

From Polynomial IOP and Commitments to Non-malleable zkSNARKs

Antonio Faonio¹, Dario Fiore², Markulf Kohlweiss³,
Luigi Russo¹, and Michal Zajac⁵

¹ EURECOM, [faonio,russol}@eurecom.fr](mailto:{faonio,russol}@eurecom.fr)

² IMDEA Software Institute, dario.fiore@imdea.org

³ University of Edinburgh and Input Output, markulf.kohlweiss@ed.ac.uk

⁴ Nethermind, michal@nethermind.io

Abstract. We study sufficient conditions for compiling simulation-extractable zkSNARKs from information-theoretic interactive oracle proofs (IOP) using a simulation-extractable commit-and-prove system for its oracles. Specifically, we define simulation extractability for opening and evaluation proofs of polynomial commitment schemes, which we then employ to prove the security of zkSNARKs obtained from polynomial IOP prove systems, such as Plonk and Marlin. To instantiate our methodology we additionally prove that KZG commitments satisfy our simulation extractability requirement, despite being naturally malleable. To this end, we design a relaxed notion of simulation extractability that matches how KZG commitments are used and optimized in real-world prove systems. Only the proof that KZG satisfies this relaxed simulation extractability property relies on the algebraic group model (AGM) and random oracle (RO). We thus isolate the use of (and thus the reliance on) these strong heuristics.

1 Introduction

Non-interactive succinct zero-knowledge arguments of knowledge (zkSNARKs) [43], are the new Swiss army knife of blockchain scalability and privacy. They effectively deliver the twin dream of probabilistically checkable proofs (PCP) [3] and zero-knowledge proofs (ZKP) [32] while also being non-interactive, short, and efficiently verifiable. These features make zkSNARK of high practical and theoretical relevance. They are an active area of research that has seen rapid progress in multiple aspects, such as efficiency [7,31,33,34], security and versatility of their setups [6,35], and proof composition [13,15].

Simulation-extractable zkSNARKs. The basic security notion of zkSNARKs is *knowledge-soundness* which, informally speaking, guarantees that, in isolation, a prover producing a valid proof must know the corresponding witness. In contrast, there exist real-world deployments and cryptographic applications of zkSNARKs that require a stronger property called *simulation extractability* (SE, for brevity) [36,46]. Intuitively, this notion considers attackers that can see proofs for some statements and may use this information in order to produce

a proof for some other statement without knowing the witness. Interestingly, simulation extractability implies that proofs are *non-malleable* [22], a relevant property in practical applications. Most zkSNARKs in the literature are only proven to be knowledge-sound. In some cases, this is due to the fact that their proofs may indeed be malleable, e.g., as in [34] (see also [4]). In other cases, the lack of SE security proof is because it is a challenging task that may require more investigation.

From polynomial commitments to SNARKs. The design of modern zkSNARKs follows the common cryptographic approach of starting with protocols that achieve information-theoretic security in idealized models and then compiling them into efficient protocols by employing a smaller computationally-secure primitive. In the world of SNARKs, the corresponding concepts are (polynomial) interactive oracle proofs F -IOP [16,17,20,26,47] and (polynomial) functional commitments F -COM [12,38,39]. An F -IOP employs two (idealized) oracles that share their state: The prover calls the first oracle *to commit* to functions $f \in F$ and the verifier calls the second *to query* the committed functions. Concretely, the F -IOP to SNARK compiler uses F -COM to replace oracles with commitments, opening proofs, and query proofs. As this only removes reliance on idealized function oracles but not interaction, the compiler additionally employs the usual Fiat-Shamir transformation for public-coin protocols to obtain the final zkSNARK. The benefits of this compilation paradigm are modularity and separation of concerns: once the compiler is proven, a line of research can address the problem of improving F -IOPs while another research line can tackle the problem of realizing F -COM schemes (e.g., with better efficiency, from different assumptions, etc.): this approach has been successfully adopted to construct several recent zkSNARKs. All this recent work, though, only shows that schemes obtained via this paradigm are knowledge-sound.

1.1 Our work

We study the simulation extractability of a broad class of zkSNARKs built through this compilation approach. In particular, our primary goal includes showing that *existing* zkSNARKs (and thus existing compilers) provide simulation extractability. This goal has a twofold motivation. On the theoretical side, we are interested in understanding sufficient conditions on F -COM to compile an F -IOP into a simulation-extractable zkSNARK. On the practical side, by capturing existing compilers we can show that existing schemes that are under deployment, e.g., Plonk [26], have already this strong security property. For this reason, in our work we focus on the popular case of the compiler where the F -IOP is a polynomial IOP (i.e., the oracle functions F are low-degree polynomials), and F -COM is a polynomial commitment. Furthermore, in terms of instantiations we are interested to cover the celebrated polynomial commitment scheme of Kate, Zaverucha and Goldberg [38] (KZG, shortly) and on a polynomial IOP framework that is flexible enough to include recent constructions, e.g., [17,20,26,41,45].

The main contributions of our work are: (i) to introduce a relaxed notion of simulation extractability for polynomial commitments; (ii) to prove that the KZG scheme satisfies our relaxed SE notion in the algebraic group model (AGM) and random oracle model (ROM); and (iii) to prove that our notion is sufficient to compile a polynomial IOP into a (full-fledged) simulation extractable zkSNARK, using the usual compilation approach. By combining these three results we obtain a simulation-extractability proof for Plonk [26] and a slight variation of Marlin [20].

Below, we elaborate more on our results and the challenges that we had to overcome along the way.

1.2 Our techniques

Background. For our work, we chose the class of Polynomial (Holographic) IOPs (PIOP) as defined by [17] as a generalization of [16].⁵ The oracle of the prover commits to low-degree polynomials over a finite field while the queries of the verifier check polynomial equations over these polynomials. These polynomial equations can depend on additional field elements sent by the prover and/or the verifier during the execution of the protocol. Slightly more in detail, the verifier can query an oracle polynomial $p(X)$ (or multiple polynomials simultaneously) by specifying polynomials G and v to test equations of the form $G(X, p(v(X))) \equiv 0$. Therefore, to be compiled, PIOPs need a commit-and-prove SNARK (CP-SNARK) for proving the validity of such equations concerning the committed polynomials. Notably, one can easily build this CP-SNARK from one for polynomial evaluations (e.g., KZG) by testing the equations on a random point chosen by the random oracle.

Simulation extractability challenges. Intuitively, the use of SE-secure CP-SNARK in the above compiler should result in a simulation-extractable zk-SNARK: The zero-knowledge simulator samples random commitments (relying either on hiding property of commitments, or the randomness in the committed functions p). It then simulates evaluations of p that satisfy the verification equation of the PIOP. The reduction to PIOP soundness extracts all committed functions from their opening proofs and the final function evaluations from the evaluation proofs. However, this approach presents two major challenges:

- The PIOP could be arbitrary. In particular, internally it could sequentially compose two PIOP sub-protocols that prove statements over different sets of polynomials. For example, consider a PIOP obtained by the sequential composition of two PIOP protocols for two independent statements. Very likely, the set of queries to the polynomials made by the two sub-protocols are independent and (unless the PIOP specifies it explicitly) the evaluation queries of the first sub-protocol may be chosen based on the verifier’s random challenges sent *before* the second sub-protocol even starts. The simulation

⁵ The choice is mainly motivated by generality as PIOPs can flexibly capture under the same hat all the most recent protocols based on the notions of [16], AHP [20], and ILDP [26].

extractability of the zkSNARK compiled from this protocol might be affected because one could strip off the second set of evaluation proofs and replace them with those for another statement⁶.

- Secondly, one needs to prove that existing, efficient, and practically deployable instantiations of polynomial commitments are simulation-extractable.

Our solutions. To solve the first challenge, we define a zkSNARK compiler that bundles all verification queries together instead of translating verification queries *on-the-fly* into calls to the commit-and-prove system. Motivated by our goal to show that existing zkSNARKs are simulation extractable, we define two (rather minimal) constraints on the PIOP: The first one is that at least one of the polynomial equations involves all the oracle polynomials and that the polynomial v chosen by the verifier (see above) is not constant.⁷ Fortunately, this constraint is natural and easy to meet in practice: Plonk naturally meets our constraint meanwhile all the other proof systems based on Aurora’s univariate sumcheck [8] can be easily (and at negligible cost) adapted by instantiating the proof of polynomial degree through an evaluation query on all the polynomials. The second constraint regards the distribution of the polynomial commitments after compilation: we require that the joint view of a transcript of the polynomial oracle protocol execution and the commitments to the oracle polynomials can be simulated. Notice, this property is obvious when we compile using hiding commitment schemes, and it is not hard to show when the hiding property is obtained through randomization of the oracle polynomials at the protocol level (for example see the proof of zero-knowledge for Plonk in [27]).

For the second challenge, unfortunately, the issue is that the most obvious candidate, the efficient and widely deployed KZG polynomial commitment scheme [38], is not simulation-extractable. Using bracket notation, KZG commitments are of the form $[p(s)]_1$ for a trapdoor secret s encoded in the parameters $([s^i]_1)_{i \in [0..d]}, [1, s]_2$, while evaluation proofs for an input x and output y are of the form $[\frac{p(s)-y}{s-x}]_1$. KZG is malleable, and thus not simulation-extractable in the usual sense: e.g., given a commitment to f anyone can compute $[f(s) + \Delta]_1$ and open it using the same proof to $(x, y + \Delta)$.

Our starting point is the observation that KZG retains a form of simulation extractability for evaluations at points that are randomly chosen *after* the commitment. Fortunately, this is the situation we encounter in the Fiat-Shamir part of the PIOP-to-SNARK compiler. The commitment forms part of the first commit-and-prove part of the statement which is hashed to determine the x of the second part of the statement. Thus, the evaluation point depends on the commitment and can be considered random in the RO model.

⁶ In particular, the adversary could have a simulated proof $\tilde{\pi} = (\pi, \pi')$ for $(\mathbb{x}, \mathbb{x}')$ and then could choose \mathbb{x}'' for which it knows a valid witness, and finally forge for $(\mathbb{x}, \mathbb{x}'')$ using (π, π'') , where π'' is honestly generated. As the simulated proof π is reused, extraction fails. Notice, this attack works even when the Fiat-Shamir challenges for π'' are derived by hashing a transcript that contains π .

⁷ This can be for example be implemented via a common random point chosen at the end of the protocol, on which all oracles are evaluated.

To formalize this important relaxation, we introduce a Φ -flexible simulation extractability (Φ -SE) notion. In the standard simulation-extractability experiment, the adversary can ask the simulator to generate proofs for statements of its choice and, eventually, must produce a *new* valid proof without knowing the witness. In Φ -SE, we consider a relaxation of the SE game in which all the simulation queries of the adversary must satisfy a predicate specified in Φ , and similarly, Φ can constrain the winning condition of the adversary. For this reason, we refer to Φ as the *policy*. One can see that Φ -SE is a generalization of existing SE notions such as true-simulation-extractability (where the adversary can only see simulated proofs on true statements) [21], or weak simulation extractability (where the adversary only wins if it provides a proof for a new statement, and contrary to (strong) SE loses if it provides a new proof for a statement previously asked to the simulation oracle). Once having defined this framework, we analyze which policies Φ are strong enough to achieve simulation extractability in the compiled zkSNARK, while at the same time being weak enough for instantiation by KZG under plausible assumptions (in the AGM [25] and RO). Specifically, we isolate the (simulation) extractability properties needed for the compiler and verify it for KZG in the AGM. This is the only part of our results where we need the AGM. Given the broad applications of KZG in the field of practical zkSNARKs and beyond, the characterization of its (non-)malleability is interesting in its own right. In fact, our policy highlights some malleability attacks that we discovered and that we needed to handle. Finally, for our Φ we prove that KZG is Φ -SE in the AGM and ROM. This proof turned out to be highly non-trivial and is one of the main technical contributions of our work.

Related work. It is hard to be exhaustive, or even representative, in discussing related work on SNARKs. For the sake of our paper, we focus on related work on simulation extractability notions. Groth and Maller [36] give a simulation-extractable zkSNARK that consists of only 3 group elements. Their construction is neither universal nor updatable. The recent work of Ganesh, Orlandi, Pancholi, Takahashi and Tschudi [29] shows that Bulletproofs [14] are non-malleable in AGM. Their work extends the framework introduced by Faust, Kohlweiss, Marson and Venturi in [23] to Fiat-Shamir applied to multi-round interactive arguments. On a similar path, the work of Ganesh, Khoshakhlagh, Kohlweiss, Nitulescu and Zajac [27] shows non-malleability for Plonk, Sonic and Marlin. Both [27,29] show that interactive *arguments* can be simulation-extractable after Fiat-Shamir transform. In contrast, our work takes a step further in the compilation pipeline showing the simulation extractability of SNARKs derived from interactive oracle protocols. In particular, [27,29] rely on properties such as *trapdoor-less zero-knowledge*, i.e. the zero-knowledge where the simulator does not rely on the SRS’s trapdoor but on the programmability of the random oracle, and on the *unique-response property*, which states that at some point of the protocol, the prover becomes a deterministic algorithm. Reliance on these two properties imposes additional effort to show that schemes are simulation-extractable. Namely, both of these properties were coined only for the purpose of showing simulation extractability, hence are usually not analyzed in papers that introduce

new protocols. Furthermore, [27] relies on rewinding-based soundness, which is a non-standard notion of soundness, and rewinding which make the simulator-extractability extractor success dependent on the probability of the adversary returning an acceptable proof; and [29] relies on the AGM. Contrary to these results, our paper’s result does not directly rely on the AGM, does not require the protocol to be trapdoor-less zero knowledge or have the unique-response property and uses more standard notion of soundness, i.e. the state restoration soundness. The work of Abdolmaleki, Ramacher and Slamanig [2] shows a generic compiler to simulation extractable SNARKs which requires key-homomorphic signatures. Their compiler produces universally-composable SNARKs (UC-SNARKs), which they prove through black-box straight-line extractor. To obtain a black-box straight-line extractor, they append to the SNARK proof an encryption of the witness, thus achieving a relaxed succinctness w.r.t. the size of the circuit describing the relation. The recent work of Ganesh, Kondi, Orlandi, Pancholi and Takahashi [28] shows how to regain full succinctness in UC-SNARKs in the ROM through Fischlin’s transform [24].

Open problems. Our framework is general enough to handle compilation from polynomial commitment schemes different than KZG. Our contribution identifies a set of properties that a polynomial commitment scheme needs to have so that the resulting SNARK is simulation extractable. We believe that FRI polynomial commitment scheme [10] readily possesses the necessary properties, which would imply the simulation extractability of STARKs [6].

Another advantage of our formalization of PIOP over previous proposals such as [16] is that it naturally supports optimization tricks in the literature [17]. As an intermediate step of our compiler, we define a CP-SNARK for polynomial evaluations based on KZG. While we capture the important use case of batched evaluations on a common point, for the sake of simplicity, we leave further extensions and optimizations for future work. In particular, we do not capture the case of proving evaluations on *arbitrary* linear combinations of committed polynomials. We believe this extension could be handled at PIOP level by extending the notion to *virtual* oracle polynomials obtained through linear combinations of other oracles (and thus using the homomorphic property of KZG). We leave as open problem to extend our result to even more flexible polynomial IOP models.

One could think that restricting the forgery on a fresh randomly-sampled evaluation point x would be a natural condition to prove simulation extractability of KZG. This restriction, however, is not sufficient to prove *strong* simulation extractability for the compiler. This is why we enlarge the class of forgeries for the PCOM: we prove that these additional forgeries are unfeasible for KZG, and the compiled SNARK is simulation extractable in the strong sense, as desired.

Recent works extend the polynomial evaluation proofs of KZG to multiple evaluation points [48,49]. Our simulation-extractability strategy for KZG can be applied partially to these schemes; however, our technique uses a clever argument to separate the realm of commitments from the realm of proofs (in KZG proofs and commitments are both of the form $[p(s)]_1$ for some polynomial p) based on their degree as polynomials. Unfortunately, the same technique does not

work when the degree of the polynomial in the proof depends on the number of evaluation points in the proved statement.

Organization of the paper. We define the framework of policy-based simulation-extractability in Section 3, we proceed with the analysis of the simulation-extractability of KZG in Section 4 and we give our generic compiler for strong simulation-extractable Universal SNARKs from (simulation-extractable) polynomial commitment and PIOP in Section 5. The thesis of the theorem on KZG polynomial commitment and the hypothesis for the theorem of the Universal SNARKs compiler do not quite match. Indeed, they could be even considered as two independent and (almost) self-contained results. We show how to fill the gaps and connect the two results in Section 6.

2 Preliminaries

A function f is negligible in λ (we write $f \in \text{negl}(\lambda)$) if it approaches zero faster than the reciprocal of any polynomial: i.e., for every $c \in \mathbb{N}$ there is an integer λ_c such that $f(\lambda) \leq \lambda^{-c}$ for all $\lambda \geq \lambda_c$. The expression $[n]$ denotes the set $\{1, 2, \dots, n\}$ for an integer $n \geq 1$. Calligraphic letters denote sets, while set sizes are written as $|\mathcal{X}|$. Lists are represented as ordered tuples, e.g. $L := (L_i)_{i \in [n]}$ is a shortcut for the list of n elements (L_1, \dots, L_n) . To get a specific value from a list, we also use the “dot” notation; e.g., we use $L.b$ to access the second element of the list $L = (a, b, c)$. The difference between lists and vectors is that elements of vectors are of the same type.

An asymmetric bilinear group \mathbb{G} is a tuple $(q, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, P_1, P_2)$, where $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T are groups of prime order q , the elements P_1, P_2 are generators of $\mathbb{G}_1, \mathbb{G}_2$ respectively, $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is an efficiently-computable non-degenerate bilinear map, and there is no efficiently computable isomorphism between \mathbb{G}_1 and \mathbb{G}_2 . Let GroupGen be some probabilistic polynomial-time (PPT) algorithm which on input 1^λ , where λ is the security parameter, returns a description of a bilinear group \mathcal{G} . Elements in $\mathbb{G}_i, i \in \{1, 2, T\}$ are denoted in implicit notation as $[a]_i := aP_i$, where $P_T := e(P_1, P_2)$. Every element in \mathbb{G}_i can be written as $[a]_i$ for some $a \in \mathbb{Z}_q$, but note that, given $[a]_i$, is in general hard to compute (discrete logarithm problem). Given $a, b \in \mathbb{Z}_q$ we distinguish between $[ab]_i$, namely the group element whose discrete logarithm base P_i is ab , and $[a]_i \cdot b$, namely the execution of the multiplication of $[a]_i$ and b , and $[a]_1 \cdot [b]_2 = [a \cdot b]_T$, namely the execution of a pairing between $[a]_1$ and $[b]_2$. We do not use the implicit notation for variables, e.g. $c = [a]_1$ indicates that c is a variable name for the group element whose logarithm is a .

Definition 1 (Algebraic algorithm, [25]). *An algorithm \mathcal{A} is called algebraic if for all group elements \mathbf{z} that \mathcal{A} outputs (either as returned by \mathcal{A} or by calling one of its oracles), it additionally provides the representation of \mathbf{z} relative to all previously received group elements. That is, if elems is the list of group elements that \mathcal{A} has received so far, then \mathcal{A} must also provide a vector \mathbf{r} such that $\mathbf{z} = \langle \mathbf{r}, \text{elems} \rangle$.*

Definition 2 (Commitment scheme). A commitment scheme with message space \mathcal{M} (and group parameters GroupGen) is a tuple of algorithms $\text{CS} = (\text{KGen}, \text{Com}, \text{VerCom})$ that works as follows:

$\text{KGen}(\text{pp}_{\mathbb{G}}) \rightarrow \text{ck}$ takes as input group parameters $\text{pp}_{\mathbb{G}} \leftarrow_{\$} \text{GroupGen}(1^\lambda)$ and outputs a commitment key ck .

$\text{Com}(\text{ck}, m) \rightarrow (c, o)$ takes the commitment key ck , and a message $m \in \mathcal{M}$, and outputs a commitment c and an opening o .

$\text{VerCom}(\text{ck}, c, m, o) \rightarrow b$ takes as input the commitment key ck , a commitment c , a message $m \in \mathcal{M}$ and an opening o , and it accepts ($b = 1$) or rejects ($b = 0$).

A commitment scheme CS satisfies correctness, binding and hiding properties.

Definition 3 (Polynomial Commitment). A polynomial commitment scheme PC is a commitment scheme (Definition 2) in which the message space \mathcal{M} is $\mathbb{F}_{\leq d}[X]$, the set of low degree polynomials over a finite field \mathbb{F} with degree bound $d \in \mathbb{N}$. The key generation algorithm might take as additional input the degree d .

Finally, we state these assumptions on distributions.

Definition 4 (Witness Sampleability, [37]). A distribution \mathcal{D} is witness sampleable if there exists a PPT algorithm $\tilde{\mathcal{D}}$ such that for any $\text{pp}_{\mathbb{G}}$ we have that the random variables $\mathbf{A} \leftarrow_{\$} \mathcal{D}(\text{pp}_{\mathbb{G}})$ and $[\tilde{\mathbf{A}}]_1$, where $\tilde{\mathbf{A}} \leftarrow_{\$} \tilde{\mathcal{D}}(\text{pp}_{\mathbb{G}})$, are equivalently distributed.

Definition 5 ($\mathcal{D}_{\ell,k}$ -Aff-MDH assumption). Given a matrix distribution $\mathcal{D}_{\ell,k}$, the Affine Diffie-Hellman Problem is given $\mathbf{A} \in \mathbb{G}_1^{\ell \times k}$, with $\mathbf{A} \leftarrow_{\$} \mathcal{D}_{\ell,k}$, find a nonzero vector $\mathbf{x} \in \mathbb{Z}_q^\ell$ and a vector $\mathbf{y} \in \mathbb{Z}_q^k$ such that $[\mathbf{x}^\top \mathbf{A}]_1 = [\mathbf{y}]_1$.

The Aff-MDH Assumption could be seen as an extension of the Kernel-MDH Assumption introduced by Morillo, Ràfols, Villar [44] since the Ker-MDH assumes $\mathbf{y} = \mathbf{0}$; Our assumption is incomparable because we also require the adversary to give $\mathbf{x} \in \mathbb{Z}_q^\ell$ “in the exponent”.

Definition 6 ((d, d') -Power Polynomial in the Exponent). The (d, d') -PEA Assumption holds for a bilinear group generator GroupGen if for every PPT adversary \mathcal{A} that receives as input $([1, \dots, s^d]_1, [1, \dots, s^{d'}]_2)$ and outputs a polynomial $p(X)$ of degree at most $\max\{d, d'\}$, and a value y , the probability that $p(s) = y$ is negligible. When $d = d'$ we use the shortcut d -PEA.

Definition 7 (d -Power Discrete Logarithm [40]). Given a degree bound $d \in \mathbb{N}$, the d -Power Discrete Logarithm (d -DL) assumption holds for a bilinear group generator GroupGen if for every PPT adversary \mathcal{A} that receives as input $([1, \dots, s^d]_1, [1, \dots, s^d]_2)$, and outputs the value s' , the probability that $s = s'$ is negligible. We also use DL as a shortcut for 1-DL.

Lemma 1 (d -DL \Rightarrow d -PEA). *We can make a reduction to the assumption that computes s . The reduction invokes the adversary, gets $p(X) - y$ of degree d , and computes s by factoring the polynomial $p(s) - y$. As $p(s) - y = 0$ we are guaranteed that s is a root.*

Lemma 2 (DL \Rightarrow $\mathcal{U}_{\ell,k}$ -Aff-MDH). *When considering the uniform random distribution $\mathcal{U}_{\ell,k}$, we can make a reduction to the assumption that computes s . The reduction samples at the exponent a uniformly random matrix $\mathbf{A} = (a_{i,j})_{i,j} \in \mathbb{Z}_q^{\ell \times k}$ and invokes the adversary on input $[(a_{i,j})_{i,j}]_1$. Finally, let $p_i(s)$ be the i -th row of $\mathbf{x}^\top \mathbf{A}$. The reduction computes s by factoring one of the k polynomials $p_i(s) - y_i$.*

3 Policy-based Simulation-Extractable NIZKs

We start by defining the basic syntax and properties for a Non-Interactive Zero-Knowledge Argument of Knowledge. Following Groth et al. [35], we define a PT relation \mathcal{R} verifying triple $(\mathbf{pp}, \mathbf{x}, \mathbf{w})$. We say that \mathbf{w} is a witness to the instance \mathbf{x} being in the relation defined by the \mathbf{pp} when $(\mathbf{pp}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}$ (equivalently, we sometimes write $\mathcal{R}(\mathbf{pp}, \mathbf{x}, \mathbf{w}) = 1$). For example, the parameters \mathbf{pp} could be the description of a bilinear group or additionally contain a commitment key for a commitment scheme or a common reference string. A NIZK for a relations \mathcal{R} (and group parameters GroupGen) is a tuple of algorithms $\Pi = (\text{KGen}, \text{Prove}, \text{Verify})$ that work as described below.

- $\text{KGen}(\mathbf{pp}_{\mathbb{G}}) \rightarrow \text{srs}$ is a probabilistic algorithm that takes as input the group parameters $\mathbf{pp}_{\mathbb{G}} \leftarrow \text{GroupGen}(1^\lambda)$ and it outputs an $\text{srs} := (\text{ek}, \text{vk}, \mathbf{pp})$ where ek is the evaluation key, vk is the verification key, and \mathbf{pp} are the parameters for the relation \mathcal{R} .
- $\text{Prove}(\text{ek}, \mathbf{x}, \mathbf{w}) \rightarrow \pi$ takes an evaluation key ek , a statement \mathbf{x} , and a witness \mathbf{w} such that $\mathcal{R}(\mathbf{pp}, \mathbf{x}, \mathbf{w})$ holds, and returns a proof π .
- $\text{Verify}(\text{vk}, \mathbf{x}, \pi) \rightarrow b$ takes a verification key, a statement \mathbf{x} , and either accepts ($b = 1$) or rejects ($b = 0$) the proof π .

Basic notions for a NIZK are completeness, (knowledge) soundness and zero-knowledge.

zkSNARKs. We consider the notion of zero-knowledge succinct argument of knowledge (zkSNARKs) which are NIZKs that are knowledge sound and succinct as defined below.

Definition 8 (Succinctness). *A NIZK Π is said succinct if the running time of Verify is $\text{poly}(\lambda + |\mathbf{x}| + \log |\mathbf{w}|)$ and the proof size is $\text{poly}(\lambda + \log |\mathbf{w}|)$.*

CP-SNARKs. Commit-and-Prove SNARKs (CP-SNARKs) are zkSNARKs whose relations verify predicates over commitments (see Campanelli, Fiore and Querol [18]). We consider the following syntax. Briefly speaking, we refer to a CP-SNARK for a relation \mathcal{R} and a commitment scheme CS as a tuple of algorithms $\text{CP} = (\text{KGen}, \text{Prove}, \text{Verify})$ where:

- $\text{KGen}(\text{ck}) \rightarrow \text{srs}$ is a probabilistic (or deterministic) algorithm that takes as input a commitment key ck for CS and it outputs $\text{srs} := (\text{ek}, \text{vk}, \text{pp})$, where ek is the evaluation key, vk is the verification key, and pp are the parameters for the relation \mathcal{R} (which include the commitment key ck).

Moreover, if we consider the key generation algorithm KGen' that upon group parameters $\text{pp}_{\mathbb{G}}$ first runs $\text{ck} \leftarrow_{\$} \text{CS.KGen}(\text{pp}_{\mathbb{G}})$, runs $\text{srs} \leftarrow_{\$} \text{CP.KGen}(\text{ck})$ and outputs srs ; then the tuple $(\text{KGen}', \text{Prove}, \text{Verify})$ defines a SNARK.

Zero-Knowledge in the SRS (and RO) model. The zero-knowledge simulator \mathcal{S} of a NIZK is a stateful PPT algorithm that can operate in three modes. $(\text{srs}, \text{st}_{\mathcal{S}}) \leftarrow \mathcal{S}(0, \text{pp}_{\mathbb{G}})$ takes care of generating the parameters and the simulation trapdoor (if necessary), $(\pi, \text{st}_{\mathcal{S}}) \leftarrow \mathcal{S}(1, \text{st}_{\mathcal{S}}, \mathbb{x})$ simulates the proof for a statement \mathbb{x} , $(a, \text{st}_{\mathcal{S}}) \leftarrow \mathcal{S}(2, \text{st}_{\mathcal{S}}, s)$ takes care of answering random oracle queries. The state $\text{st}_{\mathcal{S}}$ is shared and updated after each operation.

Similarly to [23,29], we define the following wrappers.

Definition 9 (Wrappers for NIZK Simulator). *The following algorithms are stateful and share their state $\text{st} = (\text{st}_{\mathcal{S}}, \text{coms}, \mathcal{Q}_{\text{sim}}, \mathcal{Q}_{\text{RO}}, \mathcal{Q}_{\text{aux}})$ where $\text{st}_{\mathcal{S}}$ is initially set to be the empty string, and $\mathcal{Q}_{\text{sim}}, \mathcal{Q}_{\text{RO}}$ and \mathcal{Q}_{aux} are initially set to be the empty sets.*

- $\mathcal{S}_1(\mathbb{x}, \text{aux})$ denotes an oracle that returns the first output of $\mathcal{S}(1, \text{st}_{\mathcal{S}}, \mathbb{x}, \text{aux})$ ⁸.
- $\mathcal{S}'_1(\mathbb{x}, \mathbb{w})$ denotes an oracle that first checks $(\text{pp}, \mathbb{x}, \mathbb{w}) \in \mathcal{R}$ where pp is part of srs and then runs (and returns the output of) $\mathcal{S}_1(\mathbb{x})$.
- $\mathcal{S}_1^F(\mathbb{x}, \mathbb{w})$ denotes an oracle parameterized by a function F that first checks $(\text{pp}, \mathbb{x}, \mathbb{w}) \in \mathcal{R}$ and then runs (and returns the output of) $\mathcal{S}_1(\mathbb{x}, F(\mathbb{x}, \mathbb{w}))$. As explained below, this is useful to model leaky-zero-knowledge.
- $\mathcal{S}_2(s, \text{aux})$ denotes an oracle that first checks if the query s is already present in \mathcal{Q}_{RO} and in case answers accordingly, otherwise it returns the first output a of $\mathcal{S}(2, \text{st}_{\mathcal{S}}, s)$. Additionally, the oracle updates the set \mathcal{Q}_{RO} by adding the tuple (s, aux, a) to the set.

Almost all the oracles in our definitions can take auxiliary information as additional input. We use this auxiliary information in a rather liberal form. For example, in the definition above, the auxiliary information for \mathcal{S}_1 refers to the (optional) leakage required by the simulator to work in some cases (see more in Definition 11), while the auxiliary information for \mathcal{S}_2 can contain, for example, the algebraic representations of the group elements in s (when we restrict to algebraic adversaries) or other information the security experiments might need.

Definition 10 (Zero-Knowledge). *A NIZK Π is (perfect) zero-knowledge if there exists a PPT simulator \mathcal{S} such that for all adversaries \mathcal{A} :*

$$\Pr \left[\begin{array}{l} \text{pp}_{\mathbb{G}} \leftarrow \text{GroupGen}(1^\lambda) \\ \text{srs} \leftarrow \text{KGen}(\text{pp}_{\mathbb{G}}) \\ \mathcal{A}^{\text{Prove}(\text{ek}, \cdot, \cdot)}(\text{srs}) = 1 \end{array} \right] \approx \Pr \left[\begin{array}{l} \text{pp}_{\mathbb{G}} \leftarrow \text{GroupGen}(1^\lambda) \\ (\text{srs}, \text{st}_{\mathcal{S}}) \leftarrow \mathcal{S}(0, \text{pp}_{\mathbb{G}}) \\ \mathcal{A}^{\mathcal{S}'_1(\cdot, \cdot)}(\text{srs}) = 1 \end{array} \right]$$

⁸ More often, simulators need only the first three inputs, see Definition 10; abusing notation, we assume that such simulators simply ignore the auxiliary input aux .

Zero-knowledge is a security property that is only guaranteed for valid statements in the language, hence the above definition uses \mathcal{S}'_1 as a proof simulation oracle.

We also introduce a weaker notion of zero-knowledge. A NIZK Π is F -leaky zero-knowledge if its proofs may leak some information, namely a proof leaks $F(\mathbf{x}, \mathbf{w})$, where $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$. We formalize this by giving the zero-knowledge simulator the value $F(\mathbf{x}, \mathbf{w})$, which should be interpreted as a hint for the simulation of proofs. This notion could be seen as an extension of the bounded leaky zero-knowledge property defined in [17] and tailored for CP-SNARKs. Our notion is a special case of the leakage-resilient zero-knowledge framework of Garg, Jain and Sahai [30] where the leakage of the simulator is known ahead of time.

Definition 11 (Leaky Zero-Knowledge). *A NIZK Π is F -leaky zero-knowledge if there exists a PPT simulator \mathcal{S} such that for all adversaries \mathcal{A} :*

$$\Pr \left[\begin{array}{l} \text{pp}_{\mathbb{G}} \leftarrow \text{GroupGen}(1^\lambda) \\ \text{srs} \leftarrow \text{KGen}(\text{pp}_{\mathbb{G}}) \\ \mathcal{A}^{\text{Prove}(\text{ek}, \cdot, \cdot)}(\text{srs}) = 1 \end{array} \right] \approx \Pr \left[\begin{array}{l} \text{pp}_{\mathbb{G}} \leftarrow \text{GroupGen}(1^\lambda) \\ (\text{srs}, \text{st}_{\mathcal{S}}) \leftarrow \mathcal{S}(0, \text{pp}_{\mathbb{G}}) \\ \mathcal{A}^{\mathcal{S}_1^F(\cdot, \cdot)}(\text{srs}) = 1 \end{array} \right]$$

3.1 Policy-Based Simulation Extractability

An extraction policy defines the constraints under which the extractor must extract the witness. For example, we could consider the policy that checks that the forged instance and proof were not queried/output by the zero-knowledge simulator (thus modeling the classical simulation extractability notion), or we could consider a policy that only checks that the forged instance was not queried to the zero-knowledge simulator, thus obtaining a weaker flavor of classical simulation extractability. Clearly, the more permissive the policy the stronger the security provided.

In our work, we also consider policies that constrain the behavior of the zero-knowledge simulator. For example, we could consider the policy that checks that the queried instances belong to the relation, thus obtaining a notion similar to true-simulation extractability (see Dodis et al. [21]). Looking ahead, contrary to the true-simulation extractability notion in [21], our policy-based version of the true-simulation extractability rather than disallowing certain queries, punishes the adversary at extraction time. It is not hard to see that the two definitional flavors, namely disallowing illegal queries versus punishing an adversary that made an illegal query are equivalent in the context of simulation extractability, because the adversary's goal is computational⁹.

Extraction policies. We define an extraction policy as a tuple $\Phi = (\Phi_0, \Phi_1)$ of PPT algorithms. This is used to define Φ -simulation extractability as follows. The security experiment starts by running the extraction policy algorithm Φ_0 , which generates public information pp_{Φ} . The public information may contain, for example, random values that define the constraints later checked by

⁹ Observe that for decisional tasks disallowing and punishing flavors can result in different security notions, see Bellare, Hofheinz and Kiltz [5].

Φ_1 . Therefore, we feed \mathbf{pp}_Φ to the adversary. In the case of commit-and-prove proof systems, the public information may contain commitments for which the adversary does not know openings (but on which it can still query simulated proofs).

After receiving a forgery from the adversary, the security experiment runs the extraction policy Φ_1 . The policy Φ_1 is as a predicate and takes as input:

1. The public parameter \mathbf{pp}_Φ .
2. The forged instance and proof (x, π) .
3. The view of the experiment, denoted \mathbf{view} . Such a view contains the public parameters, the set of simulated instances and proofs \mathcal{Q}_{sim} , and the set \mathcal{Q}_{RO} of queries and answers to the random oracle. Even if the given NIZK is not in the random oracle (namely neither the prover nor the verifier algorithms make random oracle queries) it still makes sense to assume the existence of the set \mathcal{Q}_{RO} . This is useful to model security for NIZK protocols that eventually are used as sub-protocols in ROM-based protocols (as Universal zkSNARKs based on Polynomial Commitments, see 5).
4. Auxiliary information \mathbf{aux}_Φ which might come along with the forged instance. We use \mathbf{aux}_Φ to provide the adversary an interface with the policy¹⁰.

Definition 12 (Simulation extractability). *Let Π be a NIZK for a relation \mathcal{R} with a simulator \mathcal{S} whose wrappers are $\mathcal{S}_1, \mathcal{S}_2$, as defined in Definition 9. Consider the experiment in Fig. 1. Π is Φ -simulation-extractable (or simply Φ -SE) if for every PPT adversary \mathcal{A} there exists an efficient extractor \mathcal{E} such that the following advantage is negligible in the security parameter:*

$$\mathbf{Adv}_{\Pi, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi\text{-se}}(\lambda) := \Pr \left[\mathbf{Exp}_{\Pi, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi\text{-se}}(\lambda) = 1 \right]$$

Below, we give a definition that explicitly considers the sub-class of PPT algebraic adversaries. To fit algebraic adversaries into our definitional framework we let the algebraic adversaries return the representation vectors (1) for any query to the oracles \mathcal{S}_1 and \mathcal{S}_2 into the auxiliary information \mathbf{aux} provided to the oracles and (2) for the forgery into the auxiliary information $\mathbf{aux}_\mathcal{E}$.

Definition 13 (Simulation extractability in the AGM). *Let Π be a NIZK for a relation \mathcal{R} with a simulator \mathcal{S} whose wrappers are $\mathcal{S}_1, \mathcal{S}_2$, as defined in Definition 9. Consider the experiment in Fig. 1. Π is Φ -simulation-extractable (or simply Φ -SE) if there exists an efficient extractor \mathcal{E} such that for every PPT algebraic adversary \mathcal{A} the advantage $\mathbf{Adv}_{\Pi, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi\text{-se}}(\lambda)$ (as defined in Definition 12) is negligible in the security parameter.*

¹⁰ For example, looking ahead, in the policy in Section 6.2 the adversary that forges a “loose” proof of opening for a commitment, additionally provides a certificate (different than the proof itself) that the commitment is indeed extractable. In this case, we require the extractor only to work for those commitments that come along with valid certificates.

<pre> Exp_{$\mathcal{A}, \mathcal{S}, \mathcal{E}$}^{$\Phi$-s$\mathcal{E}$}($\lambda$) pp_G \leftarrow \mathcal{S} GroupGen(1^λ) (srs, st_S) \leftarrow \mathcal{S}_0(pp_G) pp_{Φ} \leftarrow \mathcal{S} Φ_0(pp_G) ($\mathbb{x}, \pi, \text{aux}_{\mathcal{E}}, \text{aux}_{\Phi}$) \leftarrow $\mathcal{A}^{\mathcal{S}_1, \mathcal{S}_2}$(srs, pp_{$\Phi$}) $\mathbb{w} \leftarrow \mathcal{E}$(srs, $\mathbb{x}, \pi, \text{aux}_{\mathcal{E}}$) view \leftarrow (srs, pp_{Φ}, \mathcal{Q}_{sim}, \mathcal{Q}_{RO}, \mathcal{Q}_{aux}) if $\Phi_1((\mathbb{x}, \pi), \text{view}, \text{aux}_{\Phi}) \wedge \text{Verify}^{\mathcal{S}_2}(\text{srs}, \mathbb{x}, \pi)$ $\wedge (\text{pp}, \mathbb{x}, \mathbb{w}) \notin \mathcal{R}$ then return 1 else return 0 </pre>	<pre> $\mathcal{S}_1(\mathbb{x}, \text{aux}) :$ $\pi, \text{st}_S \leftarrow \mathcal{S}(1, \text{st}_S, \mathbb{x}, \text{aux})$ $\mathcal{Q}_{\text{sim}} \leftarrow \mathcal{Q}_{\text{sim}} \cup \{(\mathbb{x}, \text{aux}, \pi)\}$ return π $\mathcal{S}_2(s, \text{aux}) :$ if $\nexists \text{aux}, a : (s, \text{aux}, a) \in \mathcal{Q}_{\text{RO}}$: $a, \text{st}_S \leftarrow \mathcal{S}(2, \text{st}_S, s, \text{aux})$ $\mathcal{Q}_{\text{RO}} \leftarrow \mathcal{Q}_{\text{RO}} \cup \{(s, \text{aux}, a)\}$ return a </pre>
--	---

Fig. 1. The Φ -simulation extractability experiments in ROM. The extraction policy Φ takes as input the public view of the adversary view (namely, all the inputs received and all the queries and answers to its oracles). The set \mathcal{Q}_{sim} is the set of queries and answers to the simulation oracle. The set \mathcal{Q}_{RO} is the set of queries and answers to the random oracle. The set \mathcal{Q}_{aux} is the set of all the auxiliary information sent by the adversary (depending on the policy, this set might be empty or not). The wrappers \mathcal{S}_1 and \mathcal{S}_2 deal respectively with the simulation queries and the random oracle queries of \mathcal{A} in the experiment.

4 Simulation extractability of KZG in AGM

We recap the polynomial commitment scheme of Kate, Zaverucha and Goldberg [38] (KZG, shortly), **highlighting in blue** the parts related to hiding. KZG is a Polynomial Commitment scheme (see Definition 3) defined over bilinear groups $\mathbb{G} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$ as follows:

KGen($1^\lambda, d$) on input the security parameter 1^λ , and a maximum degree bound $d \in \mathbb{N}$, outputs $\text{ck} := (([s^j]_1)_{j \in [0, d]}, ([\alpha s^j]_1)_{j \in [0, d]}, [1, s]_2)$, for a random secret $s, \alpha \leftarrow \mathbb{F}_q$.

Com($\text{ck}, f(X), r(X)$) on input ck , a polynomial $f(X)$, **and a masking polynomial** $r(X)$, outputs a commitment $c := [f(s) + \alpha r(s)]_1$.

VerCom($\text{ck}, c, f(X)$) outputs 1 if $c = [f(s) + \alpha r(s)]_1$.

The above scheme is (standard) binding under the d -DL assumption (see Groth [33]), in fact, given two polynomials f and f' that evaluate to the same value on the secret point s , we can find s among the roots of the (non-zero) polynomial $f - f'$.

4.1 CP-SNARK for polynomial evaluation in AGM

We consider a CP-SNARK CP_{ev1} for the relation $\mathcal{R}_{\text{ev1}}((x, y), f) := f(x) = y$, where f is committed as $[f(s) + \alpha r(s)]_1$. The scheme constructed in this section requires one \mathbb{G}_1 element to commit to $f(X)$, one \mathbb{G}_1 **and one** \mathbb{F}_q element for the evaluation proof, and checking this proof of evaluation requires two pairings. This is taken from [17] but adapted to AGM only.

KGen_{ev1}: parse ck as $(([s^j]_1)_{j \in [0,d]}, ([\alpha s^j]_1)_{j \in [0,d]}, [1, s]_2)$ and define $\text{ek} := \text{ck}$ and $\text{vk} := [1, s]_2$, and return $\text{srs} := (\text{ek}, \text{vk})$.
Prove_{ev1}($\text{ek}, \mathbb{x} = (c, x, y), \mathbb{w} = (f, r)$): compute the polynomial $\pi(X)$ such that $\pi(X)(X-x) \equiv f(X)-y$, and the polynomial $\pi'(X)$ such that $\pi'(X)(X-x) \equiv r(X) - r(x)$ and $y' := r(x)$, and output $\pi := ([\pi(s) + \alpha \pi'(s)]_1, y')$.
Verify_{ev1}($\text{vk}, \mathbb{x} = (c, x, y), (\pi, y')$): output 1 iff $e(c - [y]_1 - [\alpha y']_1, [1]_2) = e(\pi, [s - x]_2)$.

The above CP-SNARK is knowledge extractable in the AGM [20]. The original work of [38] proves a weaker notion of security, called evaluation binding, which states that an adversary cannot find two distinct instances (with relative valid proofs) of the form $\mathbb{x} = (c, x, y)$ and $\mathbb{x}' = (c, x, y')$. The CP-SNARK supports a bounded (by $\text{deg}(r)$) number of evaluation proofs for a given commitment. One may argue that giving more than $\text{deg}(f)$ evaluations of a polynomial f on distinct points should reveal the polynomial and, thus, the zero-knowledge property would not be needed. However, there are applications in which we could give more than $\text{deg}(f)$ evaluation proofs concerning f without necessarily revealing the evaluation values: e.g., this is achieved when we only show the evaluations of linear combinations of multiple committed polynomials to known constants. Since the technique of [38] would leak information on the random masking polynomials and would therefore be usable only a limited number of times, one may use the “full-fledged” CP-SNARK of Lunar [17] for proving an unbounded number of evaluations of committed polynomials in zero-knowledge, at the cost of two additional pairings at verification time.

The scheme that we describe above is zero-knowledge for non-hiding commitments and leaky zero-knowledge (Definition 11) for hiding commitments. For the latter, given a commitment $\mathbf{c} = [f(s) + \alpha r(s)]_1$, a proof for $\mathbb{x} = (c, x, y)$ leaks $y' = r(x)$. Thus, we prove that CP_{ev1} achieves F -leaky zero-knowledge where $F(\mathbb{x} = (c, x, y), \mathbb{w} = (f, r)) := r(x)$.

We define the simulator $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1)$, where \mathcal{S}_0 outputs the trapdoor information s, α together with the srs , and \mathcal{S}_1 simulates proofs for $\mathbb{x} = (c, x, y)$ and leakage y' outputting $\pi = ((c - [y]_1 - [\alpha y']_1)(s - x)^{-1}, y')$.

Remark 1. Consider an attacker that receives a single simulated commitment \mathbf{c} and queries the ZK-simulator twice on the same evaluation point x_j with two different evaluation values y_1 and y_2 . These two queries form a linear system without solutions because we have only one variable but two linearly independent equations. These simulation queries lead to the two malleability attacks on KZG that we describe below. For sake of simplicity, we show the attacks in the non-hiding setting.

(Arbitrary Evaluation Point.) Let w_1, w_2 and c be such that $\pi_1 = [w_1]_1, \pi_2 = [w_2]_1$ and $\mathbf{c}^* = [c]_1$ we observe that:

$$(w_1 - w_2)(s - x_j) = y_2 - y_1. \quad (1)$$

For an arbitrary point $x^* \neq x_j$, the adversary sets $\mathbf{c}^* \leftarrow (\pi_2 - \pi_1), \pi^* \leftarrow \frac{\pi_1 - \pi_2}{x^* - x_j}$ and $y^* \leftarrow \frac{y_1 - y_2}{x^* - x_j}$. The tuple $(\mathbf{c}^*, x^*, y^*, \pi^*)$ is a valid forgery. Let

$$[w^*]_1 = \pi^* \text{ and } [c^*]_1 = c^*,$$

$$\begin{aligned} w^*(s - x^*) &= \frac{w_1 - w_2}{x^* - x_j}(s - x^*) = \frac{w_2 - w_1}{x^* - x_j}(x^* - x_j) + \frac{w_1 - w_2}{x^* - x_j}(s - x_j) \\ &= w_2 - w_1 - \frac{y_1 - y_2}{x^* - x_j} = c^* - y^*. \end{aligned}$$

The second equation comes directly from the definition of the elements involved, and the third equation follows because of Eq. (1).

(Same Evaluation Point.) For $x^* = x_j$, instead, by the homomorphic property of KZG we have $((\alpha + \beta)c^*, x_j, \alpha y_1 + \beta y_2)$ is a valid forgery; we can forge a proof for c^* by setting, for example, $\alpha = 2$ and $\beta = -1$.

The attacks described in Remark 1 can be mounted because of the relation between two simulation queries obtained and it can be generalized under the notion of attacks that violate the ‘‘Algebraic Consistency’’, as we explain hereafter. Since here and in the next sections we focus on CP-SNARKs in which, given a proof π for a statement \mathbb{x} , and a view view , it is possible to derive a linear system of polynomial equations $\{p_i(x_j) = y_{i,j}\}_{i,j}$, we introduce the following definition to make the policies easier to describe.

Definition 14 (Algebraic Consistency). *A view view (that might contain a set of simulated instances and proofs for CP) satisfies the algebraic consistency for a CP-SNARK CP if it is possible to derive (in a way that depends on CP) a linear system of polynomial equations that admits a solution.*

The extraction policy. We define $\Phi_{\text{ev1}}^{\text{s-adpt}} = \{\Phi_{\mathcal{D}}\}_{\mathcal{D}}$ as the family (indexed by a sampler \mathcal{D}) of *semi-adaptive* extraction policies for KZG commitment scheme. Indeed, as we show below, the evaluation points x_j for the instances for which the adversary can see simulated proofs are selectively chosen independently of the commitment key, while the evaluation values y can be adaptively chosen by the adversary. Each policy $\Phi_{\mathcal{D}}$ is a tuple of the form $(\Phi_0^{\mathcal{D}}, \Phi_1)$, as defined in Section 3.1, where $\Phi_0^{\mathcal{D}}$ outputs the parameters pp_{Φ} while Φ_1 outputs a verdict bit. In particular, $\Phi_0^{\mathcal{D}}$ on input group parameters $\text{pp}_{\mathbb{G}}$ outputs $\text{pp}_{\Phi} := (\text{coms}, \mathcal{Q}_x)$, where coms is a vector of commitments sampled from \mathcal{D} , and \mathcal{Q}_x is a set of \mathcal{Q}_x evaluation points.

For sake of clarity, we define the policy Φ_1 as the logical conjunction of a ‘‘simulator’’ policy Φ_{sim} and an ‘‘extractor’’ policy Φ_{ext} , i.e. $\Phi_1 = \Phi_{\text{sim}} \wedge \Phi_{\text{ext}}$, the first defines rules under which we can classify a simulation query *legal*, while the second defines rules under which the extractor must be able to extract a meaningful witness. We [highlight](#) the parts needed only for the hiding setting.

Definition 15. *Let Φ_{sim} be the policy that returns 1 if and only if:*

1. **Points check:** *let $(\mathbb{x}_i, \text{aux}_i, \pi_i)_i$ be all the entries of \mathcal{Q}_{sim} . Recall that an instance \mathbb{x} can be parsed as (c, x, y) . Check that $\forall i : \mathbb{x}_i.x \in \mathcal{Q}_x$.*

2. **Commitment Check:** For any $i \in [Q_{\text{sim}}]$, parse aux_i as the leakage value y'_i and the representation vectors for $\mathbb{x}_i.c$ and π_i ; in particular, let $\mathbf{r}_i = \mathbf{f}_i \parallel \mathbf{v}_i \parallel \mathbf{c}_i$ be the algebraic representation of the commitment $\mathbb{x}_i.c$. For any i check that $\langle \mathbf{f}_i \parallel \mathbf{v}_i, \text{ek} \rangle + \langle \mathbf{c}_i, \text{coms} \rangle = \mathbb{x}_i.c$ (namely, that the commitment in the instance \mathbb{x}_i is computed as a linear combination of the simulated commitments and the elements of \mathbb{G}_1 of the SRS);
3. **Algebraic Consistency:** The simulation queries satisfy the algebraic consistency for CP_{evl} . Namely, let $\mathcal{I}_j = \{i : \mathbb{x}_i.x = x_j\}$ and let $\mathbf{R}_j = (\mathbf{c}_i)_{i \in \mathcal{I}_j}$. Check that for all j : (i) the system of linear equations $\mathbf{R}_j \mathbf{z} = \mathbf{y}_j$ has at least a solution, where \mathbf{z} are the variables and $\mathbf{y}_j = (\mathbb{x}_i.y - \langle \mathbf{f}_i, (1, x_j, \dots, x_j^d) \rangle)_{i \in \mathcal{I}_j}$, and (ii) the system of linear equations $\mathbf{R}_j \mathbf{z}' = \mathbf{y}'_j$ has at least a solution, where \mathbf{z}' are the variables and $\mathbf{y}'_j = (y'_i - \langle \mathbf{v}_i, (1, x_j, \dots, x_j^d) \rangle)_{i \in \mathcal{I}_j}$.

For the sake of concreteness, we explicitly define the algebraic consistency check in the previous definition. Notice the matrix \mathbf{R}_j defines linear constraints that must hold for each of the polynomials for which the adversary asks evaluations. In the case of hiding commitments, some of the constraints are obtained by parsing adequately the inputs of the adversary to the simulation oracle.

Next, we define the policy Φ_{ext} as the logical disjunction of two policies.

First, we define some notation, let $g_c : \mathbb{G}_1 \times \{0, 1\}^* \rightarrow \{0, 1\}$ be a function that on inputs a group element c and a string s , that can be parsed as a list of group elements c_i followed by a second string \tilde{s} , outputs 1 iff $\exists i : c = c_i$.

Definition 16. Let $\Phi_{\text{ext}}, \Phi_{\text{ext}}^{\text{rnd}}$ and $\Phi_{\text{ext}}^{\text{der}}$ be predicates that parse the forgery instance $\mathbb{x}^* = (c^*, x^*, y^*)$.

- $\Phi_{\text{ext}}^{\text{rnd}}$ returns 1 if and only if there exist a query (s, aux, a) to the random oracle and aux contains a non-constant polynomial $h(X)$ such that the following conditions are satisfied:
 1. **Hashing check:** $(s, \text{aux}, a) \in \mathcal{Q}_{\text{RO}}$, note that \mathcal{Q}_{RO} is contained in view ,
 2. **Decoding check:** $g_c(c^*, s) = 1$.
 3. **Polynomial check:** $g_h(h, \text{aux}) = 1$, where $g_h : \mathbb{F}[X] \times \{0, 1\}^* \rightarrow \{0, 1\}$ is a function that on input a polynomial $h(X)$ and a string aux outputs 1 iff $h(X)$ is encoded in aux .
 4. **Computation check:** $h(a) = x^*$.
- $\Phi_{\text{ext}}^{\text{der}}$ returns 1 iff $\exists (\mathbb{x}, \cdot, \pi) \in \mathcal{Q}_{\text{sim}}$ s.t. $\mathbb{x} := (c^*, x^*, y')$ and $(y', \pi) \neq (y^*, \pi^*)$.
- Φ_{ext} returns logical disjunction of $\Phi_{\text{ext}}^{\text{rnd}}$ and $\Phi_{\text{ext}}^{\text{der}}$.

In the following theorem, we focus on the non-hiding version of KZG, we show in Section 4.2 how to handle the hiding setting.

Theorem 1. For any witness sampleable distribution \mathcal{D} that is \mathcal{D} -Aff-MDH-secure (see Definition 5), any bilinear-group generator GroupGen that samples the generator of the group \mathbb{G}_1 uniformly at random, $\forall \Phi_{\mathcal{D}} \in \Phi_{\text{evl}}^{\text{s-adpt}}$, the non-hiding KZG CP_{evl} is $\Phi_{\mathcal{D}}$ -simulation-extractable in the AGM. In particular, there exists \mathcal{E} such that for any algebraic adversary \mathcal{A} :

$$\text{Adv}_{\text{CP}_{\text{evl}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{D}}\text{-se}}(\lambda) \leq O(\epsilon_{(Q_x+d+1)\text{-DL}}(\lambda)) + O(\epsilon_{\text{Aff-MDH}}(\lambda)) + \text{poly}(\lambda)\epsilon_h$$

where d is the maximum degree supported by CP_{ev1} , $\epsilon_{(Q_x+d+1)\text{-DL}}(\lambda)$ is the maximum advantage for any algebraic PT adversary against the $(Q_x + d + 1)$ -strong Discrete-Log Assumption, $\epsilon_{\text{Aff-MDH}}(\lambda)$ is the maximum advantage for any algebraic PT adversary against the \mathcal{D} -Aff-MDH Assumption, h is the polynomial that satisfies the Polynomial check of $\Phi_{\mathcal{D}}$, and $\epsilon_h = \frac{\deg(h)}{q}$.

Before giving the proof of Theorem 1, we recall that when \mathcal{D} is the distribution \mathcal{U}_ℓ that outputs ℓ uniformly random group elements of \mathbb{G}_1 we could reduce the \mathcal{D}_ℓ -Aff-MDH Assumption to the Discrete Log (see Lemma 2). Hence we can state the following corollary, whose proof follows from the fact that $\epsilon_{\text{DL}} \leq \epsilon_{d\text{-DL}}, \forall d \geq 1$.

Corollary 1. *For any algebraic adversary \mathcal{A} , for any $\ell \in \mathbb{N}$, and for any distribution \mathcal{U}_ℓ that outputs ℓ uniformly random group elements:*

$$\text{Adv}_{\text{CP}_{\text{ev1}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{U}_\ell} \text{-se}}(\lambda) \leq O(\ell \epsilon_{(Q_x+d+1)\text{-DL}}(\lambda)) + \text{poly}(\lambda) \epsilon_h.$$

Proof intuition of Theorem 1. The proof of this theorem is rather technical. We provide here an overview to help the reader get the rationale behind our strategy. The first steps of the proof show an alternative way to simulate proofs: we give a *semi-adaptive* zero-knowledge simulator for KZG. This step allows one to move from a simulator whose trapdoor is a “secret exponent” s to a simulator whose trapdoor is a ‘tower’ of \mathbb{G}_1 -elements $[s^i]_1$. The simulated SRS seen by the adversary includes only high-degree polynomials of the form $[p(s)s^i]_1$, while the simulator keeps the low-degree monomials $[s^i]_1$ for simulation. Here, p is a polynomial that vanishes in all the points to be asked in the simulation queries (this is reminiscent of the reduction technique for Boneh-Boyen signatures [11]). Since we program the SRS based on the queries our simulator is only semi-adaptive. This first change essentially simplifies the objects involved in our analysis, from rational polynomials (with the formal variable being the trapdoor) to standard polynomials.

Next, we need to show that the adversary cannot mix the simulated commitments and the forgery material. In particular, we need to show that the forged proof is not derived as a linear combination involving simulated commitments. To show this, we use the fact that the degree of the proof must be smaller than the degrees of simulated commitments, otherwise we could break the d -DL assumption in the AGM. This intuitively comes from the fact that the verification equation *lifts* the degree of the polynomial in the forged proof (as it is multiplied by $(X - x^*)$). Similarly, we need to show that the forged instance cannot use a linear combination that involves the simulated commitments. For this, we use the Aff-MDH assumption to handle multiple evaluation proofs on different simulated commitments on the same evaluation point. In particular, we reduce the view of many simulated proofs over many commitments and many evaluation points to a view that only contains $[p(s)s^i]_1$ and (non-rational) polynomials $[p(s)/(s - x_j)]_1$. At this point, the attacker could still perform the attack if it could decide the evaluation point x^* arbitrarily:

(i) the adversary asks a simulation proof π for $\mathbf{x} = (c, x, y)$, and (ii) produces the forgery $\mathbf{x}^* = (c + \alpha\pi, x - \alpha, y, \pi)$, for any $\alpha \in \mathbb{Z}_q$. It is easy to check that the forgery satisfies the verification equation. Though, for this attack to work the attacker needs to set the commitment in the forged instance as a function of $x^* = x - \alpha$. The last part of our analysis shows that, indeed, the algebraic representation of the commitment in the forgery cannot depend on x^* and that this attack cannot be mounted when x^* is chosen after the commitment with sufficient randomness.

Proof (of Theorem 1). We stress that \mathcal{A} is algebraic (cf. Definition 1), therefore for each group element output it additionally attaches a representation \mathbf{r} of such a group element with respect to all the elements seen during the experiment (included elements in \mathbf{coms}). In particular, we assume that for each query $(\mathbf{x}, \mathbf{aux})$ to the oracle \mathcal{S}_1 we can parse the value \mathbf{aux} as $(\mathbf{r}, \mathbf{aux}')$ and \mathbf{r} is a valid representation for $\mathbf{x}.c$. Similarly, for the queries (s, \mathbf{aux}) to \mathcal{S}_2 , \mathbf{aux} includes a valid representation for all the group elements \mathbf{g}_i encoded in s , i.e. such that $g_c(\mathbf{g}_i, s) = 1$. Together with its forgery, the algebraic adversary encodes a polynomial $h(X)$ in \mathbf{aux}_ϕ , and stores in \mathbf{aux}_ε two representation vectors \mathbf{r}_{c^*} and \mathbf{r}_{π^*} for the two group elements c^* and π^* . We can parse the vectors $\mathbf{r}_\tau := \mathbf{f}_\tau \| \mathbf{c}_\tau \| \mathbf{o}_\tau$ for $\tau \in \{c^*, \pi^*\}$ where \mathbf{f}_τ is the vector of coefficients associated to group elements \mathbf{ek} , \mathbf{c}_τ is the vector of coefficients associated to group elements $\mathbf{coms} = ([c_i]_1)_{i \in [Q_c]}$, and \mathbf{o}_τ is the vector of coefficients associated to the group elements of the simulated proofs \mathbf{proofs} . Namely, we have:

$$c^* = \langle \mathbf{f}_{c^*}, \mathbf{ek} \rangle + \langle \mathbf{c}_{c^*}, \mathbf{coms} \rangle + \langle \mathbf{o}_{c^*}, \mathbf{proofs} \rangle \quad (2)$$

$$\pi^* = \langle \mathbf{f}_{\pi^*}, \mathbf{ek} \rangle + \langle \mathbf{c}_{\pi^*}, \mathbf{coms} \rangle + \langle \mathbf{o}_{\pi^*}, \mathbf{proofs} \rangle \quad (3)$$

We can assume w.l.g. that all the simulation queries and the forgery of the adversary \mathcal{A} agree with the policy $\Phi_{\mathcal{D}}$, as otherwise the adversary would automatically lose the experiment. We assume that $\mathbf{f}_{i,j} = \mathbf{0}, \forall i, j$, i.e., the adversary asks simulation queries on commitments that are a linear combination of \mathbf{coms} only: this is also w.l.g. as we briefly show below. Given a commitment $\mathbf{c}_{i,j} = \mathbf{x}_{i,j}.c$, whose representation is $\mathbf{r}_{i,j} = \mathbf{f}_{i,j} \| \mathbf{c}_{i,j}$, the adversary could compute a proof $\pi_{i,j}$ for the point x_j and the evaluation value y as follows: (i) let $y' = f_{i,j}(x_j)$, \mathcal{A} computes an “honest” proof π' for $\mathbf{c}' = (\mathbf{Com}(\mathbf{ek}, f_{i,j}(X)), x_j, y')$, (ii) asks the simulation oracle to provide a proof $\tilde{\pi}$ for the instance $(c - c', x_j, y - y')$ with representation $\mathbf{0} \| \mathbf{c}_{i,j}$ and (iii) recombines the proof $\pi_{i,j} = \pi' + \tilde{\pi}$. We define our extractor to be the *canonical* extractor that returns the polynomial $f(X) \leftarrow \langle \mathbf{f}_{c^*}, (1, X, \dots, X^d) \rangle$.

We start by proving that for any algebraic adversary \mathcal{A} whose forgery satisfies the predicate $\Phi_{\text{ext}}^{\text{der}}$, there exists an algebraic adversary \mathcal{B} whose forgery satisfies the predicate $\Phi_{\text{ext}}^{\text{rnd}}$. Let $\{\Phi'_{\mathcal{D}}\}_{\mathcal{D}}$ be the family of policies defined exactly as $\Phi_{\text{evl}}^{\text{s-adpt}}$ with the difference that the extraction policy Φ_{ext} is equal to $\Phi_{\text{ext}}^{\text{rnd}}$ (i.e., there is no logical disjunction with $\Phi_{\text{ext}}^{\text{der}}$).

Lemma 3. *For any algebraic adversary \mathcal{A} there exists an algebraic adversary \mathcal{B} such that:*

$$\text{Adv}_{\text{CP}_{\text{evl}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{D}} - \text{se}}(\lambda) = \text{Adv}_{\text{CP}_{\text{evl}}, \mathcal{B}, \mathcal{S}, \mathcal{E}}^{\Phi'_{\mathcal{D}} - \text{se}}(\lambda)$$

Proof. First, we notice that once we fix a commitment \mathbf{c} , a point x , and a value y , there is a unique proof π that can satisfy the KZG verification equation. Thus, the predicate $\Phi_{\text{ext}}^{\text{der}}$ can be simplified as requiring that an adversary outputs a valid proof π^* and a value y^* such that $\exists((c_1^*, x^*, y'), \cdot, \pi) \in \mathcal{Q}_{\text{sim}}$ and $y^* \neq y'$.

The reduction \mathcal{B} internally runs \mathcal{A} forwarding all the simulation queries, up to the forgery (\mathbf{x}^*, π^*) , where $\mathbf{x}^* = (\mathbf{c}^*, x^*, y^*)$. If the simulation queries and/or the forgery of the adversary \mathcal{A} do not agree with the policy $\Phi_{\mathcal{D}}$, i.e. \mathcal{A} automatically loses its game, \mathcal{B} aborts. Otherwise, it must be true that the forgery of \mathcal{A} either (i) satisfies the extraction predicate $\Phi_{\text{ext}}^{\text{rnd}}$ or (ii) satisfies the extraction predicate $\Phi_{\text{ext}}^{\text{der}}$. Both cases can be efficiently checked by \mathcal{B} . In case (i) \mathcal{B} would simply forward the forgery of \mathcal{A} retaining the same advantage of \mathcal{A} . Otherwise, before submitting the forgery, \mathcal{B} retrieves from \mathcal{Q}_{sim} the statement $\mathbf{x} := (\mathbf{c}^*, x^*, y')$, where $y' \neq y^*$, and the corresponding proof π output by \mathcal{S}_1 . Then \mathcal{B} produces the forgery:

$$\hat{c} \leftarrow \pi^* - \pi, \quad \hat{x} \leftarrow h(a), \quad \hat{\pi} \leftarrow \frac{\pi - \pi^*}{\hat{x} - x^*}, \quad \hat{y} \leftarrow \frac{y - y^*}{\hat{x} - x^*}$$

which satisfies the verification equation (cf. Remark 1), and the extraction predicate $\Phi_{\text{ext}}^{\text{rnd}}$ when $(\hat{c}, h, a) \in \mathcal{Q}_{\text{RO}}$. \square

Thanks to Lemma 3 we can assume that the forgery of \mathcal{A} satisfies the extraction predicate $\Phi_{\text{ext}}^{\text{rnd}}$. We let \mathbf{H}_0 be the $\mathbf{Exp}_{\mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{D}}, \text{s-e}}(\lambda)$ experiment, and we denote by ϵ_i the advantage of \mathcal{A} to win \mathbf{H}_i , i.e. $\epsilon_i := \Pr[\mathbf{H}_i = 1]$.

Hybrid \mathbf{H}_1 . Recall that \mathcal{D} is witness sampleable, thus according to Definition 4 there exists a PPT algorithm $\tilde{\mathcal{D}}$ associated with the sampler \mathcal{D} . The hybrid experiment \mathbf{H}_1 is identical to the previous one, but the group elements in coms are “sampled at exponent”, i.e. we use $\tilde{\mathcal{D}}$ to generate the field elements γ , and we let $\text{coms} \leftarrow [\gamma]_1$; we also add γ to $\text{st}_{\mathcal{S}}$. By the witness sampleability of \mathcal{D} , \mathbf{H}_0 and \mathbf{H}_1 are perfectly indistinguishable, thus $\epsilon_1 = \epsilon_0$.

Hybrid \mathbf{H}_2 . In this hybrid, we change the way we generate the SRS srs and the way in which \mathcal{S}_1 simulates the proofs.

Let $((\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e), [1]_1, [1]_2) \leftarrow_{\$} \text{GroupGen}(1^\lambda)$, sample $s \leftarrow_{\$} \mathbb{F}$ and compute $[s, \dots, s^{D+d}]_1, [1, s]_2$, where $D \leftarrow Q_x + 1$. Let $x_r \leftarrow_{\$} \mathbb{F}$, and let $p(X)$ be the vanishing polynomial in $\mathcal{Q}_x \cup \{x_r\}$, namely:

$$p(X) := (X - x_r) \prod_{x \in \mathcal{Q}_x} (X - x).$$

Let also $p_j(X) := p(X)(X - x_j)^{-1}$, for $j \in [Q_x]$. In \mathbf{H}_2 we have that:

- $\text{pp}_{\mathbb{G}} := ((\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e), [p(s)]_1, [1]_2)$,
- $\text{srs} := (\text{ek}, \text{vk})$, where $\text{ek} \leftarrow [p(s), p(s)s, \dots, p(s)s^d]_1$ and $\text{vk} \leftarrow [1, s]_2$,
- $\text{st}_{\mathcal{S}} := [1, s, \dots, s^{D+d}]_1, [1, s]_2, \gamma$.

Upon a query of the form $(\mathbf{x} = (\mathbf{c}, x_j, y_k), \text{aux} = (\mathbf{r}_c, \text{aux}'))$ to \mathcal{S}_1 , the latter outputs the proof $\pi \leftarrow [(\langle \mathbf{r}_c, \gamma \rangle - y_k) \cdot p_j(s)]_1$, and updates \mathcal{Q}_{sim} accordingly.

We now show that $\mathbf{H}_1 \equiv \mathbf{H}_2$, i.e., the view offered to the adversary \mathcal{A} is identically distributed in the two experiments.

Lemma 4. $\epsilon_2 = \epsilon_1$.

Proof. Notice that in \mathbf{H}_2 we sample from `GroupGen` the description of the group, and then we set the generator of \mathbb{G}_1 to $[p(s)]_1$ which, thanks to the random root x_r , is distributed uniformly at random even given the value s . It is not hard to verify that the simulated proofs generated by the hybrid \mathbf{H}_2 pass the verification equations, in fact, we are assuming that queried commitment \mathbf{c} are of the form $\langle \mathbf{r}_c, \text{coms} \rangle$. Additionally, since the proofs are uniquely determined given the SRS and the statements, the simulated proofs created in \mathbf{H}_2 are distributed as the simulated proofs generated by the simulator \mathcal{S}_1 in \mathbf{H}_1 . Thus the advantage of \mathcal{A} is the same in the two experiments. \square

Given an algebraic adversary \mathcal{A} we can define a new adversary, \mathcal{A}_c , that we call the *core* adversary. Whenever the adversary \mathcal{A} outputs a group element \mathbf{g} it provides a representation vector $\mathbf{r}_g := \mathbf{f}_g \parallel \mathbf{c}_g \parallel \mathbf{o}_g$ for \mathbf{g} such that:

$$\mathbf{g} = \langle \mathbf{f}_g, \text{ek} \rangle + \langle \mathbf{c}_g, \text{coms} \rangle + \langle \mathbf{o}_g, \text{proofs} \rangle.$$

The adversary \mathcal{A}_c runs internally \mathcal{A} and forwards all the queries and answers from \mathcal{A} to its simulation oracle. However, the way of simulating RO queries must ensure to not alter the result of the extractor policy, i.e. the “hash-check” for x^* . This is why we cannot simply forward the queries of \mathcal{A} to the random oracle. Therefore, we keep track of the queries made by \mathcal{A} in the list $\mathcal{Q}_{\text{RO}, \mathcal{A}}$ and the list of queries made by the core adversary in \mathcal{Q}_{RO} . More in detail, when \mathcal{A} queries the RO with (s, aux) , the adversary \mathcal{A}_c makes a “core” RO query (s_c, aux_c) such that:

1. Let s be parsed as $(\mathbf{g}_i)_{i \in [k]}$ (the group elements in s whose representations $\mathbf{r}_{\mathbf{g}_i} := \mathbf{f}_{\mathbf{g}_i} \parallel \mathbf{c}_{\mathbf{g}_i} \parallel \mathbf{o}_{\mathbf{g}_i}$ are in aux) and a string \tilde{s} . Notice, since the adversary is algebraic we can un-ambiguously parse s as such.
2. For each i , \mathcal{A}_c computes the group elements $\mathbf{g}'_i = \mathbf{g}_i - \langle \mathbf{f}_{\mathbf{g}_i}, \text{ek} \rangle$. \mathcal{A}_c encodes into the string s' the group elements $(\mathbf{g}_i, \mathbf{g}'_i)_{i \in [k]}$.
3. \mathcal{A}_c queries the RO with (s_c, aux_c) , where $s_c := s' \parallel \tilde{s}$, and aux_c contains the representations of all the group elements in s' and the same function h encoded in aux . Finally, it forwards the output to \mathcal{A} , i.e. it adds (s, aux, a) to $\mathcal{Q}_{\text{RO}, \mathcal{A}}$, and adds (s, s_c) to (the initially empty) \mathcal{Q}_s .

Eventually, the adversary \mathcal{A} outputs as forgery a string s and the tuple (c', x', y', π') , together with representation vectors $\mathbf{r}_{c'}$ and $\mathbf{r}_{\pi'}$. Let $f(X) := \langle \mathbf{f}_{c'}, (1, X, \dots, X^d) \rangle$, $y := f(x')$, and $q(X)$ be such that $q(X)(X - x') = f(X) - y$. Let \mathbf{f}_q be the vector of the coefficients of $q(X)$, namely $q(X) := \langle \mathbf{f}_q, (1, X, \dots, X^d) \rangle$. The core adversary \mathcal{A}_c returns for its forgery the string s_c such that $(s, s_c) \in \mathcal{Q}_s$, and the tuple (c^*, x', y^*, π^*) , where $y^* \leftarrow y' - f(x')$ and:

$$c^* \leftarrow c' - \underbrace{[f(s)p(s)]_1}_{\text{Com}(\text{ck}, f(X))}, \quad \pi^* \leftarrow \pi' - \underbrace{[q(s)p(s)]_1}_{\text{Com}(\text{ck}, q(X))}$$

inserting into aux_Φ the (correct) algebraic representations $(\mathbf{0} \parallel \mathbf{c}_{c'} \parallel \mathbf{o}_{c'})$ for c^* and $((\mathbf{f}_{\pi'} - \mathbf{f}_q) \parallel \mathbf{c}_{\pi'} \parallel \mathbf{o}_{\pi'})$ for π^* .

Hybrid H₃. This hybrid is exactly the same of H₂ but instead of running the experiment with the adversary \mathcal{A} we run it with the core adversary \mathcal{A}_c .

Lemma 5. $\epsilon_3 = \epsilon_2$.

Proof. First, by construction, it is easy to verify that \mathcal{A}_c is algebraic. Thus we need to show that the forgery of \mathcal{A} is valid if and only if the forgery of \mathcal{A}_c is valid. By construction, we have:

$$c^* := c' - [f(s)p(s)]_1, \quad \pi^* := \pi' - [q(s)p(s)]_1, \quad y^* := y' - f(x').$$

By the verification equation of the forgery of \mathcal{A}_c we have:

$$\begin{aligned} & e(c^* - [y^*]_1, [1]_2) - e(\pi^*, [s - x^*]_2) = \\ & e(c' - [f(s)p(s)]_1 - [y' - f(x')]_1, [1]_2) - e(\pi' - [q(s)p(s)]_1, [s - x^*]_2) = \\ & e(c' - [y']_1, [1]_2) - e(\pi', [s - x']_2) - [f(s)p(s) - f(x') - q(s)p(s)(s - x^*)]_T = \\ & e(c' - [y']_1, [1]_2) - e(\pi', [s - x']_2), \end{aligned}$$

where the last equation holds since $q(X)(X - x') = (f(X) - f(x'))$ and $x^* = x'$.

Finally, notice that a forgery is valid for \mathcal{A} if it provides a string s that satisfies the “hash check” of Φ_{ext} . We have that there exist s , \mathbf{aux} , a , and $h(X)$ such that: (i) $g_c(c^*, s) = 1$, (ii) $g_h(h, \mathbf{aux}) = 1$, (iii) $(s, \mathbf{aux}, a) \in \mathcal{Q}_{\text{RO}, \mathcal{A}}$, and (iv) $x^* = h(a)$ for the forgery of \mathcal{A} .

The way \mathcal{A}_c simulates the RO queries ensures that for the query s of \mathcal{A} to the RO, the core adversary sent the “core” RO query s_c that encodes both c' and c^* , thus we have that (i) $g_c(c^*, s_c) = 1$, (ii) $g_h(h, \mathbf{aux}_c) = 1$, (iii) $(s_c, \mathbf{aux}_c, a) \in \mathcal{Q}_{\text{RO}}$, and (iv) $x^* = h(a)$ for the forgery of \mathcal{A}_c .

Notice that if we run the canonical extractor on the outputs of the core adversary \mathcal{A}_c , the extractor sets the extracted witness to be the zero polynomial.

Hybrid H₄. The hybrid H₄ additionally checks that $\mathbf{f}_{\pi^*} \neq \mathbf{0} \vee \mathbf{c}_{\pi^*} \neq \mathbf{0}$, and if the condition holds the adversary \mathcal{A}_c loses the game.

Lemma 6. $\epsilon_3 \leq \epsilon_4 + \epsilon_{(Q_x+d+1)\text{-DL}}$

Proof. Recall that from the definition of the experiment, upon a query $(\mathbb{x}, \mathbf{aux})$ from \mathcal{A}_c to the simulation oracle of the form $\mathbb{x} = (c, x_j, y_k)$ and $\mathbf{aux} = \mathbf{r}$ where $c = \langle \mathbf{r}, \text{coms} \rangle$, the adversary receives the proof $[\pi_{\mathbf{r}, j, k}(s)]_1$ where:

$$\pi_{\mathbf{r}, j, k}(X) := (\langle \mathbf{r}, (\gamma_i)_i \rangle - y_k)p_j(X).$$

Consider the following polynomials:

$$\begin{aligned} c^*(X) &:= \sum_{i \in [Q_c]} c_{c^*, i} \cdot \gamma_i p(X) + \sum_{\mathbf{r}, j, k} o_{c^*, \mathbf{r}, j, k} \cdot \pi_{\mathbf{r}, j, k}(X) \\ \pi^*(X) &:= \sum_{i \in [Q_c]} c_{\pi^*, i} \cdot \gamma_i p(X) + \sum_{\mathbf{r}, j, k} o_{\pi^*, \mathbf{r}, j, k} \cdot \pi_{\mathbf{r}, j, k}(X) + \sum_{i \in [d+1]} f_{\pi^*, i} X^{i-1} p(X) \\ v(X) &:= c^*(X) - y^* p(X) - (X - x^*) \pi^*(X) \end{aligned}$$

By the guarantees of the AGM, we have $\mathbf{c}^* = [c^*(s)]_1$ and $\pi^* = [\pi^*(s)]_1$, moreover, if the verification equation is satisfied by the forgery of \mathcal{A}_c , then $v(s) = 0$.

Next, we show that when the forgery of the adversary is valid the probability of $\mathbf{f}_{\pi^*} \neq \mathbf{0}$ or $\mathbf{c}_{\pi^*} \neq \mathbf{0}$ is bounded by $\epsilon_{(Q_x+d+1)\text{-DL}}$.

First, notice that if the verification equation for \mathcal{A}_c holds then polynomial $v(X)$ must be equivalent to the zero polynomial with overwhelming probability. In fact, $v(s) = 0$ when the verification equation holds; if $v(X)$ is not the zero polynomial then, by Lemma 1, we can reduce \mathcal{A}_c to an adversary to the $(Q_x + d + 1)$ -DL assumption. Thus:

$$c^*(X) - y^*p(X) - (X - x^*)\pi^*(X) = v(X) = 0. \quad (4)$$

By the guarantees of the AGM, the polynomial $\pi^*(X)$ is a linear combination of elements that depend on $X^{i-1}p(X)$ for $i \in [d + 1]$ and $p_j(X)$ for $j \in [Q_x]$. However, when the verification equation holds, the degree of $\pi^*(X)$ must be strictly less than the degree of $p(X)$, because, by Eq. (4), $v(X)$ would contain a non-zero coefficient of degree $Q_x + d + 1$ which in particular implies that $v(X) \neq 0$. Then it must be the case that the forged proof $\pi^*(s)$ is a linear combination of the simulated proofs only, thus both \mathbf{f}_{π^*} and \mathbf{c}_{π^*} are null. \square

The representation of \mathbf{c}^* and π^* computed by the adversary (possibly) depends on the elements $\pi_{\mathbf{r},j,k}$ (i.e. the proof for the linear combination \mathbf{r} of the elements of **coms** with evaluation point x_j and evaluation value y_k) of **proofs**. However, it is much more convenient to give a representation that depends on the polynomials $p_j(X)$. This motivates the definition of our next hybrid.

Hybrid H₅. The hybrid **H₅** finds coefficients \mathbf{o}''_τ , for $\tau \in \{c^*, \pi^*\}$ such that:

$$\langle \mathbf{o}_\tau, \mathbf{proofs} \rangle = \langle \mathbf{o}''_\tau, ([p_j(s)]_1)_j \rangle. \quad (5)$$

Moreover, if $\mathbf{o}_{c^*} \neq \mathbf{0}$ but $\mathbf{o}''_{c^*} = \mathbf{0}$ the adversary loses the game.

Lemma 7. $\epsilon_4 \leq \epsilon_5 + \epsilon_{\text{Aff-MDH}}$

Proof. We begin by showing that the hybrid can compute such alternative representations efficiently. We proceed in steps.

Let us parse the simulated proofs $\mathbf{proofs} := (\pi_{j,\ell})_{j,\ell}$ such that $\pi_{j,\ell}$ is the ℓ -th simulated proof obtained by \mathcal{S}_1 on a query involving the j -th point x_j , i.e., $((x_j, \hat{c}_{j,\ell}, y_{j,\ell}), \mathbf{aux}_{j,\ell})$. Also, let $\mathbf{c}_{j,\ell}$ be the algebraic representation for the group element $\hat{c}_{j,\ell}$ in $\mathbf{aux}_{j,\ell}$. For every $j \in [Q_x]$, we define \mathbf{R}_j as the $Q_c \times Q_c$ matrix whose ℓ -th column is $\mathbf{c}_{j,\ell}$.

By construction of \mathcal{S}_1 in this hybrid we have that for every $j \in [Q_x]$ it holds

$$\pi_{j,\ell} := [(\mathbf{c}_{j,\ell}^\top \cdot \boldsymbol{\gamma} - y_{j,\ell}) p_j(s)]_1$$

and thus $\boldsymbol{\pi}_j := [(\mathbf{R}_j^\top \boldsymbol{\gamma} - \mathbf{y}_j) p_j(s)]_1$ with $\mathbf{y}_j := (y_{j,\ell})_\ell$.

Without loss of generality, we assume that for each x_j the adversary makes the maximum number of simulation queries (i.e., $\ell \in [Q_c]$); therefore \mathbf{R}_j is a

full rank matrix (this follows from the fact that the simulation queries of the adversary satisfy the policy Φ_{sim} , and in particular the algebraic consistency of the policy, see Item 3).

Given any vector \mathbf{o}_τ with $\tau \in \{\mathbf{c}^*, \pi^*\}$, its m -th entry $o_{\tau,m}$ corresponds to the m -th simulated proof in **proofs**. Therefore, similarly to above, we denote by $o_{\tau,j,\ell}$ the entry corresponding to proof $\pi_{j,\ell}$ and we define $\mathbf{o}_{\tau,j} := (o_{\tau,j,\ell})_\ell$.

Then, for every $j \in [Q_x]$ we define

$$\mathbf{o}'_{\tau,j} \leftarrow \mathbf{R}_j \cdot \mathbf{o}_{\tau,j} \quad (6)$$

$$\boldsymbol{\pi}'_j \leftarrow (\mathbf{R}_j^\top)^{-1} \cdot \boldsymbol{\pi}_j \quad (7)$$

from which we derive that for any j :

$$\begin{aligned} \sum_j \langle \mathbf{o}'_{\tau,j}, \boldsymbol{\pi}'_j \rangle &= \sum_j \langle \mathbf{R}_j \cdot \mathbf{o}_{\tau,j}, (\mathbf{R}_j^\top)^{-1} \cdot \boldsymbol{\pi}_j \rangle \\ &= \sum_j \mathbf{o}_{\tau,j}^\top \mathbf{R}_j^\top (\mathbf{R}_j^\top)^{-1} \cdot \boldsymbol{\pi}_j \\ &= \sum_j \langle \mathbf{o}_{\tau,j}, \boldsymbol{\pi}_j \rangle \end{aligned}$$

which is equal to $\langle \mathbf{o}_\tau, \text{proofs} \rangle$, up to a permutation of the indices j .

For all $j \in [Q_x]$ let $\mathbf{z}_j := (\mathbf{R}_j^\top)^{-1} \cdot \mathbf{y}_j$, and note that

$$\boldsymbol{\pi}'_j = [(\boldsymbol{\gamma} - \mathbf{z}_j)p_j(s)]_1,$$

namely $\boldsymbol{\pi}'_{j,i}$ is a valid proof for the instance $(\mathbf{c}_i, x_j, z_{j,i})$ w.r.t. the simulated SRS.

The hybrid \mathbf{H}_5 computes $\mathbf{o}''_{\tau,j} \leftarrow \langle \mathbf{o}'_{\tau,j}, (\boldsymbol{\gamma} - \mathbf{z}_j) \rangle$, and $\mathbf{o}''_\tau \leftarrow (o''_{\tau,j})_{j \in [Q_x]}$. By construction:

$$\sum_{j \in [Q_x]} \langle \mathbf{o}'_{\tau,j}, \boldsymbol{\pi}'_j \rangle = \sum_{j \in [Q_x]} o''_{\tau,j} \cdot [p_j(s)]_1.$$

which proves the first part of the lemma, i.e., computing $\mathbf{o}''_{\tau,j}$ satisfying Eq. (5).

In what follows, we prove that if the event that \mathbf{H}_5 outputs 0 but \mathbf{H}_4 would output 1, namely that all the conditions of \mathbf{H}_4 hold but $\mathbf{o}_{\mathbf{c}^*} \neq \mathbf{0} \wedge \mathbf{o}''_{\mathbf{c}^*} = \mathbf{0}$, then we can break the Aff-MDH assumption.

First, notice that for any j $\mathbf{o}_{\mathbf{c}^*,j} \neq \mathbf{0}$ implies that $\mathbf{o}'_{\mathbf{c}^*,j} \neq \mathbf{0}$, because the linear transformation applied to compute $\mathbf{o}'_{\mathbf{c}^*,j}$ is full rank. Second, take an index j^* such that $\mathbf{o}_{\mathbf{c}^*,j^*} \neq \mathbf{0}$ and set $\mathbf{A} \leftarrow \mathbf{o}'_{\mathbf{c}^*,j^*}$ and $\zeta \leftarrow \langle \mathbf{z}_{j^*}, \mathbf{o}'_{\mathbf{c}^*,j^*} \rangle$.

By the above definition of the values $\mathbf{o}''_{\mathbf{c}^*,j^*}$ and our assumption that the “bad event” of this hybrid is $\mathbf{o}''_{\mathbf{c}^*} = \mathbf{0}$, we have that:

$$\langle \mathbf{A}, [\boldsymbol{\gamma}]_1 \rangle = \underbrace{\langle \mathbf{o}'_{\mathbf{c}^*,j^*}, (\boldsymbol{\gamma} - \mathbf{z}_{j^*}) \rangle}_{{o''_{\mathbf{c}^*,j^*} = 0}} + \underbrace{\langle \mathbf{o}'_{\mathbf{c}^*,j^*}, \mathbf{z}_{j^*} \rangle}_\zeta = [\zeta]_1.$$

The reduction \mathcal{B} to the \mathcal{D} -Aff-MDH Assumption takes as input a distribution $[\boldsymbol{\gamma}]_1$ and runs the experiment as in \mathbf{H}_4 (it perfectly emulates \mathbf{H}_4 , and in particular

the simulation oracle, because it knows the trapdoor s “at the exponent”). Then \mathcal{B} computes the coefficients $(A_i)_{i \in [Q_c]}$ and the value ζ as described above, which is a valid \mathcal{D} -Aff-MDH solution. \square

Hybrid \mathbf{H}_6 . The hybrid \mathbf{H}_6 additionally checks that $\mathbf{r}_{c^*} \neq \mathbf{0}$, and if the condition holds the adversary \mathcal{A}_c loses the game.

Lemma 8. $\epsilon_5 \leq \epsilon_6 + \epsilon_{\text{Aff-MDH}} + 2\epsilon_{(Q_x+1+d)\text{-DL}} + \text{poly}(\lambda) \frac{\deg(h)}{q}$

Proof. We bound the probability that the adversary loses in \mathbf{H}_6 but not in \mathbf{H}_5 , namely, the probability that $\mathbf{r}_{c^*} \neq \mathbf{0}$ but the conditions of \mathbf{H}_5 hold, we show a reduction \mathcal{B} to the Aff-MDH when this event happens.

First of all, we can assume that the core adversary outputs coefficients $\mathbf{f}_{c^*} = \mathbf{f}_{\pi^*} = \mathbf{c}_{\pi^*} = \mathbf{0}$, i.e. the adversary only makes use of previous commitments $\mathbf{c}_i \in \text{coms}$ and simulated proofs $\pi_{\mathbf{r},j,k} \in \text{proofs}$ to represent \mathbf{c}^* , and only uses the simulated proofs to represent the proof π^* .

The reduction \mathcal{B} takes as input a distribution $[\gamma]_1$ and runs the experiment as in \mathbf{H}_5 . \mathcal{B} aborts if the forgery (c^*, x^*, y^*, π^*) returned by the adversary is not valid (i.e. either the extraction predicate or the verification equation is not satisfied) or $\mathbf{r}_{c^*} = \mathbf{0}$. Otherwise, we have that:

$$e(\mathbf{c}^* - [p(s)y^*]_1, [1]_2) = e(\pi^*, [s - x^*]_2) \text{ and } \mathbf{r}_{c^*} \neq \mathbf{0}$$

where $\mathbf{r}_{c^*} \neq \mathbf{0}$ if $\mathbf{o}_{c^*} \neq \mathbf{0} \vee \mathbf{c}_{c^*} \neq \mathbf{0}$.

We can then rewrite the commitment and the proof of forgery of the core adversary as a function of the coefficients \mathbf{o}_{c^*}'' and \mathbf{o}_{π^*}'' (as computed in the \mathbf{H}_5):

$$\mathbf{c}^* := \sum_{i \in [Q_c]} c_{c^*,i} [\gamma_i p(s)]_1 + \sum_{j \in [Q_x]} o_{c^*,j}'' [p_j(s)]_1, \quad \pi^* := \sum_{j \in [Q_x]} o_{\pi^*,j}'' [p_j(s)]_1$$

Since the verification equation is satisfied, and plugging in the AGM representations we have:

$$\sum_{i \in [Q_c]} c_{c^*,i} \gamma_i p(s) + \sum_{j \in [Q_x]} o_{c^*,j}'' p_j(s) - p(s)y^* = \sum_{j \in [Q_x]} o_{\pi^*,j}'' p_j(s)(s - x^*) \quad (8)$$

For all $j \in [Q_x]$, we define $\delta_j := x_j - x^*$. We can rewrite the r.h.s. of Eq. (8) as:

$$\begin{aligned} \sum_{j \in [Q_x]} o_{\pi^*,j}'' p_j(s)(s - x^*) &= \sum_{j \in [Q_x]} o_{\pi^*,j}'' p_j(s)((s - x_j) + \delta_j) \\ &= \sum_{j \in [Q_x]} o_{\pi^*,j}'' (p(s) + p_j(s)\delta_j) \end{aligned}$$

In Eq. (8), we group all the terms that depend on $p(s)$ on the left side, and we move all the terms that depend on $p_j(s)$ on the right side, thus obtaining:

$$\underbrace{\left(\sum_{i \in [Q_c]} c_{c^*,i} \gamma_i - \sum_{j \in [Q_x]} o_{c^*,j}'' - y^* \right) p(s)}_A = \sum_{j \in [Q_x]} \underbrace{\left(o_{\pi^*,j}'' \delta_j - o_{c^*,j}'' \right)}_{B_j} p_j(s) \quad (9)$$

Let $f(X) := Ap(X) - \sum_{j \in [Q_x]} B_j p_j(X)$. Notice that because of Eq. (9) we have $f(s) = 0$, thus we can assume $f(X) \equiv 0$, as otherwise we can reduce, by Lemma 1, to the $(Q_x + d + 1)$ -DL assumption. It must be the case that both $\sum_{j \in [Q_x]} B_j p_j(s) = 0$ and $A = 0$ because the degree of $p(X)$ and of $p_j(X)$ for any j are different. Moreover, the polynomials $p_j(X)$ are linearly independent, namely the only linear combination $\sum_j a_j p_j(X) = 0$ is the trivial one where the coefficients $a_j = 0$ ¹¹, thus $B_j = 0$ for all j . We have that $o''_{w^*,j} \delta_j - o''_{c^*,j} = 0, \forall j$. Thus we can rewrite the coefficients $o''_{\pi^*,j} = \frac{o''_{c^*,j}}{\delta_j}, \forall j$. Since A must be 0:

$$\sum_{i \in [Q_c]} c_{c^*,i} \gamma_i - \sum_{j \in [Q_x]} \frac{o''_{c^*,j}}{\delta_j} - y^* = 0. \quad (10)$$

\mathcal{B} can plug the definition of the coefficients $o''_{c^*,j}$ in Eq. (10) and derive:

$$\begin{aligned} 0 &= \sum_{i \in [Q_c]} c_{c^*,i} \gamma_i - \sum_{i,j} \frac{o'_{c^*,i,j} (\gamma_i - z_{j_i})}{\delta_j} - y^* \\ &= \sum_{i \in [Q_c]} c_{c^*,i} \gamma_i - \sum_i \gamma_i \sum_j \frac{o'_{c^*,i,j}}{\delta_j} + \sum_{i,j} \frac{o'_{c^*,i,j} z_{j_i}}{\delta_j} - y^* \\ &= \sum_{i \in [Q_c]} (c_{c^*,i} - \sum_j \frac{o'_{c^*,i,j}}{\delta_j}) \gamma_i + \sum_{i,j} \frac{o'_{c^*,i,j} z_{j_i}}{\delta_j} - y^*. \end{aligned}$$

Above, the second equation follows from the distributive property of the sum, while in the last step we have grouped the terms depending on γ_i . In particular, the last equation shows that \mathcal{B} can make a forgery in the Aff-MDH game since it knows $z := y^* - \sum_{i,j} \frac{o'_{c^*,i,j} z_{j_i}}{\delta_j}$ and coefficients $A_i := c_{c^*,i} - \sum_j \frac{o'_{c^*,i,j}}{\delta_j}$ such that:

$$\sum_{i \in [Q_c]} A_i [\gamma_i]_1 = [z]_1.$$

For this to be a valid solution in the Aff-MDH game, we need the existence of at least an index i such that $A_i \neq 0$. We show that this occurs with all but negligible probability, i.e., $\Pr[\exists i \in [Q_c] : A_i \neq 0] \geq 1 - \text{negl}(\lambda)$.

To this end, consider an arbitrary $\mu \in [Q_c]$, then we have $\Pr[\forall i \in [Q_c] : A_i = 0] \leq \Pr[A_\mu = 0]$. Thus, for any μ , we have:

$$\Pr[\exists i \in [Q_c] : A_i \neq 0] = 1 - \Pr[\forall i \in [Q_c] : A_i = 0] \geq 1 - \Pr[A_\mu = 0].$$

Below, we argue that $\Pr[A_\mu = 0]$ is negligible based on the randomness of x^* which is chosen by the random oracle after defining A_μ , and we make use of the assumption that $\mathbf{r}_{c^*} \neq 0$.

¹¹ To see this, $\forall x_j \in \mathcal{Q}_x$ we have that $\sum_{j'} a_{j'} p_{j'}(x_j) = a_j p_j(x_j)$ since $p_j(x_j) \neq 0$ and $p_{j'}(x_j) = 0$ for $j \neq j'$, and $a_j p_j(x_j) = 0$ iff $a_j = 0$

We claim that the value $A_\mu = c_{\mathbf{c}^*, \mu} - \sum_j \frac{o'_{\mathbf{c}^*, j, \mu}}{(x^* - x_j)}$ can be fixed before the random oracle query x^* is made. To this end, we start by showing that $\mathbf{o}'_{\mathbf{c}^*, j}$ does not depend on x^* . Let $B(j) \subseteq [Q_c]$ be the subset of indices of the simulation queries that involve x_j and that occurred before the random oracle query that returned x^* . We observe that for every $\eta \in B(j)$ it must be $o_{\mathbf{c}^*, j, \eta} = 0$ since the simulated proof $\pi_{j, \eta}$ is not in the view of the adversary. Therefore:

$$o'_{\mathbf{c}^*, j, i} = \sum_{\eta \in [Q_c]} \mathbf{R}_{j, \eta, i} \cdot o_{\mathbf{c}^*, j, \eta} = \sum_{\eta \in B(j)} \mathbf{R}_{j, \eta, i} \cdot o_{\mathbf{c}^*, j, \eta}$$

and observe that all the rows of \mathbf{R}_j belonging to $B(j)$ can all be defined before x^* is sampled. Hence, we have that A_μ depends on the values $\mathbf{c}_{\mathbf{c}^*}$, x^* , $\{x_j\}_j$, and $\mathbf{o}_{\mathbf{c}^*, j}$ which can all be defined before the random oracle query x^* is made.

Now, we bound $\Pr[A_\mu = 0]$. Recall that, since the extractor policy Φ_{ext} holds true, we have that $x^* = h(a)$ and $(s, \text{aux}, a) \in \mathcal{Q}_{\text{RO}}$ where $g_c(\mathbf{c}^*, s) = 1$ and the function h is the polynomial encoded in aux_ϕ : the adversary may want to encode up to $n \in \text{poly}(\lambda)$ different polynomials h_i into aux_ϕ to maximize its advantage, and the extractor policy does not impose any restriction on this. Moreover, by the AGM, since \mathcal{A}_c sends a query s (where \mathbf{c}^* is encoded in s) to the random oracle it also defines coefficients for \mathbf{c}^* before the value a , and therefore $x^* = h(a)$, is defined. Also, it is not hard to see that the representation vector of \mathbf{c}^* defined by \mathcal{A}_c when querying the random oracle must be the same representation vector used for the forgery. As otherwise we would break the $(Q_x + d + 1)$ -DL assumption. Thus the coefficients $\mathbf{c}_{\mathbf{c}^*}$ and $\mathbf{o}'_{\mathbf{c}^*, j}$ are defined by the adversary before seeing the random value x^* .

Notice that, once the coefficients $\mathbf{c}_{\mathbf{c}^*}$ and $\mathbf{o}'_{\mathbf{c}^*, j}$ are fixed, the coefficient A_μ can be seen as function of $x^* \in \mathbb{Z}_q$, i.e. $A_\mu = A_\mu(x^*)$, where:

$$\begin{aligned} A_\mu(X) &:= c_{\mathbf{c}^*, \mu} + \sum_j \frac{o'_{\mathbf{c}^*, j, \mu}}{X - x_j} \\ &= \frac{c_{\mathbf{c}^*, \mu} \prod_j (X - x_j) + \sum_j (o'_{\mathbf{c}^*, j, \mu} \prod_{j' \neq j} (x_{j'} - X))}{\prod_j X - x_j}. \end{aligned}$$

Notice that $A_\mu(X) (\prod_j (X - x_j))$ vanishes in at most Q_x points in $\mathbb{F} \setminus \mathcal{Q}_x$ and vanishes in the set of points \mathcal{Q}_x . Let \mathcal{R} be the set of the roots of such a polynomial, since $\forall i \in [n]$, h_i is defined before x^* is computed, and by union bound:

$$\Pr[\exists i : h_i(\text{RO}(s)) \in \mathcal{R}] \leq \sum_{r \in \mathcal{R}} \Pr[\exists i : h_i(\text{RO}(s)) = r] \leq n Q_x \frac{\max_i \deg(h_i)}{q}$$

for each string s that encodes \mathbf{c}^* . To conclude, we notice that \mathcal{A} can submit at most Q_{RO} queries to the RO with strings encoding \mathbf{c}^* , say $s_1, \dots, s_{Q_{\text{RO}}}$. Thus the probability that there exist $i \in [n]$, $j \in [Q_{\text{RO}}]$ such that $h_i(\text{RO}(s_j)) \in \mathcal{R}$ is bounded by $n Q_{\text{RO}} Q_x \frac{\max_i \deg(h_i)}{q}$. \square

Hybrid \mathbf{H}_7 . The hybrid \mathbf{H}_7 additionally checks that $y^* \neq 0$, and if the condition holds the adversary \mathcal{A}_c loses the game.

Lemma 9. $\epsilon_6 \leq \epsilon_7 + \epsilon_{(Q_x+1+d)\text{-DL}} + \text{poly}(\lambda) \frac{\deg(h)}{q}$

Proof. We reduce to the evaluation binding of KZG polynomial commitment for polynomials of maximum degree $Q_x + 1 + d$, which, in turn, can be reduced to $(Q_x + 1 + d)$ -strong Discrete Log assumption. Let \mathcal{B} be the reduction that upon input $\text{pp}_{\mathbb{G}}, \text{ck} = [1, s, \dots, s^{Q_x+d+1}]_1, [1, s]_2$ simulates experiment \mathbf{H}_4 for the adversary \mathcal{A}_c . Eventually, \mathcal{A}_c outputs its forgery (c^*, x^*, y^*, π^*) , and \mathcal{B} aborts if $y^* = 0$. The reduction sets $\tilde{f}(X) := -y^*p(X)$, sets $y := \tilde{f}(x^*)$, and computes π to be a valid KZG-proof for $([\tilde{f}(s)]_1, x^*, y)$. The forgery against evaluation binding of the reduction is set to be (y, π) and $(0, \pi^*)$ for the commitment $[\tilde{f}(s)]_1$ on the point x^* .

We need to show that (1) $([\tilde{f}(s)]_1, x^*, 0, \pi^*)$ passes the verification equation of KZG commitment where the commitment key is set to ck and that (2) $y \neq 0$. For the first item notice that, by the definition of core adversary, we have that $\mathbf{r}_{c^*} = \mathbf{0}$ thus $c^* = 0$. Therefore, by the verification equation:

$$e([0]_1 - y^* [p(s)]_1, [1]_2) = e([\tilde{f}(s)]_1 - 0 [1]_1, [1]_2) = e(\pi^*, [s]_2 - x^* [1]_2).$$

For the second item, notice that $\tilde{f}(x^*) = 0$ if and only if x^* is a root of $p(X)$, i.e. $x^* \in \mathcal{Q}_x$ or $x^* = x_r$. Thus, similarly to the previous lemma, by the assumption on the polynomials h_i and by union bound:

$$\Pr[\exists i : h_i(\text{RO}(s)) \in \mathcal{Q}_x \cup \{x_r\}] \leq nQ_{\text{RO}}(Q_x + 1) \frac{\max_i \deg(h_i)}{q}.$$

□

Finally, we have that the probability that the adversary wins in \mathbf{H}_7 is null, namely $\epsilon_7 = 0$. Indeed, the canonical extractor \mathcal{E} outputs the 0 polynomial, moreover because of the condition introduced in \mathbf{H}_6 , we have $c^* = [0]_1$, and because of the condition introduced in \mathbf{H}_7 we have $y^* = 0$, thus the witness extracted is valid for the instance $\mathfrak{x}^* = (c^* = [0]_1, x^*, y^* = 0)$. □

CP-SNARK for evaluation of multiple polynomials in AGM. We give in Appendix A a scheme $\text{CP}_{\text{m-ev1}}$ which follows from [26,42] and that allows one to prove that for all i the polynomial f_i committed in c_i evaluates to y_i on the point x . We also show how to extend the policy $\Phi_{\text{sim}}^{\text{s-adpt}}$ to the batched setting.

4.2 Simulation extractability of Hiding KZG

We extend the result of Theorem 1 to the case in which the CP_{ev1} based on KZG uses hiding commitments. In particular, we show a reduction to the simulation extractability of non-hiding KZG for uniform distributions \mathcal{U}_ℓ of the commitments.

Theorem 2. *Let \mathcal{U}_ℓ be the distribution that outputs ℓ uniformly random group elements in \mathbb{G}_1 . $\forall \ell \in \mathbb{N}$, the hiding KZG CP_{ev1} (see Section 4) is $\Phi_{\mathcal{U}_\ell}$ -simulation-extractable in the AGM. In particular, there exists \mathcal{E} such that for any algebraic adversary \mathcal{A} :*

$$\text{Adv}_{\text{CP}_{\text{ev1}}, \mathcal{A}, \mathcal{S}, \mathcal{E}, \mathcal{U}_\ell}^{\Phi\text{-se}}(\lambda) \leq O(\ell \epsilon_{q\text{-DL}}(\lambda)) + \text{poly}(\lambda) \epsilon_h(\lambda).$$

where d is the maximum degree supported by CP_{ev1} , $\epsilon_{(Q_x+d+1)\text{-DL}}(\lambda)$ is the maximum advantage for any algebraic PT adversary against the (Q_x+d+1) -strong Discrete-Log Assumption, $\epsilon_{\text{Aff-MDH}}(\lambda)$ is the maximum advantage for any algebraic PT adversary against the \mathcal{D} -Aff-MDH Assumption, h is the polynomial that satisfies the Polynomial check of $\Phi_{\mathcal{D}}$, and $\epsilon_h = \frac{\deg(h)}{q}$.

Proof. We reduce an adversary \mathcal{A} for the simulation extractability of Π_{ev1} with hiding commitments to an adversary \mathcal{B} for the simulation extractability with non-hiding commitments and a uniform distribution that generates twice as many commitments. \mathcal{B} receives as input the description of a bilinear group $\text{pp}_{\mathbb{G}}$ and an SRS $\text{srs} := (\text{ek}, \text{vk})$ where $\text{ek} = [1, s, \dots, s^d]_1$ and $\text{vk} = [1, s]_2$; \mathcal{B} samples $\alpha \leftarrow \mathbb{Z}_q$ and computes $\text{ek}_{\alpha} \leftarrow [\alpha, \alpha s, \dots, \alpha s^d]_1$ and forwards to \mathcal{A} the SRS $\text{srs}' := ((\text{ek}, \text{ek}_{\alpha}), \text{vk})$. Finally, \mathcal{B} parses the list of commitments coms as two lists of equal size $(\text{coms}_1, \text{coms}_2)$ and gives \mathcal{A} the commitments $\text{coms}' = \text{coms}_1 + \alpha \text{coms}_2$. Every time \mathcal{A} comes up with some group element \mathbf{c} , attaches a representation vector $\mathbf{r}_{\mathbf{c}}$ that we can parse as $\mathbf{f}_{\mathbf{c}} \parallel \mathbf{v}_{\mathbf{c}} \parallel \mathbf{c}_{\mathbf{c}} \parallel \mathbf{o}_{\mathbf{c}}$ where $\mathbf{f}_{\mathbf{c}}$ (resp. $\mathbf{v}_{\mathbf{c}}$) is the vector of coefficients associated to group elements ek (resp. ek_{α}), $\mathbf{c}_{\mathbf{c}}$ is the vector of coefficients associated to group elements coms' , and $\mathbf{o}_{\mathbf{c}}$ is the vector of coefficients associated to the group elements of the simulated proofs proofs' . The strategy of \mathcal{B} is to forward the simulation queries of \mathcal{A} whose representation vectors are given w.r.t. srs' , attaching the representation vectors w.r.t. srs .

On input a simulation query of the form $\mathbf{x} = (\mathbf{c}, x_j, y)$ and leakage y' , with $\mathbf{r}_{\mathbf{c}} = \mathbf{f}_{\mathbf{c}} \parallel \mathbf{v}_{\mathbf{c}} \parallel \mathbf{c}_{\mathbf{c}}$ as representation vector for \mathbf{c} , \mathcal{B} invokes the simulation oracle \mathcal{S}_1 first on input $\mathbf{x}_1 = (\langle \mathbf{f}_{\mathbf{c}} + \alpha \mathbf{v}_{\mathbf{c}}, \text{ek} \rangle + \langle \mathbf{c}_{\mathbf{c}} \parallel \mathbf{0}, \text{coms} \rangle, x_j, y)$ and then on $\mathbf{x}_2 = (\langle \mathbf{0} \parallel \mathbf{c}_{\mathbf{c}}, \text{coms} \rangle, x_j, y')$, obtaining the two proofs π_1 and π_2 . Finally returns to \mathcal{A} the proof $\pi := \pi_1 + \alpha \pi_2$ which satisfies the verification equation:

$$e(\mathbf{c} - [y]_1 - [\alpha y']_1, [1]_2) = e(\pi, [s - x_j]_2)$$

and adds π to proofs' . When \mathcal{A} makes the forgery $(\mathbf{c}^*, x^*, y^*, y'^*, \pi^*)$, with associated representation vectors $\mathbf{r}'_{\tau} = \mathbf{f}_{\tau} \parallel \mathbf{v}_{\tau} \parallel \mathbf{c}_{\tau} \parallel \mathbf{o}_{\tau}$, for $\tau \in \{\mathbf{c}^*, \pi^*\}$, \mathcal{B} forwards the forgery $(\mathbf{c}, x^*, y^* + \alpha y'^*, \pi^*)$. \mathcal{B} needs to attach representation vectors for the group elements \mathbf{c}^*, π^* w.r.t. srs, coms and proofs ; to do that, \mathcal{B} applies the following transformation and computes $\mathbf{r}_{\tau} = \mathbf{f}_{\tau} + \alpha \mathbf{v}_{\tau} \parallel \mathbf{c}_{\tau} \parallel \alpha \mathbf{c}_{\tau} \parallel (\mathbf{o}_{\tau, i} \parallel \alpha \mathbf{o}_{\tau, i})_i$. This choice, indeed, is such that:

$$\langle \mathbf{r}'_{\tau}, (\text{ek} \parallel \text{ek}_{\alpha} \parallel \text{coms}' \parallel \text{proofs}') \rangle = \langle \mathbf{r}_{\tau}, (\text{ek} \parallel \text{coms} \parallel \text{proofs}) \rangle$$

If the forgery of \mathcal{A} is valid, so is the forgery of \mathcal{B} and, unless with negligible probability, the canonical extractor \mathcal{E} extracts the valid witness $f(X) + \alpha v(X)$.

5 Simulation-Extractable Universal zkSNARKs

We extend the notion of relations to indexed relations [20]. We define a PT indexed relation \mathcal{R} verifying tuple $(\text{pp}, \mathbf{i}, \mathbf{x}, \mathbf{w})$. We say that \mathbf{w} is a witness to the instance \mathbf{x} being in the relation defined by the pp and index \mathbf{i} when $(\text{pp}, \mathbf{i}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}$ (equivalently, we sometimes write $\mathcal{R}(\text{pp}, \mathbf{i}, \mathbf{x}, \mathbf{w}) = 1$).

Definition 17 (Universal NIZKs). A universal NIZK for an indexed relation \mathcal{R} is a tuple of algorithms $\Pi = (\text{KGen}, \text{Derive}, \text{Prove}, \text{Verify})$ where

- Derive is a deterministic algorithm that takes as input an srs (which includes relation parameters pp) produced by KGen , an index \mathbf{i} , and outputs specialized SRS $\text{srs}_{\mathbf{i}} = (\text{ek}_{\mathbf{i}}, \text{vk}_{\mathbf{i}})$.
- Consider the relation \mathcal{R}' such that $\mathcal{R}'(\text{pp}, (\mathbf{i}, \mathbf{x}), \mathbf{w}) \iff \mathcal{R}(\text{pp}, \mathbf{i}, \mathbf{x}, \mathbf{w})$, the tuple of algorithms $(\text{KGen}, \text{Prove}, \text{Verify}')$ is a NIZK for the relation \mathcal{R}' and Verify' is the algorithm that upon input srs, instance (\mathbf{i}, \mathbf{x}) and a proof π , first runs Derive on srs and index \mathbf{i} , then runs $\text{Verify}(\text{vk}_{\mathbf{i}}, \mathbf{x}, \pi)$.

Moreover, we say Π is a universal SNARK if the length of $\text{vk}_{\mathbf{i}}$ is $\text{poly}(\lambda, \log |\mathbf{i}|)$, the proof size is $\text{poly}(\lambda, \log |\mathbf{w}|)$.

5.1 Polynomial Holographic Interactive Oracle Proofs

Definition 18 (Polynomial Holographic IOP). Let \mathcal{F} be a family of finite fields and let \mathcal{R} be an indexed relation. A (public-coin non-adaptive) Polynomial Holographic IOP over \mathcal{F} for \mathcal{R} is a tuple $\text{PIOP} = (r, n, m, D, n_e, \text{RE}, \mathcal{P}, \mathcal{V})$ where $r, n, m, D, n_e: \{0, 1\}^* \rightarrow \mathbb{N}$ are polynomial-time computable functions, and $\text{RE}, \mathcal{P}, \mathcal{V}$ are three algorithms for the encoder, prover and verifier respectively, that work as follows.

Offline phase: The encoder $\text{RE}(\mathbb{F}, \mathbf{i})$ is executed on input a field $\mathbb{F} \in \mathcal{F}$ and a relation description \mathbf{i} , and it returns $n(0)$ polynomials $\{p_{0,j}\}_{j \in [n(0)]}$ encoding the relation \mathbf{i} .

Online phase: The prover $\mathcal{P}(\mathbb{F}, \mathbf{i}, \mathbf{x}, \mathbf{w})$ and the verifier $\mathcal{V}^{\text{RE}(\mathbb{F}, \mathbf{i})}(\mathbb{F}, \mathbf{x})$ are executed for $r(|\mathbf{i}|)$ rounds; the prover has a tuple $(\mathbb{F}, \mathbf{i}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}$ and the verifier has an instance \mathbf{x} and oracle access to the polynomials encoding \mathbf{i} .

In the i -th round, \mathcal{P} sends $m(i)$ messages $\{\pi_{i,j} \in \mathbb{F}\}_{j \in [m(i)]}$, and $n(i)$ oracle polynomials $\{p_{i,j} \in \mathbb{F}[X]\}_{j \in [n(i)]}$ of degree at most $D := D(|\mathbf{i}|)$, while \mathcal{V} replies (except for the last round) with a uniformly random message $\rho_i \in \mathbb{F}$.

Decision phase: After the $r := r(|\mathbf{i}|)$ -th round, let $n_e := n_e(|\mathbf{i}|)$, the verifier $\mathcal{V}(\mathbb{F}, \mathbf{x}, \boldsymbol{\rho})$, on input the description of the field \mathbb{F} , the input \mathbf{x} and all the random messages of the verifier $\boldsymbol{\rho} := (\rho_1, \dots, \rho_{r-1})$, outputs tuples $(G^{(k)}, v_1^{(k)}, \dots, v_n^{(k)})_{k \in [n_e]}$ which define the following algebraic checks. Let $n := \sum_{k=0}^r n(k)$, let $m := \sum_{k=1}^r m(k)$, and denote by (p_1, \dots, p_n) all the oracle polynomials (including the $n(0)$ ones from the encoder) and by (π_1, \dots, π_m) all the messages sent by the prover. For any $k \in [n_e]$ we have $G^{(k)} \in \mathbb{F}[X, X_1, \dots, X_n, Y_1, \dots, Y_m]$ and $v_j^{(k)} \in \mathbb{F}[X]$. Then a tuple $(G^{(k)}, v_1^{(k)}, \dots, v_n^{(k)})$ is satisfied if and only if $F^{(k)}(X) \equiv 0$ where

$$F^{(k)}(X) := G^{(k)}(X, \{p_i(v_i^{(k)}(X))\}_{i \in [n]}, \{\pi_i\}_{i \in [m]})$$

The verifier accepts if and only if all the checks are satisfied.

A list $\mathcal{L} = \{(i_1, y_1), \dots\}$ is (\mathbf{b}, C) -bounded where $\mathbf{b} \in \mathbb{N}^n$ and C is a PT algorithm if $\forall i \in [n] : |\{(i, y) : (i, y) \in \mathcal{L}\}| \leq \mathbf{b}_i$ and $\forall (i, y) \in \mathcal{L} : C(i, y) = 1$.

Definition 19 (b-bounded Zero-Knowledge). A *PIOP* PIOP is **b-Zero-Knowledge** if there exists a checker C such that $\Pr[C(i, x) = 0] \leq \text{negl}(\lambda)(|\mathbb{F}|)$ over random x such that for every index i , and $(\mathbf{pp}, \mathbf{i}, \mathbf{x}, \mathbf{w}) \in \mathcal{R}$, and every (\mathbf{b}, C) -bounded list \mathcal{L} , the following random variables are within ϵ statistical distance: $(\text{view}(\mathcal{P}(\mathbb{F}, \mathbf{i}, \mathbf{x}, \mathbf{w}), \mathcal{V}^{\text{RE}(\mathbb{F}, \mathbf{i})}(\mathbb{F}, \mathbf{x})), (p_i(y))_{(i, y) \in \mathcal{L}}) \approx_\epsilon \mathcal{S}(\mathbb{F}, \mathbf{i}, \mathbf{x}, \mathcal{L})$. where p_1, \dots, p_n are the polynomials returned by the prover \mathcal{P} and \mathcal{S} is a PPT simulator. Moreover, PIOP is special honest-verifier **b-bounded zero-knowledge** if \mathcal{S} first samples uniformly at random the verifier's messages $\rho_1, \dots, \rho_{r-1}$ and then computes the full transcript. Namely, \mathcal{S} can take as additional input the verifier's messages. PIOP is independent leakage if \mathcal{S} can be divided in two algorithms $(\mathcal{S}_0, \mathcal{S}_1)$ where \mathcal{S}_0 outputs the simulated transcript, and \mathcal{S}_1 the leakage. Namely $\mathcal{S}(\mathbb{F}, \mathbf{i}, \mathbf{x}, \mathcal{L}; r) = \mathcal{S}_0(\mathbb{F}, \mathbf{i}, \mathbf{x}; r), \mathcal{S}_1(\mathbb{F}, \mathbf{i}, \mathbf{x}, \mathcal{L}; r)$.

Remark 2. Our notion of PIOP is slightly different from that of [17]: first, we consider a verifier that checks polynomial equations between the oracle polynomials but we do not explicitly check the degree of the polynomials; second, we let the prover send the first message. This is mainly for simplicity in the presentation of the compiler since, given a PIOP according to the definition of [17], we can define a PIOP in our model: \mathcal{P} sends an extra (initial) message which is ignored, and for each poly p_i where we check degree d_i the PIOP additionally sends $p'_i(X) = p(X) \cdot X^{D-d_i}$ where D is the maximum degree, and \mathcal{V} additionally checks the equation $\sum \rho^i \cdot (X^{D-d_i} p_i(X) - p'_i(X)) \equiv 0$, for a randomizer ρ .

Definition 20. A *PIOP* PIOP is (almost) instance-independent verifier queries if there exists an alternative deterministic PT algorithm $\tilde{\mathcal{V}}$ such for any \mathbf{i} and \mathbf{x} we have $(\tilde{v}_1^{(j)}, \dots, \tilde{v}_n^{(j)})_{j \in [n_e]} \leftarrow \tilde{\mathcal{V}}(\mathbb{F}, \mathbf{i})$ where for any $j \in [n]$ and $k \in [n_e]$ we have $\tilde{v}_j^{(k)} \in \mathbb{F}[X, X_i, \dots, X_{r-1}]$. Moreover, for any field elements $\rho_1, \dots, \rho_{r-1}$ let $(G^{(k)}, v_1^{(j)}, \dots, v_n^{(j)})_{j \in [n_e]} \leftarrow \mathcal{V}(\mathbb{F}, \mathbf{i}, \mathbf{x}, \rho_1, \dots, \rho_{r-1})$ we have for any $i \in [r(|\mathbf{i}|) - 1], j \in [n_{i-1}, n_i], k \in [n_e]$:

$$\tilde{v}_j^{(k)}(X, \rho_i, \dots, \rho_{r-1}) = v_j^{(k)}(X).$$

State restoration. We introduce state restoration attacks on PIOPs , extending the notion of state-restoration for IOPs introduced in [9]. Informally, a state-restoring prover receives in each round a verifier message ρ_i as well as a complete verifier state cvs , and then sends to the verifier $m(i)$ messages $\pi_{i,j}$ and $n(i)$ oracle polynomials $p_{i,j}$, and a previously-seen complete verifier state which sets the verifier to that state; this forms a state restoration attack on the verifier. The interaction between a state-restoring prover $\tilde{\mathcal{P}}$ and the verifier \mathcal{V} is mediated through the experiment described below. The experiment considers a straight-line extractor \mathcal{E} run after the interaction of the state-restoring prover and verifier. We consider the experiment $\mathbf{Exp}_{\tilde{\mathcal{P}}, \text{PIOP}, \mathcal{E}}^{sr}(\mathbb{F})$:

1. The challenger initializes the list `SeenStates` to be empty.
2. Repeat the following until the challenger halts:

- (a) $\tilde{\mathcal{P}}$ either (1) chooses a complete verifier state cvs in the list SeenStates or (2) sends a new fresh tuple $(\mathbf{i}, \mathbf{x}, \{\pi_{1,j}\}_j, \{p_{1,j}\}_j)$ to the challenger.
 - (b) If (1) the challenger sets the verifier to cvs :
 - i. if $\text{cvs} = (\mathbf{i}, \mathbf{x}, \{\pi_{1,j}\}_j, \{p_{1,j}\}_j \parallel \rho_1 \parallel \dots \parallel \{\pi_{i,j}\}_j, \{p_{i,j}\}_j)$ and $i < r(\mathbf{x})$: $\tilde{\mathcal{P}}$ outputs $\{\pi_{i-1,j}\}_j, \{p_{i-1,j}\}_j$; \mathcal{V} samples ρ_i and sends it to $\tilde{\mathcal{P}}$; the game appends $\text{cvs}' := (\text{cvs} \parallel \{\pi_{i-1,j}\}_j \parallel \{p_{i-1,j}\}_j \parallel \rho_i)$ to the list SeenStates ;
 - ii. if $\text{cvs} = (\mathbf{i}, \mathbf{x}, \{\pi_{1,j}\}_j, \{p_{1,j}\}_j \parallel \rho_1 \parallel \dots \parallel \rho_{r-1})$: $\tilde{\mathcal{P}}$ outputs $\{\pi_{r,j}\}_j, \{p_{r,j}\}_j$; the challenger runs $\text{RE}(\mathbb{F}, \mathbf{i})$ and \mathcal{V} performs the decision phase of the PIOP. The challenger sets cvs to be the final cvs , sets the decision bit d as the output of the verifier \mathcal{V} and halts.
 - (c) If (2) the verifier samples ρ_1 and sends it to $\tilde{\mathcal{P}}$; the game appends $\text{cvs}' = (\mathbf{i}, \mathbf{x}, \{\pi_{1,j}\}_j, \{p_{1,j}\}_j \parallel \rho_1)$ to the list SeenStates .
3. The game computes the extraction bit $b \stackrel{\text{def}}{=} (\mathbf{i}, \mathbf{x}, \mathcal{E}(\mathbf{i}, \mathbf{x}, p_1, \dots, p_n)) \in \mathcal{R}$ where the instance \mathbf{x} and the polynomials p_1, \dots, p_n are the ones generated by RE and the ones included in the final cvs . The game returns $(d \wedge \neg b)$ (namely, the malicious prover convinces the verifier but the extractor fails).

Definition 21 (State-restoration (straight-line) proof of knowledge).

A PIOP PIOP is state-restoration (straight-line) proof of knowledge if there exists an extractor \mathcal{E} such that for any $\tilde{\mathcal{P}}$ and any \mathbb{F} :

$$\Pr \left[\mathbf{Exp}_{\tilde{\mathcal{P}}, \text{PIOP}, \mathcal{E}}^{sr}(\mathbb{F}) = 1 \right] \leq \text{negl}(|\mathbb{F}|)$$

We consider PIOPs with the following structural condition.

Definition 22. A PIOP PIOP is compiler-safe if for any \mathbf{i}, \mathbf{x} and $\boldsymbol{\rho} := \rho_1, \dots, \rho_{r-1}$ and any tuple $(G^{(k)}, v_1^{(j)}, \dots, v_n^{(j)})_{j \in [n_e]} \leftarrow \mathcal{V}(\mathbb{F}, \mathbf{i}, \mathbf{x}, \boldsymbol{\rho})$ there exists an index k such that for all j the polynomials $v_j^{(k)}$ are of degree at least one.

This is the formalization of the technical condition mentioned in the introduction that allows us to “bundle together” all the oracles and solve the first challenge.

We notice that some PIOPs have this technical condition (an example is [26]) while other PIOPs that internally run different sub-protocols might not have this property. However, at PIOP level we can always obtain this property by adding an extra degree of randomness to the polynomials (obtaining $\mathbf{b} + \mathbf{1}$ -bounded zero-knowledge) and adding an extra trivial equation over the polynomials with the v_j set to be the identity function.

Commitments Simulator. We abstract how our compiler handles the simulation of the commitments for the oracles sent during a PIOP protocol’s execution.

Definition 23 (Commitment Simulator for PIOP). Let PIOP be a PIOP with \mathbf{b} -Zero-Knowledge simulator \mathcal{S} and checker \mathcal{C} and consider a two-stage PPT algorithm where $\mathcal{S}_{\text{Com}}(0, \text{ck})$ outputs a vector coms of commitments, and $\mathcal{S}_{\text{Com}}(1, \mathbf{i}, \mathbf{x})$ outputs a matrix $\mathbf{M}_{\mathbf{i}, \mathbf{x}}$ of linearly independent rows. \mathcal{S}_{Com} is a commitments simulator for a CP (defined over the same commitment scheme) and

PIOP, if for every (\mathbf{b}, C) -bounded list \mathcal{L} , \mathbb{F}, \mathfrak{i} and \mathfrak{x} the following distributions are computationally indistinguishable:

$$\begin{aligned} & (\text{ck}, \text{view}(\mathcal{P}(\mathbb{F}, \mathfrak{i}, \mathfrak{x}, \mathfrak{w}), \mathcal{V}), (p_j(x))_{(j,x) \in \mathcal{L}}, (\text{Com}(\text{ck}, p_i))_{i \in [n]}) \approx_c \\ & (\text{ck}, \mathcal{S}(\mathbb{F}, \mathfrak{i}, \mathfrak{x}, \mathcal{L}), (\mathbf{M}_{\mathfrak{i}, \mathfrak{x}} \cdot \text{coms} \parallel \text{ck})) \end{aligned}$$

Moreover, we have that the view $(\mathcal{L}, \mathcal{S}(\mathbb{F}, \mathfrak{i}, \mathfrak{x}, \mathcal{L}), \mathfrak{x}, \mathbf{M}_{\mathfrak{i}, \mathfrak{x}})$ satisfies the algebraic consistency for the CP-SNARK CP (as in Definition 14).

The definition is in two stages so to fit our result of Theorem 1. For KZG the first part handles the generation of instance-independent commitments according to a \mathcal{D}_k -Aff-MDH Assumption, while, the second part acts over the formerly sampled commitments adapting them to the instance at hand. If the commitment scheme is hiding then the task of defining the commitment simulator \mathcal{S}_{Com} for an arbitrary PIOP is trivial (the distribution of $\mathcal{S}_{\text{Com}}(0, \text{ck})$ is uniformly over the commitment space), thus this property is interesting for the case of non-hiding commitments. Similarly, this notion can be applied to commitment schemes that are not homomorphic, in that case, we must require $\mathbf{M}_{\mathfrak{i}, \mathfrak{x}}$ to be the identity.

5.2 The Compilation-Ready CP-SNARK

Consider the relation $\mathcal{R}_{\text{poly}}$ that upon relation parameters ck and an instance $(\text{msg}, \mathbf{c}, (\bar{G}^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$ where msg is an arbitrary string and with $\mathbf{c} = (c_j)_{j \in [n]}$ and witness \mathbf{p}, \mathbf{o} outputs 1 if and only if:

$$\forall k \in [n_e] : \bar{G}^{(k)}(X, p_1(v_1^{(k)}(X)), \dots, p_n(v_n^{(k)}(X))) \equiv 0 \quad \forall j : \text{VerCom}(\text{ck}, c_j, p_j, o_j) = 1$$

Notice that a simulation-extractable CP-SNARK for the relation $\mathcal{R}_{\text{poly}}$ forms *signature of knowledge* [19] for the message msg and an instance $(\mathbf{c}, (\bar{G}^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$. Consider the relation for multi-instance opening of commitments defined below:

$$\mathcal{R}_{\text{opn}}(\text{ck}, \mathbf{c}, (\mathbf{p}, \mathbf{o})) := (\forall j \in [|\mathbf{c}|] : \text{VerCom}(\text{ck}, c_j, p_j, o_j)).$$

The second ingredient for our compiler is a CP-SNARK that can simultaneously prove the polynomial evaluation relation and the knowledge of the openings. We consider the *join* of the two relations $\mathcal{R}_{\text{poly}}$ and \mathcal{R}_{opn} :

$$\hat{\mathcal{R}} \stackrel{\text{def}}{=} \{\text{ck}, (\text{poly}, \mathfrak{x}), \mathfrak{w} : \mathcal{R}_{\text{poly}}(\text{ck}, \mathfrak{x}, \mathfrak{w})\} \cup \{\text{ck}, (\text{opn}, \mathbf{c}), \mathfrak{w} : \mathcal{R}_{\text{opn}}(\text{ck}, \mathbf{c}, \mathfrak{w})\} \quad (11)$$

Notice that if we have a CP-SNARK for $\mathcal{R}_{\text{poly}}$ and a CP-SNARK for \mathcal{R}_{opn} we can easily define a single proof system that proves the relation $\hat{\mathcal{R}}$. In fact, the relation $\hat{\mathcal{R}}$ could be seen as a join of the two relations plus some syntactic sugar.

The reason to use a single CP-SNARK for $\hat{\mathcal{R}}$ (instead of one for each relation) is rather technical and it has to do with the security guarantees when we compose protocols in the AGM [1]. In fact, we need a CP-SNARK for $\mathcal{R}_{\text{poly}}$ that is simulation extractable in the AGM even in presence of the simulated proofs for a CP-SNARK for the \mathcal{R}_{opn} . In particular, the simulated proofs for the latter

CP-SNARK could contain group elements in \mathbb{G}_1 that might interfere with the security proved for the former CP-SNARK (and vice versa). In other words, even if the first CP-SNARK is simulation extractable (in the AGM), it could potentially be insecure when we use it in combination with the second CP-SNARK (if we are using the same group to instantiate the proof of opening).

We consider a set of policies $\hat{\Phi}$ where $\hat{\Phi} = (\hat{\Phi}_0, \hat{\Phi}_1) \in \hat{\Phi}$ such that $\hat{\Phi}_0(\text{pp}_{\mathbb{G}})$ outputs parameters $\text{pp}_{\hat{\Phi}}$ that contain an arbitrary set of vectors of polynomials $\mathcal{Q}_v = \{\mathbf{v}^{(i)} = (v_1^{(i)}, \dots, v_n^{(i)})\}_{i \in \text{poly}(\lambda)}$ and a list of commitments sampled $\text{coms} = (\mathbf{c}^{(i)})_{i \in [g]}$ from a sampler \mathcal{D} where the \mathcal{D} -Aff-MDH assumption holds, while $\hat{\Phi}_1$ is such that:

- **Semi-adaptive w.r.t. poly-queries.** Given the set of simulation queries \mathcal{Q}_{sim} , define the projections of the set to the **poly**-simulation queries and **opn**-simulation queries. Let the i -th **poly**-query to the simulation oracle be $(\mathbb{x}, \text{aux}, \pi) \in \mathcal{Q}_{\text{sim}}$, parse \mathbb{x} as $(\text{msg}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$. (1) $\forall k : \mathbf{v}^{(k)} \in \mathcal{Q}_v$ and (2) parse aux two matrices \mathbf{C}, \mathbf{F} the tuple $(\mathbf{c}, \mathbf{C}, \mathbf{F})$ satisfies the commitment check, namely $\mathbf{c} = \mathbf{C} \cdot \mathbf{c}^{(i)} + \mathbf{F} \cdot \text{ck}$. (3) the view defined by set of all **poly**-queries satisfies the algebraic consistency for CP (cf. Definition 14).
- **Extractor policy for opn-forgery.** If the forgery of the adversary (\mathbb{x}^*, π^*) is of the form $\mathbb{x}^* := (\text{opn}, \mathbf{c}^*)$, parse $\text{aux}_{\hat{\Phi}}$ as (yet another) forgery $\tilde{\mathbb{x}} := (\text{poly}, \text{msg}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{i \in [n_e]})$, $\tilde{\pi}$ and check that:
 1. All the commitments \mathbf{c}^* are in the vector of commitments \mathbf{c} ,
 2. The proof $\tilde{\pi}$ is valid according to the extractor policy for **poly**-forgeries below.
- **Extractor policy for poly-forgery.** If the forgery of the adversary (\mathbb{x}^*, π^*) is of the form $\mathbb{x}^* := (\text{poly}, \mathbb{x}')$, then check that \mathbb{x}^* was never queried to the simulation oracle and there exists a $k \in [n_e]$ such that for all j the polynomial $v_j^{(k)}$ has degree at least 1 (and degree $\text{poly}(\lambda)$).

Intuitively, the extractor policy for **opn**-forgery means that a proof of opening for a commitment \mathbf{c} can be extracted (given some auxiliary information) only when the adversary can exhibit a proof of evaluation that involves such a commitment.

In Section 6.1 we describe a simple and unoptimized CP-SNARK for $\mathcal{R}_{\text{poly}}$ for KZG, in Section 6.2 we show how to extend it to $\hat{\mathcal{R}}$. The security analysis is given in the AGM. The CP-SNARK uses classical random-point evaluation for the polynomial equations and a *vacuous proof of opening* that can be extracted given the algebraic representation. The extraction of the latter is successful, namely the representation depends only on the element in the commitment key, only if the adversary exhibits a proof of polynomial equation. Both the proof of polynomial equation and of opening are extracted through the algebraic representations. Thus, unless the binding property is broken, when the proof of evaluation holds the proof of opening can be extracted correctly.

5.3 The Universal zkSNARK

Let Φ_{SE} be the standard (strong) simulation extractability policy. Namely, the policy checks that the forgery of the adversary is a tuple $(\mathbb{x}_{\Pi}, \pi_{\Pi}) \notin \mathcal{Q}_{\text{sim}}$.

```

Π.Derive(srs, i) :
-----
p0 ← RE(℔, i);
for j ∈ [n(|i|, 0)] do c0,j ← Com(ck, p0,j; 0) // Deterministic commitments
return (eki, vki) where eki ← p0, vki ← (c0,i)i ∈ [n(|i|, 0)]

Π.Prove(srs, eki, x, w) :
-----
for i ∈ [r(|i|)] do :
  (pi, πi) ← P(℔, i, x, w, ρ1, ..., ρi-1) // Get polynomials and messages from PIOP prover
  for j ∈ [n(i)] do (ci,j, oi,j) ← Com(ck, pi,j)
  xo,i ← (opn, ci), ci := (ci,j)j, oi := (oi,j)j, πopn,i ← CP.Prove(ck, xo,i, (pi, oi)), π̄i := (ci, πopn,i, πi)
  ρi ← RO(vki, x, π̄1, ..., π̄i) // Fiat-Shamir transform
  (G(k), v(k))k ∈ [ne] ← V(℔, x, ρ1, ..., ρr-1)
  for k ∈ [ne] do : G(k)(X, X1, ..., Xn) ← G(k)(X, X1, ..., Xn, π)
  trns ← (vki, x, π̄1, ..., π̄r), xpoly ← (poly, trns, c, (G(k), v(k))k ∈ [ne]), πpoly ← CP.Prove(srs, xpoly, (p, o))
  return (c, π, (πopn,i)i ∈ [r], πpoly)

Π.Verify(vki, x, ππ) :
-----
compute (π̄1, ..., π̄r)
for i ∈ [r(|i|) - 1] do : ρi ← RO(vki, x, π̄1, ..., π̄i-1) // Fiat-Shamir transform
(G(k), v(k))k ∈ [ne] ← V(℔, x, ρ1, ..., ρr-1)
for k ∈ [ne] do : G(k)(X, X1, ..., Xn) ← G(k)(X, X1, ..., Xn, π)
trns ← (vki, x, π̄1, ..., π̄r), xpoly ← (poly, trns, c, (G(k), v(k))k ∈ [ne])
return ⋀i ∈ [r] CP.Verify(ck, xopn,i, πopn,i) ∧ CP.Verify(srs, xpoly, πpoly)

```

Fig. 2. The Compiler to Universal zkSNARKs.

Theorem 3. *Let CP be a compilation-ready CP-SNARK, namely a $\hat{\Phi}$ -simulation-extractable CP-SNARK for the joint relation $\hat{\mathcal{R}}$ (cf. Eq. (11)). Let PIOP be a PIOP for an indexed relation \mathcal{R} that is state-restoration straight-line extractable (cf. Definition 21), and bounded special honest-verifier zero-knowledge, where the technical condition of Definition 22 holds and equipped with a commitment simulator for CP (cf. Definition 23). Let Π be the zkSNARK compiled from PIOP using the compiler from Fig. 2. Then Π is Φ_{SE} -simulation-extractable.*

The proof relies mostly on the simulation extractability of compiler-ready scheme CP. By inspection of the policy set $\hat{\Phi}$, we notice that a technical condition similar to Definition 22 is already required to prove the simulation extractability of the compilation-ready scheme. In fact, without the technical condition, the mix-and-match malleability attack mentioned in the introduction would already apply for two disjoint sets of equations for CP over disjoint sets of committed polynomials. Given an adversary \mathcal{A} for Π , we can derive a malicious prover $\hat{\mathcal{P}}$ for the state-restoration game of the PIOP. The idea is to translate all the RO queries of \mathcal{A} to messages of the malicious prover. The prover additionally needs to send polynomial oracles, we extract them using the extractor for the proof of opening. Notice this extractor is quite weak, we have an assurance that its extractions are valid only if the adversary \mathcal{A} later can provide a proof of

polynomial evaluation. Luckily for us, the adversary must provide the latter proof if it wants to win its security game. Notice, this avoids rewinding, since the extraction can be performed at the same moment when the adversary sends the proof of opening (through a RO call).

The bounded special honest-verifier zero-knowledge property, together with the property of the commitment simulator for the PIOP (see Definition 23), guarantee that the simulated proofs in our reduction to the security of CP match the simulation policy of the latter.

Proof. We prove the theorem assuming the commitments are not hiding. The adaption to hiding commitments is straightforward. We start showing the zero-knowledge simulator for Π . The simulator is in Fig. 3. For the description of the simulator, we assume we can couple the leakage function F_{poly} , which defines the F_{poly} -leaky zero-knowledge of the CP-SNARK, with a function \tilde{F}_{poly} that upon input the instance outputs the set $\mathcal{L}_{\mathbf{x}}$ of evaluation points necessary to compute a simulated proof. The zero-knowledge guarantees of the simulator come from the zero-knowledge property of the PIOP, the zero-knowledge property of the CP and the simulator commitment property of \mathcal{S}_{Com} (see Definition 23). Notice that \mathcal{S}_1 might abort if $(\text{vk}_i, \mathbf{x}, \bar{\pi}_1, \dots, \bar{\pi}_i)$ for some i was already queried to the random oracle. However, by simply assuming that the first message of \mathcal{P} has $\omega(\log |\mathbb{F}|)$ min-entropy then the event that the simulator aborts happens with negligible probability.

We define the extractor for Π for a given adversary \mathcal{A}_Π . We make some simplifying assumptions on the behavior of \mathcal{A}_Π : (1) the adversary always queries first the RO on a string that can be parsed as (\mathbf{i}, \mathbf{x}) before querying the simulation oracle on the same string, (2) the auxiliary string $\text{aux}_{\mathcal{E}}$ output by \mathcal{A}_Π can be parsed as a list of strings $(s_i, \text{aux}_i, \text{st}_i)_i$ and a string $\text{aux}'_{\mathcal{E}}$ where for any i we have $(s_i, \text{aux}_i, \text{st}_i)$ string is identical to the auxiliary input output at the i -th query of the adversary and (3) for any i the auxiliary input of the i -th random oracle query of the adversary contains the full internal state of the adversary. These assumptions are w.l.g. In fact, given an adversary \mathcal{A}_Π that does not respect these rules we can always define another adversary that runs internally \mathcal{A}_Π , collects all the necessary information to comply with (2) and (3), and moreover follows the rule (1).

Let \mathcal{B}_i be a reduction to the simulation extractability of CP that runs \mathcal{A}_Π (simulating the oracles for \mathcal{A}_Π using its own oracles and the code defined in Fig. 3) and sets its output to be the i -th random oracle query of \mathcal{A}_Π . For any $i \in [q]$ let \mathcal{E}_i be the extractor associated to \mathcal{B}_i (thanks to the $\hat{\Phi}$ -simulation extractability of CP we can associate to \mathcal{B}_i an extractor, in Lemma 11 we describe \mathcal{B}_i with more details and we show that it complies with the $\hat{\Phi}$ policy). Similarly, let $\mathcal{E}_{\text{poly}}$ be the extractor for the adversary $\mathcal{B}_{\text{poly}}$ that runs \mathcal{A}_Π (simulating the oracles of \mathcal{A}_Π using its own oracles and the code defined in Fig. 3) until completion, and isolates the instance $\mathbf{x}_{\text{poly}} := (\text{poly}, \text{trns}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$ and the proof π_{poly} in the forgery of \mathcal{A}_Π , and outputs all the auxiliary outputs that \mathcal{A}_Π does.

We first set some notation:

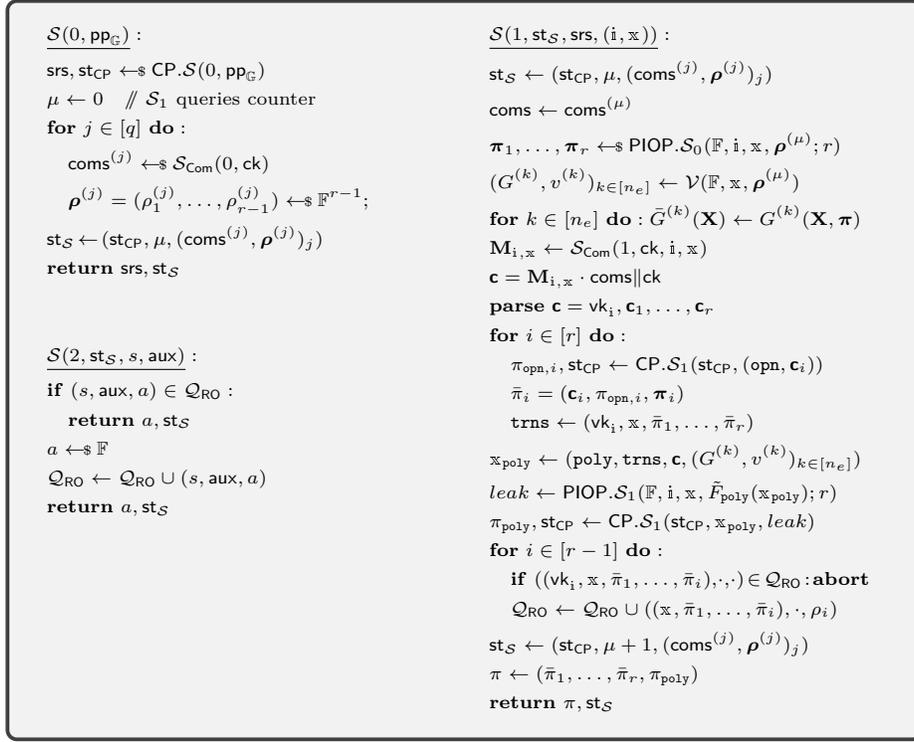


Fig. 3. The simulator \mathcal{S} for Π .

- Let \mathcal{P}_i be the indexes of the polynomials sent at the i -th round by the PIOP.
- Given a proof π for Π we define *the RO-queries set of π* be the set of string $\{(\text{vk}_i, \mathbf{x}), \dots, (\text{vk}_i, \mathbf{x}, \bar{\pi}_1, \dots, \bar{\pi}_r)\}$.
- We say that a proof π *shares a simulated commitment* \mathbf{c} of proof π' simulated by \mathcal{S} if in st_S one can find $\mathbf{m}_1 \neq \mathbf{0}$ and \mathbf{m}_2 such that $\mathbf{c} = \mathbf{m}_1 \parallel \mathbf{m}_2 \cdot \text{coms} \parallel \text{ck}$.

The extractor $\mathcal{E}_\Pi(\mathbf{x}_\Pi, \pi_\Pi, \text{aux}_\mathcal{E})$:

1. Parse $\text{aux}_\mathcal{E}$ as the concatenation of a list $(s_i, \text{aux}_i, \text{st}_i)_i$ and $\text{aux}'_\mathcal{E}$, where $(s_i, \text{aux}_i, \text{st}_i)$ is the output of \mathcal{A}_Π at the i -th query to the ROM and $\text{aux}'_\mathcal{E}$ the remaining auxiliary information given by the adversary (namely, the auxiliary information associated with its last output).
2. From π_Π derive the messages $\bar{\pi}_1, \dots, \bar{\pi}_r$ and find the indexes q_i such that $s_{q_i} = (\text{vk}_i, \mathbf{x}, \bar{\pi}_1, \dots, \bar{\pi}_i)$.
3. If π_Π shares a simulated commitment with one of the simulated proofs then return \perp .
4. For $i \in [r]$ run $\mathbf{w}_{i, \text{opn}} \leftarrow \mathcal{E}_{q_i}(\text{srs}, \mathbf{c}_i, \pi_{\text{opn}, i}, \text{aux}_i)$ where $\bar{\pi}_i$ contains both the instance and the proof of opening and $\mathbf{w}_{i, \text{opn}} = (p'_j)_{j \in \mathcal{P}_i}$.
5. Run $(p_j)_{j \in [n]} \leftarrow \mathcal{E}_{\text{poly}}(\text{srs}, \mathbf{x}_{\text{poly}}, \pi_{\text{poly}}, \text{aux}'_\mathcal{E})$.
6. If $\exists i : p'_i \neq p_i$ then return \perp .

7. If $\exists i \in [n]$ and $j \in \mathcal{P}_i : \text{VerCom}(\text{ck}, \mathbf{c}_j, p'_j) \neq 1$ then return \perp .
8. If $\exists k : G^{(k)}(X, p_1(v_1^{(k)}(X)), \dots, p_n(v_n^{(k)}(X)), \pi_1, \dots, \pi_r) \neq 0$ then return \perp .
9. Return $\mathcal{E}_{\text{PIOP}}(\mathbf{i}, \mathbf{x}_\Pi, (p_j)_{j \in [n]})$.

To analyze the success of the extractor we define a series of hybrid games. We start from the first hybrid that is the $\mathbf{Exp}_{\mathcal{A}_{\text{PIOP}}, \text{PIOP}}^{sr}(\mathbb{F})$ experiment for PIOP (see Definition 21) for an adversary $\mathcal{A}_{\text{PIOP}}$ that we define next.

$\mathcal{A}_{\text{PIOP}}$:

1. Run simulator $\text{srs}, \text{st}_S \leftarrow \mathcal{S}(0, \text{pp}_{\mathbb{G}})$ and set $\mathcal{Q}_{\text{RO}}, \mathcal{Q}_{\text{sim}}$ empty sets.
2. Run $\mathcal{A}_\Pi(\text{srs})$ and answer all the simulation query of \mathcal{A}_Π with \mathcal{S}_1 .
3. Upon i -th query (s_i, aux_i) from \mathcal{A}_Π to \mathcal{S}_2 :
 - (a) if s_i is in the RO-queries set of a simulated proof in \mathcal{Q}_{sim} then run \mathcal{S}_2 on input s_i .
 - (b) Else parse s_i as a (partial) transcript $\text{trns} = (\text{vk}_i, \mathbf{x}, \bar{\pi}_1, \dots, \bar{\pi}_{r'})$ for $r' \in [r]$; parse $\bar{\pi}_j = (\mathbf{c}_j, \pi_{\text{opn}, j}, \boldsymbol{\pi}_j)$, compute $\mathbf{w} \leftarrow \mathcal{E}_i(\text{srs}, \mathbf{c}_{r'}, \pi_{\text{opn}, r'}, \text{aux}_i)$, parse \mathbf{w} as some polynomials, find in SeenStates the state $\text{cvs} = (\mathbf{i}, \mathbf{x}, \boldsymbol{\pi}_1, \{p_{1,i}\}_i \|\rho_1\| \dots \|\boldsymbol{\pi}_{r'-1}, \{p_{r'-1,i}\}_i \|\rho_{r'-1}\|)$, set the verifier state to cvs , and send the message $(\mathbf{w}, \boldsymbol{\pi}_{r'})$ to the PIOP verifier. Receive the challenge $\rho_{r'}$, and program the random oracle adding $(s_i, \text{aux}_i, \rho_{r'})$ to \mathcal{Q}_{RO} .
4. Eventually the adversary outputs a (valid) forgery $(\mathbf{x}_\Pi, \pi_\Pi)$. From π_Π derive the PIOP transcript $\text{trns} := (\mathbf{i}, \mathbf{x}, \bar{\pi}_1, \rho_1, \dots, \bar{\pi}_r)$, and act as if \mathcal{A} has queried \mathcal{S}_1 with (trns, \cdot) : i.e., let i be the index of RO query of the partial transcript $(\mathbf{i}, \mathbf{x}, \bar{\pi}_1, \rho_1, \dots, \bar{\pi}_{r-1})$; as described in the previous step, find the cvs in SeenStates associated with s_i , set the verifier state to cvs , extract the (last) witness \mathbf{w} and send $(\mathbf{w}, \boldsymbol{\pi}_r)$ to the verifier. cvs and the last messages $(\mathbf{w}, \boldsymbol{\pi}_r)$ define a full transcript: this would trigger the verifier to perform the decision phase of the PIOP and set the decision bit d of the game.

Let \mathbf{H}_0 be the $\mathbf{Exp}_{\mathcal{A}_{\text{PIOP}}, \text{PIOP}}^{sr}(\mathbb{F})$. By the (state-restoration) knowledge extractability of PIOP:

$$\Pr[\mathbf{H}_0] \in \text{negl}(|\mathbb{F}|).$$

Consider \mathbf{H}_1 that additionally extracts $(p_j)_{j \in [n]} \leftarrow \mathcal{E}_{\text{poly}}(\text{srs}, \mathbf{x}_{\text{poly}}, \pi_{\text{poly}}, \text{aux}')$ and returns 1 if $\exists k : G^{(k)}(X, p_1(v_1^{(k)}(X)), \dots, p_n(v_n^{(k)}(X)), \pi_1, \dots, \pi_r) \neq 0$ or $\exists j : \text{VerCom}(\text{ck}, \mathbf{c}_j, p_j) = 0$.

Lemma 10. $\Pr[\mathbf{H}_1] \leq \Pr[\mathbf{H}_0] + \epsilon_{\text{CP}}$

Proof. We reduce to $\hat{\Phi}$ -simulation extractability of CP. We define with more details the adversary $\mathcal{B}_{\text{poly}}$ that runs \mathcal{A}_Π , and we define a policy $\hat{\Phi}_0$ that samples parameters for $\mathcal{B}_{\text{poly}}$.

$$\begin{array}{l} \text{Policy } \hat{\Phi}_0(\text{pp}_{\mathbb{G}}) \\ 1. \text{ Run } \mathcal{S}(0, \text{pp}_{\mathbb{G}}) \end{array}$$

2. Parse $\text{st}_{\mathcal{S}} = (\text{st}_{\text{CP}}, 0, (\text{coms}^{(j)})_j, (\boldsymbol{\rho}^{(j)})_j)$.
3. Let $(\tilde{v}_j^{(k)})_{j,k}$ be the polynomials defined by the instance-independent verifier queries of PIOP (see Definition 20).
4. Set $\mathcal{Q}_v = \{\tilde{v}_j^{(k)}(X, \boldsymbol{\rho}^i) : i \in [q], j \in [n], k \in [n_e]\}$

Next, we define the reduction $\mathcal{B}_{\text{poly}}$.

Reduction $\mathcal{B}_{\text{poly}}(\text{srs}, \text{pp}_{\hat{\phi}})$

1. Run $\mathcal{A}_{\Pi}(\text{srs})$.
2. Upon query $((i, \mathbb{x}), \text{aux})$ to the simulation oracle, run the same strategy of \mathcal{S}_1 defined in Fig. 3 where the calls to $\text{CP}.\mathcal{S}_1$ are forwarded to the simulation oracle of $\mathcal{B}_{\text{poly}}$.
3. Given the forgery $((i, \mathbb{x}_{\Pi}), \pi_{\Pi})$ output by \mathcal{A}_{Π} , define the instance $\mathbb{x}_{\text{poly}} = (\text{poly}, \text{trns}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$ and the proof π_{poly} .
4. Return the forgery $(\mathbb{x}_{\text{poly}}, \pi_{\text{poly}})$ and set the auxiliary input $\text{aux}_{\mathcal{E}}$ as the adversary \mathcal{A}_{Π} does.

By inspection, if the forgery of \mathcal{A}_{Π} passes the verification then forgery of $\mathcal{B}_{\text{poly}}$ passes the verification too. Moreover, by the (strong) simulation extractability game of \mathcal{A}_{Π} we have that $((i, \mathbb{x}_{\Pi}), \pi_{\Pi})$ is not in the set of simulation proofs $\mathcal{Q}_{\text{sim}, \Pi}$ of \mathcal{A}_{Π} , thus the pair $(\mathbb{x}_{\text{poly}}, \pi_{\text{poly}})$ is not in the simulation proofs $\mathcal{Q}_{\text{sim}, \text{poly}}$ of $\mathcal{B}_{\text{poly}}$. In particular, there are three cases:

1. If \mathcal{A} never queried (i, \mathbb{x}_{Π}) to its simulation oracle, i.e. \mathbb{x}_{Π} is a “fresh” instance for Π , then $\mathcal{B}_{\text{poly}}$ never queried the simulator with \mathbb{x}_{poly} (notice that \mathbb{x}_{poly} contains \mathbb{x}_{Π} in the message trns).
2. Otherwise, \mathcal{A} queried (i, \mathbb{x}_{Π}) to its simulation oracle:
 - (a) For any simulated proof π'_{Π} for (i, \mathbb{x}_{Π}) the transcript of π_{Π} is different from the simulated one. This implies that $\mathbb{x}_{\text{poly}} \neq \mathbb{x}'_{\text{poly}}$ (for the same reason of above).
 - (b) There exists a simulated proof π'_{Π} for (i, \mathbb{x}_{Π}) such that the transcripts of the forgery and of the simulated proof are equal. Since the forgery of \mathcal{A}_{Π} is not in $\mathcal{Q}_{\text{sim}, \Pi}$ then the proofs $\pi_{\text{poly}} \neq \pi'_{\text{poly}}$.

By the property of the PIOP we have that there exists an index k such that for all j we have $v_j^{(k)}$ is not constant polynomial, thus the forgery of $\mathcal{B}_{\text{poly}}$ meets the extractor policy.

Finally, by the definition of the simulator \mathcal{S} in Fig. 3, the query to the simulation oracle of $\mathcal{B}_{\text{poly}}$ respects the simulator policy of $\hat{\phi}_1$. In fact, the simulator policy of $\hat{\phi}$ matches the constraint in Definition 23. Thus the reduction $\mathcal{B}_{\text{poly}}$ follows the policy $\hat{\phi}$. On the other hand, the distinguish event implies that the extractor fails, thus this would break the $\hat{\phi}$ -simulation extractability of CP. \square

Let \mathbf{H}_2 additionally return 1 if $\exists i, j : j \in \mathcal{P}_i \wedge \text{VerCom}(\text{ck}, \mathbf{c}_j, p'_j) \neq 1$, where q_i is the index of the random oracle query such that $(s_{q_i}, \text{aux}_{q_i}, \rho_i) \in \mathcal{Q}_{\text{RO}}$, and $w_{q_i} = (p'_j)_{j \in \mathcal{P}_i}$ are the polynomials extracted by \mathcal{E}_{q_i} .

Lemma 11. $\Pr[\mathbf{H}_2] \leq \Pr[\mathbf{H}_1] + r \cdot \epsilon_{\text{CP}}$

Proof. We prove through a series of r hybrids. Let $\mathbf{H}_{1,0}$ be the same as \mathbf{H}_1 ; moreover, let $\mathbf{H}_{1,i}$ be the same as $\mathbf{H}_{1,i-1}$ but that additionally returns 1 if $\exists j \in \mathcal{P}_i : \text{VerCom}(\text{ck}, c_j, p'_j) \neq 1$, where q_i is the index of the random oracle query such that $(s_{q_i}, \text{aux}_{q_i}, \rho_i) \in \mathcal{Q}_{\text{RO}}$, and $w_{q_i} = (p'_j)_{j \in \mathcal{P}_i}$ are the polynomials extracted by \mathcal{E}_{q_i} . Clearly, $\mathbf{H}_{1,r} \equiv \mathbf{H}_2$ (where r is the number of rounds of the PIOP).

We reduce again to $\hat{\Phi}$ -simulation extractability of CP. We define the adversary \mathcal{B} that runs \mathcal{A}_Π . The adversary \mathcal{B} corresponds to the reduction \mathcal{B}_{q_i} mentioned before. Additionally, we need to define a policy $\hat{\Phi}_0$ that samples the parameters for \mathcal{B} , we set the policy identical to the one used in Lemma 10.

Reduction $\mathcal{B}(\text{srs}, \text{pp}_{\hat{\Phi}})$

1. Run $\mathcal{A}_\Pi(\text{srs})$.
2. Upon query $((i, \mathbf{x}), \text{aux})$ to simulation oracle, run the same strategy of \mathcal{S}_1 defined in Fig. 3 with the only difference that instead of calling $\text{CP}.\mathcal{S}_1$ it makes a query to its simulator.
3. Parse the q_i -th query $(s_{q_i}, \text{aux}_{q_i})$ to the RO of \mathcal{A}_Π as a partial transcript where $\bar{\pi}_i := (\mathbf{c}_i, \pi_{\text{opn},i}, \boldsymbol{\pi}_i)$.
4. Continue running the adversary, parse the forgery output by \mathcal{A}_Π as $(\bar{\pi}'_1, \dots, \bar{\pi}'_r, \pi_{\text{poly}})$. Set the instance $\tilde{\mathbf{x}} := (\text{poly}, \text{trns}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$ as the verifier would do.
5. Return the forgery $((\text{opn}, \mathbf{c}_i), \pi_{\text{opn},i})$ with auxiliary information $\text{aux}_\mathcal{E} \leftarrow \text{aux}_{q_i}$ and $\text{aux}_\Phi \leftarrow (\tilde{\mathbf{x}}, \pi_{\text{poly}})$.

By inspection, if the forgery of \mathcal{A}_Π pass the verification (both for poly -instance in the final proof and for the opn -instance in the q_i -th query) then \mathcal{B} 's forgery passes the verification too. Moreover, by the (strong) simulation extractability game of \mathcal{A}_Π we have that $((i, \mathbf{x}_\Pi), \pi_\Pi)$ is not in the set of simulations \mathcal{Q}_{sim} of \mathcal{A}_Π , thus the pair $(\tilde{\mathbf{x}}, \pi_{\text{poly}})$ is not in the simulations of \mathcal{B} (cf. Lemma 10). By the property of the PIOP, there exists an index k such that for all j we have $v_j^{(k)}$ is not constant polynomial, thus the forgery of \mathcal{B} meets the extractor policy.

Finally, by the definition of the simulator \mathcal{S} in Fig. 3, the query to the simulation oracle of \mathcal{B} respects the simulator policy of $\hat{\Phi}_1$. In fact, the simulator policy of $\hat{\Phi}$ matches with the constraint in Definition 23. Notice we are running the experiment for \mathcal{B} with the same extractor of \mathcal{B}_{q_i} . However, the first three outputs of \mathcal{A} and \mathcal{B} are distributed equivalently (\mathcal{B} additionally outputs aux_Φ), thus the same extractor \mathcal{E}_{q_i} works either for \mathcal{B}_{q_i} or for \mathcal{B} . \square

Let \mathbf{H}_3 additionally return 1 if $\exists i : p'_i \neq p_i$ where $(p'_j)_{j \in \mathcal{P}_i}$ are the polynomials extracted by \mathcal{E}_{q_i} and $(p_j)_{j \in [n]}$ are the polynomials extracted by $\mathcal{E}_{\text{poly}}$.

Lemma 12. $\Pr[\mathbf{H}_3] \leq \Pr[\mathbf{H}_2] + \epsilon_{\text{CP}}$.

Proof. The distinguishing event implies that we can define an adversary \mathcal{B} that runs the $\hat{\Phi}$ -simulation extractability experiment for CP and finds a commitment c and two different polynomials p, p' such that $\text{VerCom}(\text{ck}, c, p) = \text{VerCom}(\text{ck}, c, p')$ =

¹²Let D be the maximum degree allowed by the commitment key, let $x \in \mathbb{F}$ such that $p(x) \neq p'(x)$ and consider the following (false) statement for $\mathcal{R}_{\text{poly}}$ (and in turn of $\hat{\mathcal{R}}$):

- For $j \in [D + 1]$ set $G^{(j)}(X, X_1, X_j) = (X_1 - p(j)) - (X - j)X_j$. Namely, the polynomial associated with formal variable X_1 evaluates on $p(j)$ at point j , and X_j is the *witness*.
- Set $G^{(D+2)}(X, X_1, X_{D+2}) = (X_1 - p'(x)) - (X - x)X_{D+2}$. Namely, the polynomial associated to formal variable X_1 evaluates on $p'(x)$ at point x , and X_{D+2} is its *witness*.

We can create a valid proof for such a statement using the prover of CP: the first $d + 1$ equations are proved using p and the polynomials $p_j(X) = (p(X) - p(j))/(X - j)$, while to prove the last equation we use the witness p' . Here, we use the hypothesis that CP uses internally a CP-SNARK for $\mathcal{R}_{\text{m-ev1}}$. The proof is for a statement that is not in the language, thus we break the simulation extractability of CP.

Let \mathbf{H}_4 additionally return 1 if the forged proof shares a simulated commitment with a simulated proof in \mathcal{Q}_{sim} .

Lemma 13. $\Pr[\mathbf{H}_4] \leq \Pr[\mathbf{H}_3] + \text{negl}(\lambda)$

Proof. We need to bound the probability that \mathcal{A}_Π returns a forgery that shares a simulated commitment with a proof in \mathcal{Q}_{sim} .

Let \mathcal{B} be the same reduction described in Lemma 10: as shown before, if \mathcal{A}_Π satisfies the simulation-extractability policy then the reduction \mathcal{B} satisfies the policy $\hat{\Phi}$. We notice that the extractor of \mathcal{B} extracts a witness for all the commitments in \mathbf{c} in the forgery. Let assume that the forgery shares a commitment with the i -th simulated proof for some $i \in [q]$ and $\text{coms} \leftarrow \mathbf{M}_{i,x} \cdot (\text{coms}^{(i)} \parallel \text{ck})$ (in case the commitments are not homomorphic we assume that $\mathbf{M}_{i,x}$ is the identity matrix). Since $\text{coms} \leftarrow \mathbf{M}_{i,x} \cdot (\text{coms}^{(i)} \parallel \text{ck})$ there exists a row \mathbf{m} of the matrix $\mathbf{M}_{i,x}$ such that $c_j = \mathbf{m} \cdot \text{coms}^{(i)}$. Moreover, we have a valid opening for c_j so $\text{VerCom}(\text{ck}, c_j, p)$. Notice that a commitment function defines a one-way function; moreover, the simulated commitment is sampled from a distribution that is computationally indistinguishable from a distribution that samples *random* commitments. Thus, such an extractor would break the one-way property of the commitment function. \square

Lemma 14. $\Pr[\mathbf{H}_4] = \text{Adv}_{\mathcal{A}_\Pi, \mathcal{S}, \mathcal{E}_\Pi}^{\Phi_{\text{SE-se}}}(\lambda)$

Proof. We now show that the probability that \mathbf{H}_4 outputs 1 is equal to the probability that the adversary \mathcal{A}_Π wins the $\Phi_{\text{SE-se}}$ experiment against \mathcal{E}_Π . First, notice that the srs is generated by $\mathcal{S}(0, \text{pp}_G)$ in both experiments. Upon (valid) forgery $((i, \mathbf{x}_\Pi), \pi_\Pi)$, we notice that:

¹² Intuitively, this breaks the binding property of the commitment scheme, however, the binding property does not assume that the adversary can see simulated proofs thus we cannot reduce to it. Here, instead, we show how forge a proof for an instance that is not in the language.

- π_Π cannot share a simulated commitment with one of the simulated proofs (see Item 3) because of the check introduced in Lemma 13. Thus all the RO queries of \mathcal{A}_Π that constitute π_Π (namely the queries q_1, \dots, q_r) were forwarded to the challenger in Item 3b (in particular, any of the query was already answered by the programming of the RO by the simulator \mathcal{S}_1). This implies that the complete transcript sent by $\mathcal{A}_{\text{PIOP}}$ is in the list `SeenStates`. On the other hand, the extractor \mathcal{E}_Π does not abort at Item 3.
- $\forall j \in [n]$, the polynomial p_j extracted by \mathcal{E}_{q_i} is equal to p'_j extracted from $\mathcal{E}_{\text{poly}}$ and $\text{VerCom}(\text{ck}, \mathbf{c}_j, p_j) = 1$; this is ensured by the checks introduced in Lemma 11 and Lemma 12. This implies that the extractor $\mathcal{E}_{\text{PIOP}}$ in the $\text{Exp}_{\mathcal{A}_{\text{PIOP}}, \text{PIOP}}^{sr}$ in \mathbf{H}_3 is fed with the same polynomials extracted by \mathcal{E}_Π .
- Finally, the polynomial check on \mathbb{x}_{poly} must be satisfied by the extracted polynomials p_j because of the check introduced in Lemma 10. Thus, if the proof π_{PIOP} is valid, the decision bit in the state-restoration game is 1.

6 The Glue

In this section, we bridge together the results from Sections 4 and 5. In Section 5 we defined the notion of *compilation-ready* CP-SNARK as the (simulation-extractable) CP-SNARK for the joint relationship $\hat{\mathcal{R}}$ (see Eq. (11) in Section 5.2) that proves both knowledge of the opening and polynomial equations. In the next sections, we first give a (simulation extractable) CP-SNARK for polynomial equations in the AGM based on our result in Section 4 and then we show a CP-SNARK for $\hat{\mathcal{R}}$ leveraging the extractability of the AGM.

6.1 CP-SNARK for polynomial equations in the AGM

We present a simple and un-optimized simulation-extractable CP-SNARK for the relation of polynomial equations in the random oracle model (and AGM). We recall the CP-SNARK should be a *signature of knowledge* [19] for an arbitrary message msg and an instance of the following relation $\mathcal{R}_{\text{poly}}$ that upon parameters ck , instance $((\mathbf{c}, (\bar{G}^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]}))$, and witness (\mathbf{p}, \mathbf{o}) outputs 1 if and only if:

$$\forall k \in [n_e] : \bar{G}^{(k)}(X, p_1(v_1^{(k)}(X)), \dots, p_n(v_n^{(k)}(X))) \equiv 0$$

To make the scheme a signature of knowledge we extend the relation by adding to the instance the message msg , and requiring simulation extractability with respect to the extended instance. Our scheme is based on the CP-SNARK for $\mathcal{R}_{\text{m-ev1}}$ for Hiding KZG from Section 4. We describe the CP-SNARK CP_{poly} :

- $\text{Prove}(\text{ck}, \mathbb{x} = (\text{msg}, \mathbf{c}, (\bar{G}^{(k)}, v_1^{(k)}, \dots, v_n^{(k)})_{k \in [n_e]}), \mathbb{w} = (\mathbf{p}, \mathbf{o}))$:
 1. Compute $x^* \leftarrow \text{RO}(\mathbb{x})$.
 2. Compute for any j, k the value $x_j^{(k)} \leftarrow v_j^{(k)}(x^*)$.
 3. Let $\mathcal{X} := \{x_1^*, \dots, x_m^*\}$ be the set of the points $x_j^{(k)}$ for all k and j computed at the previous step, and let $\mathcal{P}_i := \{j : v_j^{(k)}(x^*) = x_i^*\}$.

4. For $i \in [m]$, compute $\pi_{\text{m-evil},i} \leftarrow \mathcal{P}_{\text{m-evil}}(\text{ck}, \mathbb{x}_{\text{m-evil}}, (p_j, o_j)_{j \in \mathcal{P}_i})$ for the instance $\mathbb{x}_{\text{m-evil}} = (x_i^*, (\mathbf{c}_j, p_j(x_i^*))_{j \in \mathcal{P}_i})$.
 5. Output $(\pi_{\text{m-evil},i}, (p_j(x_i^*))_{j \in \mathcal{P}_i})_{i \in [m]}$.
- **Verify**(ck, \mathbf{c} , $\mathbb{x} = (\text{msg}, \bar{G}^{(k)}, v_1^{(k)}, \dots, v_n^{(k)})_{k \in [n_e]}, \pi$)
1. Compute $x^* \leftarrow \text{RO}(\mathbb{x})$.
 2. Compute for any j, k the value $x_j^{(k)} \leftarrow v_j^{(k)}(x^*)$.
 3. Let $\mathcal{X} := \{x_1^*, \dots, x_m^*\}$ be the set of the points $x_j^{(k)}$ computed at the previous step, and let $\mathcal{P}_i := \{j : v_j^{(k)}(x^*) = x_i^*\}$.
 4. Parse π as $(\pi_{\text{m-evil},i}, (y_{i,j})_{j \in \mathcal{P}_i})_{i \in [m]}$. For any j, k let $\bar{y}_j^{(k)}$ be the (claimed) evaluation of the polynomial committed in \mathbf{c}_j on point $v_j^{(k)}(x^*) \in \mathcal{X}$, this value is equivalent to $y_{i,j}$ for the index i such that $v_j^{(k)}(x^*) = x_i^*$.
 5. Output 1 iff
 - (a) $\forall i \in [m] : \mathcal{V}_{\text{m-evil}}(\text{ck}, \mathbb{x}_{\text{m-evil},i}, \pi_{\text{m-evil},i}) = 1$, where $\mathbb{x}_{\text{m-evil},i} = (x_i^*, (\mathbf{c}_j, y_{i,j})_{j \in \mathcal{P}_i})$
 - (b) $\forall k \in [n_e] : \bar{G}^{(k)}(x^*, \bar{y}_1^{(k)}, \dots, \bar{y}_n^{(k)}) = 0$.

Consider the family of policies Φ_{poly} , such that for any $(\Phi_0, \Phi_1) \in \Phi_{\text{poly}}$ we have $\Phi_0(\text{pp}_{\mathbb{G}})$ outputs parameters pp_{Φ} that contain an arbitrary set $\mathcal{Q}_v = \{\mathbf{v}^{(i)} = (v_1^{(i)}, \dots, v_n^{(i)})\}_{i \in \text{poly}(\lambda)}$ and a list of commitments sampled from a sampler \mathcal{D} where the \mathcal{D} -Aff-MDH assumption holds. On the other hand, Φ_1 returns 1 iff:

- **Semi-adaptive simulation queries.** Let \mathcal{Q}_{sim} be the set of simulation queries made by the adversary and parse coms as a list of vectors of commitments $(\mathbf{c}^{(i)})_{i \in [\mathcal{Q}_{\text{sim}}]}$ where $\mathbf{c}^{(i)} \in \mathbb{G}_1^n$. Let the i -th query to the simulation oracle be $(\mathbb{x}, \text{aux}, \pi) \in \mathcal{Q}_{\text{sim}}$, we can parse \mathbb{x} as $(\text{msg}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$ where (1) $\forall k : \mathbf{v}^{(k)} \in \mathcal{Q}_v$ and (2) we can find in aux two matrices \mathbf{C}, \mathbf{F} such that $\mathbf{c} = \mathbf{C} \cdot \mathbf{c}^{(i)} + \mathbf{F} \cdot \text{ck}$. Moreover (3) the algebraic consistency for CP_{poly} holds¹³.
- **Forgery.** The forgery $(\mathbb{x}, \pi) \notin \mathcal{Q}_{\text{sim}}$, i.e. π is “fresh” and not produced by the simulation oracle on input \mathbb{x} . Moreover, there exists a $k \in [n_e]$ such that for all j the polynomial $v_j^{(k)}$ has a degree at least 1 (and degree polynomial in the security parameter).

Let $\text{CP}_{\text{m-evil}}$ be $F_{\text{m-evil}}$ -leaky zero-knowledge, consider the following leakage function:

Function $F_{\text{poly}}(\mathbb{x}, \mathbb{w})$:

1. Parse \mathbb{w} as (\mathbf{p}, \mathbf{o}) , leak $\{(j, p_j(x_j^{(k)}))\}_{j,k}$
2. For any $i \in [m]$ compute $\mathbb{x}_{\text{m-evil},i}$ from \mathbb{x} and the leaked points (as in the scheme above), leak points $F_{\text{m-evil}}(\mathbb{x}_{\text{m-evil},i}, (p_j, o_j)_{j \in \mathcal{P}_i})$.

¹³ Specifically, as in the protocol specification, if we derive from \mathbb{x} the m different instances $(\mathbb{x}_{\text{m-evil},i})_{i \in [m]}$ and let \mathbf{C}_i (resp \mathbf{F}_i) be the sub-matrix of \mathbf{C} (resp. \mathbf{F}) in which the set of row indices is \mathcal{P}_i . Then the following consistency opening check must hold: $\forall i$ the system of linear equations $\mathbf{C}_i \cdot \mathbf{z} = \mathbf{y}_i$ has at least a solution, where \mathbf{z} are the variables and $\mathbf{y}_i = \mathbb{x}_i \cdot \mathbf{y} + \mathbf{F}_i \cdot (1, x_i^*, \dots, x_i^{*d})^\top$.

Theorem 4. *If the CP-SNARK $\text{CP}_{m\text{-ev1}}$ is $F_{m\text{-ev1}}$ -leaky zero-knowledge then the CP-SNARK CP_{poly} is F_{poly} -leaky zero-knowledge.*

Let $\epsilon_{m\text{-ev1}}$ be the maximum winning probability of a PT adversary against the $\Phi_{m\text{-ev1}}$ -simulation extractability. Let \mathcal{E} be the canonical extractor in the AGM, let \mathcal{S} be the leaky-ZK simulators for CP_{poly} . For every $\Phi_{\text{poly}} \in \Phi_{\text{poly}}$, for every adversary \mathcal{A} that makes at most Q_{sim} simulation queries,

$$\text{Adv}_{\text{CP}_{\text{poly}}, \mathcal{S}, \mathcal{E}, \mathcal{A}}^{\Phi_{\text{poly}}\text{-se}}(\lambda) \leq (Q_{\text{sim}} + 1)\epsilon_{m\text{-ev1}} + \epsilon_{\text{Aff-MDH}} + \frac{\text{poly}(d, \lambda)}{|\mathbb{F}|}.$$

Proof. We first prove leaky zero-knowledge and then simulation extractability.

Leaky zero-knowledge. We rely on the leaky zero knowledge simulator $\mathcal{S}' = (\mathcal{S}'_0, \mathcal{S}'_1, \mathcal{S}'_2)$ of the scheme $\text{CP}_{m\text{-ev1}}$. In particular, we define the simulator $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2)$, where \mathcal{S}_0 (resp. \mathcal{S}_2) is simply a wrapper of \mathcal{S}'_0 (resp. \mathcal{S}'_2). \mathcal{S}_1 upon input the statement \mathbb{x}_{poly} and the leakage $(j, \tilde{y}_{i,j})_{i \in [m], j \in \mathcal{P}_i}, y'_1, \dots, y'_m \leftarrow F_{\text{poly}}(\mathbb{x}_{\text{poly}}, \mathbb{W}_{\text{poly}})$ derives, for any i , the statement $\mathbb{x}_{m\text{-ev1}, i}$ from \mathbb{x}_{poly} and $(j, \tilde{y}_{i,j})_{i \in [m]}$, then it runs \mathcal{S}'_1 on the derived statement and leakage y'_i . The indistinguishability of the simulated view from the real view can be proved with an hybrid argument where at each step we use the leaky zero-knowledge of $\text{CP}_{m\text{-ev1}}$.

Simulation Extractability. We consider the *canonical* extractor given by the AGM guarantees, namely the extractor that parses and outputs the polynomials in the instance as derived by the coefficients in the algebraic representations.

We define an hybrid experiment \mathbf{H}_1 that is equivalent to the Φ_{poly} -simulation extractability experiment but additionally the hybrid experiment outputs 0 if the instance in the forgery $(\mathbb{x}_{\text{poly}}^*, \pi_{\text{poly}}^*)$ of the adversary is in the set of simulation queries and the canonical extractor fails to extract a valid witness.

Lemma 15. $\text{Adv}_{\text{CP}_{\text{poly}}, \mathcal{S}, \mathcal{E}, \mathcal{A}}^{\Phi_{\text{poly}}\text{-se}}(\lambda) \leq \Pr[\mathbf{H}_1] + \epsilon_{m\text{-ev1}} + \epsilon_{\text{Aff-MDH}}$

Proof. The distinguishing event between the original experiment and \mathbf{H}_1 is that \mathcal{A} returns a valid “fresh” proof for a statement for which has seen a simulated proof which the extractor cannot extract.

We reduce to $\Phi_{m\text{-ev1}}$ -simulation extractability showing an adversary that produces a *fresh* proof for a statement $\mathbb{x}_{m\text{-ev1}}$ for which has seen a simulated proof.

Consider the policy $\Phi_{m\text{-ev1}} = (\bar{\Phi}_0, \bar{\Phi}_1) \in \Phi_{m\text{-ev1}}$ where $\bar{\Phi}_0(\text{pp}_{\mathbb{G}})$ does the following:

1. Runs the policy $\bar{\Phi}_0(\text{pp}_{\mathbb{G}})$ and obtains the set of vectors \mathcal{Q}_v .
2. Samples random $\tilde{x}_1, \dots, \tilde{x}_q$ from \mathbb{F} and computes $\mathcal{Q}_x := \{v_j(\tilde{x}_i) : \mathbf{v} \in \mathcal{Q}_v\}_{i,j}$.
3. Runs for $i \in [Q_{\text{sim}}]$ the sampler $\mathbf{c}^{(i)} \leftarrow_{\mathcal{S}} \mathcal{D}(\text{pp}_{\mathbb{G}})$, where Q_{sim} is the maximum number of simulation queries made by \mathcal{A} , and defines $\text{coms} := (\mathbf{c}^{(i)})_i$.
4. Outputs $(\mathcal{Q}_x, \text{coms})$.

We can assume w.l.g. that \mathcal{A} queries the random oracle on \mathbb{x} before querying the simulation oracle on such an instance.

Consider the adversary \mathcal{B} for the $\Phi_{m\text{-ev1}}$ -simulation-extractability game that:

1. Run the adversary \mathcal{A} on parameters ck and $\text{pp}_{\Phi} := (\mathcal{Q}_v, \text{coms})$

2. Parse \mathbf{coms} as $(\mathbf{c}^{(i)})_{i \in [Q_{\text{sim}}]}$ and keep a list \mathcal{Q}'_{RO} (initially empty) of the random oracle call of \mathcal{A} .
3. At the i -th random-oracle query on input (s, \mathbf{aux}) , store $(s, \mathbf{aux}, \tilde{x}_i)$ in \mathcal{Q}'_{RO} , and forward \tilde{x}_i to the adversary \mathcal{A} .
4. Upon simulation query with a tuple $(\mathbb{x}_{\text{poly}}, \mathbf{aux}_{\text{poly}})$ from \mathcal{A} :
 - Find in $\mathbf{aux}_{\text{poly}}$ the leakage-input for the simulator $(y_{i,j})_{i \in [m], j \in \mathcal{P}_i}, (y'_i)_{i \in [m]}$.
 - Following the specification of the prover, query the simulation oracle with instances $\mathbb{x}_{\text{m-ev1},i} = ((\mathbf{c}_j)_{j \in \mathcal{P}_i}, x_i^*, (y_{i,j})_{i,j \in \mathcal{P}_i})$, using y'_i as the leakage for the i -th instance. Parse them as a proof for poly .
5. Upon forgery $(\mathbb{x}_{\text{poly}}^*, \mathbf{aux}_{\mathcal{E}}^*, \pi^* = (\pi_{\text{m-ev1},i}^*, (y_{i,j})_i))$:
 - (a) **Abort** if \mathcal{A} never queried \mathcal{S}_2 with $\mathbb{x}_{\text{poly}}^*$.
 - (b) Let $\tilde{\pi}_{\text{poly}} := (\tilde{\pi}_{\text{m-ev1},i}, (\tilde{y}_{i,j})_i)$ be the first simulated proof for $\mathbb{x}_{\text{poly}}^*$ and let i^* be the index such that either $(y_{i^*,j}) \neq (\tilde{y}_{i^*,j})$ or $\tilde{\pi}_{\text{m-ev1},i^*} \neq \pi_{\text{m-ev1},i^*}$, return the forgery $(\mathbb{x}_{\text{m-ev1},i^*}^*, \mathbf{aux}_{\mathcal{E}}^*, \pi_{\text{m-ev1},i^*}^*)$.

We need to show that if Φ_{poly} holds for \mathcal{A} then the policy $\Phi_{\text{m-ev1}}$ holds for \mathcal{B} . Notice that, by condition (1) of the semi-adaptive simulation queries property of Φ_{poly} and the definition of Φ_0 the reduction \mathcal{B} calls its own simulator on points in \mathcal{Q}_x , moreover (2) and (3) (of the semi-adaptive simulation queries property of Φ_{poly}) easily imply respectively the Commitment Check and the Algebraic Consistency of $\Phi_{\text{m-ev1}}$ (cf. Definition 24).

By inspection the forgery, if the reduction \mathcal{B} does not abort and there exists only one $\tilde{\pi}_{\text{poly}}$ associated with the forged instance then \mathcal{B} matches the predicate $\Phi_{\text{ext}}^{\text{der}}$ of the $\Phi_{\text{m-ev1}}$ policy (cf. Definition 25). We show that, because of condition (2) and the algebraic consistency check of the semi-adaptive simulation queries property of Φ_{poly} , that there exists indeed only one simulated proof $\tilde{\pi}_{\text{poly}}$ for each queried instance. In fact, from an adversary that asks twice the same instance to the simulator, let say at query i and i' , we can derive that $\mathbf{C}^{(i)} \cdot \mathbf{c}^{(i)} + \mathbf{F}^{(i)} \cdot \text{ck} = \mathbf{C}^{(i')} \cdot \mathbf{c}^{(i')} + \mathbf{F}^{(i')} \cdot \text{ck}$. If the matrices $\mathbf{C}^{(i)}$ or $\mathbf{C}^{(i')}$ are non-zero then we can break the \mathcal{D} -Aff-MDDH assumption. On the other hand, if both matrices are the zero matrix, we have that the commitments are fully defined given the ck and thus by the algebraic consistency check the leakage for the simulated proofs must be the same (and thus the proofs are the same). \square

The next hybrid experiment \mathbf{H}_2 is equivalent to \mathbf{H}_1 but additionally it outputs 0 if the instance in the forgery $(\mathbb{x}_{\text{poly}}^*, \pi_{\text{poly}}^*)$ of the adversary is valid and the canonical extractor fails to extract valid witnesses for the m-ev1 -instances derived by $\mathbb{x}_{\text{poly}}^*$.

Lemma 16. $\Pr[\mathbf{H}_1] \leq \Pr[\mathbf{H}_2] + \frac{\epsilon_{\text{m-ev1}}}{Q_{\text{sim}}}$

Proof. We reduce again to $\Phi_{\text{m-ev1}}$ -simulation-extractability. We consider the same $\bar{\Phi}_0(\text{pp}_{\mathbb{G}})$ as in the previous lemma.

Consider the adversary \mathcal{B} that:

1. Sample an index $q^* \leftarrow_{\$} [Q_{\text{sim}}]$.
2. Run the adversary \mathcal{A} on parameters ck and $\text{pp}_{\Phi} := (Q_v, \text{coms})$

3. Parse \mathbf{coms} as $(\mathbf{c}^{(i)})_{i \in [Q_{\text{sim}}]}$ and keep a list \mathcal{Q}'_{RO} (initially empty) of the random oracle call of \mathcal{A} .
4. At the i -th random-oracle query on input (s, \mathbf{aux}) :
 - if $i \neq q^*$ then store $(s, \mathbf{aux}, \tilde{x}_i)$ in \mathcal{Q}'_{RO} , and forward \tilde{x}_i to \mathcal{A}
 - else parse s as $(\mathbf{msg}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_k)$, find k^* such that $\forall j : \deg(v_j^{(k^*)}) \geq 1$, set $\mathbf{aux}' := (\mathbf{aux}, (v_j^{(k^*)})_{j \in [n]})$, and send the query (s, \mathbf{aux}') to \mathcal{S}_2 , and forward a to the adversary \mathcal{A}
5. Upon simulation query with a tuple $(\mathbb{x}_{\text{poly}}, \mathbf{aux}_{\text{poly}})$ from \mathcal{A} (as in the previous lemma):
 - find in \mathcal{Q}'_{RO} the tuple $(\mathbb{x}_{\text{poly}}, \cdot, \tilde{x})$
 - find in $\mathbf{aux}_{\text{poly}}$ the leakage-input for the simulator $(y_{i,j})_{i \in [m], j \in \mathcal{P}_i}, (y'_i)_{i \in [m]}$
 - following the specification of the prover, query the simulation oracle with instances $\mathbb{x}_{\text{m-ev1}, i} = ((\mathbf{c}_j)_{j \in \mathcal{P}_i}, x_i^*, (y_{i,j})_{i,j \in \mathcal{P}_i})$, using y'_i as the leakage for the i -th instance
6. Upon forgery $(\mathbb{x}_{\text{poly}}^*, \mathbf{aux}_{\mathcal{E}}^*, \pi^* = (\pi_{\text{m-ev1}, i}^*, (y_{i,j})_j)_i)$
 - (a) **abort** if $\mathbb{x}_{\text{poly}}^*$ was not queried at the q^* -th random oracle query by \mathcal{A}
 - (b) **abort** if $\forall j \in [n]$ c_j can be extracted as $p_j(X)$ and $\forall i : p_j(x_i^*) = y_{i,j}$,
 - (c) Let j^*, i be such that $p_{j^*}(x_i^*)$ the previous check does not hold. Return the forgery $(\mathbb{x}_{\text{m-ev1}, i^*}^*, \mathbf{aux}_{\mathcal{E}}^*, \pi_{\text{m-ev1}, i^*}^*)$.

Let **Abort** be the event that \mathcal{B} aborts. We notice that \mathcal{B} could abort if one of two distinct events happens. The first event **Abort**₁ in Item 6a, and the second event **Abort**₂ is the condition in Item 6b. Notice that the distinguishing event between the two hybrids implies that **Abort**₂ does not happen. Moreover, notice that when the reduction does not abort, it returns a valid proof that either the canonical extractor cannot extract or such that the extracted polynomial $p_j(x_i^*) \neq y_{i^*, j^*}$; thus, if the policy $\Phi_{\text{m-ev1}}$ is valid the reduction wins the $\Phi_{\text{m-ev1}}$ -simulation extractability experiment. Namely,

$$\Pr[\bar{\Phi}_1 \wedge \neg \text{Abort} | b = 0] \leq \epsilon_{\text{m-ev1}}$$

The event **Abort**₁ only depends on the uniformly random index q^* which is independent of the view of the adversary \mathcal{A} , i.e., $\Pr[\neg \text{Abort}_1] = \frac{1}{Q_{\text{sim}}}$. Notice that \mathcal{A} wins the Φ_{poly} -simulation extractability experiment when $\neg \text{Abort}_2$ happens and Φ_1 holds. Thus, we only need to show that, conditioned on $\neg \text{Abort}$, when the policy Φ_1 holds w.r.t. the view of \mathcal{A} then the policy $\bar{\Phi}_1$ holds w.r.t. the view of \mathcal{B} . By construction of \mathcal{Q}_x and the answers to the random oracle queries to \mathcal{A} , if the simulation query of \mathcal{A} has $\mathbf{v}^{(k)} \in \mathcal{Q}_v$ then the evaluation points derived are in \mathcal{Q}_x . Moreover, it is not hard to see that the consistency check of the two policies match: in particular, the conditions on matrix \mathbf{C}_i (resp \mathbf{F}_i) are equivalent to the consistent opening check described in Item 3 of the extraction policy (the adversary is considered valid if its queries form a solvable linear system of equations defined by the coefficients of the linear combinations of commitments and the evaluation values). As for the forgery, in case of $\neg \text{Abort}_1$, the entry $(\mathbb{x}_{\text{poly}}^*, \mathbf{aux}', a^*)$ is indeed in the list of random oracle checked by the m-ev1 -policy: in fact, it is the only query. And, by construction, \mathbf{aux}' contains the polynomial $v_{j^*}^{(k^*)}$. Putting things together we have the statement of the lemma. \square

The next hybrid experiment \mathbf{H}_3 is equivalent to \mathbf{H}_2 but additionally it outputs 0 if the instance in the forgery $(\mathbb{x}_{\text{poly}}^*, \pi_{\text{poly}}^*)$ of the adversary is valid and the canonical extractor fails to extract valid witness, namely:

$$\exists k : \bar{G}^{(k)}(X, p_1(v_n^{(k)}(X)), \dots, p_n(v_n^{(k)}(X))) \neq 0$$

Lemma 17. $\Pr[\mathbf{H}_3] \leq \frac{\text{poly}(d, \lambda)}{|\mathbb{F}|}$.

Proof. First notice that by the changes introduced in the previous hybrid the canonical extractor must extract valid polynomials and $p_j(x_i^*) = y_{i,j}$ for any i and $j \in \mathcal{P}_i$. Also x^* is sampled after all the polynomials extracted by the canonical extractor are defined. Let d the degree of $\bar{G}^{(k)}(X, p_1(v_n^{(k)}(X)), \dots, p_n(v_n^{(k)}(X)))$, such value is polynomial in the security parameter and the maximum degree for the commitments. If the verification passes we have:

$$\bar{G}^{(k)}(x^*, p_1(v_n^{(k)}(x^*)), \dots, p_n(v_n^{(k)}(x^*))) = 0$$

which, by Swartz-Zippel lemma happens with probability $\frac{d}{|\mathbb{F}|}$. \square

6.2 The Compilation-Ready CP-SNARK in the AGM

We define a CP-SNARK CP for $\hat{\mathcal{R}}$ that uses the CP-SNARK CP_{poly} from Section 6.1. We set the CP = (KGen, Prove, Verify) where KGen is the same as in CP_{poly} and Prove upon poly -instances runs $\text{CP}_{\text{poly}}.\text{Prove}$, while upon opn -instances simply outputs an empty string. (The verifier algorithm is the obvious one.)

We recall the policy set Φ defined in Section 5.2. For every $(\Phi_0, \Phi_1) \in \Phi_{\text{poly}}$, $(\hat{\Phi}_0, \hat{\Phi}_1) \in \Phi$ where $\hat{\Phi}_1$ satisfies the 3 properties:

- **Semi-adaptive w.r.t. poly-queries.** Given the set of simulation queries \mathcal{Q}_{sim} , define the projections of the set to the poly -simulation queries and opn -simulation queries. Let the i -th poly -query to the simulation oracle be $(\mathbb{x}, \text{aux}, \pi) \in \mathcal{Q}_{\text{sim}}$, parse \mathbb{x} as $(\text{msg}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{k \in [n_e]})$. (1) $\forall k : \mathbf{v}^{(k)} \in \mathcal{Q}_v$ and (2) parse aux two matrices \mathbf{C}, \mathbf{F} the tuple $(\mathbf{c}, \mathbf{C}, \mathbf{F})$ satisfies the commitment check, namely $\mathbf{c} = \mathbf{C} \cdot \mathbf{c}^{(i)} + \mathbf{F} \cdot \text{ck}$. (3) the view defined by set of all poly -queries satisfies the algebraic consistency for CP (cf. Definition 14).
- **Extractor policy for opn-forgery.** If the forgery of the adversary (\mathbb{x}^*, π^*) is of the form $\mathbb{x}^* := (\text{opn}, \mathbf{c}^*)$, parse aux_{Φ} as (yet another) forgery $\tilde{\mathbb{x}} := (\text{poly}, \text{msg}, \mathbf{c}, (G^{(k)}, \mathbf{v}^{(k)})_{i \in [n_e]})$, $\tilde{\pi}$ and check that:
 1. All the commitments \mathbf{c}^* are in the vector of commitments \mathbf{c} ,
 2. The proof $\tilde{\pi}$ is valid according to the extractor policy for poly -forgeries below.
- **Extractor policy for poly-forgery.** If the forgery of the adversary (\mathbb{x}^*, π^*) is of the form $\mathbb{x}^* := (\text{poly}, \mathbb{x}')$, then check that \mathbb{x}^* was never queried to the simulation oracle and there exists a $k \in [n_e]$ such that for all j the polynomial $v_j^{(k)}$ has degree at least 1 (and degree $\text{poly}(\lambda)$).

Theorem 5. *The CP-SNARK CP for $\hat{\mathcal{R}}$ is $\hat{\Phi}$ -simulation extractable.*

Proof. The proof is mostly syntactical as we need to match the interface of an adversary of CP with the interface of an adversary of CP_{poly} .

We reduce to the Φ_{poly} -simulation extractability of CP_{poly} . In particular, given an adversary \mathcal{A} for CP consider the reduction that runs \mathcal{A} and forwards all its poly-simulation queries, moreover, the reduction answers all the opn-simulation queries by trivially outputting the empty string. If \mathcal{A} outputs as a forgery a opn-instance then the reduction outputs as a forgery the instance and valid proof in aux_{Φ} .

We assume that the representations of the commitments in the opn-forgery and the representations of the same commitments in the auxiliary output aux_{Φ} are the same. Notice, if we have two distinct representations for the same commitment then we break the binding of KZG. Thus the canonical extractor would output the same opening both when extracting the opn-forgery (using the representations in $\text{aux}_{\mathcal{E}}$) and when extracting from aux_{Φ} .

References

1. M. Abdalla, M. Barbosa, J. Katz, J. Loss, and J. Xu. Algebraic adversaries in the universal composability framework. In M. Tibouchi and H. Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 311–341. Springer, Heidelberg, Dec. 2021.
2. B. Abdolmaleki, S. Ramacher, and D. Slamanig. Lift-and-shift: Obtaining simulation extractable subversion and updatable SNARKs generically. In J. Ligatti, X. Ou, J. Katz, and G. Vigna, editors, *ACM CCS 2020*, pages 1987–2005. ACM Press, Nov. 2020.
3. S. Arora and S. Safra. Probabilistic checking of proofs; A new characterization of NP. In *33rd FOCS*, pages 2–13. IEEE Computer Society Press, Oct. 1992.
4. K. Bagheri, M. Kohlweiss, J. Siim, and M. Volkhov. Another look at extraction and randomization of groth’s zk-snark. In N. Borisov and C. Díaz, editors, *Financial Cryptography and Data Security - 25th International Conference, FC 2021, Virtual Event, March 1-5, 2021, Revised Selected Papers, Part I*, volume 12674 of *Lecture Notes in Computer Science*, pages 457–475. Springer, 2021.
5. M. Bellare, D. Hofheinz, and E. Kiltz. Subtleties in the definition of IND-CCA: When and how should challenge decryption be disallowed? *Journal of Cryptology*, 28(1):29–48, Jan. 2015.
6. E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev. Scalable zero knowledge with no trusted setup. In A. Boldyreva and D. Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 701–732. Springer, Heidelberg, Aug. 2019.
7. E. Ben-Sasson, A. Chiesa, D. Genkin, E. Tromer, and M. Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 90–108. Springer, Heidelberg, Aug. 2013.
8. E. Ben-Sasson, A. Chiesa, M. Riabzev, N. Spooner, M. Virza, and N. P. Ward. Aurora: Transparent succinct arguments for R1CS. In Y. Ishai and V. Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 103–128. Springer, Heidelberg, May 2019.

9. E. Ben-Sasson, A. Chiesa, and N. Spooner. Interactive oracle proofs. In M. Hirt and A. D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 31–60. Springer, Heidelberg, Oct. / Nov. 2016.
10. E. Ben-Sasson, L. Goldberg, S. Kopparty, and S. Saraf. DEEP-FRI: Sampling outside the box improves soundness. In T. Vidick, editor, *ITCS 2020*, volume 151, pages 5:1–5:32. LIPIcs, Jan. 2020.
11. D. Boneh and X. Boyen. Short signatures without random oracles. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 56–73. Springer, Heidelberg, May 2004.
12. D. Boneh, J. Drake, B. Fisch, and A. Gabizon. Efficient polynomial commitment schemes for multiple points and polynomials. Cryptology ePrint Archive, Report 2020/081, 2020. <https://eprint.iacr.org/2020/081>.
13. D. Boneh, J. Drake, B. Fisch, and A. Gabizon. Halo infinite: Proof-carrying data from additive polynomial commitments. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 649–680, Virtual Event, Aug. 2021. Springer, Heidelberg.
14. B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
15. B. Bünz, A. Chiesa, P. Mishra, and N. Spooner. Recursive proof composition from accumulation schemes. In R. Pass and K. Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 1–18. Springer, Heidelberg, Nov. 2020.
16. B. Bünz, B. Fisch, and A. Szepieniec. Transparent SNARKs from DARK compilers. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Heidelberg, May 2020.
17. M. Campanelli, A. Faonio, D. Fiore, A. Querol, and H. Rodríguez. Lunar: A toolbox for more efficient universal and updatable zkSNARKs and commit-and-prove extensions. In M. Tibouchi and H. Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 3–33. Springer, Heidelberg, Dec. 2021.
18. M. Campanelli, D. Fiore, and A. Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In L. Cavallaro, J. Kinder, X. Wang, and J. Katz, editors, *ACM CCS 2019*, pages 2075–2092. ACM Press, Nov. 2019.
19. M. Chase and A. Lysyanskaya. On signatures of knowledge. In C. Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 78–96. Springer, Heidelberg, Aug. 2006.
20. A. Chiesa, Y. Hu, M. Maller, P. Mishra, N. Vesely, and N. P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In A. Canteaut and Y. Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Heidelberg, May 2020.
21. Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs. Efficient public-key cryptography in the presence of key leakage. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 613–631. Springer, Heidelberg, Dec. 2010.
22. D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography (extended abstract). In *23rd ACM STOC*, pages 542–552. ACM Press, May 1991.
23. S. Faust, M. Kohlweiss, G. A. Marson, and D. Venturi. On the non-malleability of the Fiat-Shamir transform. In S. D. Galbraith and M. Nandi, editors, *INDOCRYPT 2012*, volume 7668 of *LNCS*, pages 60–79. Springer, Heidelberg, Dec. 2012.
24. M. Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In V. Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 152–168. Springer, Heidelberg, Aug. 2005.

25. G. Fuchsbauer, E. Kiltz, and J. Loss. The algebraic group model and its applications. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Heidelberg, Aug. 2018.
26. A. Gabizon, Z. J. Williamson, and O. Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>.
27. C. Ganesh, H. Khoshakhlagh, M. Kohlweiss, A. Nitulescu, and M. Zajac. What makes fiat-shamir zk-snarks (updatable SRS) simulation extractable? In C. Galdi and S. Jarecki, editors, *Security and Cryptography for Networks, SCN 2022*, volume 13409 of *Lecture Notes in Computer Science*, pages 735–760. Springer, 2022.
28. C. Ganesh, Y. Kondi, C. Orlandi, M. Pancholi, A. Takahashi, and D. Tschudi. Witness-succinct universally-composable snarks. Cryptology ePrint Archive, Paper 2022/1618, 2022. <https://eprint.iacr.org/2022/1618>.
29. C. Ganesh, C. Orlandi, M. Pancholi, A. Takahashi, and D. Tschudi. Fiat-shamir bulletproofs are non-malleable (in the algebraic group model). In O. Dunkelman and S. Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 397–426. Springer, Heidelberg, May / June 2022.
30. S. Garg, A. Jain, and A. Sahai. Leakage-resilient zero knowledge. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 297–315. Springer, Heidelberg, Aug. 2011.
31. R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic span programs and succinct NIZKs without PCPs. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
32. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In *17th ACM STOC*, pages 291–304. ACM Press, May 1985.
33. J. Groth. Short pairing-based non-interactive zero-knowledge arguments. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 321–340. Springer, Heidelberg, Dec. 2010.
34. J. Groth. On the size of pairing-based non-interactive arguments. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
35. J. Groth, M. Kohlweiss, M. Maller, S. Meiklejohn, and I. Miers. Updatable and universal common reference strings with applications to zk-SNARKs. In H. Shacham and A. Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 698–728. Springer, Heidelberg, Aug. 2018.
36. J. Groth and M. Maller. Snarky signatures: Minimal signatures of knowledge from simulation-extractable SNARKs. In J. Katz and H. Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 581–612. Springer, Heidelberg, Aug. 2017.
37. C. S. Jutla and A. Roy. Shorter quasi-adaptive NIZK proofs for linear subspaces. In K. Sako and P. Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 1–20. Springer, Heidelberg, Dec. 2013.
38. A. Kate, G. M. Zaverucha, and I. Goldberg. Constant-size commitments to polynomials and their applications. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, Dec. 2010.
39. J. Lee. Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. In K. Nissim and B. Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 1–34. Springer, Heidelberg, Nov. 2021.

40. H. Lipmaa. Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In R. Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 169–189. Springer, Heidelberg, Mar. 2012.
41. M. Maller, S. Bowe, M. Kohlweiss, and S. Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In L. Cavallaro, J. Kinder, X. Wang, and J. Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, Nov. 2019.
42. M. Maller, S. Bowe, M. Kohlweiss, and S. Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updateable structured reference strings. Cryptology ePrint Archive, Report 2019/099, 2019. <https://eprint.iacr.org/2019/099>.
43. S. Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, Nov. 1994.
44. P. Morillo, C. Ràfols, and J. L. Villar. The kernel matrix Diffie-Hellman assumption. In J. H. Cheon and T. Takagi, editors, *ASIACRYPT 2016, Part I*, volume 10031 of *LNCS*, pages 729–758. Springer, Heidelberg, Dec. 2016.
45. C. Ràfols and A. Zapico. An algebraic framework for universal and updatable SNARKs. In T. Malkin and C. Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 774–804, Virtual Event, Aug. 2021. Springer, Heidelberg.
46. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th FOCS*, pages 543–553. IEEE Computer Society Press, Oct. 1999.
47. A. Szepieniec. Polynomial IOPs for linear algebra relations. Cryptology ePrint Archive, Report 2020/1022, 2020. <https://eprint.iacr.org/2020/1022>.
48. A. Tomescu, I. Abraham, V. Buterin, J. Drake, D. Feist, and D. Khovratovich. Aggregatable subvector commitments for stateless cryptocurrencies. In C. Galdi and V. Kolesnikov, editors, *SCN 20*, volume 12238 of *LNCS*, pages 45–64. Springer, Heidelberg, Sept. 2020.
49. A. Zapico, V. Buterin, D. Khovratovich, M. Maller, A. Nitulescu, and M. Simkin. Caulk: Lookup arguments in sublinear time. In H. Yin, A. Stavrou, C. Cremers, and E. Shi, editors, *ACM CCS 2022*, pages 3121–3134. ACM Press, Nov. 2022.

A CP-SNARK for evaluation of multiple polynomials in AGM

Here we generalize the scheme described in Section 4.1 to the batched setting, **highlighting** the parts of the construction that are not needed if hiding is not desired. In particular, this scheme CPm-ev1 allows one to prove that for all i the polynomial f_i committed in \mathbf{c}_i evaluates to y_i on the point x . This batched version, which is given in the ROM, follows from [26,42] and relies on the linearity of the polynomials and the homomorphic properties of KZG. Our contribution is to prove its policy-based simulation extractability.

$\text{KGen}_{\text{m-ev1}}$: parse ck as $(([s^j]_1)_{j \in [0,d]}, ([\alpha s^j]_1)_{j \in [0,d]}, [1, s]_2)$ and define $\text{ek} := \text{ck}$ and $\text{vk} := [1, s]_2$, and return $\text{srs} := (\text{ek}, \text{vk})$.

$\text{Prove}_{\text{m-ev1}}(\text{ek}, \mathbf{x} = (x, (\mathbf{c}_i, y_i)_i), \mathbf{w} = (f_i, r_i)_i)$: for all i compute the polynomials $\pi_i(X)$ such that $\pi_i(X)(X - x) \equiv f_i(X) - y_i$, **the polynomials $\pi_i(X)$ such that $\pi_i(X)(X - x) \equiv r_i(X) - r(x)$** , $\rho \leftarrow \text{RO}(\text{batch} \parallel \mathbf{x})$, and output

$$\left(\sum_i \rho^{i-1} [\pi_i(s) + \alpha \pi'_i(s)]_1, \sum_i \rho^{i-1} r_i(x) \right)$$

$\text{Verify}_{\text{m-ev1}}(\text{vk}, \mathbf{x} = (x, (\mathbf{c}_i, y_i)_i), (\pi, \mathbf{y}'))$: compute $\rho \leftarrow \text{RO}(\text{batch} \parallel \mathbf{x})$ and output 1 iff

$$e\left(\sum_i \rho^{i-1} \mathbf{c}_i - \left[\sum_i \rho^{i-1} y_i\right]_1 - [\alpha \mathbf{y}']_1, [1]_2\right) = e(\pi, [s - x]_2).$$

Similarly to CPev1 , we can prove that CPm-ev1 achieves *F-leaky* zero-knowledge (see Definition 11) **where $F(\mathbf{x} = (x, (\mathbf{c}_i, y_i)_i), \mathbf{w} = (f_i, r_i)_i) = \sum_i \rho^{i-1} r_i(x)$ and $\rho = \text{RO}(\text{batch} \parallel \mathbf{x})$** .

We define the simulator $\mathcal{S} = (\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2)$, where \mathcal{S}_0 outputs the trapdoor information s, α together with the srs , \mathcal{S}_2 simulates the random oracle on any input via lazy sampling, and \mathcal{S}_1 simulates proofs for $\mathbf{x} = (x, (\mathbf{c}_i, y_i)_i)$ **and leakage \mathbf{y}'** by outputting $\pi = ((\mathbf{c} - [\mathbf{y}]_1 - [\alpha \mathbf{y}']_1)(s - x)^{-1}, \mathbf{y}')$, where $\mathbf{c} \leftarrow \sum_i \rho^{i-1} \mathbf{c}_i$, $\mathbf{y} \leftarrow \sum_i \rho^{i-1} y_i$ for $\rho \leftarrow \mathcal{S}_2(\text{batch} \parallel \mathbf{x})$.

The extraction policy. The extraction policy $\Phi_{\text{m-ev1}}^{\text{s-adpt}}$ is naturally extended from the single point case of CPev1 . The evaluation points x_j for the instances for which the adversary can see simulated proofs are selectively chosen independently of the commitment key, while the evaluation values $(y_i)_i$ can be arbitrarily chosen by the adversary. Each policy $\Phi_{\mathcal{D}}$ in the family is a tuple of the form $(\Phi_0^{\mathcal{D}}, \Phi_1)$. Each policy $\Phi_{\mathcal{D}}$ is a tuple of the form $(\Phi_0^{\mathcal{D}}, \Phi_1)$, as defined in Section 3.1, where $\Phi_0^{\mathcal{D}}$ outputs the parameters pp_{Φ} while Φ_1 outputs a verdict bit. In particular, $\Phi_0^{\mathcal{D}}$ on input group parameters $\text{pp}_{\mathbb{G}}$ outputs $\text{pp}_{\Phi} := (\text{coms}, \mathcal{Q}_x)$, where coms is a vector of commitments sampled from \mathcal{D} , and \mathcal{Q}_x is a set of \mathcal{Q}_x evaluation points.

For sake of clarity, we define the policy Φ_1 as the logical conjunction of a “simulator” policy Φ_{sim} and an “extractor” policy Φ_{ext} , i.e. $\Phi_1 = \Phi_{\text{sim}} \wedge \Phi_{\text{ext}}$, the

first defines rules under which we can classify a simulation query *legal*, while the second defines rules under which the extractor must be able to extract a meaningful witness. We [highlight](#) the parts needed only for the hiding setting.

Definition 24. Let Φ_{sim} be the policy that returns 1 if and only if:

1. **Points check:** let $(\mathbb{x}_i, \text{aux}_i, \pi_i)_i$ be all the entries of \mathcal{Q}_{sim} . Recall that an instance \mathbb{x} can be parsed as $(x, (\mathbf{c}, \mathbf{y}))$, check that $\forall i : \mathbb{x}_i.x \in \mathcal{Q}_x$.
2. **Commitment Check:** For any $i \in [\mathcal{Q}_{\text{sim}}]$, parse aux_i as *the leakage value y'_i and the representation vectors for $\mathbb{x}_i.\mathbf{c}$ and π_i* ; in particular, let $\mathbf{M}_i = \mathbf{F}_i \parallel \mathbf{V}_i \parallel \mathbf{C}_i$ be the algebraic representation of the commitments $\mathbb{x}_i.\mathbf{c}$. For any i check that $(\mathbf{F}_i \parallel \mathbf{V}_i) \cdot \text{ek} + \mathbf{C}_i \cdot \text{coms} = \mathbb{x}_i.\mathbf{c}$ (namely, that the commitments in the instance \mathbb{x}_i are computed as a linear combination of the simulated commitments and the elements of \mathbb{G}_1 of the SRS)
3. **Algebraic Consistency:** The simulation queries satisfy the algebraic consistency for $\text{CP}_{\text{m-evl}}$. Namely, let $\mathcal{I}_J = \{i : \mathbb{x}_i.x = x_j\}$ and let $\mathbf{R}_j = (\mathbf{C}_i)_{i \in \mathcal{I}_j}$. Check that for all j : (i) the system of linear equations $\mathbf{R}_j \cdot \mathbf{z} = \mathbf{y}_j$ has at least a solution, where \mathbf{z} are the variables and $\mathbf{y}_j = (\mathbb{x}_i.\mathbf{y} + \mathbf{F}_i \cdot (1, x_j, \dots, x_j^d)^\top)_{i \in \mathcal{I}_j}$, and (ii) the system of linear equations $\mathbf{R}_j \cdot \mathbf{z}' = \mathbf{y}'_j$ has at least a solution, where \mathbf{z}' are the variables and $\mathbf{y}'_j = (y'_i + (1, \rho_i, \dots, \rho_i^m) \cdot \mathbf{V}_i \cdot (1, x_j, \dots, x_j^d)^\top)_{i \in \mathcal{I}_j}$, and $\rho_i = \text{RO}(\text{batch} \parallel \mathbb{x}_i)$.

Next, we define the policy Φ_{ext} as the logical disjunction of two policies. We recall and extend the notation introduced in Section 4.1: let $g_c : \mathbb{G}_1^* \times \{0, 1\}^* \rightarrow \{0, 1\}$ be a function that on inputs a list of group elements c_i and a string s , that can be parsed as a list of group elements \hat{c}_i followed by a second string \tilde{s} , outputs 1 iff $\forall i, \exists j : c_i = \hat{c}_j$.

Definition 25. Let $\Phi_{\text{ext}}, \Phi_{\text{ext}}^{\text{rnd}}$ and $\Phi_{\text{ext}}^{\text{der}}$ be predicates that parse the forgery instance $\mathbb{x}^* = (x^*, \mathbf{c}^*, \mathbf{y}^*)$.

- $\Phi_{\text{ext}}^{\text{rnd}}$ returns 1 if and only if there exists a query (s, aux, a) to the random oracle and aux contains a non-constant polynomial $h(X)$ such that the following conditions are satisfied:
 1. **Hashing check:** $(s, \text{aux}, a) \in \mathcal{Q}_{\text{RO}}$
 2. **Decoding check:** $g_c(\mathbf{c}^*, s) = 1$.
 3. **Polynomial check:** $g_h(h, \text{aux}) = 1$, where $g_h : \mathbb{F}[X] \times \{0, 1\}^* \rightarrow \{0, 1\}$ is a function that on input a polynomial $h(X)$ and a string aux outputs 1 iff $h(X)$ is encoded in aux .
 4. **Computation check:** $h(a) = x^*$.
- $\Phi_{\text{ext}}^{\text{der}}$ returns 1 iff $\exists (\mathbb{x}, \cdot, \pi) \in \mathcal{Q}_{\text{sim}}$ s.t. $\mathbb{x} := (x^*, \mathbf{c}^*, \mathbf{y}')$ and $(\mathbf{y}', \pi) \neq (\mathbf{y}^*, \pi^*)$.
- Φ_{ext} returns logical disjunction of $\Phi_{\text{ext}}^{\text{rnd}}$ and $\Phi_{\text{ext}}^{\text{der}}$.

Theorem 6. For any witness sampleable distribution \mathcal{D} that is \mathcal{D} -Aff-MDH-secure (see Definition 5), any bilinear-group generator GroupGen that samples the generator of the group \mathbb{G}_1 uniformly at random, $\forall \Phi_{\mathcal{D}} \in \Phi_{\text{m-evl}}^{\text{s-adpt}}$, the scheme $\text{CP}_{\text{m-evl}}$ is $\Phi_{\mathcal{D}}$ -simulation-extractable in the AGM. In particular, there exists \mathcal{E} such that for any algebraic adversary \mathcal{A} :

$$\text{Adv}_{\text{CP}_{\text{m-evl}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{D}}^{\text{se}}}(\lambda) \leq O(\epsilon_{(\mathcal{Q}_x + d + 1)\text{-DL}}(\lambda)) + O(\epsilon_{\text{Aff-MDH}}(\lambda)) + \text{poly}(\lambda)\epsilon_h$$

where d is the maximum degree supported by CPm-ev1 , $\epsilon_{(Q_x+d+1)\text{-DL}}(\lambda)$ is the maximum advantage for any algebraic PT adversary against the $(Q_x + d + 1)$ -strong Discrete-Log Assumption, $\epsilon_{\text{Aff-MDH}}(\lambda)$ is the maximum advantage for any algebraic PT adversary against the \mathcal{D} -Aff-MDH Assumption, h is the polynomial that satisfies the Polynomial check of $\Phi_{\mathcal{D}}$, and $\epsilon_h = \frac{\deg(h)}{q}$.

Proof (of Theorem 6). The goal of the proof is to show that for any algebraic \mathcal{A} there exists an extractor \mathcal{E} such that $\Pr[\mathbf{Exp}_{\Pi_{\text{m-ev1}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}}^{\Phi_{\mathcal{D}}\text{-se}}(\lambda) = 1]$ is negligible.

Since \mathcal{A} is algebraic (cf. Definition 1), for each group element \mathbf{g} that it produces, it additionally attaches a representation vector $\mathbf{r}_{\mathbf{g}} := \mathbf{f}_{\mathbf{g}} \parallel \mathbf{c}_{\mathbf{g}} \parallel \mathbf{o}_{\mathbf{g}}$ such that:

$$\mathbf{g} = \langle \mathbf{f}_{\mathbf{g}}, \text{ek} \rangle + \langle \mathbf{c}_{\mathbf{g}}, \text{coms} \rangle + \langle \mathbf{o}_{\mathbf{g}}, \text{proofs} \rangle.$$

where ek are all the group elements of the SRS, $\text{coms} = ([c_i]_1)_{i \in [Q_c]}$ are the simulated commitments, and proofs are the simulated proofs.

The algebraic assumption applies to each query $(\mathbf{x} = (x, (\mathbf{c}_i, y_i)_i), \mathbf{aux})$ to the oracle \mathcal{S}_1 , in which case we can parse \mathbf{aux} as $((\mathbf{r}_{\mathbf{c}_i})_i, \mathbf{aux}')$ such that $\mathbf{r}_{\mathbf{c}_i}$ is the representation of \mathbf{c}_i , as well as to each query (s, \mathbf{aux}) to \mathcal{S}_2 , in which case \mathbf{aux} includes a valid representation for all the group elements \mathbf{g}_i encoded in s , i.e., such that $g_c(\mathbf{g}_i, s) = 1$.

Furthermore, when returning its forgery $\mathbf{x}^* = (x^*, (\mathbf{c}_i^*, y_i)_i), [\pi]_1^*$, the algebraic adversary encodes a polynomial $h(X)$ in \mathbf{aux}_{ϕ} , and stores in $\mathbf{aux}_{\mathcal{E}}$ the representation vectors $\{\mathbf{r}_{\mathbf{c}_i^*}\}_i$ and \mathbf{r}_{π^*} for the group elements $(\mathbf{c}_i^*)_i$ and π^* .

We can assume w.l.o.g. that all the simulation queries and the forgery of the adversary \mathcal{A} agree with the policy $\Phi_{\mathcal{D}}$, as otherwise the adversary would automatically lose the experiment. In particular, we assume that the forgery satisfies the extraction predicate $\Phi_{\text{ext}}^{\text{rnd}}$: we show in Lemma 27 why this is without loss of generality. For every query $(\mathbf{x} = (x, (\mathbf{c}_i, y_i)_i), \mathbf{aux})$ to the oracle \mathcal{S}_1 , recall that this means that $\mathbf{o}_{\mathbf{c}_i} = \mathbf{0}$ for all i .

We define the extractor \mathcal{E} to be the *canonical* extractor that, given $\mathbf{aux}_{\mathcal{E}}$, returns the polynomials $(f_i(X))_i$ where:

$$f_i(X) = \langle \mathbf{f}_{\mathbf{c}_i^*}, (1, X, \dots, X^d) \rangle$$

We let $\mathbf{H}_0(\mathcal{A})$ be the $\mathbf{Exp}_{\Pi_{\text{m-ev1}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}}^{\Phi_{\mathcal{D}}\text{-se}}(\lambda)$ experiment (executed with adversary \mathcal{A}).

Hybrid \mathbf{H}_1 . This is a modification of \mathbf{H}_0 in which the adversary is allowed to only do queries to \mathcal{S}_1 of the form $(\mathbf{x} = (x, \mathbf{c}, y), \mathbf{aux})$ (i.e., for single evaluations) and such that $\mathbf{f}_{\mathbf{c}} = \mathbf{0}$. This means that an adversary not meeting this requirement loses, i.e., the hybrid returns 0. As we show in the following lemma, we can perfectly simulate adversaries in \mathbf{H}_0 using an adversary in \mathcal{A}_1 . Such a change is useful as it allows us to later consider “simpler” adversaries.

Lemma 18. *For any \mathcal{A}_0 there exists \mathcal{A}_1 such that $\Pr[\mathbf{H}_1(\mathcal{A}_1)] = \Pr[\mathbf{H}_0(\mathcal{A}_0)]$.*

Proof. We build \mathcal{A}_1 runs \mathcal{A}_0 providing its own inputs and forwarding all the random oracle queries, except for \mathcal{S}_1 queries that are handled as follows. Whenever

\mathcal{A}_0 makes a simulation query $(\mathbf{x} = (x, (\mathbf{c}_i, y_i)_i), \mathbf{aux})$, \mathcal{A}_1 proceeds as described below.

- For every i , parse $\mathbf{r}_{\mathbf{c}_i} := \mathbf{f}_{\mathbf{c}_i} \parallel \mathbf{c}_{\mathbf{c}_i} \parallel \mathbf{o}_{\mathbf{c}_i}$ and define $f_i(X) := \langle \mathbf{f}_{\mathbf{c}_i}, (1, X, \dots, X^d) \rangle$.
- Compute $y'_i \leftarrow f_i(x)$, $\pi'_i(X) \leftarrow (f_i(X) - y'_i)/(X - x)$ and $\pi'_i \leftarrow [\pi'_i(s)]_1$, which is essentially a “honest” evaluation proof w.r.t. a commitment to the polynomial $f_i(X)$, input x and output y'_i .
- For every i , ask the query $(x, \mathbf{c}_i - [f_i(s)]_1, y_i - y'_i)$ to its simulation oracle. Notice that the queried commitments satisfy the condition $\mathbf{f} = \mathbf{0}$. Let $\tilde{\pi}_i$ be the obtained proofs.
- “Merge” these proofs by computing $\pi_i \leftarrow \pi'_i + \tilde{\pi}_i$. One can verify that π_i is a valid evaluation proof for (x, \mathbf{c}_i, y_i) .
- Aggregate all the proofs by computing $\pi \leftarrow \sum_i \rho^{i-1} \pi_i$, where $\rho \leftarrow \mathcal{S}_2(\mathbf{x})$.

Hybrid \mathbf{H}_2 . Recall that \mathcal{D} is witness sampleable, thus according to Definition 4 there exists a PPT algorithm $\tilde{\mathcal{D}}$ associated with the sampler \mathcal{D} . The hybrid experiment \mathbf{H}_2 is identical to the previous one, but the group elements in \mathbf{coms} are “sampled at the exponent”, i.e., we use $\tilde{\mathcal{D}}$ to first generate the field elements γ , we compute $\mathbf{coms} \leftarrow [\gamma]_1$, and we also add γ to $\mathbf{st}_{\mathcal{S}}$. By the witness sampleability of \mathcal{D} , \mathbf{H}_1 and \mathbf{H}_2 are perfectly indistinguishable, i.e., for any \mathcal{A} $\Pr[\mathbf{H}_2(\mathcal{A})] = \Pr[\mathbf{H}_1(\mathcal{A})]$.

Hybrid \mathbf{H}_3 . In this hybrid, we change the way we generate the SRS \mathbf{srs} with \mathcal{S}_0 and the way in which \mathcal{S}_1 simulates proofs.

Let $((\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e), [1]_1, [1]_2) \leftarrow \mathcal{S} \text{ GroupGen}(1^\lambda)$, sample $s \leftarrow \mathcal{S} \mathbb{F}$ and compute $[s, \dots, s^{D+d}]_1, [1, s]_2$, where $D \leftarrow Q_x + 1$. Let $x_r \leftarrow \mathcal{S} \mathbb{F}$, and let $p(X)$ be the vanishing polynomial in $\mathcal{Q}_x \cup \{x_r\}$, namely:

$$p(X) := (X - x_r) \prod_{x \in \mathcal{Q}_x} (X - x).$$

Let also $p_j(X) := p(X)(X - x_j)^{-1}$, for $j \in [Q_x]$. In this game \mathcal{S}_0 computes:

- $\mathbf{pp}_{\mathbb{G}} := ((\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e), [p(s)]_1, [1]_2)$,
- $\mathbf{srs} := (\mathbf{ek}, \mathbf{vk})$, where $\mathbf{ek} \leftarrow [p(s), p(s)s, \dots, p(s)s^d]_1$ and $\mathbf{vk} \leftarrow [1, s]_2$,
- $\mathbf{coms} := [p(s)\gamma]_1$,
- $\mathbf{st}_{\mathcal{S}} := [1, s, \dots, s^{D+d}]_1, [1, s]_2, \gamma$.

Upon a query of the form $(\mathbf{x} = (x, \mathbf{c}, y), \mathbf{aux} = (\mathbf{r}_{\mathbf{c}}, \mathbf{aux}'))$ to \mathcal{S}_1 such that $x = x_j \in \mathcal{Q}_x$, we proceed as follows.

- Parse $\mathbf{r}_{\mathbf{c}} := \mathbf{f}_{\mathbf{c}} \parallel \mathbf{c}_{\mathbf{c}} \parallel \mathbf{o}_{\mathbf{c}}$ (recall, $\mathbf{o}_{\mathbf{c}} = \mathbf{0}$ by simulation policy and $\mathbf{f}_{\mathbf{c}} = \mathbf{0}$ by the change introduced earlier).
- Compute $\pi \leftarrow [(\langle \mathbf{c}_{\mathbf{c}}, \gamma \rangle - y) \cdot p_j(s)]_1$, which is a valid simulated proof w.r.t. commitment $[\langle \mathbf{c}_{\mathbf{c}}, \gamma \rangle p(s)]_1$, input x and output y .

We now show that the view offered to the adversary \mathcal{A} in the two experiments is statistically close.

Lemma 19. For any \mathcal{A} $\Pr[\mathbf{H}_3(\mathcal{A})] \leq \Pr[\mathbf{H}_2(\mathcal{A})] + \frac{Q_x+1+d}{q}$.

Proof. Let us first look at the distribution of the input provided to \mathcal{A} . Notice that in \mathbf{H}_3 we sample from **GroupGen** the description of the group, and then we set the generator of \mathbb{G}_1 to be $[p(s)]_1$ which, thanks to the random root x_r , is distributed uniformly at random even given the value s . This however holds except with the negligible probability $\frac{Q_x+1+d}{p}$ that s is one of the roots of $p(X)$.

Second, it is not hard to verify that the simulated proofs generated in the hybrid \mathbf{H}_3 pass the verification equations. Since the proofs are unique, given the SRS and the statements, the simulated proofs created in \mathbf{H}_3 are distributed identically to the simulated proofs generated by the simulator \mathcal{S}_1 in \mathbf{H}_2 . \square

Hybrid \mathbf{H}_4 and core adversaries. This is a modification of \mathbf{H}_3 in which we provide the adversary with a slightly different random oracle interface and in which the adversary loses if in its forgery there is a commitment c_i^* and associated representation vector such that $\mathbf{f}_{c_i^*} \neq \mathbf{0}$. The random oracle is changed only in the sense that each input queried to the random oracle in \mathbf{H}_4 may contain more elements than the same query in \mathbf{H}_3 . We call an adversary that satisfies this additional requirement of \mathbf{H}_4 , i.e., $\mathbf{f}_{c_i^*} = \mathbf{0}$ for all i , a *core* adversary.

To continue the proof, we show how to perfectly simulate any adversary \mathcal{A} playing in \mathbf{H}_3 using a core adversary \mathcal{A}_c that runs in hybrid \mathbf{H}_4 .

Lemma 20. For any \mathcal{A} there exists \mathcal{A}_c such that $\Pr[\mathbf{H}_4(\mathcal{A}_c)] = \Pr[\mathbf{H}_3(\mathcal{A})]$.

Proof. We define \mathcal{A}_c as follows. Whenever the adversary \mathcal{A} outputs a group element \mathbf{g} it provides a representation vector $\mathbf{r}_g := \mathbf{f}_g \parallel \mathbf{c}_g \parallel \mathbf{o}_g$ for \mathbf{g} such that:

$$\mathbf{g} = \langle \mathbf{f}_g, \mathbf{ek} \rangle + \langle \mathbf{c}_g, \mathbf{coms} \rangle + \langle \mathbf{o}_g, \mathbf{proofs} \rangle.$$

The adversary \mathcal{A}_c runs internally \mathcal{A} and forwards all the queries and answers from \mathcal{A} to its simulation oracle. However, the way of simulating RO queries must ensure to not alter the result of the extractor policy, i.e. the “hash-check” for x^* . For this reason, we cannot simply forward the queries of \mathcal{A} to the random oracle. Therefore, we keep track of the queries made by \mathcal{A} in the list $\mathcal{Q}_{\text{RO}, \mathcal{A}}$ and the list of queries made by the core adversary in \mathcal{Q}_{RO} . More in detail, when \mathcal{A} queries the RO with (s, \mathbf{aux}) , the adversary \mathcal{A}_c makes a “core” RO query (s_c, \mathbf{aux}_c) such that:

1. Let s be parsed as $(\mathbf{g}_i)_{i \in [k]}$ (the group elements in s whose representations $\mathbf{r}_{\mathbf{g}_i} := \mathbf{f}_{\mathbf{g}_i} \parallel \mathbf{c}_{\mathbf{g}_i} \parallel \mathbf{o}_{\mathbf{g}_i}$ are in \mathbf{aux}) and a string \tilde{s} . Notice, since the adversary is algebraic we can un-ambiguously parse s as such.
2. For each i , \mathcal{A}_c computes the group elements $\mathbf{g}'_i = \mathbf{g}_i - \langle \mathbf{f}_{\mathbf{g}_i}, \mathbf{ek} \rangle$. \mathcal{A}_c encodes into the string s' the group elements $(\mathbf{g}_i, \mathbf{g}'_i)_{i \in [k]}$.
3. \mathcal{A}_c queries its RO with (s_c, \mathbf{aux}_c) , where $s_c := s' \parallel \tilde{s}$, and \mathbf{aux}_c contains the representations of all the group elements in s' and the same function h encoded in \mathbf{aux} . Finally, it forwards the output to \mathcal{A} , i.e. it adds (s, \mathbf{aux}, a) to $\mathcal{Q}_{\text{RO}, \mathcal{A}}$, and adds (s, s_c) to (the initially empty) \mathcal{Q}_s .

Eventually, the adversary \mathcal{A} outputs as forgery a string s and the tuple $(x', (c'_i, y'_i)_i, \pi')$, together with representation vectors $(\mathbf{r}_{c'_i})_i$ and $\mathbf{r}_{\pi'}$. For every i , let $f_i(X) := \langle \mathbf{f}_{c'_i}, (1, X, \dots, X^d) \rangle$, $y_i := f_i(x')$, and $q_i(X)$ be such that $q_i(X)(X - x') = f_i(X) - y_i$. Let \mathbf{f}_{q_i} be the vector of the coefficients of $q_i(X)$, namely $q_i(X) := \langle \mathbf{f}_{q_i}, (1, X, \dots, X^d) \rangle$. The core adversary \mathcal{A}_c returns for its forgery the string s_c such that $(s, s_c) \in \mathcal{Q}_s$, and the tuple $(x^*, (c_i^*, y_i^*)_i, \pi^*)$, where:

$$\begin{aligned} c_i^* &\leftarrow c'_i - \text{Com}(\text{ck}, f_i(X)) = c'_i - [f_i(s)p(s)]_1 \\ x^* &\leftarrow x' \\ y_i^* &\leftarrow y'_i - y_i = y'_i - f_i(x') \\ \pi^* &\leftarrow \pi' - \sum_i \rho^{i-1} \text{Com}(\text{ck}, q_i(X)) = \pi' - \left[p(s) \sum_i \rho^{i-1} q_i(s) \right]_1 \end{aligned}$$

where ρ is the value obtained by \mathcal{A} as the result of the random oracle query on $\text{batch}\|(x', (c_i, y_i)_i)$. \mathcal{A}_c inserts into aux_Φ the (correct) algebraic representations $(\mathbf{0} \| \mathbf{c}_{c'_i} \| \mathbf{o}_{c'_i})$ for each c'_i and $((\mathbf{f}_{\pi'} - \sum_i \rho^{i-1} \mathbf{f}_{q_i}) \| \mathbf{c}_{\pi'} \| \mathbf{o}_{\pi'})$ for π^* .

First, by construction, it is easy to verify that \mathcal{A}_c is algebraic if so is \mathcal{A} . Thus we need to show that the forgery of \mathcal{A} is valid if and only if the forgery of \mathcal{A}_c is valid too. By construction, we have:

$$c_i^* := c'_i - [f_i(s)p(s)]_1, \quad \pi^* := \pi' - \left[p(s) \sum_i \rho^{i-1} q_i(s) \right]_1, \quad y'_i := y_i^* - f_i(x^*).$$

By the verification equation of the forgery of \mathcal{A}_c we have:

$$\begin{aligned} &e\left(\sum_i \rho^{i-1} c_i^* - \left[\sum_i \rho^{i-1} y_i^* \right]_1, [1]_2\right) - e(\pi^*, [s - x^*]_2) = \\ &e\left(\sum_i \rho^{i-1} c'_i - \left[\sum_i \rho^{i-1} f_i(s)p(s) \right]_1 - [y'_i - f_i(x')]_1, [1]_2\right) - e\left(\pi' - \left[\sum_i \rho^{i-1} q_i(s)p(s) \right]_1, [s - x^*]_2\right) = \\ &e\left(\sum_i \rho^{i-1} c'_i - \left[\sum_i \rho^{i-1} y'_i \right]_1, [1]_2\right) - e(\pi', [s - x']_2) - [p(s) \sum_i \rho^{i-1} (f_i(s) - f_i(x') - q_i(s)(s - x'))]_T = \\ &e\left(\sum_i \rho^{i-1} c'_i - \left[\sum_i \rho^{i-1} y'_i \right]_1, [1]_2\right) - e(\pi', [s - x']_2) \end{aligned}$$

where the last equation holds because for every i we have $q_i(X)(X - x') = (f_i(X) - f_i(x'))$ and $x^* = x'$ by definition.

Finally, notice that a forgery is valid for \mathcal{A} if it provides a string s that satisfies the ‘‘hash check’’ of Φ_{ext} . We have that there exist s , aux , a , and $h(X)$ such that: (i) $g_c((c_i^*)_i, s) = 1$, (ii) $g_h(h, \text{aux}) = 1$, (iii) $(s, \text{aux}, a) \in \mathcal{Q}_{\text{RO}, \mathcal{A}}$, and (iv) $x^* = h(a)$ for the forgery of \mathcal{A} .

The way \mathcal{A}_c simulates the RO queries ensures that for the query s of \mathcal{A} to the RO, the core adversary sent the ‘‘core’’ RO query s_c that encodes both vectors

$(c'_i)_i$ and $(c_i^*)_i$, thus we have that (i) $g_c((c_i^*)_i, s_c) = 1$, (ii) $g_h(h, \text{aux}_c) = 1$, (iii) $(s_c, \text{aux}_c, a) \in \mathcal{Q}_{\text{RO}}$, and (iv) $x^* = h(a)$ for the forgery of \mathcal{A}_c . \square

Notice that if we run the canonical extractor on the outputs of a core adversary \mathcal{A}_c in \mathbf{H}_4 , the extractor sets the extracted witness to be the zero polynomials.

Hybrid \mathbf{H}_5 . The hybrid \mathbf{H}_5 additionally checks that $\mathbf{f}_{\pi^*} \neq \mathbf{0} \vee \mathbf{c}_{\pi^*} \neq \mathbf{0}$, and if the condition holds the adversary \mathcal{A}_c loses the game.

Lemma 21. *For any \mathcal{A}_c it holds $\Pr[\mathbf{H}_5(\mathcal{A}_c)] \leq \Pr[\mathbf{H}_4(\mathcal{A}_c)] + \epsilon_{(Q_x+d+1)\text{-DL}}$*

Proof. Recall that from the definition of the experiment, upon the ℓ -th query (\mathbb{x}, aux) from \mathcal{A}_c to the simulation oracle \mathcal{S}_1 of the form $\mathbb{x} = (x_j, \mathbf{c}, y)$ and $\text{aux} = \mathbf{r}_c = (\mathbf{0} \parallel \mathbf{c}_c \parallel \mathbf{0})$ such that $\mathbf{c} = \langle \mathbf{c}_c, \text{coms} \rangle$, the adversary receives the proof $[\pi_{\mathbb{x}}(s)]_1$ where:

$$\pi_\ell(X) := (\langle \mathbf{c}_c, \boldsymbol{\gamma} \rangle - y)p_j(X).$$

Given the forgery, consider the following polynomials:

$$\begin{aligned} c_i^*(X) &:= \langle \mathbf{c}_{c_i^*}, \boldsymbol{\gamma} \rangle p(X) + \sum_{\ell=1}^{|\mathcal{Q}_{\text{sim}}|} o_{c_i^*, \ell} \cdot \pi_\ell(X) \\ \pi^*(X) &:= \langle \mathbf{f}_{\pi^*}, (1, X, \dots, X^d) \rangle p(X) + \langle \mathbf{c}_{\pi^*}, \boldsymbol{\gamma} \rangle p(X) + \sum_{\ell=1}^{|\mathcal{Q}_{\text{sim}}|} o_{\pi^*, \ell} \cdot \pi_\ell(X) \\ v(X) &:= \sum_i \rho^{i-1} (c_i^*(X) - y_i^* p(X)) - (X - x^*) \pi^*(X) \end{aligned}$$

By the guarantees of the AGM, we have $c_i^* = [c_i^*(s)]_1$ and $\pi^* = [\pi^*(s)]_1$, moreover, if the verification equation is satisfied by the forgery of \mathcal{A}_c , then $v(s) = 0$.

Next, we show that when the forgery of the adversary is valid the probability of $\mathbf{f}_{\pi^*} \neq \mathbf{0}$ or $\mathbf{c}_{\pi^*} \neq \mathbf{0}$ is bounded by $\epsilon_{(Q_x+d+1)\text{-DL}}$.

First, notice that if the verification equation for \mathcal{A}_c holds then polynomial $v(X)$ must be equivalent to the zero polynomial with overwhelming probability. In fact, if $v(X) \not\equiv 0$ and $v(s) = 0$ then, by Lemma 1, we can reduce \mathcal{A}_c to an adversary against the $(Q_x + d + 1)$ -DL assumption. To conclude the proof we show that, whenever $\mathbf{f}_{\pi^*} \neq \mathbf{0}$ or $\mathbf{c}_{\pi^*} \neq \mathbf{0}$, we have $v(X) \not\equiv 0$.

Assume by contradiction that:

$$\sum_i \rho^{i-1} (c_i^*(X) - y_i^* p(X)) - (X - x^*) \pi^*(X) = v(X) = 0. \quad (12)$$

For this equation to hold, it must be the case that the degree of $(X - x^*) \pi^*(X)$ is less than or equal to the degree of $\sum_i \rho^{i-1} (c_i^*(X) - y_i^* p(X))$. However, by the definition of $\pi^*(X)$ above, this cannot occur if $\mathbf{f}_{\pi^*} \neq \mathbf{0}$ or $\mathbf{c}_{\pi^*} \neq \mathbf{0}$ as in such a case the degree of $(X - x^*) \pi^*(X)$ would be $> Q_x + 1$ while the degree of $\sum_i \rho^{i-1} (c_i^*(X) - y_i^* p(X))$ is $\leq Q_x + 1$. Therefore we can conclude that the forged proof $\pi^*(s)$ is a linear combination of the simulated proofs only, i.e., both \mathbf{f}_{π^*} and \mathbf{c}_{π^*} are null. \square

The representation of the commitments $(\mathbf{c}_i^*)_i$ and the proof π^* computed by the adversary in the forgery (possibly) depends on the simulated proofs **proofs** returned by \mathcal{S}_1 . This dependence is characterized by the vector of coefficients \mathbf{o}_τ for $\tau \in \{\mathbf{c}_i^*\}_i \cup \{\pi^*\}$. However, to proceed with the proof it is more convenient to obtain an equivalent representation that depends on the polynomials $p_j(X)$. This motivates the definition of our next hybrid.

Hybrid \mathbf{H}_6 . The hybrid \mathbf{H}_6 finds coefficients \mathbf{o}'_τ , for $\tau \in \{\mathbf{c}_i^*\}_i \cup \{\pi^*\}$ such that:

$$\langle \mathbf{o}_\tau, \mathbf{proofs} \rangle = \langle \mathbf{o}'_\tau, ([p_j(s)]_1)_j \rangle. \quad (13)$$

Moreover, if there is an index k such that $\mathbf{o}_{\mathbf{c}_k^*} \neq \mathbf{0}$ but $\mathbf{o}'_{\mathbf{c}_k^*} = \mathbf{0}$ the adversary loses the game.

Lemma 22. *For any PPT \mathcal{A}_c we have $\Pr[\mathbf{H}_6(\mathcal{A}_c)] \leq \Pr[\mathbf{H}_5(\mathcal{A}_c)] + \epsilon_{\text{Aff-MDH}}$*

Proof. We begin by showing that the hybrid can compute such alternative representations efficiently. We proceed in steps.

Let us parse the simulated proofs $\mathbf{proofs} := (\pi_{j,\ell})_{j,\ell}$ such that $\pi_{j,\ell}$ is the ℓ -th simulated proof obtained by \mathcal{S}_1 on a query involving the j -th point x_j , i.e., $((x_j, \hat{\mathbf{c}}_{j,\ell}, y_{j,\ell}), \mathbf{aux}_{j,\ell})$. Also, let $\mathbf{c}_{j,\ell}$ be the algebraic representation for the group element $\hat{\mathbf{c}}_{j,\ell}$ in $\mathbf{aux}_{j,\ell}$. For every $j \in [Q_x]$, we define $\mathbf{\Gamma}_j$ as the $Q_c \times Q_c$ matrix whose ℓ -th column is $\mathbf{c}_{j,\ell}$.

By construction of \mathcal{S}_1 in this hybrid we have that for every $j \in [Q_x]$ it holds

$$\pi_{j,\ell} := [(\mathbf{c}_{j,\ell}^\top \cdot \boldsymbol{\gamma} - y_{j,\ell}) p_j(s)]_1$$

and thus $\boldsymbol{\pi}_j := [(\mathbf{\Gamma}_j^\top \boldsymbol{\gamma} - \mathbf{y}_j) p_j(s)]_1$ with $\mathbf{y}_j := (y_{j,\ell})_\ell$.

Without loss of generality, we assume that for each x_j the adversary makes the maximum number of simulation queries (i.e., $\ell \in [Q_c]$); therefore $\mathbf{\Gamma}_j$ is a full rank matrix (this follows from the fact that the simulation queries of the adversary satisfy the policy Φ_{sim} , and in particular the *consistent opening checks* of the policy, see Item 3).

Given any vector \mathbf{o}_τ with $\tau \in \{\mathbf{c}_i^*\}_i \cup \{\pi^*\}$, its entry m -th entry $o_{\tau,m}$ corresponds to the m -th simulated proof in **proofs**. Therefore, similarly to above, we denote by $o_{\tau,j,\ell}$ the entry corresponding to proof $\pi_{j,\ell}$ and we define $\mathbf{o}_{\tau,j} := (o_{\tau,j,\ell})_\ell$.

Then, for every $j \in [Q_x]$ we define

$$\mathbf{o}'_{\tau,j} \leftarrow \mathbf{\Gamma}_j \cdot \mathbf{o}_{\tau,j} \quad (14)$$

$$\boldsymbol{\pi}'_j \leftarrow (\mathbf{\Gamma}_j^\top)^{-1} \cdot \boldsymbol{\pi}_j \quad (15)$$

from which we derive that for any j :

$$\begin{aligned} \sum_j \langle \mathbf{o}'_{\tau,j}, \boldsymbol{\pi}'_j \rangle &= \sum_j \langle \boldsymbol{\Gamma}_j \cdot \mathbf{o}_{\tau,j}, (\boldsymbol{\Gamma}_j^\top)^{-1} \cdot \boldsymbol{\pi}_j \rangle \\ &= \sum_j \mathbf{o}_{\tau,j}^\top \boldsymbol{\Gamma}_j^\top (\boldsymbol{\Gamma}_j^\top)^{-1} \cdot \boldsymbol{\pi}_j \\ &= \sum_j \langle \mathbf{o}_{\tau,j}, \boldsymbol{\pi}_j \rangle. \end{aligned}$$

Note that the value above is equal to $\langle \mathbf{o}_\tau, \text{proofs} \rangle$, up to a permutation of the indices j .

For all $j \in [Q_x]$ let $\mathbf{z}_j := (\boldsymbol{\Gamma}_j^\top)^{-1} \cdot \mathbf{y}_j$, and note that

$$\boldsymbol{\pi}'_j = [(\boldsymbol{\gamma} - \mathbf{z}_j)p_j(s)]_1$$

namely $\boldsymbol{\pi}'_{j,i}$ is a valid proof for the instance $(\mathbf{c}_i, x_j, z_{j,i})$ w.r.t. the simulated SRS.

The hybrid computes $o''_{\tau,j} \leftarrow \langle \mathbf{o}'_{\tau,j}, (\boldsymbol{\gamma} - \mathbf{z}_j) \rangle$ and sets $\mathbf{o}''_\tau := (o''_{\tau,j})_{j \in [Q_x]}$. By construction, it holds:

$$\sum_{j \in [Q_x]} \langle \mathbf{o}'_{\tau,j}, \boldsymbol{\pi}'_j \rangle = \sum_{j \in [Q_x]} o''_{\tau,j} \cdot [p_j(s)]_1.$$

which proves the first part of the lemma, i.e., computing $\mathbf{o}''_{\tau,j}$ satisfying Eq. (13).

In what follows, we prove that if the event that \mathbf{H}_6 outputs 0 but \mathbf{H}_5 would output 1, namely that all the conditions of \mathbf{H}_5 hold but there is an index k such that $\mathbf{o}_{\mathbf{c}_k^*} \neq \mathbf{0} \wedge \mathbf{o}''_{\mathbf{c}_k^*} = \mathbf{0}$, then we can break the Aff-MDH assumption.

First, notice that for any j and k $\mathbf{o}_{\mathbf{c}_k^*,j} \neq \mathbf{0}$ implies that $\mathbf{o}'_{\mathbf{c}_k^*,j} \neq \mathbf{0}$, because the linear transformation applied to compute $\mathbf{o}'_{\mathbf{c}_k^*,j}$ is full rank.

Second, take an index j^* such that $\mathbf{o}_{\mathbf{c}_k^*,j^*} \neq \mathbf{0}$ and set $\mathbf{A} \leftarrow \mathbf{o}'_{\mathbf{c}_k^*,j^*}$ and $\zeta \leftarrow \langle \mathbf{z}_{j^*}, \mathbf{o}'_{\mathbf{c}_k^*,j^*} \rangle$.

By the above definition of the values $o''_{\mathbf{c}_k^*,j^*}$ and our assumption that the “bad event” of this hybrid is $\mathbf{o}''_{\mathbf{c}_k^*} = \mathbf{0}$, we have that:

$$\langle \mathbf{A}, [\boldsymbol{\gamma}]_1 \rangle = \left[\underbrace{\langle \mathbf{o}'_{\mathbf{c}_k^*,j^*}, (\boldsymbol{\gamma} - \mathbf{z}_{j^*}) \rangle}_{o''_{\mathbf{c}_k^*,j^*} = 0} + \underbrace{\langle \mathbf{o}'_{\mathbf{c}_k^*,j^*}, \mathbf{z}_{j^*} \rangle}_{\zeta} \right]_1 = [\zeta]_1.$$

The reduction \mathcal{B} to the \mathcal{D} -Aff-MDH Assumption takes as input a distribution $[\boldsymbol{\gamma}]_1$ and runs the experiment as in \mathbf{H}_5 (it perfectly emulates \mathbf{H}_5 , and in particular the simulation oracle, because it knows the trapdoor s “at the exponent”). Then \mathcal{B} computes the coefficients $(A_i)_{i \in [Q_c]}$ and the value ζ as described above, which is a valid \mathcal{D} -Aff-MDH solution. \square

Hybrid \mathbf{H}_7 . This hybrid proceeds as the previous one except that we abort in the case the representation vectors of the commitments in the forgery are different from the representations provided for the same commitments when querying

the batching random oracle. More precisely, let \mathbf{x}^* be the statement returned by the adversary at the end of its execution and along with the representation vectors $\{\mathbf{r}_{c_i^*}\}_i$ of the commitments $(c_i^*)_i$. Look for the tuple $(\text{batch} \parallel \mathbf{x}^*, \text{aux}, a) \in \mathcal{Q}_{\text{RO}}$ (this tuple must exist w.l.o.g. since otherwise we can always define another adversary that makes this query at the very last) and let $\tilde{\mathbf{r}}_{c_i^*}$ be the representation vectors of the commitments in aux . If there exists an index i such that $\mathbf{r}_{c_i^*} \neq \tilde{\mathbf{r}}_{c_i^*}$ the adversary loses the game.

Lemma 23. *For any PPT \mathcal{A}_c , $\Pr[\mathbf{H}_7(\mathcal{A}_c)] \leq \Pr[\mathbf{H}_6(\mathcal{A}_c)] + \epsilon_{(Q_x+d+1)\text{-DL}}(\lambda) + \epsilon_{\text{Aff-MDH}}$.*

Proof. Let c^* be a commitment of the forgery with two different representations $\mathbf{r}_{c^*} \neq \tilde{\mathbf{r}}_{c^*}$. By definition of core adversary we have that $\mathbf{f}_{c^*} = \tilde{\mathbf{f}}_{c^*} = \mathbf{0}$. Thus we have

$$\langle (\mathbf{c}_{c^*} - \tilde{\mathbf{c}}_{c^*}), \gamma p(s) \rangle + \langle (\mathbf{o}_{c^*} - \tilde{\mathbf{o}}_{c^*}), (p_j(s))_j \rangle = \langle \mathbf{\Delta}_c, \gamma p(s) \rangle + \langle \mathbf{\Delta}_0, (p_j(s))_j \rangle = 0$$

Similarly to previous cases, under the $(Q_x + d + 1)$ -DL assumption we can write the equation above symbolically as

$$\langle \mathbf{\Delta}_c, \gamma \rangle p(X) + \langle \mathbf{\Delta}_0, (p_j(X))_j \rangle = 0$$

which, by looking at the degree of $p(X)$ and $p_j(X)$ and the linear independence of the latter, implies that $\langle \mathbf{\Delta}_c, \gamma \rangle = 0$ and $\mathbf{\Delta}_0 = \mathbf{0}$. However, the event that $\langle \mathbf{\Delta}_c, \gamma \rangle = 0$ for a nonzero $\mathbf{\Delta}_c$ can be reduced to the Aff-MDH assumption.

Hybrid \mathbf{H}_8 . The hybrid \mathbf{H}_8 additionally checks if there exists an index $k \in [m]$ such that $\mathbf{r}_{c_k^*} \neq \mathbf{0}$, and if the condition holds the adversary \mathcal{A}_c loses the game.

Lemma 24. *For any PPT \mathcal{A}_c ,*

$$\Pr[\mathbf{H}_8(\mathcal{A}_c)] \leq \Pr[\mathbf{H}_7(\mathcal{A}_c)] + \epsilon_{\text{Aff-MDH}} + 2\epsilon_{(Q_x+1+d)\text{-DL}} + \text{poly}(\lambda) \frac{\text{deg}(h)}{q}.$$

Proof. We bound the probability that the adversary loses in \mathbf{H}_8 but not in \mathbf{H}_7 , namely, the probability that $\exists k : \mathbf{r}_{c_k^*} \neq \mathbf{0}$ but the conditions of \mathbf{H}_7 hold. The main idea is to do a reduction \mathcal{B} to the Aff-MDH when this event happens.

First of all, notice that by the changes introduced earlier we can assume that the core adversary outputs commitments $(c_\ell^*)_l$ and proof π whose representations are such that $\mathbf{f}_{c_\ell^*} = \mathbf{f}_{\pi^*} = \mathbf{c}_{\pi^*} = \mathbf{0}$, i.e. the adversary only makes use of previous commitments in coms and simulated proofs in proofs to represent each c_ℓ^* , and only uses the simulated proofs to represent the proof π^* .

The reduction \mathcal{B} takes as input a distribution $[\gamma]_1$ and runs the experiment as in \mathbf{H}_7 . \mathcal{B} aborts if the forgery $\mathbf{x}^* = (x^*, (c_\ell^*, y_\ell^*)_{\ell \in [m]}, \pi^*)$ returned by the adversary is not valid (i.e. either the extraction predicate or the verification equation is not satisfied) or if $\mathbf{r}_{c_\ell^*} = \mathbf{0}$ for all ℓ . Otherwise, we have that:

$$e\left(\sum_{\ell} \rho^{\ell-1} c_\ell^* - \left[\sum_i \rho^{\ell-1} y_\ell^*\right]_1, [1]_2\right) = e(\pi^*, [s - x^*]_2) \text{ and } \exists k : \mathbf{r}_{c_k^*} \neq \mathbf{0}$$

where $\mathbf{r}_{\mathbf{c}_k^*} \neq \mathbf{0}$ if $\mathbf{o}_{\mathbf{c}_k^*} \neq \mathbf{0} \vee \mathbf{c}_{\mathbf{c}_k^*} \neq \mathbf{0}$.

We can rewrite each commitment and the proof in the forgery as functions of the coefficients $\mathbf{o}_{\mathbf{c}_\ell^*}''$ and \mathbf{o}_{π^*}'' (as computed in the previous hybrid):

$$\begin{aligned}\mathbf{c}_\ell^* &:= \left[\langle \mathbf{c}_{\mathbf{c}_\ell^*}, \gamma p(s) \rangle + \langle \mathbf{o}_{\mathbf{c}_\ell^*}'', (p_j(s))_j \rangle \right]_1 \\ \pi^* &:= \left[\langle \mathbf{o}_{\pi^*}'', (p_j(s))_j \rangle \right]_1\end{aligned}$$

Since the verification equation is satisfied, and plugging in the AGM representations we have:

$$\begin{aligned}\left\langle \sum_{\ell} \rho^{\ell-1} \cdot \mathbf{c}_{\mathbf{c}_\ell^*}, \gamma p(s) \right\rangle + \left\langle \sum_{\ell} \rho^{\ell-1} \cdot \mathbf{o}_{\mathbf{c}_\ell^*}'', (p_j(s))_j \right\rangle - p(s) \sum_{\ell} \rho^{\ell-1} \cdot y_\ell^* \\ = \langle \mathbf{o}_{\pi^*}'', (p_j(s))_j \rangle (s - x^*)\end{aligned}\quad (16)$$

For all $j \in [Q_x]$, let $\delta_j := x_j - x^* \neq 0$. We can rewrite the r.h.s. of Eq. (16) as:

$$\begin{aligned}\langle \mathbf{o}_{\pi^*}'', (p_j(s)(s - x^*))_j \rangle &= \langle \mathbf{o}_{\pi^*}'', (p_j(s)((s - x_j) + \delta_j))_j \rangle \\ &= \langle \mathbf{o}_{\pi^*}'', (p(s) + p_j(s)\delta_j)_j \rangle\end{aligned}$$

In Eq. (16), we group all the terms that depend on $p(s)$ on the left side, and we move all the terms that depend on $p_j(s)$ on the right side, thus obtaining:

$$H \cdot p(s) = \sum_{j \in [Q_x]} L_j \cdot p_j(s)\quad (17)$$

where

$$\begin{aligned}H &:= \left\langle \sum_{\ell} \rho^{\ell-1} \cdot \mathbf{c}_{\mathbf{c}_\ell^*}, \gamma \right\rangle - \sum_j o_{\pi^*,j}'' - \sum_{\ell} \rho^{\ell-1} \cdot y_\ell^* \\ L_j &:= o_{\pi^*,j}'' \delta_j - \sum_{\ell} \rho^{\ell-1} o_{\mathbf{c}_\ell^*,j}''.\end{aligned}$$

Let $v(X) := Hp(X) - \sum_{j \in [Q_x]} L_j p_j(X)$. Notice that because of Eq. (17) we have $v(s) = 0$, thus we can assume $v(X) \equiv 0$, as otherwise we can reduce, by Lemma 1, to the $(Q_x + d + 1)$ -DL assumption. It must be the case that both $\sum_{j \in [Q_x]} L_j p_j(X) = 0$ and $H = 0$ because the degrees of $p(X)$ and $p_j(X)$, for all j , are different. Moreover, the polynomials $p_j(X)$ are linearly independent, namely the only linear combination $\sum_j a_j p_j(X) = 0$ is the trivial one where the coefficients $a_j = 0$ ¹⁴, thus $L_j = 0$ for all j . From this, we have that:

$$\forall j : o_{\pi^*,j}'' = \frac{-\sum_{\ell} \rho^{\ell-1} o_{\mathbf{c}_\ell^*,j}''}{\delta_j}.$$

¹⁴ To see this, $\forall x_j \in \mathcal{Q}_x$ we have that $\sum_{j'} a_{j'} p_{j'}(x_j) = a_j p_j(x_j)$ since $p_j(x_j) \neq 0$ and $p_{j'}(x_j) = 0$ for $j \neq j'$, and $a_j p_j(x_j) = 0$ iff $a_j = 0$

Since H must be 0, we have that:

$$\sum_{\ell} \rho^{\ell-1} \sum_{i \in [Q_c]} c_{c_{\ell}^*, i} \gamma_i - \sum_{j \in [Q_x]} \frac{\sum_{\ell} \rho^{\ell-1} o_{c_{\ell}^*, j}''}{\delta_j} - \sum_{\ell} \rho^{\ell-1} \cdot y_{\ell}^* = 0. \quad (18)$$

\mathcal{B} can derive from Eq. (18) that:

$$\begin{aligned} 0 &= \sum_{i \in [Q_c]} \sum_{\ell} \rho^{\ell-1} c_{c_{\ell}^*, i} \gamma_i - \sum_{j \in [Q_x]} \frac{\sum_{\ell} \rho^{\ell-1} o_{c_{\ell}^*, j}''}{\delta_j} - \sum_{\ell} \rho^{\ell-1} y_{\ell}^* \\ &= \sum_{i \in [Q_c]} \sum_{\ell} \rho^{\ell-1} c_{c_{\ell}^*, i} \gamma_i - \sum_{i, j, \ell} \frac{\rho^{\ell-1} o_{c_{\ell}^*, j, i}'(\gamma_i - z_{j, i})}{\delta_j} - \sum_{\ell} \rho^{\ell-1} y_{\ell}^* \\ &= \sum_{i \in [Q_c]} \underbrace{\left(\sum_{\ell} \rho^{\ell-1} \underbrace{\left(c_{c_{\ell}^*, i} - \sum_j \frac{o_{c_{\ell}^*, j, i}'}{\delta_j} \right)}_{A_{i, \ell}} \right)}_{A_i} \gamma_i + \underbrace{\sum_{i, j, \ell} \frac{\rho^{\ell-1} o_{c_{\ell}^*, i, j} z_{j, i}}{\delta_j}}_{-z} - \sum_{\ell} \rho^{\ell-1} y_{\ell}^* \end{aligned}$$

Above, the second equation comes directly from the definition of the coefficients $o_{\tau, j}''$, and in the last step we have grouped the terms depending on γ_i . In particular, the last equation shows that \mathcal{B} can make a forgery in the Aff-MDH game since it knows z and all the coefficients A_i such that:

$$\sum_{i \in [Q_c]} A_i [\gamma_i]_1 = [z]_1.$$

For this to be a valid solution in the Aff-MDH game we need the existence of at least an index i such that $A_i \neq 0$. We show that this occurs with all but negligible probability, i.e., there is a negligible value ν such that $\Pr[\exists i \in [Q_c] : A_i \neq 0] \geq 1 - \nu$.

To this end, consider an arbitrary $\mu \in [Q_c]$, then we have $\Pr[\forall i \in [Q_c] : A_i = 0] \leq \Pr[A_{\mu} = 0]$. Thus, for any μ , we have

$$\Pr[\exists i \in [Q_c] : A_i \neq 0] = 1 - \Pr[\forall i \in [Q_c] : A_i = 0] \geq 1 - \Pr[A_{\mu} = 0].$$

Next, let k be the index such that $\mathbf{r}_{c_k^*} \neq 0$. We have

$$\Pr[A_{\mu} = 0] \leq \Pr[A_{\mu} = 0 | A_{\mu, k} \neq 0] + \Pr[A_{\mu, k} = 0]$$

Below, we argue that $\Pr[A_{\mu} = 0 | A_{\mu, k} \neq 0]$ is negligible based on the randomness of ρ , which can be shown to be chosen by the random oracle after all the coefficients $A_{\mu, \ell}$ are defined, and that $\Pr[A_{\mu, k} = 0]$ is negligible based on the randomness of x^* , which also in this case is chosen by the random oracle after defining $A_{\mu, k}$. Furthermore, for the latter, we make use of the assumption that $\mathbf{r}_{c_k^*} \neq 0$.

For every ℓ , we claim that the values $A_{\mu,\ell} = c_{\mathbf{c}_\ell^*,\mu} - \sum_j \frac{o'_{\mathbf{c}_\ell^*,j,\mu}}{(x^* - x_j)}$ can be fixed before the random oracle query $\text{batch}||_{\mathbf{x}^*}$ is made. To this end, we start by showing that that $\mathbf{o}'_{\mathbf{c}_\ell^*,j}$ does not depend on x^* . Let $B(j) \subseteq [Q_c]$ be the subset of indices of the simulation queries that involve x_j and that occurred before the random oracle query that returned x^* . We observe that for every $\eta \in B(j)$ it must be $o_{\mathbf{c}_\ell^*,j,\eta} = 0$ since the simulated proof $\pi_{j,\eta}$ is not in the view of the adversary. Therefore, we have

$$o'_{\mathbf{c}_\ell^*,j,i} = \sum_{\eta \in [Q_c]} \Gamma_{j,\eta,i} \cdot o_{\mathbf{c}_\ell^*,j,\eta} = \sum_{\eta \in B(j)} \Gamma_{j,\eta,i} \cdot o_{\mathbf{c}_\ell^*,j,\eta}$$

and observe that all the rows of Γ_j belonging to $B(j)$ can all be defined before x^* is sampled.

Hence, we have that each $A_{\mu,\ell}$ depends on the values $\mathbf{c}_{\mathbf{c}_\ell^*}, x^*, \{x_j\}_j$, and $\mathbf{o}_{\mathbf{c}_\ell^*,j}$ which can all be defined before the random oracle query $\text{batch}||_{\mathbf{x}^*}$ is made.

Therefore, assuming the existence of at least a nonzero $A_{\mu,k} \neq 0$ we have that $A_\mu = \sum_{\ell=1}^m \rho^{\ell-1} A_{\mu,\ell} = 0$ is at most m/q over the random choice of ρ .

Next, we bound $\Pr[A_{\mu,k} = 0]$.

Recall that, since the extractor policy Φ_{ext} holds true, we have that $x^* = h(a)$ and $(s, \text{aux}, a) \in \mathcal{Q}_{\text{RO}}$ where $= g_c(\mathbf{c}_k^*, s) = 1$ and the function h is the polynomial encoded in aux_ϕ . Moreover, by the AGM, since \mathcal{A}_c sends a query s (where \mathbf{c}_k^* is encoded in s) to the random oracle it also defines coefficients for \mathbf{c}_k^* before the value a , and therefore $x^* = h(a)$, is defined. Also, it is not hard to see that the representation vector of \mathbf{c}_k^* defined by \mathcal{A}_c when querying the random oracle must be the same representation vector used for the forgery. As otherwise we would break the $(Q_x + d + 1)$ -DL assumption. Thus the coefficients $\mathbf{c}_{\mathbf{c}_k^*}$ and $\mathbf{o}'_{\mathbf{c}_k^*,j}$ are defined by the adversary before seeing the random value x^* .

Notice that, once the coefficients $\mathbf{c}_{\mathbf{c}_k^*}$ and $\mathbf{o}'_{\mathbf{c}_k^*,j}$ are fixed, the coefficient $A_{\mu,k}$ can be seen as function of $x^* \in \mathbb{Z}_q$, i.e. $A_{\mu,k} = A_{\mu,k}(x^*)$, where

$$\begin{aligned} A_{\mu,k}(X) &:= c_{\mathbf{c}_k^*,\mu} + \sum_j \frac{o'_{\mathbf{c}_k^*,j,\mu}}{X - x_j} \\ &= \frac{c_{\mathbf{c}_k^*,\mu} \prod_j (X - x_j) + \sum_j (o'_{\mathbf{c}_k^*,j,\mu} \prod_{j' \neq j} (x_{j'} - X))}{\prod_j (X - x_j)}. \end{aligned}$$

Notice that $A_{\mu,k}(X)(\prod_j (X - x_j))$ vanishes in at most Q_x points in $\mathbb{F} \setminus \mathcal{Q}_x$ and vanishes in the set of points \mathcal{Q}_x . Let \mathcal{R} be the set of the roots of such a polynomial, since h is defined before x^* is computed, and by union bound:

$$\Pr[h(\text{RO}(s)) \in \mathcal{R}] \leq \sum_{r \in \mathcal{R}} \Pr[h(\text{RO}(s)) = r] \leq Q_x \frac{\deg(h)}{q}$$

for each string s that encodes \mathbf{c}_k^* . To conclude, we notice that \mathcal{A} can submit at most Q_{RO} queries to the RO with strings encoding \mathbf{c}_k^* , say $s_1, \dots, s_{Q_{\text{RO}}}$. Thus the probability that there exists $i \in [Q_{\text{RO}}]$ such that $h(\text{RO}(s_i)) \in \mathcal{R}$ is bounded by $Q_{\text{RO}} Q_x \frac{\deg(h)}{q}$. \square

Hybrid \mathbf{H}_9 . This hybrid proceeds as the previous one except that the adversary loses if the forgery triggers the following event:

$$\sum_{\ell \in [m]} \rho^{\ell-1} y_\ell^* = 0 \text{ and } \exists k : y_k^* \neq 0$$

Lemma 25. *For any PPT \mathcal{A}_c we have $\Pr[\mathbf{H}_9(\mathcal{A}_c)] \leq \Pr[\mathbf{H}_8(\mathcal{A}_c)] + m/q$.*

The proof follows by observing that, since all the outputs $(y_\ell^*)_\ell$ are fixed in the input of the query $(\text{batch} \parallel \mathbb{x}^*)$ that returns ρ , we can apply the Schwartz-Zippel lemma.

Hybrid \mathbf{H}_{10} . The hybrid \mathbf{H}_{10} additionally checks if there exists k such that $y_k^* \neq 0$, and if the condition holds the adversary \mathcal{A}_c loses the game.

Lemma 26. *For any PPT \mathcal{A}_c we have $\Pr[\mathbf{H}_{10}(\mathcal{A}_c)] \leq \Pr[\mathbf{H}_9(\mathcal{A}_c)] + \epsilon_{(Q_x+1+d)\text{-DL}} + \text{poly}(\lambda) \frac{\text{deg}(d)}{q}$*

Proof. We reduce to the evaluation binding of KZG polynomial commitment for polynomials of maximum degree $Q_x + 1 + d$, which, in turn, can be reduced to $(Q_x + 1 + d)$ -strong Discrete Log assumption. Let \mathcal{B} be the reduction that upon input $\text{pp}_{\mathbb{G}}, \text{ck} = [1, s, \dots, s^{Q_x+d+2}]_1, [1, s]_2$ simulates experiment \mathbf{H}_9 for the adversary \mathcal{A}_c . Eventually, \mathcal{A}_c outputs its forgery $(x^*, (c_\ell^*, y_\ell^*)_\ell, \pi^*)$, and \mathcal{B} aborts if $y_\ell^* = 0$ for all ℓ . Else, the reduction sets $y^* \leftarrow \sum_{\ell \in [m]} \rho^{\ell-1} y_\ell^*$, $f(X) := -y^* p(X)$, $y := \tilde{f}(x^*)$, and computes π to be a valid KZG-proof for $([\tilde{f}(s)]_1, x^*, y)$. The forgery against evaluation binding of the reduction is set to be (y, π) and $(0, \pi^*)$ for the commitment $[\tilde{f}(s)]_1$ on the point x^* .

We need to show that (1) $([\tilde{f}(s)]_1, x^*, 0, \pi^*)$ passes the verification equation of KZG commitment where the commitment key is set to ck and that (2) $y \neq 0$. For the first item notice that, by the definition of core adversary, we have that $\mathbf{r}_{c_k^*} = \mathbf{0}$ thus $c^* = [0]_1$. Therefore, by the verification equation:

$$e([0]_1 - y^* [p(s)]_1, [1]_2) = e([\tilde{f}(s)]_1 - 0 [1]_1, [1]_2) = e(\pi^*, [s]_2 - x^* [1]_2).$$

For the second item, first notice that by the change in the previous hybrid we have that $y^* \neq 0$. Second, notice that $\tilde{f}(x^*) = 0$ if and only if x^* is a root of $p(X)$, i.e. $x^* \in \mathcal{Q}_x$ or $x^* = x_r$. Thus, similarly to the previous lemma, by the assumption on h and by union bound:

$$\Pr[h(\text{RO}(s)) \in \mathcal{Q}_x \cup \{x_r\}] \leq Q_{\text{RO}}(Q_x + 1) \frac{\text{deg}(d)}{q}.$$

□

Finally, we have that the probability that the adversary wins in \mathbf{H}_{10} is null. Indeed, the canonical extractor \mathcal{E} outputs the 0 polynomial, moreover because of the condition introduced in \mathbf{H}_8 , we have $\mathbf{c}^* = [0]_1$, and because of the condition introduced in \mathbf{H}_{10} we have $y^* = 0$, thus the witness extracted is valid for the instance $\mathbb{x}^* = (x^*, ([0]_1, \mathbf{0}))$. □

Finally, we prove that for any algebraic adversary \mathcal{A} whose forgery satisfies the predicate $\Phi_{\text{ext}}^{\text{der}}$, there exists an algebraic adversary \mathcal{B} whose forgery satisfies the predicate $\Phi_{\text{ext}}^{\text{nd}}$.

Lemma 27. *For any algebraic adversary \mathcal{A} there exists an algebraic adversary \mathcal{B} such that:*

$$\mathbf{Adv}_{\text{CP}_{\text{m-evil}}, \mathcal{A}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{D}}\text{-se}}(\lambda) = \mathbf{Adv}_{\text{CP}_{\text{m-evil}}, \mathcal{B}, \mathcal{S}, \mathcal{E}}^{\Phi_{\mathcal{D}}'\text{-se}}(\lambda)$$

Proof. The reduction \mathcal{B} internally runs \mathcal{A} forwarding all the queries, up to the forgery (\mathbb{x}^*, π^*) , where $\mathbb{x}^* = (x^*, (c_i^*, y_i^*)_i)$. If the simulation queries and/or the forgery of the adversary \mathcal{A} do not agree with the policy $\Phi_{\mathcal{D}}$, i.e. \mathcal{A} automatically loses its game, \mathcal{B} aborts. Otherwise, it must be true that the forgery of \mathcal{A} either (i) satisfies the extraction predicate $\Phi_{\text{ext}}^{\text{nd}}$ or (ii) satisfies the extraction predicate $\Phi_{\text{ext}}^{\text{der}}$. Both cases can be efficiently checked by \mathcal{B} . In case (i) \mathcal{B} would simply forward the forgery of \mathcal{A} retaining the same advantage of \mathcal{A} . Otherwise, before submitting the forgery, \mathcal{B} retrieves from \mathcal{Q}_{sim} the statement $\mathbb{x} = (x^*, (c_i^*, y_i)_i)$, where there exists an index i such that $y_i \neq y_i^*$. Let π_i be a proof for the “single-eval” statement $\mathbb{x}_i = (x, (c_i, y_i))$: these additional simulation queries submitted to \mathcal{S}_1 do not alter the consistency check, because they are consistent with the simulation for the statement \mathbb{x} . Also, let $\rho \leftarrow \text{RO}(\text{batch} \parallel \mathbb{x}^*)$ be the coefficient used for batch-verify the proof π^* , and let $\pi := \sum_i \rho^{i-1} \pi_i$. We can define:

$$\begin{aligned} \hat{c} &\leftarrow \pi^* - \pi \\ \hat{x} &\leftarrow h(a) \\ \hat{y} &\leftarrow \frac{\sum_i \rho^{i-1} (y_i - y_i^*)}{\hat{x} - x} \\ \hat{\pi} &\leftarrow \frac{\pi - \pi^*}{\hat{x} - x} \end{aligned}$$

where $(\hat{c}, h, a) \in \mathcal{Q}_{\text{RO}}$. We need to prove that $\hat{\mathbb{x}} := (\hat{x}, \hat{c}, \hat{y})$ and $\hat{\pi}$ satisfy the (single-eval) verification equation:

$$\hat{c} - [\hat{y}]_1 = \hat{\pi}([s - x]_2)$$

We notice that the probability that $\exists i : y_i \neq y_i^*$ but $\hat{y}_1 := \sum_i \rho^{i-1} (y_i - y_i^*) = 0$ is only negligible; this is because ρ is chosen uniformly at random and after the values y_i^* and the values y_i are determined by the simulation queries made up to the forgery. Then, the correctness of the above forgery follows from Remark 1, where for the same commitment $\sum_i \rho^{i-1} c_i$ we have two valid openings $(\sum_i \rho^{i-1} y_i, \pi)$ and $(\sum_i \rho^{i-1} y_i^*, \pi^*)$ on the point x^* .