in \mathsf{poly}(\lambda)$, this algorithm outputs a public key pk, n private key shares $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ and verification keys $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$.

SetupABO$(\lambda, t, n, \mathsf{tag}^\star)$**:** takes as input a security parameter $\lambda \in \mathbb{N}$, integers $t, n \in \mathsf{poly}(\lambda)$ and a tag $\mathsf{tag}^\star$. It outputs a public key pk, n private key shares $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$, the corresponding verification keys $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ as well as a simulation trapdoor τ. It is important that τ be independent of $\{\mathsf{sk}_i\}_{i=1}^n$.

Sample(pk)**:** is a probabilistic algorithm that takes as input a public key pk. It draws random coins $r \xleftarrow{R} \mathcal{R}$ and outputs an element $\Phi \in \mathcal{V}$ along with the random coins r that will serve as a witness explaining Φ as an element of $\mathcal{V}$.

Prove$(\mathsf{pk}, \mathsf{tag}, r, \Phi)$**:** takes in a public key pk, a tag tag, an element $\Phi \in \mathcal{V}$ and the random coins $r \in \mathcal{R}$ that were used to sample Φ. It generates a non-interactive proof $\pi_\mathcal{V}$ that $\Phi \in \mathcal{V}$.

SimProve$(\mathsf{pk}, \tau, \mathsf{tag}, \Phi)$**:** takes as input a public key pk and a simulation trapdoor τ produced by SetupABO$(\lambda, t, n, \mathsf{tag}^\star)$, a tag tag and an element $\Phi \in \mathcal{C}$. If $\mathsf{tag} \neq \mathsf{tag}^\star$, the algorithm outputs $\perp$. If $\mathsf{tag} = \mathsf{tag}^\star$, the algorithm produces a simulated NIZK proof $\pi_\mathcal{V}$ that $\Phi \in \mathcal{V}$.

Verify$(\mathsf{pk}, \mathsf{tag}, \Phi, \pi_\mathcal{V})$**:** takes as input a public key pk, a tag tag, an element $\Phi \in \mathcal{C}$ and a purported proof $\pi_\mathcal{V}$. It outputs 1 if and only if $\pi_\mathcal{V}$ is deemed as a valid proof that $\Phi \in \mathcal{V} \subset \mathcal{C}$.

PubEval(pk, r, Φ)**:** takes as input a public key pk, an element $\Phi \in \mathcal{V}$ and the random coins $r \in_R \mathcal{R}$ such that $(r, \Phi) \leftarrow$ Sample(pk). It outputs a value $K \in \mathcal{K}$, which is called *public evaluation* of Φ.

SharePrivEval$(\mathsf{pk}, \mathsf{sk}_i, \Phi)$**:** is a deterministic algorithm that takes in a public key pk, a private key share sk_i and an element $\Phi \in \mathcal{C}$. It outputs a value $K_i \in \mathcal{K}'$, called *private evaluation share* and a proof π_{K_i} that K_i was evaluated correctly.

ShareEvalVerify$(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i})$**:** given a public key pk, a verification key vk_i, an element $\Phi \in \mathcal{C}$, a private evaluation share $K_i \in \mathcal{K}'$ and its proof π_{K_i}, this algorithm outputs 1 if π_{K_i} is considered as a valid proof of the correct evaluation of K_i. Otherwise, it outputs 0.

Combine$(\mathsf{pk}, \Phi, \{(K_i, \pi_{K_i})\}_{i \in S})$**:** takes as input a public key pk, an element $\Phi \in \mathcal{C}$ and a set of t pairs $\{(K_i, \pi_{K_i})\}_{i \in S}$, where $S \subset \{1, \ldots, n\}$, each one of which consists of a private evaluation share $K_i \in \mathcal{K}'$ and its proof π_{K_i}. If ShareEvalVerify$(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i}) = 0$ for some $i \in S$, it outputs $\perp$. Otherwise, it outputs a value $K \in \mathcal{K}$.

We also define this algorithm which is implied by the above ones but will be convenient to use.

PrivEval$(\mathsf{pk}, \{\mathsf{sk}_i\}_{i \in S}, \Phi)$**:** given a public key pk, a set of private key shares $\{\mathsf{sk}_i\}_{i \in S}$ where S is an arbitrary t-subset of $\{1, \ldots, n\}$, and an element $\Phi \in \mathcal{C}$, this algorithm outputs the result of Combine$(\mathsf{pk}, \Phi, \{(K_i, \pi_{K_i})\}_{i \in S})$ where $(K_i, \pi_{K_i}) \leftarrow$ SharePrivEval$(\mathsf{pk}, \mathsf{sk}_i, \Phi)$ for each $i \in S$.

The following properties are required from these algorithms and the sets $(\mathcal{C}, \mathcal{V}, \mathcal{K}, \mathcal{K}', \mathcal{R})$.

(Setup indistinguishability): For any integers (λ, t, n) such that $1 \le t \le n$ and any tag $\mathsf{tag}^\star$, the output of $\mathsf{SetupSound}(\lambda, t, n)$ and the outputs $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n)$ of $\mathsf{SetupABO}(\lambda, t, n, \mathsf{tag}^\star)$ are computationally indistinguishable.

(Correctness and Public Evaluability on $\mathcal{V}$): For any $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n)$ returned by $\mathsf{SetupSound}$ or $\mathsf{SetupABO}$, if $(r, \Phi) \xleftarrow{R} \mathsf{Sample}(\mathsf{pk})$ (and thus $\Phi \in \mathcal{V}$), it holds that:

1. For any $i \in \{1, \dots, n\}$, if $(K_i, \pi_{K_i}) \leftarrow \mathsf{SharePrivEval}(\mathsf{pk}, \mathsf{sk}_i, \Phi)$, then the private evaluation share $K_i \in \mathcal{K}'$ is *uniquely* determined by $(\mathsf{pk}, \mathsf{vk}_i)$ and Φ. Moreover, the proof π_{K_i} satisfies the verification test: $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i}) = 1$.

2. For any t-subset $S \subset \{1, \dots, n\}$, combining the corresponding private evaluation shares allows recomputing the public evaluation:
$$\mathsf{PubEval}(\mathsf{pk}, r, \Phi) = \mathsf{PrivEval}(\mathsf{pk}, \{\mathsf{sk}_i\}_{i \in S}, \Phi).$$

(Universality): For any $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n)$ produced by $\mathsf{SetupSound}$ or $\mathsf{SetupABO}$ and any $\Phi \in \mathcal{C} \backslash \mathcal{V}$, for any subset $\bar{S} \subset \{1, \dots, n\}$ of size $|\bar{S}| = t - 1$, the statistical distance
$$\Delta\left[\left(\mathsf{pk}, \{\mathsf{vk}_i\}_{i=1}^n, \{\mathsf{sk}_i\}_{i \in \bar{S}}, \Phi, \mathsf{PrivEval}(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^t, \Phi)\right), \left(\mathsf{pk}, \{\mathsf{vk}_i\}_{i=1}^n, \{\mathsf{sk}_i\}_{i \in \bar{S}}, \Phi, K\right)\right],$$
where $K \xleftarrow{R} \mathcal{K}$, should be negligible.

(All-But-One Soundness): For all integers (λ, t, n) such that $1 \le t \le n$, any tag $\mathsf{tag}^\star$ and any outputs $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n, \tau)$ of $\mathsf{SetupABO}(\lambda, t, n, \mathsf{tag}^\star)$, these conditions are satisfied.

1. For any $\mathsf{tag} \ne \mathsf{tag}^\star$, proofs are always perfectly sound. Namely, if a proof $\pi_\mathcal{V}$ satisfies $\mathsf{Verify}(\mathsf{pk}, \mathsf{tag}, \Phi, \pi_\mathcal{V}) = 1$ for some $\Phi \in \mathcal{C}$, then it necessarily holds that $\Phi \in \mathcal{V}$.

2. For any $\Phi \in \mathcal{C}$, the trapdoor τ allows simulating a proof $\pi_\mathcal{V} \leftarrow \mathsf{SimProve}(\mathsf{pk}, \tau, \mathsf{tag}^\star, \Phi)$ such that $\mathsf{Verify}(\mathsf{pk}, \mathsf{tag}^\star, \Phi, \pi_\mathcal{V}) = 1$ (note that $\pi_\mathcal{V}$ is a proof for a false statement if $\Phi \in \mathcal{C} \backslash \mathcal{V}$). Moreover, if $\Phi \in \mathcal{V}$, the simulated proof $\pi_\mathcal{V}$ should be perfectly indistinguishable from a real proof (*i.e.*, that would be generated by Prove using a witness $r \in \mathcal{R}$ of the fact that $\Phi \in \mathcal{V}$).

(Simulatability of Share Proofs): For all integers (λ, t, n) such that $1 \le t \le n$, any tag $\mathsf{tag}^\star$, any outputs $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n, \tau)$ of $\mathsf{SetupABO}(\lambda, t, n, \mathsf{tag}^\star)$ and any $\Phi \in \mathcal{C}$, the proofs π_{K_i} produced by $(K_i, \pi_{K_i}) \leftarrow \mathsf{SharePrivEval}(\mathsf{pk}, \mathsf{sk}_i, \Phi)$ should be simulatable using the trapdoor τ instead of $\{\mathsf{sk}_i\}_{i=1}^n$. Using τ and public values $(\mathsf{pk}, \{\mathsf{vk}_i\}_{i=1}^n, \Phi)$, an efficient algorithm $\mathcal{S}$ should be able to produce simulated proofs π_{K_i} that are perfectly indistinguishable from real proofs.

(Consistency): For all integers (λ, t, n) such that $1 \le t \le n$, any output $(\mathsf{pk}, \{(\mathsf{vk}_i, \mathsf{sk}_i)\}_{i=1}^n)$ of $\mathsf{SetupSound}(\lambda, t, n)$, given $(\mathsf{pk}, \{(\mathsf{vk}_i, \mathsf{sk}_i)\}_{i=1}^n)$, it should be computationally infeasible to come up with a triple $(\mathsf{tag}, \Phi, \pi_\mathcal{V})$ as well as two distinct t-sets $\Gamma = \{(K_{i_1}, \pi_{K_{i_1}}), \dots, (K_{i_t}, \pi_{K_{i_t}})\}$ and $\Gamma' = \{(K'_{j_1}, \pi'_{K_{j_1}}), \dots, (K'_{j_t}, \pi'_{K_{j_t}})\}$, with $i_k, j_k \in \{1, \dots, n\}$ for each $k \in \{1, \dots, t\}$, such that: (i) $\mathsf{Verify}(\mathsf{pk}, \mathsf{tag}, \Phi, \pi_\mathcal{V}) = 1$; (ii) for each $k \in \{1, \dots, t\}$, $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_{i_k}, \Phi, K_{i_k}, \pi_{K_{i_k}}) = 1$ and $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_{j_k}, \Phi, K'_{j_k}, \pi'_{K_{j_k}}) = 1$; (iii) $\mathsf{Combine}(\mathsf{pk}, \Phi, \Gamma) \ne \mathsf{Combine}(\mathsf{pk}, \Phi, \Gamma')$.

(Subset Membership Hardness): membership in $\mathcal{C}$ should be easy to check but membership in $\mathcal{V}$ should not. Moreover, this should hold *even* if τ is given. Namely, for all integers (λ, t, n) such that $1 \le t \le n$, any tag $\mathsf{tag}^\star$ and any outputs $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n, \tau)$ of $\mathsf{SetupABO}(\lambda, t, n, \mathsf{tag}^\star)$, for any PPT distinguisher $\mathcal{D}$, it must hold that:
$$\mathbf{Adv}^{\mathrm{SM}}(\mathcal{D}) = |\Pr[\mathcal{D}(\mathcal{C}, \mathcal{V}, C_1, \tau) = 1 | C_1 \xleftarrow{R} \mathcal{C} \backslash \mathcal{V}] - \Pr[\mathcal{D}(\mathcal{C}, \mathcal{V}, C_0, \tau) = 1 | C_0 \xleftarrow{R} \mathcal{V}]| \in \mathsf{negl}(\lambda).$$

In the definition of the subset membership hardness property, the trapdoor τ should not carry any side information helping the distinguisher. For this reason, the latter receives τ as part of its input.

4 Adaptively Secure Robust Non-Interactive CCA2-Secure Threshold Cryptosystems from All-But-One Perfectly Sound Threshold Hash Proof Systems

Let $\Pi^{\mathsf{ABO\text{-}THPS}} = (\mathsf{SetupSound}, \mathsf{SetupABO}, \mathsf{Sample}, \mathsf{Prove}, \mathsf{SimProve}, \mathsf{Verify}, \mathsf{PubEval}, \mathsf{SharePrivEval},$ $\mathsf{ShareEvalVerify}, \mathsf{Combine})$ be an all-but-one perfectly sound threshold hash proof system for sets $(\mathcal{C}, \mathcal{V}, \mathcal{K}, \mathcal{K}', \mathcal{R})$ that satisfy the conditions specified in Section 3. We assume that messages are in $\mathcal{K}$. The generic construction of CCA2-secure threshold cryptosystem goes as follows.

Keygen(λ, t, n)**:** given integers $\lambda, t, n \in \mathbb{N}$, choose a one-time signature scheme $\Sigma = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$, generate $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n) \leftarrow \mathsf{SetupSound}(\lambda, t, n)$ and output $(PK, \mathbf{SK}, \mathbf{VK})$, where the vectors of private key shares and verification keys are defined as $\mathbf{SK} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ and $\mathbf{VK} = (\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$, respectively. The public key is $PK = (\mathsf{pk}, \Sigma)$.

Encrypt(M, PK)**:** to encrypt a message $M \in \mathcal{K}$ using $PK = (\mathsf{pk}, \Sigma)$,

1. Generate a one-time signature key pair $(\mathsf{SSK}, \mathsf{SVK}) \leftarrow \Sigma.\mathsf{Gen}(\lambda)$.
2. Choose $r \xleftarrow{R} \mathcal{R}$, compute $(r, \Phi) \leftarrow \mathsf{Sample}(\mathsf{pk}, r)$ as well as $C_0 = M \odot \mathsf{PubEval}(\mathsf{pk}, r, \Phi)$.
3. Generate a proof $\pi_{\mathcal{V}} \leftarrow \mathsf{Prove}(\mathsf{pk}, \mathsf{SVK}, r, \Phi)$ that $\Phi \in \mathcal{V}$ with respect to the tag SVK.
4. Output the ciphertext $C = (\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$, where $\sigma = \Sigma.\mathsf{Sig}(\mathsf{SSK}, (C_0, \Phi, \pi_{\mathcal{V}}))$.

Ciphertext-Verify(PK, C)**:** parse C as $C = (\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$ and PK as (pk, Σ). Return 1 if $\Sigma.\mathsf{Ver}(\mathsf{SVK}, (C_0, \Phi, \pi_{\mathcal{V}}), \sigma) = 1$ and $\mathsf{Verify}(\mathsf{pk}, \mathsf{SVK}, \Phi, \pi_{\mathcal{V}}) = 1$. Otherwise, return 0.

Share-Decrypt(SK_i, C)**:** given the private key share $SK_i = \mathsf{sk}_i$ and $C = (\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$, return $(i, \perp)$ if it turns out that **Ciphertext-Verify**$(PK, C) = 0$. Otherwise, compute a pair $(K_i, \pi_{K_i}) \leftarrow \mathsf{SharePrivEval}(\mathsf{pk}, \mathsf{sk}_i, \Phi)$ and return $\mu_i = (i, \hat{\mu}_i)$ where $\hat{\mu}_i = (K_i, \pi_{K_i})$.

Share-Verify$(PK, VK_i, C, (i, \hat{\mu}_i))$**:** parse the ciphertext C as $(\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$. If the decryption share $\hat{\mu}_i$ is such that $\hat{\mu}_i = \perp$ or if it cannot be properly parsed as a pair (K_i, π_{K_i}), return 0. Otherwise, return 1 if $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i}) = 1$. In any other situation, return 0.

Combine$(PK, \mathbf{VK}, C, \{(i, \hat{\mu}_i)\}_{i \in S})$**:** parse C as $(\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$. Return $\perp$ if there exists $i \in S$ such that **Share-Verify**$(PK, C, (i, \hat{\mu}_i)) = 0$ or if **Ciphertext-Verify**$(PK, C) = 0$. Otherwise, compute $K = \mathsf{Combine}(\mathsf{pk}, \Phi, \{(K_i, \pi_{K_i})\}_{i \in S}) \in \mathcal{K}$, which unveils $M = C_0 \odot K^{-1}$.

We observe that there is no need to bind the one-time verification key SVK to the ciphertext components $(C_0, \Phi, \pi_{\mathcal{V}})$ in any other way than by using it as a tag to compute the non-interactive proof $\pi_{\mathcal{V}}$. Indeed, if the adversary attempts to re-use parts $(C_0^\star, \Phi^\star, \pi_{\mathcal{V}}^\star)$ of the challenge ciphertext and simply replaces the one-time verification key $\mathsf{SVK}^\star$ by a verification key SVK of its own, it will be forced to compute a proof $\pi_{\mathcal{V}}$ that correspond to the same $\Phi^\star$ as in the challenge phase but under the *new* tag SVK. Our security proof shows that this is infeasible as long as $\Pi^{\mathsf{ABO\text{-}THPS}}$ satisfies the properties of setup indistinguishability and all-but-one soundness.

The consistency property of the scheme is trivially implied by that of $\Pi^{\mathsf{ABO\text{-}THPS}}$ and we focus on proving its IND-CCA security. In the threshold setting, adaptive security is achieved by taking advantage of the fact that, in security reductions using hash proof systems, the simulator typically knows the private key and can thus answer adaptive queries at will. At the same time, invalid ciphertexts are harmless as they are made publicly recognizable due to the use of non-interactive proofs of validity: as long as these proofs are perfectly sound in all decryption queries, the simulator is guaranteed not to leak too much information about the particular private key it is using.

The main problem to solve is thus to make sure that *only* the simulator can simulate a fake proof in the challenge phase and this is where the all-but-one soundness property is handy.

Theorem 1. *The above threshold cryptosystem is IND-CCA secure against adaptive corruptions assuming that: (i) $\Pi^{\mathsf{ABO\text{-}THPS}}$ is an all-but-one perfectly sound hash proof system; (ii) Σ is a strongly unforgeable one-time signature.*

Proof. The proof uses of a sequence of games starting with the real attack game and ending with a game where the adversary $\mathcal{A}$ has no advantage. For each i, S_i is the event that $\mathcal{A}$ wins in Game_i.

Game_1: is the real attack game. In details, the adversary is given the public key PK and the set of verification keys $\mathbf{VK} = (\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ and starts making adaptive queries. At each corruption query $i \in \{1, \ldots, n\}$, the challenger $\mathcal{B}$ reveals the queried private key share $SK_i = \mathsf{sk}_i$ and, at each decryption query, $\mathcal{B}$ runs the real shared decryption algorithm. In the challenge phase, the adversary $\mathcal{A}$ chooses messages $M_0, M_1 \in \mathcal{K}$ and obtains $C^\star = (\mathsf{SVK}^\star, C_0^\star, \Phi^\star, \pi_{\mathcal{V}}^\star, \sigma^\star)$ which is an encryption of M_β, for some random coin $\beta \stackrel{R}{\leftarrow} \{0,1\}$ internally flipped by $\mathcal{B}$. For simplicity, we assume that the one-time signature key pair $(\mathsf{SSK}^\star, \mathsf{SVK}^\star)$ is chosen by $\mathcal{B}$ at the outset of the game. In the second phase, $\mathcal{A}$ makes more queries under the restriction of not asking for a partial decryption of $C^\star$ or for more than $t-1$ private key shares throughout the entire game. Eventually, $\mathcal{A}$ halts and outputs β'. We denote by S_1 the event that $\beta = \beta'$.

Game_2: we change the distribution of the public key $PK = (\mathsf{pk}, \Sigma)$. Namely, instead of generating $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n)$ as per $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n) \leftarrow \mathsf{SetupSound}(\lambda, t, n)$, the challenger $\mathcal{B}$ runs the all-but-one setup algorithm $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n, \tau) \leftarrow \mathsf{SetupABO}(\lambda, t, n, \mathsf{SVK}^\star)$, discards τ and uses $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n)$ as in Game_1. We note that, after this change, the one-time verification key $\mathsf{SVK}^\star$ may not be completely independent of $\mathcal{A}$'s view before the challenge phase. However, due to the setup indistinguishability property of $\Pi^{\mathsf{ABO\text{-}THPS}}$, this modification cannot significantly affect $\mathcal{A}$'s behavior. This implies $|\Pr[S_2] - \Pr[S_1]| \in \mathsf{negl}(\lambda)$.

Game_3: we introduce a failure event F_3 and let the challenger $\mathcal{B}$ halt and output a random bit if this event occurs. We call F_3 the event that $\mathcal{A}$ makes a decryption query involving a valid ciphertext $C = (\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$ such that $\mathsf{SVK} = \mathsf{SVK}^\star$. We note that Game_3 and Game_2 are identical until F_3 occurs and argue that $|\Pr[S_3] - \Pr[S_2]| \leq \Pr[F_3] \in \mathsf{negl}(\lambda)$. Indeed, if F_3 occurs before the challenge phase, it means that $\mathcal{A}$ was able to forge a valid one-time signature even before seeing a signature. If F_3 comes about in a post-challenge query, $\mathcal{A}$ must have been able to break the strong unforgeability of the one-time signature.

Game_4: we modify the generation of the challenge ciphertext $C^\star$. Namely, the challenger still picks $\Phi^\star \in \mathcal{V}$ as per $(r^\star, \Phi^\star) \leftarrow \mathsf{Sample}(\mathsf{pk})$, using random coins $r^\star \stackrel{R}{\leftarrow} \mathcal{R}$, as in previous games. However, $C_0^\star$ is now computed as $C_0^\star = M_\beta \odot \mathsf{PrivEval}(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^t, \Phi^\star)$ (instead of $C_0^\star = M_\beta \odot \mathsf{PubEval}(\mathsf{pk}, r^\star, \Phi^\star)$). As long as $\Pi^{\mathsf{ABO\text{-}THPS}}$ satisfies the property of correctness and public evaluability on $\mathcal{V}$, $\mathcal{A}$'s view does not change since $C_0^\star$ has the same distribution either way. We thus have $\Pr[S_4] = \Pr[S_3]$.

Game_5: we modify again the generation of the challenge ciphertext $C^\star$. We observe that the proof $\pi_{\mathcal{V}}^\star$ must be generated w.r.t. the tag $\mathsf{SVK}^\star$ which, due to the modification introduced in Game_2, is the tag for which $\mathcal{B}$ can generate simulated NIZK proofs using the trapdoor τ. To construct the ciphertext $C^\star$, the challenger $\mathcal{B}$ chooses $\Phi^\star \in \mathcal{V}$ as in Game_4 and sets

$$C_0^\star = M_\beta \odot \mathsf{PrivEval}(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^t, \Phi^\star), \qquad \pi_{\mathcal{V}}^\star = \mathsf{SimProve}(\mathsf{pk}, \tau, \mathsf{SVK}^\star, \Phi^\star). \tag{1}$$

Note that, with this modification, $\pi_{\mathcal{V}}^\star$ is now independent of $\{\mathsf{sk}_i\}_{i=1}^n$ as these are independent of τ. Since $(C_0^\star, \Phi^\star, \pi_{\mathcal{V}}^\star)$ have the same distribution as in Game_4, we have $\Pr[S_5] = \Pr[S_4]$.

Game_6: is as Game_5 but we change the treatment of decryption queries $C = (\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$. More precisely, whenever $\mathcal{B}$ runs $\mathsf{SharePrivEval}(\mathsf{pk}, \mathsf{sk}_i, \Phi)$ in order to answer decryption queries, to obtain a private evaluation share K_i and a proof π_{K_i} of its validity, the latter is generated using the simulator $\mathcal{S}$ and the simulation trapdoor τ instead of the private key share sk_i. The property that we called "simulatability of share proofs" guarantees the existence of such an efficient simulator $\mathcal{S}$ and that simulated proofs π_{K_i} will be distributed as real proofs. Hence, we can write $\Pr[S_6] = \Pr[S_5]$.

Game_7: we bring one last change in the generation of the challenge ciphertext. Instead of computing $(C_0^\star, \pi_{\mathcal{V}}^\star)$ as per (1) using a random $\Phi^\star \in \mathcal{V}$, the value $\Phi^\star$ is randomly chosen in $\mathcal{C} \backslash \mathcal{V}$. Under the subset membership hardness assumption in $(\mathcal{C}, \mathcal{V})$, this modification cannot be noticed by $\mathcal{A}$ and we must have $|\Pr[S_7] - \Pr[S_6]| \leq \mathbf{Adv}^{\mathrm{SM}}(\mathcal{A}) \in \mathsf{negl}(\lambda)$ for any PPT adversary $\mathcal{A}$.

In Game_7, we have $\Pr[S_7] \approx 1/2$ so that $\mathcal{A}$'s advantage is statistically negligible. To see this, we observe that, for any valid decryption query $C = (\mathsf{SVK}, C_0, \Phi, \pi_{\mathcal{V}}, \sigma)$ such that $\mathsf{SVK} \neq \mathsf{SVK}^\star$, the proof $\pi_{\mathcal{V}}$ is perfectly sound since it is generated for a tag $\mathsf{SVK} \neq \mathsf{SVK}^\star$ and this guarantees that $\Phi \in \mathcal{V}$ (as even an unbounded $\mathcal{A}$ would be unable to generate a convincing proof $\pi_{\mathcal{V}}$ otherwise). Consequently, for each revealed decryption share $\hat{\mu}_i = (i, (K_i, \pi_{K_i}))$, it holds that: (1) K_i does not reveal any more information about sk_i than $(\mathsf{pk}, \mathsf{vk}_i)$ since it is uniquely determined by $(\mathsf{pk}, \mathsf{vk}_i, \Phi)$; (2) the distribution of π_{K_i} does not depend on sk_i thanks to the modification introduced in Game_6.

The universality property of $\Pi^{\mathsf{ABO\text{-}THPS}}$ tells us that, for any $(t-1)$-subset $\bar{S} \subset \{1, \ldots, n\}$, the distribution $\left(\mathsf{pk}, \{\mathsf{vk}_i\}_{i=1}^n, \{\mathsf{sk}_i\}_{i \in \bar{S}}, \Phi^\star, \mathsf{PrivEval}(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^t, \Phi^\star)\right)$ is statistically indistinguishable from the distribution $\left(\mathsf{pk}, \{\mathsf{vk}_i\}_{i=1}^n, \{\mathsf{sk}_i\}_{i \in \bar{S}}, \Phi^\star, K\right)$, where $K \xleftarrow{R} \mathcal{K}$. In other words, $C_0^\star$ statistically hides M_β and $\Pr[S_7]$ is negligibly far apart from $1/2$, as claimed. □

5 Instantiations

5.1 Construction in Groups of Composite Order $N = p_1 p_2$

The construction relies on a hash proof system in a group $\mathbb{G}$ of composite order $N = p_1 p_2$ and it is conceptually close to the one in [33] (notably because it builds on a $\log p_2$-entropic hash proof system, as defined in [37]). The public key includes group elements $(g, X = g^x)$ in the subgroup $\mathbb{G}_{p_1}$ of order p_1 and the sets $\mathcal{C}$ and $\mathcal{V}$ are defined to be $\mathbb{G}$ and $\mathbb{G}_{p_1}$, respectively. The sampling algorithm returns $\Phi = g^r \in \mathbb{G}_{p_1}$ for a randomly chosen exponent $r \xleftarrow{R} \mathbb{Z}_N$, which allows publicly evaluating $H(X^r) = H(\Phi^x)$ using a pairwise independent hash function $H : \mathbb{G} \to \{0, 1\}^\ell$. Since the public key is independent of $x \bmod p_2$, for any $\Phi \in \mathbb{G}$ that has a non-trivial component of order p_2, the "hash value" Φ^x has exactly $\log p_2$ bits of min-entropy and the leftover hash lemma implies that $H(\Phi^x)$ is statistically close to the uniform distribution in $\{0, 1\}^\ell$ when ℓ is sufficiently small.

In order to turn the scheme into an all-but-one perfectly sound threshold HPS, we need a mechanism that proves membership in the subgroup $\mathbb{G}_{p_1}$ and guarantees the perfect soundness of proofs of membership for all tags $\mathsf{tag} \in \mathbb{Z}_N$ such that $\mathsf{tag} \neq \mathsf{tag}^\star$. To this end, we use additional public parameters $(u, v) \in \mathbb{G}^2$ and a tag-dependent group element $u^{\mathsf{tag}} \cdot v$ will serve as a common reference string to generate a non-interactive proof that $\Phi \in \mathbb{G}_{p_1}$. Membership in $\mathbb{G}_{p_1}$ can be non-interactively proved using a technique that can be traced back to [30]. The proof consists of a group element $\pi_{\mathrm{SD}} \in \mathbb{G}$ satisfying the equality $e(\Phi, u^{\mathsf{tag}} \cdot v) = e(g, \pi_{\mathrm{SD}})$, which ensures that $\Phi \in \mathbb{G}_{p_1}$ as

long as $u^{\mathsf{tag}} \cdot v$ has a $\mathbb{G}_{p_2}$ component. In the public parameters produced by $\mathsf{SetupABO}$, the value $u^{\mathsf{tag}} \cdot v$ thus has to be in $\mathbb{G} \backslash \mathbb{G}_{p_1}$ for any $\mathsf{tag} \neq \mathsf{tag}^\star$ in such a way that generating fake proofs that $\Phi \in \mathbb{G}_{p_1}$ is impossible. At the same time, $u^{\mathsf{tag}^\star} \cdot v$ should be in $\mathbb{G}_{p_1}$ so that fake proofs can be generated for $\mathsf{tag}^\star$.

SetupSound(λ, t, n)**:** choose a group $\mathbb{G}$ of composite order $N = p_1 p_2$ for large primes $p_i > 2^{l(\lambda)}$ for each $i \in \{1, 2\}$ and for some polynomial $l : \mathbb{N} \to \mathbb{N}$. Then, conduct the following steps

1. Pick $g \xleftarrow{R} \mathbb{G}_{p_1}$, $u, v \xleftarrow{R} \mathbb{G}$, $x \xleftarrow{R} \mathbb{Z}_N$ and set $X = g^x \in \mathbb{G}_{p_1}$.
2. Choose a random polynomial $P[Z] \in \mathbb{Z}_N[Z]$ of degree $t-1$ such that $P(0) = x$. For each $i \in \{1, \ldots, n\}$, compute $Y_i = g^{P(i)} \in \mathbb{G}_{p_1}$.
3. Select a pairwise independent hash function $H : \mathbb{G} \to \{0,1\}^\ell$, where $\ell \leq l(\lambda) - 2\lambda$. Note that the range $\mathcal{K} = \{0,1\}^\ell$ of H forms a group for the bitwise exclusive OR operation $\odot = \oplus$.
4. Define private key shares $(\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ as $\mathsf{sk}_i = P(i) \in \mathbb{Z}_N$ for each $i = 1$ to n. The vector $(\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ is defined as $\mathsf{vk}_i = Y_i \in \mathbb{G}_{p_1}$ for each i and the public key consists of $\mathsf{pk} = \big((\mathbb{G}, \mathbb{G}_T), N, g, X, u, v, H \big)$. In addition, we have $(\mathcal{C}, \mathcal{V}, \mathcal{K}, \mathcal{K}', \mathcal{R}) = (\mathbb{G}, \mathbb{G}_{p_1}, \{0,1\}^\ell, \mathbb{G}, \mathbb{Z}_N)$.

SetupABO$(\lambda, t, n, \mathsf{tag}^\star)$**:** is identical to $\mathsf{SetupSound}$ with the difference that, instead of being chosen uniformly in $\mathbb{G}$, v is defined as $v = u^{-\mathsf{tag}^\star} \cdot g^\alpha$ for some random $\alpha \xleftarrow{R} \mathbb{Z}_N$. The algorithm also outputs the simulation trapdoor $\tau = \alpha \in \mathbb{Z}_N$.

Sample(pk)**:** parse pk as $\big((\mathbb{G}, \mathbb{G}_T), N, g, X, u, v, H \big)$. Choose $r \xleftarrow{R} \mathbb{Z}_N$, compute $\Phi = g^r \in \mathbb{G}_{p_1}$ and output the pair $(r, \Phi) \in \mathbb{Z}_N \times \mathbb{G}_{p_1}$.

Prove$(\mathsf{pk}, \mathsf{tag}, r, \Phi)$**:** parse pk as $\big((\mathbb{G}, \mathbb{G}_T), N, g, X, u, v, H \big)$ and return $\perp$ if $\Phi \neq g^r$. Otherwise, compute and return $\pi_{\mathrm{SD}} = (u^{\mathsf{tag}} \cdot v)^r$.

SimProve$(\mathsf{pk}, \tau, \mathsf{tag}, \Phi)$**:** return $\perp$ if $\mathsf{tag} \neq \mathsf{tag}^\star$ or if $\Phi \notin \mathbb{G}$. Otherwise, use the simulation trapdoor $\tau = \alpha \in \mathbb{Z}_N$ to compute and output $\pi_{\mathrm{SD}} = \Phi^\alpha$.

Verify$(\mathsf{pk}, \mathsf{tag}, \Phi, \pi_{\mathrm{SD}})$**:** return 1 if and only if $(\Phi, \pi_{\mathrm{SD}}) \in \mathbb{G}^2$ and $e(\Phi, u^{\mathsf{tag}} \cdot v) = e(g, \pi_{\mathrm{SD}})$.

PubEval(pk, r, Φ)**:** on input of $\mathsf{pk} = \big((\mathbb{G}, \mathbb{G}_T), N, g, X, u, v, H \big)$, return $\perp$ if $(r, \Phi) \notin \mathbb{Z}_N \times \mathbb{G}$. Otherwise, compute and return $K = H(X^r) \in \{0,1\}^\ell$.

SharePrivEval$(\mathsf{pk}, \mathsf{sk}_i, \Phi)$**:** return $\perp$ if $\Phi \notin \mathbb{G}$. Otherwise, compute and return (K_i, π_{K_i}), where $K_i = \Phi^{\mathsf{sk}_i} = \Phi^{P(i)}$ and $\pi_{K_i} = \varepsilon$ is simply the empty string.

ShareEvalVerify$(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i})$**:** if $K_i \notin \mathbb{G}$, $\mathsf{vk}_i \notin \mathbb{G}$ or $\pi_{K_i} \neq \varepsilon$, return 0. Otherwise, return 1 if $e(g, K_i) = e(\Phi, \mathsf{vk}_i)$. In any other situation, return 0 (the proof π_{K_i} is completely ignored in this instantiation since, given $\mathsf{vk}_i = Y_i$, the private evaluation share K_i is directly verifiable).

Combine$(\mathsf{pk}, \Phi, \{(K_i, \pi_{K_i})\}_{i \in S})$**:** return $\perp$ if $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i}) = 0$ for some $i \in S$. Otherwise, compute $K = H(\prod_{i \in S} K_i^{\Delta_{i,S}(0)}) = H(\Phi^x) \in \mathcal{K}$ by interpolation in the exponent.

Theorem 2. *The above construction is an all-but-one perfectly sound threshold hash proof system if the SD assumption holds in $\mathbb{G}$. (The proof is given in appendix C.1).*

When the above all-but-one perfectly sound threshold hash proof system is plugged into the generic construction of Section 4, the resulting threshold cryptosystem bears resemblance with the scheme in [39], which makes use of groups whose order is a product of three primes. However, it is more efficient and its security proof is completely different as the dual system encryption approach [50] is not used here.

5.2 Construction from the Decision Linear Assumption in Prime Order Groups

This section presents an all-but-one threshold hash proof system based on the DLIN assumption in prime order bilinear groups. The public key comprises elements $(g, g_1, g_2, X_1, X_2) \in \mathbb{G}^5$, where $X_1 = g_1^{x_1} \cdot g^z$, $X_2 = g_2^{x_2} \cdot g^z$ and (x_1, x_2, z) are part of the private key. The sets $\mathcal{C}$ and $\mathcal{V} \subset \mathcal{C}$ consist of $\mathcal{C} = \mathbb{G}^3$ and $\mathcal{V} = \{(\Phi_1, \Phi_2, \Phi_3) = (g_1^{\theta_1}, g_2^{\theta_2}, g^{\theta_1 + \theta_2}) \mid \theta_1, \theta_2 \in \mathbb{Z}_p\}$, respectively. For any $\Phi = (\Phi_1, \Phi_2, \Phi_3) \in \mathcal{V}$, the public evaluation algorithm computes $X_1^{\theta_1} \cdot X_2^{\theta_2}$, which can be privately evaluated as $\Phi_1^{x_1} \cdot \Phi_2^{x_2} \cdot \Phi_3^z$.

As in the previous instantiation, we append to elements $\Phi \in \mathcal{V}$ a non-interactive proof of their membership of $\mathcal{V}$ (*i.e.*, a proof that $(g, g_1, g_2, \Phi_1, \Phi_2, \Phi_3)$ is a linear tuple) and, in this case, the proof is obtained using the Groth-Sahai techniques (which are recalled in appendix B). However, we cannot simply combine them with a DLIN-based hash proof system in the obvious way. The reason is that, using parameters produced by SetupABO and under the special tag tag*, SimProve must be able to compute a fake non-interactive proof of the statement $\Phi \in \mathcal{V}$ for an element $\Phi \notin \mathcal{V}$. At the same time, we should make sure that, for any tag such that tag $\neq$ tag*, it will be impossible to simulate such proofs. To solve this problem, we need a form of one-time simulation soundness [46] which can be possibly obtained from Groth's simulation-sound non-interactive proofs [29] or a more efficient variant suggested by Katz and Vaikuntanathan [35]. However, the specific language that we consider allows for even more efficient constructions: it is actually possible to build on the Groth-Sahai proofs essentially without any loss of efficiency.

The solution is as follows. After having sampled a linear tuple $\Phi = (\Phi_1, \Phi_2, \Phi_3) \in \mathcal{V}$, the sampler generates his proof using a Groth-Sahai CRS that depends on tag. Algorithm SetupABO produces parameters in the fashion of the all-but-one technique [7]: the tag-based CRS is perfectly WI on the special tag tag* (which allows generating NIZK proofs for this tag) and perfectly sound for any other tag, which makes it impossible to convincingly prove false statements on tags tag $\neq$ tag*. Malkin, Teranishi, Vahlis and Yung [42] used a similar idea of message-dependent CRS in the context of signatures. A difference with [42] is that we do not need to extract witnesses from adversarially-generated proofs and only use them as proofs of membership.

Interestingly, the same technique can be applied to have a more efficient simulation-sound proof of plaintext equality in the Naor-Yung-type [43] cryptosystem in [35][Section 3.2.2]: the proof can be reduced from 60 to 22 group elements and the ciphertext size is decreased by more than 50%.

SetupSound(λ, t, n)**:** Choose a group $\mathbb{G}$ of prime order $p > 2^\lambda$ with generators $g, g_1, g_2, f_1, f_2 \stackrel{R}{\leftarrow} \mathbb{G}$.

1. Choose $x_1, x_2, z \stackrel{R}{\leftarrow} \mathbb{Z}_p$ and set $X_1 = g_1^{x_1} g^z$, $X_2 = g_2^{x_2} g^z$. Define the vectors $\vec{g}_1 = (g_1, 1, g)$ and $\vec{g}_2 = (1, g_2, g)$. Then, pick $\xi_1, \xi_2 \stackrel{R}{\leftarrow} \mathbb{Z}_p$ and define $\vec{g}_3 = \vec{g}_1^{\xi_1} \cdot \vec{g}_2^{\xi_2}$.

2. Choose $\phi_1, \phi_2 \stackrel{R}{\leftarrow} \mathbb{Z}_p$ and define $\vec{f}_1 = (f_1, 1, g)$, $\vec{f}_2 = (1, f_2, g)$ and $\vec{f}_3 = \vec{f}_1^{\phi_1} \cdot \vec{f}_2^{\phi_2} \cdot (1, 1, g)$.

3. Choose random polynomials $P_1[Z], P_2[Z], P[Z] \in \mathbb{Z}_p[Z]$ of degree $t-1$ such that $P_1(0) = x_1$, $P_2(0) = x_2$ and $P(0) = z$. For each $i = 1$ to n, compute $Y_{i,1} = g_1^{P_1(i)} g^{P(i)}$, $Y_{i,2} = g_2^{P_2(i)} g^{P(i)}$.

4. Define private key shares $\mathbf{SK} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ as $\mathsf{sk}_i = (P_1(i), P_2(i), P(i)) \in (\mathbb{Z}_p)^3$ for each $i \in \{1, \ldots, n\}$. Verification keys $\mathbf{VK} = (\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ are defined as $\mathsf{vk}_i = (Y_{i,1}, Y_{i,2}) \in \mathbb{G}^2$ for each $i \in \{1, \ldots, n\}$ and the public key is defined to be

$$\mathsf{pk} = \Big((\mathbb{G}, \mathbb{G}_T),\ g,\ \vec{g}_1,\ \vec{g}_2,\ \vec{g}_3,\ \vec{f}_1,\ \vec{f}_2,\ \vec{f}_3,\ X_1,\ X_2 \Big).$$

As for the sets $(\mathcal{C}, \mathcal{K}, \mathcal{K}', \mathcal{R})$, they are defined as $\mathcal{C} = \mathbb{G}^3$, $\mathcal{K} = \mathcal{K}' = \mathbb{G}$ and $\mathcal{R} = (\mathbb{Z}_p)^2$, respectively. The subset $\mathcal{V} \subset \mathcal{C}$ consists of the language $(\Phi_1, \Phi_2, \Phi_3) \in \mathbb{G}^3$ for which there

exists $\theta_1, \theta_2 \in \mathbb{Z}_p$ such that $\Phi_1 = g_1^{\theta_1}$, $\Phi_2 = g_2^{\theta_2}$ and $\Phi_3 = g^{\theta_1 + \theta_2}$.

SetupABO$(\lambda, t, n, \mathsf{tag}^\star)$**:** is identical to SetupSound with the following differences.

1. In step 1, $\vec{g_3}$ is set as $\vec{g_3} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2} \cdot (1, 1, g)^{-\mathsf{tag}^\star}$ instead of being chosen in $\mathrm{span}(\vec{g_1}, \vec{g_2})$.
2. In step 2, the vectors $(\vec{f_1}, \vec{f_2}, \vec{f_3})$ are chosen so as to have $\vec{f_3} = \vec{f_1}^{\phi_1} \cdot \vec{f_2}^{\phi_2}$.
3. The algorithm additionally outputs the trapdoor $\tau = (\xi_1, \xi_2, \phi_1, \phi_2) \in (\mathbb{Z}_p)^4$.

Sample(pk)**:** choose $\theta_1, \theta_2 \xleftarrow{R} \mathbb{Z}_p$, compute $\Phi = (\Phi_1, \Phi_2, \Phi_3) = (g_1^{\theta_1}, g_2^{\theta_2}, g^{\theta_1 + \theta_2})$ and output $\big((\theta_1, \theta_2), \Phi\big)$.

Prove$\big(\mathsf{pk}, \mathsf{tag}, (\theta_1, \theta_2), \Phi\big)$**:** parse pk as $\big((\mathbb{G}, \mathbb{G}_T), g, \vec{g_1}, \vec{g_2}, \vec{g_3}, \vec{f_1}, \vec{f_2}, \vec{f_3}, X_1, X_2\big)$ and Φ as (Φ_1, Φ_2, Φ_3). Construct[1] a vector $\vec{g}_{\mathsf{tag}} = \vec{g_3} \cdot (1, 1, g)^{\mathsf{tag}}$ and use $\mathbf{g}_{\mathsf{tag}} = (\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}})$ as a Groth-Sahai CRS to generate a NIZK proof that $(g, g_1, g_2, \Phi_1, \Phi_2, \Phi_3)$ is a linear tuple. More precisely, generate commitments $\vec{C}_{\theta_1}, \vec{C}_{\theta_2}$ to exponents $\theta_1, \theta_2 \in \mathbb{Z}_p$ (in other words, compute $\vec{C}_{\theta_i} = \vec{g}_{\mathsf{tag}}^{\theta_i} \cdot \vec{g_1}^{r_i} \cdot \vec{g_2}^{s_i}$ with $r_i, s_i \xleftarrow{R} \mathbb{Z}_p$ for each $i \in \{1, 2\}$) and a proof $\pi_{(\theta_1, \theta_2)}$ that they satisfy

$$\Phi_1 = g_1^{\theta_1}, \qquad \Phi_2 = g_2^{\theta_2}, \qquad \Phi_3 = g^{\theta_1 + \theta_2}. \tag{2}$$

The whole proof π_{LIN} for (2) consists of $\vec{C}_{\theta_1}, \vec{C}_{\theta_2}$ and $\pi_{(\theta_1, \theta_2)}$ (see appendix E.1 for details about the generation of this proof) and requires 12 elements of $\mathbb{G}$.

SimProve$(\mathsf{pk}, \tau, \mathsf{tag}, \Phi)$**:** parses pk as above, τ as $(\xi_1, \xi_2, \phi_1, \phi_2) \in (\mathbb{Z}_p)^4$ and Φ as $(\Phi_1, \Phi_2, \Phi_3) \in \mathbb{G}^3$. If $\mathsf{tag} \neq \mathsf{tag}^\star$, return $\bot$. Otherwise, the commitments $\vec{C}_{\theta_1}, \vec{C}_{\theta_2}$ and the proof π_{LIN} must be generated for the CRS $\mathbf{g}_{\mathsf{tag}^\star} = (\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}^\star})$, where $\vec{g}_{\mathsf{tag}^\star} = \vec{g_3} \cdot (1, 1, g)^{\mathsf{tag}^\star} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$, which is a Groth-Sahai CRS for the witness indistinguishability setting (as recalled in appendix B).

1. Using the trapdoor (ξ_1, ξ_2), simulate proofs for multi-exponentiation equations (see appendix E.1 for details as to how such proofs can be simulated). That is, generate $\vec{C}_{\theta_1}, \vec{C}_{\theta_2}$ as commitments to 0 and compute $\pi_{(\theta_1, \theta_2)}$ as a simulated proof that relations (2) hold.
2. Output $\pi_{\mathrm{LIN}} = (\vec{C}_{\theta_1}, \vec{C}_{\theta_2}, \pi_{(\theta_1, \theta_2)})$ that consists of perfectly hiding commitments and simulated NIZK proofs which, on the CRS $(\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}^\star})$, are distributed as real proofs.

Verify$(\mathsf{pk}, \mathsf{tag}, \Phi, \pi_{\mathrm{LIN}})$**:** parse pk and Φ as above and π_{LIN} as $(\vec{C}_{\theta_1}, \vec{C}_{\theta_2}, \pi_{(\theta_1, \theta_2)}) \in \mathbb{G}^{12}$. Then, compute $\vec{g}_{\mathsf{tag}} = \vec{g_3} \cdot (1, 1, g)^{\mathsf{tag}}$ and use $\mathbf{g}_{\mathsf{tag}} = (\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}})$ as a Groth-Sahai CRS to verify the proof π_{LIN}. If the latter is deemed as a valid proof for the relations (2), return 1. Otherwise, return 0.

PubEval$\big(\mathsf{pk}, (\theta_1, \theta_2), \Phi\big)$**:** parse pk and Φ as above. Return $\bot$ if $(\Phi_1, \Phi_2, \Phi_3) \neq (g_1^{\theta_1}, g_2^{\theta_2}, g^{\theta_1 + \theta_2})$. Otherwise, compute and return $K = X_1^{\theta_1} \cdot X_2^{\theta_2} \in \mathcal{K}$.

SharePrivEval$(\mathsf{pk}, \mathsf{sk}_i, \Phi)$**:** parse sk_i as $(P_1(i), P_2(i), P(i)) \in (\mathbb{Z}_p)^3$ and return $\bot$ if $\Phi \notin \mathbb{G}^3$. Otherwise, compute and return a pair (K_i, π_{K_i}), where $K_i = \Phi_1^{P_1(i)} \cdot \Phi_2^{P_2(i)} \cdot \Phi_3^{P(i)} \in \mathcal{K}'$ and $\pi_{K_i} = (\vec{C}_{P_1}, \vec{C}_{P_2}, \vec{C}_P, \pi'_{K_i}) \in \mathbb{G}^{15}$ is a proof consisting of commitments $\vec{C}_{P_1}, \vec{C}_{P_2}, \vec{C}_P$ to exponents $P_1(i), P_2(i), P(i) \in \mathbb{Z}_p$ and a proof π'_{K_i} that these satisfy the equations

$$K_i = \Phi_1^{P_1(i)} \cdot \Phi_2^{P_2(i)} \cdot \Phi_3^{P(i)}, \qquad Y_{i,1} = g_1^{P_1(i)} g^{P(i)}, \qquad Y_{i,2} = g_2^{P_2(i)} g^{P(i)}. \tag{3}$$

The perfectly binding commitments $\vec{C}_{P_1}, \vec{C}_{P_2}, \vec{C}_P$ and the proof π'_{K_i} are generated using the vectors $\mathbf{f} = (\vec{f_1}, \vec{f_2}, \vec{f_3})$ as a Groth-Sahai CRS (in such a way that $\vec{C}_{P_1} = \vec{f_3}^{P_1(i)} \cdot \vec{f_1}^{r_{P_1}} \cdot \vec{f_2}^{s_{P_1}}$, for some $r_{P_1}, s_{P_1} \xleftarrow{R} \mathbb{Z}_p$, for example).

---

[1] We assume that tags are non-zero. This can be enforced by having Prove and Verify output $\bot$ when $\mathsf{tag} = 0$.

ShareEvalVerify$(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i})$: parse vk_i as $(Y_{i,1}, Y_{i,2}) \in \mathbb{G}^2$ and return $\perp$ if (K_i, π_{K_i}) cannot be parsed as a tuple in $\mathbb{G} \times \mathbb{G}^{15}$. Otherwise, parse π_{K_i} as $\pi_{K_i} = (\vec{C}_{P_1}, \vec{C}_{P_2}, \vec{C}_P, \pi'_{K_i}) \in \mathbb{G}^{15}$ and return 1 if π'_{K_i} is a valid proof for equations (3). In any other situation, return 0.

Combine$(\mathsf{pk}, \Phi, \{(K_i, \pi_{K_i})\}_{i \in S})$: return $\perp$ if $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_i, \Phi, K_i, \pi_{K_i}) = 0$ for some $i \in S$. Otherwise, compute $K = \prod_{i \in S} K_i^{\Delta_{i,S}(0)} = \Phi_1^{x_1} \cdot \Phi_2^{x_2} \cdot \Phi_3^z \in \mathcal{K}$.

Theorem 3. *The above construction is an all-but-one perfectly sound threshold hash proof system assuming that the DLIN assumption holds in $\mathbb{G}$.* (The proof is given in appendix C.2.)

The proof π_{LIN} takes 6 group elements whereas commitments $\vec{C}_{\theta_1}, \vec{C}_{\theta_2}$ require 3 group elements each. If the scheme is instantiated using Groth's one-time signature [29] (which relies on the discrete logarithm assumption), SVK and σ demand 3 and 2 group elements, respectively. The whole ciphertext C thus consists of 21 group elements. Concretely, if each element has a representation of 512 bits, at the 128-bit security level, the ciphertext overhead amounts to 10240 bits.

From a computational standpoint, assuming that a multi-exponentiation with two base elements has roughly the same cost as a single-base exponentiation, the sender has to compute 19 exponentiations in $\mathbb{G}$ (we include the cost of generating SVK which incurs three exponentiations in Groth's one-time signature [29]). As for the verifier's workload, the validity of a ciphertext can be checked by computing a product of 12 pairings (which is significantly more efficient than naively evaluating 12 individual pairings) using probabilistic batch verification techniques as in [6].

In appendix D, we show an even more efficient instantiation based on the Symmetric eXternal Diffie-Hellman assumption in prime order groups: only 6 pairing evaluations suffice to check $\pi_{\mathcal{V}}$.

Acknowledgements

References

1. M. Abe. Robust Distributed Multiplicaton with out Interaction. In *Crypto'99*, *LNCS* 1666, pp. 130–147, 1999.
2. M. Abe, S. Fehr. Adaptively Secure Feldman VSS and Applications to Universally-Composable Threshold Cryptography. In *Crypto'04*, *LNCS* 3152, pp. 317–334, 2004.
3. J. Almansa, I. Damgård, J.-B. Nielsen. Simplified Threshold RSA with Adaptive and Proactive Security. In *Eurocrypt'06*, *LNCS* 4004, pp. 593–611, 2006.
4. P. Barreto, M. Naehrig. Pairing-Friendly Elliptic Curves of Prime Order. In *SAC'05*, *LNCS* 3897, pp. 319–331, 2005.
5. M. Bellare, P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS*, pp. 62–73, 1993.
6. O. Blazy, G. Fuchsbauer, M. Izabachène, A. Jambert, H. Sibert, D. Vergnaud. Batch Groth-Sahai. In *Applied Cryptography and Network Security (ACNS'10)*, *LNCS* 6123, pp. 218–235, 2010.
7. D. Boneh, X. Boyen. Efficient Selective-ID Secure Identity-Based Encryption Without Random Oracles. In *Eurocrypt'04*, *LNCS* 3027, pp. 223–238, 2004.
8. D. Boneh, X. Boyen, S. Halevi. Chosen Ciphertext Secure Public Key Threshold Encryption Without Random Oracles. In *CT-RSA'06*, *LNCS* 3860, pp. 226–243, 2006.
9. D. Boneh, X. Boyen, H. Shacham. Short group signatures. In *Crypto'04*, LNCS 3152, pp. 41–55, 2004.
10. D. Boneh, M. Franklin. Identity-Based Encryption from the Weil Pairing. In *SIAM* J. of Computing 32(3), pp. 586–615, 2003. Earlier version in *Crypto'01*.
11. D. Boneh, E.-J. Goh, K. Nissim. Evaluating 2-DNF formulas on ciphertexts. In *Theory of Cryptography Conference – TCC 2005*, *LNCS* 3378, pp. 325–341. Springer, 2005.

12. C. Boyd. Digital Multisignatures. In *Cryptography and Coding* (H.J. Beker and F.C. Piper Eds.), Oxford University Press, pp. 241–246, 1989.

13. X. Boyen, Q. Mei, B. Waters. Direct Chosen Ciphertext Security from Identity-Based Techniques. in *ACM CCS'05*, pp. 320–329, 2005.

14. R. Canetti, R. Gennaro, S. Jarecki, H. Krawczyk, T. Rabin. Adaptive Security for Threshold Cryptosystems. In *Crypto'99*, *LNCS* 1666, pp. 98–115, 1999.

15. R. Canetti, S. Goldwasser. An Efficient Threshold Public Key Cryptosystem Secure Against Adaptive Chosen Ciphertext Attack. In *Eurocrypt'99*, *LNCS* 1592, pp. 90–106, 1999.

16. R. Cramer, V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *Crypto'98*, LNCS 1462, pp. 13–25, 1998.

17. R. Cramer, V. Shoup. Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In *Eurocrypt'02*, *LNCS* 2332, pp. 45–64, 2002.

18. R. Canetti, S. Halevi, J. Katz. Chosen-Ciphertext Security from Identity-Based Encryption. In *Eurocrypt'04*, *LNCS* 3027, pp. 207–222, 2004.

19. R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, T. Rabin. Efficient Multi-Party Computations Secure Against an Adaptive Adversary. In *Eurocrypt'99*, *LNCS* 1592, pp. 311–326, 1999.

20. R. Cramer, I. Damgård, Y. Ishai. Share Conversion, Pseudorandom Secret-Sharing and Applications to Secure Computation. In *TCC'05*, *LNCS* 3378, pp. 342–362, 2005.

21. I. Damgård. Towards Practical Public Key Systems Secure Against Chosen Ciphertext Attacks. In *Crypto'91*, *LNCS* 576, pp. 445–456 1991.

22. Y. Desmedt. Society and Group Oriented Cryptography: A New Concept. In *Crypto'87*, *LNCS* 293, pp. 120–127, 1987.

23. Y. Desmedt, Y. Frankel. Threshold Cryptosystems. In *Crypto'89*, *LNCS* 435, pp. 307–315, 1989.

24. Y. Dodis, J. Katz. Chosen-Ciphertext Security of Multiple Encryption. In *TCC'05*, *LNCS* 3378, pp. 188–209, 2005.

25. P.-A. Fouque, D. Pointcheval. Threshold Cryptosystems Secure against Chosen-Ciphertext Attacks. In *Asiacrypt'01*, *LNCS* 2248, pp. 351–368, 2001.

26. Y. Frankel, P. MacKenzie, M. Yung. Adaptively-Secure Distributed Public-Key Systems. In *ESA'99*, *LNCS* 1643, pp. 4–27, 1999.

27. Y. Frankel, P. MacKenzie, M. Yung. Adaptively-Secure Optimal-Resilience Proactive RSA. In *Asiacrypt'99*, *LNCS* 1716, pp. 180–194, 1999.

28. D. Freeman. Converting Pairing-Based Cryptosystems from Composite-Order Groups to Prime-Order Groups. In *Eurocrypt'10*, *LNCS* 6110, pp. 44–61, 2010.

29. J. Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In *Asiacrypt 2006*, *LNCS* 4284, pp. 444–459, 2006.

30. J. Groth, R. Ostrovsky, A. Sahai. Perfect non-interactive zero knowledge for NP. In *Eurocrypt'06*, volume 4004 of *Lecture Notes in Computer Science*, pages 339–358. Springer, 2006.

31. J. Groth, A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt'08*, LNCS 4965, pp. 415–432, 2008.

32. J. Håstad, R. Impagliazzo, L. Levin, M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, vol. 28(4), pp. 1364–1396, 1999.

33. D. Hofheinz, E. Kiltz. The Group of Signed Quadratic Residues and Applications. In *Crypto'09*, *LNCS* 5677, pp. 637–653, 2009.

34. S. Jarecki, A. Lysyanskaya. Adaptively Secure Threshold Cryptography: Introducing Concurrency, Removing Erasures. In *Eurocrypt'00*, *LNCS* 1807, pp. 221–242, 2000.

35. J. Katz, V. Vaikuntanathan. Round-Optimal Password-Based Authenticated Key Exchange. In *TCC'11*, LNCS 6597, pp. 293–310, 2011.

36. E. Kiltz. Chosen-ciphertext security from tag-based encryption. In *TCC'06*, LNCS 3876, pp. 581–600, 2006.

37. E. Kiltz, K. Pietrzak, M. Stam, M. Yung. A New Randomness Extraction Paradigm for Hybrid Encryption. In *Eurocrypt'09*, *LNCS* 5479, pp. 590–609, 2009.

38. A. Lewko, B. Waters. New Techniques for Dual System Encryption and Fully Secure HIBE with Short Ciphertexts. In *TCC 2010*, LNCS 5978, pp. 455–479, 2010.

39. B. Libert, M. Yung. Adaptively Secure Non-Interactive Threshold Cryptosystems. In *ICALP 2011*, LNCS 6756, pp. 588–600, 2011.

40. P. MacKenzie. An Efficient Two-Party Public Key Cryptosystem Secure against Adaptive Chosen Ciphertext Attack. In *PKC'03*, LNCS 2567, pp. 47–61, 2003.

41. P. MacKenzie, M. Reiter, K. Yang. Alternatives to non-malleability: Definitions, constructions, and applications. In *TCC'04*, *LNCS* 2951, pp. 171–190. Springer, 2004.
42. T. Malkin, I. Teranishi, Y. Vahlis, M. Yung. Signatures resilient to continual leakage on memory and computation. In *TCC'11*, *LNCS* 6597, pp. 89–106, 2011.
43. M. Naor, M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC'90*, ACM Press, 1990.
44. R. Ostrovsky, M. Yung. How to Withstand Mobile Virus Attacks. In 10^{th} *ACM Symp. on Principles of Distributed Computing (PODC'91)*, 1991.
45. C. Rackoff, D. Simon. Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In *Crypto'91*, *LNCS* 576, pp. 433–444, 1991.
46. A. Sahai. Non-Malleable Non-Interactive Zero Knowledge and Adaptive Chosen-Ciphertext Security. In *FOCS'99*, pp. 543–553, 1999.
47. M. Scott. Authenticated ID-based Key Exchange and remote log-in with simple token and PIN number. Cryptology ePrint Archive: Report 2002/164.
48. A. Shamir. Identity-Based Cryptosystems and Signature Schemes. In *Crypto'84*, *LNCS* 196, pp. 47–53, 1984.
49. V. Shoup, R. Gennaro. Securing Threshold Cryptosystems against Chosen Ciphertext Attack. In *J. of Cryptology*, 15(2), pp. 75–96, 2002. Earlier version in *Eurocrypt'98*, *LNCS* 1403, pp. 1–16, 1998.
50. B. Waters. Dual System Encryption: Realizing Fully Secure IBE and HIBE under Simple Assumptions. In *Crypto'09*, *LNCS* 5677, pp. 619–636, 2009.
51. H. Wee. Efficient Chosen-Ciphertext Security via Extractable Hash Proofs. In *Crypto'10*, *LNCS* 6223, pp. 314–332, 2010.
52. H. Wee. Threshold and Revocation Cryptosystems via Extractable Hash Proofs. In *Eurocrypt'11*, *LNCS* 6632, pp. 589–609, 2011.

A One-time Signatures

A one-time signature scheme is a triple of algorithms $\Sigma = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ such that, on input of a security parameter λ, $\mathcal{G}$ generates a one-time key pair $(\mathsf{SSK}, \mathsf{SVK}) \leftarrow \Sigma.\mathcal{G}(\lambda)$ while, for any message M, $\Sigma.\mathsf{Ver}(\mathsf{SVK}, M, \sigma)$ outputs 1 whenever $\sigma = \Sigma.\mathsf{Sig}(\mathsf{SSK}, M)$ and 0 otherwise.

As in [18, 24], we need strongly unforgeable one-time signatures: no PPT adversary can be able to create a new signature for a previously signed message.

Definition 4. $\Sigma = (\mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ *is a strongly unforgeable one-time signature if the probability*

$$
\begin{aligned}
\mathbf{Adv}^{\mathsf{OTS}} = \Pr\big[\ & (\mathsf{SSK}, \mathsf{SVK}) \leftarrow \mathcal{G}(\lambda); (M, St) \leftarrow \mathcal{F}(\mathsf{SVK}); \\
& \sigma \leftarrow \Sigma.\mathsf{Sig}(\mathsf{SSK}, M); (M', \sigma') \leftarrow \mathcal{F}(M, \sigma, \mathsf{SVK}, St) : \\
& \Sigma.\mathsf{Ver}(\sigma', \mathsf{SVK}, M') = 1 \wedge (M', \sigma') \neq (M, \sigma)\ \big],
\end{aligned}
$$

where St denotes the state information maintained by $\mathcal{F}$ between stages, is negligible for any PPT forger $\mathcal{F}$.

B Groth-Sahai Proof Systems

In the following notations, for equal-dimension vectors $\vec{A}$ and $\vec{B}$ containing group elements, $\vec{A} \cdot \vec{B}$ stands for their component-wise product.

In their instantiation based on the DLIN assumption in symmetric bilinear groups $(\mathbb{G}, \mathbb{G}_T)$, Groth-Sahai (GS) proofs [31] use a common reference string comprising vectors $\vec{g_1}, \vec{g_2}, \vec{g_3} \in \mathbb{G}^3$, where $\vec{g_1} = (g_1, 1, g)$, $\vec{g_2} = (1, g_2, g)$ for some $g_1, g_2 \in \mathbb{G}$. A commitment to $X \in \mathbb{G}$ is obtained as $\vec{C} = (1, 1, X) \cdot \vec{g_1}^{\,r} \cdot \vec{g_2}^{\,s} \cdot \vec{g_3}^{\,t}$ with $r, s, t \xleftarrow{R} \mathbb{Z}_p^*$. When proofs should be perfectly sound, $\vec{g_3}$ is set as

$\vec{g_3} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$, with $\xi_1, \xi_2 \xleftarrow{R} \mathbb{Z}_p^*$, so that $\vec{C} = (g_1^{r+\xi_1 t}, g_2^{s+\xi_2 t}, X \cdot g^{r+s+t(\xi_1+\xi_2)})$ is a Boneh-Boyen-Shacham (BBS) encryption [9] that can be decrypted using $\alpha_1 = \log_g(g_1)$, $\alpha_2 = \log_g(g_2)$. In the witness indistinguishability (WI) setting, $\vec{g_1}, \vec{g_2}, \vec{g_3}$ are linearly independent and $\vec{C}$ is a perfectly hiding commitment. Under the DLIN assumption, the two settings are indistinguishable.

To commit to an exponent $x \in \mathbb{Z}_p$, one computes $\vec{C} = \vec{\varphi}^x \cdot \vec{g_1}^r \cdot \vec{g_2}^s$, with $r, s \xleftarrow{R} \mathbb{Z}_p^*$, using a CRS comprising vectors $\vec{\varphi}, \vec{g_1}, \vec{g_2}$. In the soundness setting $\vec{\varphi}, \vec{g_1}, \vec{g_2}$ are linearly independent vectors (typically, one chooses $\vec{\varphi} = \vec{g_3} \cdot (1, 1, g)$ where $\vec{g_3} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$) whereas, in the WI setting, choosing $\vec{\varphi} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$ gives a perfectly hiding commitment since $\vec{C}$ is always a BBS encryption of $1_{\mathbb{G}}$. On a perfectly sound CRS (where $\vec{g_3} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$ and $\vec{\varphi} = \vec{g_3} \cdot (1, 1, g)$), commitments to exponents are not fully extractable since the trapdoor (α_1, α_2) only allows recovering g^x from $\vec{C} = \vec{\varphi}^x \cdot \vec{g_1}^r \cdot \vec{g_2}^s$.

To prove that committed variables satisfy certain relations, the techniques of [31] require one commitment per variable and one proof element per relation. Such efficient proofs notably exist for multi-exponentiation equations which are equations of the form

$$\prod_{i=1}^m \mathcal{A}_i^{y_i} \cdot \prod_{j=1}^n \mathcal{X}_j^{b_j} \cdot \prod_{i=1}^m \cdot \prod_{j=1}^n \mathcal{X}_j^{y_i \gamma_{ij}} = T,$$

for variables $\mathcal{X}_1, \ldots, \mathcal{X}_n \in \mathbb{G}$, $y_1, \ldots, y_m \in \mathbb{Z}_p$ and constants $T, \mathcal{A}_1, \ldots, \mathcal{A}_m \in \mathbb{G}$, $b_1, \ldots, b_n \in \mathbb{Z}_p$ and $\gamma_{ij} \in \mathbb{G}$, for $i \in \{1, \ldots, m\}, j \in \{1, \ldots, n\}$.

Multi-exponentiation equations admit zero-knowledge proofs at no additional cost. On a simulated CRS (prepared for the WI setting), the trapdoor (ξ_1, ξ_2) makes it possible to simulate proofs without knowing witnesses, and simulated proofs are perfectly indistinguishable from real proofs.

For linear equations (*i.e.*, when $\gamma_{ij} = 0$ for all i, j) depends on the form of the considered equation. Namely, linear multi-exponentiation equations of the type $\prod_{j=1}^n \mathcal{X}_j^{b_j} = T$ (resp. $\prod_{i=1}^m \mathcal{A}_i^{y_i} = T$) demand 3 (resp. 2) group elements.

The Groth-Sahai techniques can also be instantiated in groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ with an asymmetric bilinear map $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$, where $\mathbb{G} \neq \hat{\mathbb{G}}$. In this case, they rely on the Symmetric eXternal Diffie-Hellman assumption according to which the DDH problem is hard in both $\mathbb{G}$ and $\hat{\mathbb{G}}$. In this setting, we only use them to prove multi-exponentiation equations of the form $\prod_{i=1}^m \mathcal{A}_i^{y_i} = T$, for constants $\mathcal{A}_1, \ldots, \mathcal{A}_m, T \in \mathbb{G}$ and variables $y_1, \ldots, y_m \in \mathbb{Z}_p$. To this end, commitments to exponents $x \in \mathbb{Z}_p$ have to be computed in $\hat{\mathbb{G}}$. The common reference string includes vector $\vec{u}, \vec{u_1} \in \hat{\mathbb{G}}^2$ and commitments are calculated as per $\vec{C} = \vec{u}^x \cdot \vec{u_1}^r$, with $r \xleftarrow{R} \mathbb{Z}_p$. It is easy to see that the commitment $\vec{C}$ is perfectly hiding if $(\vec{u}, \vec{u_1})$ are linearly independent and perfectly hiding if $\vec{u} \in \text{span}(\vec{u_1})$. The corresponding proof for the equation $\prod_{i=1}^m \mathcal{A}_i^{y_i} = T$ will be perfectly sound if $\vec{u} \notin \text{span}(\vec{u_1})$ and perfectly WI if $\vec{u} \in \text{span}(\vec{u_1})$. In either case, the proof consists of a single element of $\mathbb{G}$.

C Deferred Proofs

C.1 Proof of Theorem 2

The theorem is proved by demonstrating that, under the Subgroup Decision assumption, the scheme provides all the properties required from an all-but-one perfectly sound HPS.

The subset membership hardness property is straightforward as it is exactly the Subgroup Decision assumption in this instantiation. The simulatability of share proofs is also trivial to verify since no non-interactive proof is needed to check the validity of private evaluation shares. We thus

focus on remaining properties.

To prove the universality property, we rely on the leftover hash lemma [32].

Lemma 1. *Let $X \in \mathcal{X}$ be a random variable such that $H_\infty(X) \geq k$ and let $\mathcal{H}$ be a family of pairwise independent hash functions with domain $\mathcal{X}$ and range $\{0,1\}^\ell$. Then, if $H \overset{R}{\leftarrow} \mathcal{H}$, we have*

$$\Delta((H, H(X)), (H, U_\ell)) \leq 1/2^{(k-\ell)/2},$$

where U_ℓ denotes the uniform distribution over $\{0,1\}^\ell$.

SETUP INDISTINGUISHABILITY. The only difference between the outputs $(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^n, \{\mathsf{vk}_i\}_{i=1}^n)$ of $\mathsf{SetupSound}(\lambda, t, n)$ and $\mathsf{SetupABO}(\lambda, t, n, \mathsf{tag}^\star)$ is the distribution of $v \in \mathbb{G}$ which is uniform in $\mathbb{G}$ in the former case and equals $v = u^{-\mathsf{tag}^\star} \cdot g^\alpha$, where $\alpha \overset{R}{\leftarrow} \mathbb{Z}_N$, when it is returned by $\mathsf{SetupABO}$. In the latter situation, v can be seen a Boneh-Goh-Nissim encryption [11] of $-\mathsf{tag}^\star$ whereas a uniformly random $v \in_R \mathbb{G}$ can be interpreted as a BGN encryption of a random plaintext. Consequently, the public outputs of $\mathsf{SetupSound}$ and $\mathsf{SetupABO}$ cannot be told apart if the Subgroup Decision assumption (which is equivalent to the semantic security of the BGN cryptosystem) holds.

CORRECTNESS AND PUBLIC EVALUABILITY ON $\mathcal{V}$. Since the public values $\mathsf{pk} = g^x$ and $\mathsf{vk}_i = g^{P(i)}$ uniquely determine $P[X] \bmod p_1$ as well as $\mathsf{sk}_i \bmod p_1$, for any $\Phi \in \mathbb{G}_{p_1}$, there is only one possible value $\mathsf{SharePrivEval}(\mathsf{pk}, \mathsf{sk}_i, \Phi) = (K_i, \varepsilon) = (\Phi^{P(i)}, \varepsilon)$. The second condition is immediate to verify.

UNIVERSALITY. Let Φ be a random element of order N in $\mathbb{G}$. For any $(t-1)$-subset $\bar{S} \subset \{1, \ldots, n\}$, if we consider the min-entropy of Φ^x given Φ, g^x, $\{\mathsf{vk}_i = g^{P(i)}\}_{i=1}^n$ and $\{\mathsf{sk}_i = P(i)\}_{i \in \bar{S}}$, we have

$$H_\infty\big(\Phi^x | (g, \Phi, g^x, \{\mathsf{vk}_i\}_{i=1}^n, \{\mathsf{sk}_i\}_{i \in \bar{S}})\big) = H_\infty\big(\Phi^x \mid (g, \Phi, g^x)\big) = H_\infty(x \bmod N \mid x \bmod p_1)$$
$$= H_\infty(x \bmod p_2 \mid x \bmod p_1) = H_\infty(x \bmod p_2) = \log p_2.$$

Lemma 1 tells us that the statistical distance

$$\Delta\left[\big((g^x, \{\mathsf{vk}_i\}_{i=1}^n, \{sk_i\}_{i \in \bar{S}}, \Phi), H, H(\Phi^x)\big), \big((g^x, \{\mathsf{vk}_i\}_{i=1}^n, \{sk_i\}_{i \in \bar{S}}, \Phi), H, U_\ell\big)\right],$$

is smaller than $1/2^{(\log p_2 - \ell)/2} < 1/2^{(l(\lambda) - \ell)/2}$, which is negligible as long as $l(\lambda) \geq 2\lambda + \ell$.

ALL-BUT-ONE SOUNDNESS. Since $\mathsf{SetupABO}$ chooses u at random in $\mathbb{G}$, u has a non-trivial $\mathbb{G}_{p_2}$ component with overwhelming probability. We know that $u^{\mathsf{tag}} \cdot v = u^{\mathsf{tag} - \mathsf{tag}^\star} \cdot g^\alpha$ has a non-trivial $\mathbb{G}_{p_2}$ component whenever $\mathsf{tag} \neq \mathsf{tag}^\star$. The equality $e(\Phi, u^{\mathsf{tag}} \cdot v) = e(g, \pi_{\mathrm{SD}})$ – more precisely, the fact that its right-hand-side member has order p_1 – thus guarantees that $\Phi \in \mathbb{G}_{p_1}$ as long as $\mathsf{tag} \neq \mathsf{tag}^\star$. At the same time, $u^{\mathsf{tag}^\star} \cdot v = g^\alpha$ has order p_1, and the trapdoor allows simulating proofs that $\Phi \in \mathbb{G}_{p_1}$. When Φ is really in the subgroup $\mathbb{G}_{p_1}$, $\pi_{\mathrm{SD}} = \Phi^\alpha$ equals the proof that would be produced using the real witness $r = \log_g(\Phi)$. When, $\Phi \in_R \mathbb{G}$, $\pi_{\mathrm{SD}} = \Phi^\alpha$ still satisfies the equality $e(\Phi, u^{\mathsf{tag}^\star} \cdot v) = e(g, \pi_{\mathrm{SD}})$ and can thus serve as a simulated proof that $\Phi \in \mathbb{G}_{p_1}$.

CONSISTENCY. Let us assume that a PPT adversary $\mathcal{A}$ can break the consistency property of the all-but-one HPS with non-negligible probability. We show that $\mathcal{A}$ implies a distinguisher $\mathcal{B}$ for the Subgroup Decision assumption. The distinguisher $\mathcal{B}$ receives $(g \in \mathbb{G}_{p_1}, \ h \in \mathbb{G})$ and $\eta \in \mathbb{G}$ with the aim of deciding if η has a non-trivial $\mathbb{G}_{p_2}$ component. This is done by generating the public key pk using $g \in \mathbb{G}_{p_1}$ and $h \in \mathbb{G}$ and by choosing $\{(\mathsf{sk}_i, \mathsf{vk}_i)\}_{i=1}^n$ as in the specification of the scheme.

The only way for the adversary to break the consistency property is to output $(\mathsf{tag}, \Phi, \pi_{\mathrm{SD}})$ (note that Φ's membership in $\mathbb{G}_{p_1}$ is guaranteed by the perfectly sound proof π_{SD}) and two sets of

decryption shares where at least one share is of the form $K_i = \Phi^{P(i)} \cdot R_{2,i}$, for some $R_{2,i} \in \mathbb{G}_{p_2}^*$. Since $\mathcal{B}$ knows $\mathsf{sk}_i = P(i)$, it can compute $R_{2,i} = K_i/\Phi^{P(i)} \in \mathbb{G}_{p_2}$, which allows deciding whether $\eta \in \mathbb{G}_{p_1}$ by testing whether the equality $e(\eta, R_{2,i}) = 1_{\mathbb{G}_T}$ (which only holds if $\eta \in \mathbb{G}_{p_1}$) holds. $\quad\square$

C.2 Proof of Theorem 3

We show that, under the Decision Linear assumption, the scheme meets all the requirements of all-but-one perfectly sound threshold HPS. The subset membership hardness property is trivially implied by the DLIN assumption and we thus focus on remaining properties.

SETUP INDISTINGUISHABILITY. The difference between the public outputs of $\mathsf{SetupSound}(\lambda, t, n)$ and $\mathsf{SetupABO}(\lambda, t, n, \mathsf{tag}^\star)$ is in the distributions of vectors $\vec{g_3}$ and $\vec{f_3}$ since $\mathsf{SetupSound}$ chooses $\vec{g_3}$ in $\mathsf{span}(\vec{g_1}, \vec{g_2})$ and $\vec{f_3} \notin \mathsf{span}(\vec{f_1}, \vec{f_2})$ whereas $\mathsf{SetupABO}$ proceeds the other way around.

We first prove the indistinguishability of the two possible distributions for $\vec{g_3}$ (the case of $\vec{f_3}$ can be handled in a completely analogous way). To this end, we define an intermediate setup procedure $\mathsf{SetupInt}$ which produces vectors of the form $(\vec{g_1}, \vec{g_2}, \vec{g_3})$ with $\vec{g_1} = (g_1, 1, g)$, $\vec{g_2} = (1, g_2, g)$, $\vec{g_3} = (g_{3,1}, g_{3,2}, g_{3,3})$ where $g_1, g_2, g_{3,1}, g_{3,2}, g_{3,3} \overset{R}{\leftarrow} \mathbb{G}$. The result is obtained by combining the following two claims, the proofs of which are straightforward but given for completeness.

Claim 1. If DLIN holds, no PPT adversary can distinguish the output of $\mathsf{SetupSound}$ from $\mathsf{SetupInt}$.

Proof. We show a distinguisher $\mathcal{B}$ that takes in $(g, g_1, g_2, g_1^{\delta_1}, g_2^{\delta_2}, \chi)$, for some $\delta_1, \delta_2 \overset{R}{\leftarrow} \mathbb{Z}_p$, with the purpose of deciding if $\chi = g^{\delta_1 + \delta_2}$ or $\chi \in_R \mathbb{G}$. To this end, $\mathcal{B}$ defines $\vec{g_1} = (g_1, 1, g)$, $\vec{g_2} = (1, g_2, g)$. As for $\vec{g_3}$, $\mathcal{B}$ defines it as $\vec{g_3} = (g_1^{\delta_1}, g_2^{\delta_2}, \chi)$. It is clear that, if $\chi = g^{\delta_1 + \delta_2}$, $(\vec{g_1}, \vec{g_2}, \vec{g_3})$ is distributed as an output of $\mathsf{SetupSound}$ whereas, if $\chi \in_R \mathbb{G}$, it is an output of $\mathsf{SetupInt}$. $\quad\blacksquare$

Claim 2. If DLIN holds, no PPT adversary can distinguish the outputs of $\mathsf{SetupInt}$ and $\mathsf{SetupABO}$.

Proof. Consider a distinguisher $\mathcal{B}$ that takes as input $(g, g_1, g_2, g_1^{\delta_1}, g_2^{\delta_2}, \chi)$ and decides if $\chi = g^{\delta_1 + \delta_2}$ or $\chi \in_R \mathbb{G}$. To do so, $\mathcal{B}$ defines $\vec{g_1} = (g_1, 1, g)$ and $\vec{g_2} = (1, g_2, g)$. As for the third vector $\vec{g_3}$, $\mathcal{B}$ and computes $\vec{g_3} = (g_1^{\delta_1}, g_2^{\delta_2}, \chi \cdot g^{-\mathsf{tag}^\star})$. If $\chi \in_R \mathbb{G}$, the vector $\vec{g_3}$ has the same distribution no matter if χ is multiplied by $g^{-\mathsf{tag}^\star}$ or not and its distribution corresponds to that of an output of $\mathsf{SetupInt}$. If $\chi = g^{\delta_1 + \delta_2}$, $\vec{g_3}$ is distributed as in parameters produced by $\mathsf{SetupABO}$. $\quad\blacksquare$

CORRECTNESS AND PUBLIC EVALUABILITY ON $\mathcal{V}$. This property is implied by the public evaluability of the underlying standard hash proof system. Namely, for any element $\Phi \in \mathcal{V}$, which is a triple of the form $(\Phi_1, \Phi_2, \Phi_3) = \big(g_1^{\theta_1}, g_2^{\theta_2}, g^{\theta_1 + \theta_2}\big)$ and for each i, the value $K_i = \Phi_1^{P_1(i)} \cdot \Phi_2^{P_2(i)} \cdot \Phi_3^{P(i)}$ equals $Y_{i,1}^{\theta_1} \cdot Y_{i,2}^{\theta_2}$ and is uniquely defined by pk and $\mathsf{vk}_i = (Y_{i,1}, Y_{i,2})$. It is also immediate that combining any t values $K_i = \Phi_1^{P_1(i)} \cdot \Phi_2^{P_2(i)} \cdot \Phi_3^{P(i)}$ allows recovering $X_1^{\theta_1} X_2^{\theta_2}$.

UNIVERSALITY. Let $\Phi = (\Phi_1, \Phi_2, \Phi_3)$ be a random triple in $\mathbb{G}^3$. With overwhelming probability, we have $\Phi_3 \neq g^{\theta_1 + \theta_2}$, where $\theta_1 = \log_{g_1}(\Phi_1)$ and $\theta_2 = \log_{g_2}(\Phi_2)$. For any $(t-1)$-subset $\bar{S} \subset \{1, \ldots, n\}$, given shares $\{\mathsf{sk}_i = (P_1(i), P_2(i), P(i))\}_{i \in \bar{S}}$ and public elements $X_1 = g_1^{x_1} \cdot g^z$, $X_2 = g_2^{x_2} \cdot g^z$, $\{\mathsf{vk}_i = (Y_{i,1}, Y_{i,2}) = \big(g_1^{P_1(i)} g^{P(i)}, g_2^{P_2(i)} g^{P(i)}\big)\}_{i=1}^n$, the value $z = P(0)$ is completely undetermined. Since Φ can be written $(\Phi_1, \Phi_2, \Phi_3) = (g_1^{\theta_1}, g_2^{\theta_2}, g^{\theta_1 + \theta_2 + \theta})$ for some non-zero $\theta \in_R \mathbb{Z}_p^*$, its private evaluation can be expressed as

$$\mathsf{PrivEval}(\mathsf{pk}, \{\mathsf{sk}_i\}_{i=1}^t, \Phi) = \Phi_1^{x_1} \cdot \Phi_2^{x_2} \cdot \Phi_3^z = X_1^{\theta_1} \cdot X_2^{\theta_2} \cdot g^{\theta \cdot z},$$

which is uniformly random since z is itself random and independent of publicly available elements.

ALL-BUT-ONE SOUNDNESS. Algorithm SetupABO chooses $(\vec{g_1}, \vec{g_2}, \vec{g_3})$ in such a way that, for any tag $\neq$ tag*, the vector $\vec{g}_{\mathsf{tag}} = \vec{g_3} \cdot (1,1,g)^{\mathsf{tag}} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2} \cdot (1,1,g)^{\mathsf{tag}-\mathsf{tag}^\star}$ is not in $\mathrm{span}(\vec{g_1}, \vec{g_2})$ and $(\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}})$ forms a Groth-Sahai CRS for the perfect soundness setting. Consequently, for any tag tag $\neq$ tag*, even an unbounded adversary would be unable to produce a convincing proof π_{LIN} for an element $\Phi \notin \mathcal{V}$. At the same time, $\vec{g}_{\mathsf{tag}^\star} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$ is such that $(\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}^\star})$ is a Groth-Sahai CRS for the perfect WI setting, and the trapdoor (ξ_1, ξ_2) makes it possible to generate simulated proofs π_{LIN} for elements $\Phi = (\Phi_1, \Phi_2, \Phi_3) \in \mathbb{G}^3$ that can be outside the language $\mathcal{V}$ of linear tuples. Whenever $(g, g_1, g_2, \Phi_1, \Phi_2, \Phi_3)$ is actually a linear tuple, simulated proofs (see appendix E.1 for details on how to construct them) are distributed exactly as the proofs that would be produced using real witnesses.

SIMULATABILITY OF SHARE PROOFS. In the public parameters produced by SetupABO, the vectors $\mathbf{f} = (\vec{f_1}, \vec{f_2}, \vec{f_3})$ are chosen in such a way that $\vec{f_3} = \vec{f_1}^{\phi_1} \cdot \vec{f_2}^{\phi_2}$. This means that (ϕ_1, ϕ_2) can be used as a trapdoor to generate simulated NIZK proofs $\pi_{K_i} = (\vec{C}_{P_1}, \vec{C}_{P_2}, \vec{C}_P, \pi'_{K_i})$ that committed exponents $(P_1(i), P_2(i), P(i))$ satisfy the multi-exponentiation equations

$$Y_{i,1} = g_1^{P_1(i)} \cdot g^{P(i)}, \qquad Y_{i,2} = g_2^{P_2(i)} g^{P(i)}, \qquad K_i = \Phi_1^{P_1(i)} \cdot \Phi_2^{P_2(i)} \cdot \Phi_3^{P(i)}. \tag{4}$$

Namely, $\vec{C}_{P_1}, \vec{C}_{P_2}, \vec{C}_P$ are generated as commitments to 0 and the proof for (4) is simulated using (ϕ_1, ϕ_2). The resulting proof π'_{K_i} – which is a simulated proof for a true statement – has the same distribution as a real proof.

CONSISTENCY. This property holds unconditionally. This is implied by the perfect soundness of Groth-Sahai proofs. Namely, SetupSound produces common reference strings $\mathbf{g}_{\mathsf{tag}} = (\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}})$ and $\mathbf{f} = (\vec{f_1}, \vec{f_2}, \vec{f_3})$ that are always perfectly sound. This guarantees the impossibility of producing a convincing proof π_{LIN} for an element $\Phi = (\Phi_1, \Phi_2, \Phi_3)$ such that $(g, g_1, g_2, \Phi_1, \Phi_2, \Phi_3)$ is not a linear tuple. Moreover, thanks to the perfect soundness of proofs π_{K_i} for the CRS $\mathbf{f} = (\vec{f_1}, \vec{f_2}, \vec{f_3})$, invalid private evaluation shares K_i can never be accepted by the ShareEvalVerify algorithm. Consequently, there is no way for two distinct sets of acceptable private evaluation shares to yield two distinct private evaluations for a valid $\Phi \in \mathcal{V}$. □

D Instantiation from the SXDH Assumption in Prime Order Groups

The construction of Section 5.2 relies on a well-established assumption in prime order groups and it is described in terms of symmetric pairings for simplicity. However, it readily extends to asymmetric pairing configurations.

Further efficiency improvements can be obtained if we choose to rely on asymmetric pairings $e : \mathbb{G} \times \hat{\mathbb{G}} \to \mathbb{G}_T$ and the Symmetric eXternal Diffie-Hellman assumption (SXDH), which posits that the DDH problem is hard in $\mathbb{G}$ and $\hat{\mathbb{G}}$ when $\mathbb{G} \neq \hat{\mathbb{G}}$ and no isomorphism is efficiently computable between $\mathbb{G}$ and $\hat{\mathbb{G}}$.

In this case, the public key comprises group elements $(g_1, g_2, X) \in \mathbb{G}$ with $X = g_1^{x_1} g_2^{x_2}$ and where $(x_1, x_2) \xleftarrow{R} (\mathbb{Z}_p)^2$ is part of the private key. The public key also includes vectors $(\vec{u}_1, \vec{u}_2)$, where $\vec{u}_1 = (\hat{g}, \hat{h}) \in \hat{\mathbb{G}}^2$ and $\vec{u}_2 = \vec{u}_1^{\rho_u} = (\hat{g}^{\rho_u}, \hat{h}^{\rho_u}) \in \hat{\mathbb{G}}^2$, for some $\rho_u \xleftarrow{R} \mathbb{Z}_p$. It finally contains vectors $(\vec{v}_1, \vec{v}_2)$, where $\vec{v}_1 = (v_{1,1}, v_{1,2}) \in \hat{\mathbb{G}}^2$ and $\vec{v}_2 = \vec{v}_1^{\rho_v} \cdot (1, \hat{g})$, for some $\rho_v \xleftarrow{R} \mathbb{Z}_p$. These vectors $(\vec{v}_1, \vec{v}_2)$ are the counterpart of $(\vec{f_1}, \vec{f_2}, \vec{f_3})$ in Section 5.2 and they form the CRS that allows generating proofs of well-formedness for private evaluation shares.

SetupSound(λ, t, n)**:** Choose a configuration of asymmetric bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ with generators $g_1, g_2 \overset{R}{\leftarrow} \mathbb{G}$ and $\hat{g}, \hat{h} \overset{R}{\leftarrow} \hat{\mathbb{G}}$.

1. Choose $x_1, x_2 \overset{R}{\leftarrow} \mathbb{Z}_p$ and set $X = g_1^{x_1} g_2^{x_2}$. Define vectors $(\vec{u}_1, \vec{u}_2)$, where $\vec{u}_1 = (\hat{g}, \hat{h}) \in \hat{\mathbb{G}}^2$ and $\vec{u}_2 = \vec{u}_1^{\rho_u} = (\hat{g}^{\rho_u}, \hat{h}^{\rho_u}) \in \hat{\mathbb{G}}^2$, for some $\rho_u \overset{R}{\leftarrow} \mathbb{Z}_p$.
2. Choose $v_{1,1}, v_{1,2} \overset{R}{\leftarrow} \hat{\mathbb{G}}$ and define the vectors $(\vec{v}_1, \vec{v}_2)$, where $\vec{v}_1 = (v_{1,1}, v_{1,2}) \in \hat{\mathbb{G}}^2$ and $\vec{v}_2 = \vec{v}_1^{\rho_v} \cdot (1, \hat{g})$, for some $\rho_v \overset{R}{\leftarrow} \mathbb{Z}_p$.
3. Choose random polynomials $P_1[Z], P_2[Z] \in \mathbb{Z}_p[Z]$ of degree $t-1$ such that $P_1(0) = x_1$ and $P_2(0) = x_2$. For each $i \in \{1, \ldots, n\}$, compute $Y_i = g_1^{P_1(i)} g_2^{P_2(i)}$.
4. Define private key shares $\mathbf{SK} = (\mathsf{sk}_1, \ldots, \mathsf{sk}_n)$ as $\mathsf{sk}_i = (P_1(i), P_2(i)) \in (\mathbb{Z}_p)^2$ for each $i \in \{1, \ldots, n\}$. Verification keys $\mathbf{VK} = (\mathsf{vk}_1, \ldots, \mathsf{vk}_n)$ are set as $\mathsf{vk}_i = Y_i \in \mathbb{G}$ for each i and the public key is defined as

$$\mathsf{pk} = \Big((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T), \ g, \ \vec{u}_1, \ \vec{u}_2, \ \vec{v}_1, \ \vec{v}_2, \ X \Big).$$

The sets $(\mathcal{C}, \mathcal{K}, \mathcal{K}', \mathcal{R})$, they are defined as $\mathcal{C} = \mathbb{G}^2$, $\mathcal{K} = \mathcal{K}' = \mathbb{G}$ and $\mathcal{R} = \mathbb{Z}_p$, respectively. The subset $\mathcal{V} \subset \mathcal{C}$ consists of the language $(\Phi_1, \Phi_2) \in \mathbb{G}^2$ for which there exists $\theta \in \mathbb{Z}_p$ such that $\Phi_1 = g_1^\theta$ and $\Phi_2 = g_2^\theta$.

SetupABO$(\lambda, t, n, \mathsf{tag}^\star)$**:** is identical to **SetupSound** with the following differences.

1. In step 1, $\vec{u}_2$ is set as $\vec{u}_2 = \vec{u}_1^{\rho_u} \cdot (1, \hat{g})^{-\mathsf{tag}^\star}$ instead of being chosen in $\mathrm{span}(\vec{g_1}, \vec{g_2})$.
2. In step 2, the vectors $(\vec{v}_1, \vec{v}_2)$ are chosen so as to have $\vec{v}_2 = \vec{v}_1^{\rho_v}$.
3. The algorithm additionally outputs the trapdoor $\tau = (\rho_u, \rho_v) \in (\mathbb{Z}_p)^2$.

Sample(pk)**:** choose $\theta \overset{R}{\leftarrow} \mathbb{Z}_p$ and compute a pair $\Phi = (\Phi_1, \Phi_2) = (g_1^\theta, g_2^\theta)$. Then, output (θ, Φ).

Prove$(\mathsf{pk}, \mathsf{tag}, (\theta_1, \theta_2), \Phi)$**:** parse pk as $\big((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T), g, \vec{u}_1, \vec{u}_2, \vec{v}_1, \vec{v}_2, X \big)$ and Φ as $(\Phi_1, \Phi_2) \in \mathbb{G}^2$. Construct a vector $\vec{u}_{\mathsf{tag}} = \vec{u}_2 \cdot (1, \hat{g})^{\mathsf{tag}}$ and use $\mathbf{u}_{\mathsf{tag}} = (\vec{u}_1, \vec{u}_{\mathsf{tag}})$ as a Groth-Sahai CRS to generate a NIZK proof that $(g, g_1, g_2, \Phi_1, \Phi_2)$ is a Diffie-Hellman tuple. More precisely, generate a commitment $\vec{C}_\theta$ to $\theta \in \mathbb{Z}_p$ (in other words, compute $\vec{C}_\theta = \vec{u}_{\mathsf{tag}}^\theta \cdot \vec{u}_1^r$ with $r \overset{R}{\leftarrow} \mathbb{Z}_p$ for each $i \in \{1, 2\}$) and a proof π_{DH} that it satisfies

$$\Phi_1 = g_1^\theta, \qquad\qquad \Phi_2 = g_2^\theta. \tag{5}$$

The entire proof π_{DH} for (5) consists of $\vec{C}_\theta$ and π_θ (see appendix E.2 for details about the generation of this proof) and requires 2 elements of $\hat{\mathbb{G}}$ and 2 elements of $\mathbb{G}$.

SimProve$(\mathsf{pk}, \tau, \mathsf{tag}, \Phi)$**:** parses pk as above, τ as $(\rho_u, \rho_v) \in (\mathbb{Z}_p)^2$ and Φ as $(\Phi_1, \Phi_2) \in \mathbb{G}^2$. If $\mathsf{tag} \neq \mathsf{tag}^\star$, return $\bot$. Otherwise, the commitment $\vec{C}_\theta$ and the proof π_{DH} must be generated for the CRS $\mathbf{u}_{\mathsf{tag}^\star} = (\vec{u}_1, \vec{u}_{\mathsf{tag}^\star})$, where $\vec{u}_{\mathsf{tag}^\star} = \vec{u}_2 \cdot (1, \hat{g})^{\mathsf{tag}^\star} = \vec{u}_1^{\rho_u}$, which is a Groth-Sahai CRS for the perfect WI setting. The algorithm thus proceeds as follows.

1. Using the trapdoor ρ_u, simulate proofs for multi-exponentiation equations (see appendix E.2 for details). That is, generate $\vec{C}_\theta$ as a commitment to 0 and compute π_θ as a simulated proof for relations (5).
2. Output $\pi_{\mathrm{DH}} = (\vec{C}_\theta, \pi_\theta)$ that consists of perfectly hiding commitments and simulated NIZK proofs which, on the CRS $\mathbf{u}_{\mathsf{tag}^\star} = (\vec{u}_1, \vec{u}_{\mathsf{tag}^\star})$, have the same distribution as real proofs.

Verify$(\mathsf{pk}, \mathsf{tag}, \Phi, \pi_{\mathrm{DH}})$**:** parse pk and Φ as above and π_{DH} as $(\vec{C}_\theta, \pi_\theta) \in \hat{\mathbb{G}}^2 \times \mathbb{G}^2$. Then, compute the vector $\vec{u}_{\mathsf{tag}} = \vec{u}_2 \cdot (1, \hat{g})^{\mathsf{tag}}$ and use $\mathbf{u}_{\mathsf{tag}} = (\vec{u}_1, \vec{u}_{\mathsf{tag}})$ as a Groth-Sahai CRS to verify the proof π_{DH}. If the latter is deemed as a valid proof for relations (5), return 1. Otherwise, return 0.

PubEval$\big(\mathsf{pk}, (\theta_1, \theta_2), \varPhi\big)$**:** parse pk and $\varPhi$ as above. Return $\perp$ if $(\varPhi_1, \varPhi_2) \neq (g_1^\theta, g_2^\theta)$. Otherwise, compute and return $K = X^\theta \in \mathcal{K}$.

SharePrivEval$(\mathsf{pk}, \mathsf{sk}_i, \varPhi)$**:** parse sk_i as $(P_1(i), P_2(i)) \in (\mathbb{Z}_p)^2$ and return $\perp$ if $\varPhi \notin \mathbb{G}^2$. Otherwise, return (K_i, π_{K_i}), where $K_i = \varPhi_1^{P_1(i)} \cdot \varPhi_2^{P_2(i)} \in \mathcal{K}'$ and $\pi_{K_i} = (\vec{C}_{P_1}, \vec{C}_{P_2}, \pi'_{K_i}) \in \hat{\mathbb{G}}^4 \times \mathbb{G}^4$ is a proof consisting of commitments $\vec{C}_{P_1}, \vec{C}_{P_2}$ to exponents $P_1(i), P_2(i) \in \mathbb{Z}_p$ and a proof π'_{K_i} that they satisfy the equations

$$K_i = \varPhi_1^{P_1(i)} \cdot \varPhi_2^{P_2(i)}, \qquad Y_i = g_1^{P_1(i)} \cdot g_2^{P_2(i)}. \tag{6}$$

The perfectly binding commitments $\vec{C}_{P_1}, \vec{C}_{P_2}$ and the proof π'_{K_i} are generated using the vectors $\mathbf{v} = (\vec{v}_1, \vec{v}_2)$ as a Groth-Sahai CRS .

ShareEvalVerify$(\mathsf{pk}, \mathsf{vk}_i, \varPhi, K_i, \pi_{K_i})$**:** parse vk_i as $Y_i \in \mathbb{G}$ and return $\perp$ in the event that (K_i, π_{K_i}) cannot be parsed as a sequence of elements in $\mathbb{G} \times \hat{\mathbb{G}}^4 \times \mathbb{G}^4$. Otherwise, parse the proof π_{K_i} as $\pi_{K_i} = (\vec{C}_{P_1}, \vec{C}_{P_2}, \pi'_{K_i}) \in \hat{\mathbb{G}}^4 \times \mathbb{G}^4$ and return 1 if π'_{K_i} is a valid proof for equations (6). In any other situation, return 0.

Combine$(\mathsf{pk}, \varPhi, \{(K_i, \pi_{K_i})\}_{i \in S})$**:** return $\perp$ if $\mathsf{ShareEvalVerify}(\mathsf{pk}, \mathsf{vk}_i, \varPhi, K_i, \pi_{K_i}) = 0$ for some $i \in S$. Otherwise, compute $K = \prod_{i \in S} K_i^{\Delta_{i,S}(0)} = \varPhi_1^{x_1} \cdot \varPhi_2^{x_2} \in \mathcal{K}$.

The proof of the following theorem is completely similar to the proof of theorem 3 and omitted.

Theorem 4. *The above construction is an all-but-one perfectly sound hash proof system assuming that the SXDH assumption holds in $(\mathbb{G}, \hat{\mathbb{G}})$.*

When the generic construction of Section 4 is instantiated with the above all-but-one hash proof system, the resulting cryptosystem can be seen as a combination between Damgård's ElGamal encryption scheme [21] (as it is described in [37]) with a non-interactive one-time simulation-sound proof of validity of the ciphertext. The latter makes it possible to publicly verify the validity of ciphertexts so as to achieve security in the threshold setting.

As detailed in appendix E.2, the proof π_{DH} consists of 2 elements of $\mathbb{G}$ and the commitment $\vec{C}_\theta$ requires 2 elements of $\hat{\mathbb{G}}$ (each one of which has a representation as large as two elements of $\mathbb{G}$ with the choice of parameters suggested in [4]). The ciphertext overhead now amounts to the length of 13 elements of $\mathbb{G}$ if the one-time signature Σ is instantiated using [29]. On Barreto-Naehrig curves [4], if each element of $\mathbb{G}$ has a 256-bit representation (as recommended at the 128-bit security level), this overhead reduces to 3328 bits.

From a computational point of view, if we assume that a multi-exponentiation with two base elements has roughly the same cost as a single exponentiation, the sender has to compute 8 exponentiations in $\mathbb{G}$ and 2 exponentiations in $\hat{\mathbb{G}}$. The validity of a ciphertext can be verified using only 6 pairing evaluations in a batch-verification process.

E Construction of Non-Interactive Proofs for Schemes in Prime Order Groups

E.1 Construction of Proof Elements for the DLIN-based Instantiation

In the following notations, we define a coordinate-wise pairing $E : \mathbb{G} \times \mathbb{G}^3 \to \mathbb{G}_T^3$ such that, for any element $h \in \mathbb{G}$ and any vector $\vec{g} = (g_1, g_2, g_3)$, we have $E(h, \vec{g}) = \big(e(h, g_1), e(h, g_2), e(h, g_3)\big)$.

To construct the proof π_{LIN} that $\varPhi = (\varPhi_1, \varPhi_2, \varPhi_3) = \big(g_1^{\theta_1}, g_2^{\theta_2}, g^{\theta_1 + \theta_2}\big)$, for some $(\theta_1, \theta_2) \in (\mathbb{Z}_p)^2$,

the sender first computes commitments $\vec{C}_{\theta_i} = \vec{g}_{\mathsf{tag}}^{\theta_i} \cdot \vec{g_1}^{r_i} \cdot \vec{g_2}^{s_i} = \left(g_{\mathsf{tag},1}^{\theta_i} \cdot g_1^{r_i}, \ g_{\mathsf{tag},2}^{\theta_i} \cdot g_2^{s_i}, g_{\mathsf{tag},3}^{\theta_i} \cdot g^{r_i+s_i} \right)$, for each $i \in \{1,2\}$, with $r_1, r_2, s_1, s_2 \xleftarrow{R} \mathbb{Z}_p$ and where $\vec{g}_{\mathsf{tag}} = (g_{\mathsf{tag},1}, g_{\mathsf{tag},2}, g_{\mathsf{tag},3}) \in \mathbb{G}^3$. Then, he generates the proof $\pi_{(\theta_1, \theta_2)}$ as

$$\pi_{(\theta_1,\theta_2)} = (\pi_1, \pi_2, \pi_3, \pi_4, \pi_5, \pi_6) = \left(g_1^{r_1}, \ g_1^{s_1}, \ g_2^{r_2}, \ g_2^{s_2}, \ g^{r_1+r_2}, \ g^{s_1+s_2} \right)$$

which satisfies the verification equations

$$
\begin{aligned}
E\left(g_1, \vec{C}_{\theta_1}\right) &= E\left(\Phi_1, \vec{g}_{\mathsf{tag}}\right) \cdot E\left(\pi_1, \vec{g_1}\right) \cdot E\left(\pi_2, \vec{g_2}\right) \\
E\left(g_2, \vec{C}_{\theta_2}\right) &= E\left(\Phi_2, \vec{g}_{\mathsf{tag}}\right) \cdot E\left(\pi_3, \vec{g_1}\right) \cdot E\left(\pi_4, \vec{g_2}\right) \\
E\left(g, \vec{C}_{\theta_1} \cdot \vec{C}_{\theta_2}\right) &= E\left(\Phi_3, \vec{g}_{\mathsf{tag}}\right) \cdot E\left(\pi_5, \vec{g_1}\right) \cdot E\left(\pi_6, \vec{g_2}\right).
\end{aligned}
\tag{7}
$$

When the above verifications are performed in the naive way, they require to evaluate 30 pairings altogether. However, using randomized batch verification techniques (which, as illustrated in [6], can provide substantial savings in the context of Groth-Sahai proofs), they can be more efficiently processed by computing a product of 12 pairings at the expense of a tiny probability of accepting an invalid ciphertext.

On a CRS $(\vec{g_1}, \vec{g_2}, \vec{g}_{\mathsf{tag}^\star})$ for the WI setting (*i.e.*, where $\vec{g}_{\mathsf{tag}^\star} = \vec{g_1}^{\xi_1} \cdot \vec{g_2}^{\xi_2}$ for some $\xi_1, \xi_2 \in_R \mathbb{Z}_p$), the proof π_{LIN} can be simulated as follows. First, commitments $\vec{C}_{\theta_1}, \vec{C}_{\theta_2}$ are computed as commitments to 0 (say $\vec{C}_{\theta_i} = \vec{g_1}^{r_i} \cdot \vec{g_2}^{s_i}$ for each $i \in \{1,2\}$ with $r_1, r_2, s_1, s_2 \xleftarrow{R} \mathbb{Z}_p$). Then, proof elements $\pi_{(\theta_1,\theta_2)} = (\pi_1, \pi_2, \pi_3, \pi_4, \pi_5, \pi_6)$ satisfying (7) can be obtained as per

$$
\begin{aligned}
\pi_1 &= g_1^{r_1} \cdot \Phi_1^{-\xi_1} & \pi_3 &= g_2^{r_2} \cdot \Phi_2^{-\xi_1} & \pi_5 &= g^{r_1+r_2} \cdot \Phi_3^{-\xi_1} \\
\pi_2 &= g_1^{s_1} \cdot \Phi_1^{-\xi_2} & \pi_4 &= g_2^{s_2} \cdot \Phi_2^{-\xi_2} & \pi_6 &= g^{s_1+s_2} \cdot \Phi_3^{-\xi_2}.
\end{aligned}
$$

E.2 Construction of Proof Elements for the SXDH-based instantiation

Here, our notations use a coordinate-wise pairing $E : \mathbb{G} \times \hat{\mathbb{G}}^2 \to \mathbb{G}_T^2$ such that, for any element $h \in \mathbb{G}$ and any vector $\vec{g} = (\hat{g}_1, \hat{g}_2) \in \hat{\mathbb{G}}^2$, we have $E\left(h, \vec{g}\right) = \left(e(h, \hat{g}_1), e(h, \hat{g}_2)\right)$.

To construct the non-interactive proof π_{DH} that $(\Phi_1, \Phi_2) = (g_1^\theta, g_2^\theta)$, for some $\theta \in_R \mathbb{Z}_p$, the sender first computes a commitment $\vec{C}_\theta = \vec{u}_{\mathsf{tag}}^\theta \cdot \vec{u_1}^r = \left(\hat{u}_{\mathsf{tag},1}^\theta \cdot \hat{g}^r, \ \hat{u}_{\mathsf{tag},2}^\theta \cdot \hat{h}^r \right)$, using a randomly drawn $r \xleftarrow{R} \mathbb{Z}_p$ and where $\vec{u}_{\mathsf{tag}} = (\hat{u}_{\mathsf{tag},1}, \hat{u}_{\mathsf{tag},2}) \in \hat{\mathbb{G}}^2$. Then, he generates the proof π_θ as

$$\pi_\theta = (\pi_1, \pi_2) = \left(g_1^r, g_2^r \right) \in \mathbb{G}^2$$

which satisfies the verification equations

$$
\begin{aligned}
E\left(g_1, \vec{C}_\theta\right) &= E\left(\Phi_1, \vec{u}_{\mathsf{tag}}\right) \cdot E\left(\pi_1, \vec{u_1}\right) \\
E\left(g_2, \vec{C}_\theta\right) &= E\left(\Phi_2, \vec{u}_{\mathsf{tag}}\right) \cdot E\left(\pi_2, \vec{u_1}\right).
\end{aligned}
\tag{8}
$$

Instead of naively verifying equations (8) separately, the verifier can choose $\omega \xleftarrow{R} \mathbb{Z}_p$ and test whether

$$E(g_1 \cdot g_2^\omega, \vec{C}_\theta) = E(\Phi_1 \cdot \Phi_2^\omega, \vec{u}_{\mathsf{tag}}) \cdot E(\pi_1 \cdot \pi_2^\omega, \vec{u_1}),$$

which fails with overwhelming probability when one of the two equations (8) is not satisfied. With further optimizations (when coordinate-wise equalities are simultaneously batch-verified), the verifier only needs to compute a product of 6 pairings.

On a CRS $(\vec{u}_{\mathsf{tag}}, \vec{u}_1)$ for the perfect WI setting (*i.e.*, where $\vec{u}_{\mathsf{tag}} = \vec{u}_1^{\rho_u}$ for some $\rho_u \in_R \mathbb{Z}_p$), a NIZK proof π_{DH} can be simulated by computing $\vec{C}_\theta$ as a commitment to 0 (say $\vec{C}_\theta = \vec{u}_1^r$ for some $r \xleftarrow{R} \mathbb{Z}_p$) and the assignment

$$\pi_1 = g_1^r \cdot \Phi_1^{-\rho_u} \qquad\qquad \pi_2 = g_2^r \cdot \Phi_2^{-\rho_u}$$

is easily seen to satisfy the verification equations (8).