Knowledge Soundness of Polynomial Commitments in the Algebraic Group Model Does Not Guarantee Extractability

Petr Chmel[1], Pavel Hubáček[1,2], and Dominik Stejskal[1]

[1] Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic
chmel@iuuk.mff.cuni.cz
[2] Institute of Mathematics, Czech Academy of Sciences, Prague, Czech Republic
hubacek@math.cas.cz

Abstract. The Algebraic Group Model (AGM) has become a standard framework for analyzing the knowledge soundness of group-based polynomial commitment schemes. In this work, we formally establish inherent limitations of this methodology. For many PCSs analyzed in the AGM, the algebraic explanation directly determines a polynomial and any required randomness that together reproduce the commitment. We call this structural property AGM-clarity. The extraction proof then only needs to show that the accepted evaluation agrees with this polynomial. For AGM-clear schemes, evaluation binding alone implies knowledge soundness in the AGM. If the commitment algorithm is itself algebraic, knowledge soundness in the AGM is in fact equivalent to *weak* evaluation binding, a relaxation of evaluation binding introduced by Catalano, Fiore, and Tucker (ASIACRYPT 2022).
We then show that this guarantee implies neither standard-model extraction nor evaluation binding. Under the discrete-logarithm and SBDH assumptions, we give a succinct modification of the KZG polynomial commitment that is evaluation binding and, hence, knowledge-sound in the AGM. It admits a prover that is accepted on every evaluation point, yet no efficient extractor recovers an opening of its commitment, even given the prover's code and state. For a non-succinct variant, no polynomial consistent with the transcript exists at all, so extraction fails even for unbounded extractors. A minor modification of those PCSs destroys evaluation binding while leaving knowledge soundness in the AGM intact, so the AGM notion lies strictly below evaluation binding. Finally, we show that the common formulation, which also requires algebraic explanations of evaluation proofs, is strictly weaker than the variant characterized by weak evaluation binding. Our separating schemes satisfy even the stronger variant, so no property implied by either formulation guarantees extraction in the standard model.

Keywords: Polynomial Commitment Schemes · Extractability · Knowledge Soundness · Evaluation Binding · Algebraic Group Model · AGM.

1 Introduction

Polynomial commitment schemes (PCSs) are a central building block in modern zk-SNARKs, offering succinct openings and fast verification. The shortest proofs in practice, at a few hundred bytes, come from group-based constructions, whether pairing-based [KZG10, PST13] or inner-product-based [BBB$^+$18]. For pairing-based constructions, security guarantees are historically the hardest to establish.

In their seminal work, Kate, Zaverucha, and Goldberg [KZG10] showed that their scheme satisfies *evaluation binding* under the strong Diffie-Hellman assumption in bilinear groups (SBDH), i.e., that no efficient prover can produce accepting evaluation proofs for two distinct values at a single evaluation point for the same commitment. Applications to zk-SNARKs motivate the stronger property of *knowledge soundness*, i.e., that from every prover producing an accepting evaluation proof one can extract a polynomial of bounded degree that opens the commitment and agrees with the claimed evaluation; we call such a polynomial *transcript-consistent*. Establishing knowledge soundness for group-based schemes in the standard model under falsifiable assumptions became a major open problem, addressed only recently by Lipmaa, Parisella, and Siim [LPS24]. Those proofs, and the ones that followed [BDH$^+$25, LPS25, CGKY26, BHKM26], cover specific schemes and

rely on new assumptions such as ARSDH. Knowledge soundness under better-studied assumptions remains open.

Besides the direction initiated by Lipmaa et al. [LPS24], the de facto standard route to arguing extractability for group-based schemes goes through the *Algebraic Group Model* (AGM) of Fuchsbauer, Kiltz, and Loss [FKL18]. In the AGM, the adversary is restricted to being *algebraic*: for every group element it outputs, it must also output an *algebraic explanation*, i.e., the coefficients of a linear combination of its inputs producing the output. For the schemes analyzed in this model, the algebraic explanation of a commitment *is* the coefficient vector of the committed polynomial. This crucial syntactic property was leveraged in all AGM proofs of knowledge soundness for PCSs.

Notably, Chiesa et al. [CHM+20] put this to use in a more high-level manner. They gave two variants of an extractable PCS. In the standard-model variant, commitments carry a second group element to facilitate an extractability argument under the non-falsifiable power knowledge of exponent (PKoE) assumption. As a heuristic, they proposed that, in the AGM, the second group element can be dropped, because the algebraic explanation supplies the opening instead. They showed that extractability in the AGM follows from evaluation binding alone [CHM+19, Lemma B.14]. The authors present this as an efficiency improvement: one group element per commitment rather than two. Gabizon, Williamson, and Ciobotaru [GWC19, Remark 3.2] had noted the same convenience earlier, observing that the AGM permits modeling binding and knowledge soundness in a single game, with the extractor returning a polynomial immediately after the commitment. These observations explain how to construct the extractor. They leave open which schemes the argument applies to, what security property it establishes, and whether that property suffices for extraction in the standard model.

The recent standard-model proofs of knowledge soundness under falsifiable assumptions have not displaced the AGM. The practice dates to the first generation of KZG-based universal zk-SNARKs [GWC19, MBKM19, CHM+20], and new constructions of pairing-based polynomial commitments still lead with an AGM analysis as their argument for extractability [GPS25, EG25, Zha26]. We therefore ask the following question.

What does knowledge soundness in the AGM add over evaluation binding?

Prior Critiques of the AGM. One line of criticism concerns the *limited expressivity* of the AGM. Lipmaa, Parisella, and Siim [LPS23] argue that the AGM forbids *oblivious sampling* of group elements without knowing their linear representations, a capability real adversaries possess, e.g., via admissible encodings in elliptic curves. They introduced the *AGM with Oblivious Sampling (AGMOS)* to remedy this, motivated by the observation that assumptions like SpurKE (the spurious knowledge assumption of immediate knowledge of exponent) are false in reality despite being provable in the AGM. While valuable, this line of work critiques the *model*. Our motivation is different. We accept the AGM as a model but study the guarantees of the *definition* of knowledge soundness within it. Our contribution is *descriptive* rather than prescriptive: we characterize what knowledge soundness in the AGM guarantees for practical PCSs, without proposing a replacement notion.

1.1 Our Results and Techniques

Our motivating question has a surprising answer: nothing. When the algebraic explanation of a commitment determines a polynomial and any required randomness that reproduce the commitment, the extraction proof only needs to show that the accepted evaluation agrees with this polynomial. This argument gives no method for recovering a polynomial from a prover that supplies no algebraic explanation.

We isolate the structural property behind this and call it *AGM-clarity*; it holds for the pairing-based constructions previously analyzed in the AGM whose commitment key spans exactly the class of polynomials committed to, as detailed in Appendix A. For AGM-clear schemes, evaluation binding implies knowledge soundness in the AGM, so the AGM notion excludes no scheme that evaluation binding already admits. For

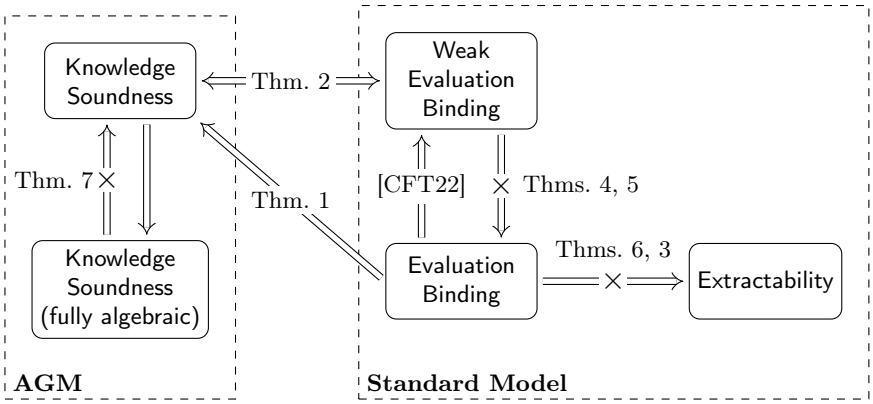

Fig. 1. Visual summary of the relationships between security notions for AGM-clear PCSs with algebraic commitment algorithms. The upper knowledge soundness box refers to Definition 4, which requires algebraic explanations of commitments and permits unrestricted evaluation proofs; the lower node in the AGM box is the common formulation, in which the adversary must explain its evaluation proof as well. The dashed boxes distinguish these formulations from the standard-model notions. The crossed arrows depict our three separations. Theorems 3 and 6 show that AGM knowledge soundness, even together with evaluation binding, does not imply standard-model extractability, the former for a succinct scheme and against efficient extractors, the latter against unbounded ones. Theorems 4 and 5 show that it does not imply evaluation binding either, for the coefficient-wise and the succinct scheme respectively, and Theorem 7 shows that the common fully algebraic formulation does not recover it.

AGM-clear schemes with an algebraic commitment algorithm,[3] knowledge soundness in the AGM is in fact equivalent to *weak evaluation binding* (WEB) [CFT22, CCC$^+$25], a standard-model relaxation of evaluation binding. By this equivalence, the AGM guarantee for these schemes is exactly weak evaluation binding. We show that this guarantee implies neither extraction in the standard model nor evaluation binding.

First, knowledge soundness in the AGM does not guarantee extraction. We give a modification of the KZG scheme that is evaluation binding, hence knowledge-sound in the AGM, yet admits a prover that is accepted on every challenge but from which no efficient extractor recovers an opening. A non-succinct variant with coefficient-wise commitment rules out unbounded extraction as well, since no consistent polynomial exists at all. Second, it does not guarantee evaluation binding either, since a minor modification of the verifier leaves the scheme knowledge-sound in the AGM while destroying evaluation binding outright, so *the AGM notion lies strictly below evaluation binding.*

Finally, our equivalence with WEB uses the formulation in which the adversary explains its commitment but may produce arbitrary evaluation proofs. We show that the common formulation, which also requires explanations of evaluation proofs, is strictly weaker. The schemes we use to rule out standard-model extraction satisfy both formulations, so no property implied by either formulation guarantees such extraction. Figure 1 summarizes these relations.

The AGM Notion Does Not Yield Extraction. By Theorem 2 below, WEB is all that the AGM notion certifies, which leaves open whether it is enough for extraction. Külaots, Lipmaa, Parisella, and Siim [KLPS26] showed that it is not, constructing a PCS that is evaluation binding, and hence WEB, but not extractable in the sense of computational special soundness, with succinct KZG-based variants under ARSDH. However, their separating schemes are not AGM-clear in that the explanation of a commitment need not supply an opening, so weak evaluation binding does not make them knowledge-sound in the AGM, and the question remains open for the schemes to which AGM analyses are applied. We settle it with a scheme that is itself knowledge-sound in the AGM.

[3] That is, the honest commitment algorithm can return an algebraic explanation along with the commitment; see Remark 2.

The scheme we modify is KZG, which is defined over bilinear groups. We write $[x]$ for x times the generator of its first source group, and assume one can obliviously hash into that group. More formally, there is an admissible encoding E from a set of strings to the group, for which recovering the discrete logarithm of $E(s)$ for a uniform seed s is hard under the discrete-logarithm assumption. Fixing a public, uniformly random seed $s^\star$ and writing $E(s^\star) = [x^\star]$, we define the *distinguished commitment* $C^\star = E(s^\star)$, a group element that everyone can compute and whose discrete logarithm $x^\star$ nobody knows.

As our separating PCS, we define $\Pi_{\mathsf{KZG}}[\{0\}]$ to be the ordinary KZG with two changes. First, the commitment key carries the seed $s^\star$ alongside the usual powers of τ, as an ordinary string rather than a further group element. Second, for the distinguished commitment $C^\star$, the verifier additionally accepts the *empty* proof $\pi = \perp$ whenever the claimed value is zero, at any evaluation point. Commitment, opening and the ordinary verification equation are those of KZG, so commitments and honest evaluation proofs remain single group elements.

Theorem 6 (Informal). *Under the discrete-logarithm and SBDH assumptions, $\Pi_{\mathsf{KZG}}[\{0\}]$ is evaluation binding, and hence knowledge-sound in the AGM, yet a deterministic standard-model prover is accepted on every challenge while no PPT extractor with rewinding access outputs any polynomial that opens its commitment, even when given the accepting transcript, the prover's code and state.*

The prover $\mathsf{P}^\star$ establishing the theorem is the obvious one. It commits via $C^\star$, answers every challenge with the value zero and the empty proof, and is accepted throughout. No efficient extractor can output a polynomial that opens $C^\star$, since a reduction that samples the key's trapdoor itself turns any such opening into $x^\star$, which the public seed does not reveal. That remains true however much access to $\mathsf{P}^\star$ the extractor is granted. The AGM never has to account for $\mathsf{P}^\star$, because an algebraic adversary must explain $C^\star$ against the powers in the key, and such an explanation is itself an opening.

A coefficient-wise variant strengthens the failure from efficient to unbounded extraction. Committing to f as $([f_0], \ldots, [f_d])$ makes the commitment injective, and the distinguished commitment becomes $C^\star = (E(s^\star), [0], \ldots, [0])$. It has exactly one opening, the constant polynomial $x^\star$, which is nonzero with overwhelming probability and so never evaluates to the accepted value zero. No transcript-consistent polynomial exists at all, and no extractor succeeds regardless of running time (Theorem 3). For KZG, such a polynomial does exist when $d \geq 1$ and $\alpha \neq \tau$ (see Remark 6), which is why the above separation is only computational. Both results rule out extraction under every notion requiring transcript consistency, computational special soundness included.

The separating prover computes a group element without knowing its discrete logarithm, the capability that Lipmaa, Parisella, and Siim [LPS23] identify as missing from the AGM. This is forced by the question rather than chosen: for an AGM-clear scheme, an algebraic adversary that outputs a commitment with a valid explanation thereby knows an opening, so knowledge soundness in the AGM would apply to any prover the model can express. A witness to the failure of standard-model extraction must lie outside the AGM adversary class.

AGM-Clarity and What AGM Knowledge Soundness Certifies. We first define a structural property satisfied by existing pairing-based polynomial commitments, which we term *AGM-clarity*. A polynomial commitment scheme is AGM-clear if, given the algebraic explanation of a commitment C, one can efficiently construct a polynomial f consistent with C. This property is immediate for practical proposals such as KZG [KZG10], its multivariate variants [PST13], KZG-based Gemini [BCHO22] or Mercury [EG25] since, in these constructions, the commitment algorithm combines exactly the group elements of the commitment key, so an algebraic explanation of a commitment already is the coefficient vector of an opening. AGM-clarity is not a security property, however, and it alone does not suffice for knowledge soundness in the AGM.

Our initial observation is that, for AGM-clear schemes, evaluation binding already implies knowledge soundness. The extractor can simply output the polynomial f reconstructed from the adversary's algebraic explanation of C. To argue that f is consistent with a claimed evaluation (α, η) for which the adversary produced an accepting proof π, suppose instead that $f(\alpha) \neq \eta$. Then one could use f with C to construct another accepting proof for $(\alpha, f(\alpha))$, contradicting evaluation binding.

Theorem 1 (Informal). *Every AGM-clear polynomial commitment scheme that is evaluation binding is also knowledge-sound in the AGM.*

The contribution of Theorem 1 is not the argument, which is familiar from AGM proofs in the literature [CHM+20, GWC19], but how far it reaches across the pairing-based constructions previously analyzed in the AGM. In Appendix A, we check this case by case and discuss technicalities that arise for particular schemes.

Take the original KZG scheme. It is AGM-clear, and Kate, Zaverucha, and Goldberg [KZG10] proved it evaluation binding under SBDH in the standard model. By Theorem 1, their theorem alone already establishes knowledge soundness in the AGM, under the same assumption and with no further argument. The AGM notion therefore imposes no requirement on KZG beyond evaluation binding. What it still certifies once the extractor is trivial is a different question, and the one we answer next.

Characterization via weak evaluation binding. We show that the precise standard-model guarantee of AGM knowledge soundness is weak evaluation binding (WEB, Definition 6), introduced by Catalano, Fiore, and Tucker [CFT22]. In the WEB game, the adversary supplies a polynomial f and commitment randomness, together with an accepting evaluation proof at α that contradicts $f(\alpha)$ for the resulting commitment. By contrast, in the standard evaluation binding game, the adversary can pick an arbitrary commitment and need only provide two accepting evaluation proofs for distinct values at some α, without ever revealing an opening. Catalano et al. [CFT22] already showed that evaluation binding implies WEB via a reduction similar to the one establishing our Theorem 1. Evaluation binding therefore implies both WEB for any scheme and knowledge soundness in the AGM for AGM-clear ones. This is not an accident. For AGM-clear schemes with an algebraic commitment algorithm, the two notions are the same. The formulation in Definition 4 requires an explanation only of the commitment, whereas the customary fully algebraic formulation requires explanations of the evaluation proof as well [LPS23, Definition 6]. Leaving the evaluation proof unrestricted admits more adversaries, so the equivalence characterizes the stronger requirement.

Theorem 2 (Informal). *Every AGM-clear polynomial commitment scheme with an algebraic commitment algorithm satisfies weak evaluation binding if and only if it is knowledge-sound in the AGM in the sense of Definition 4.*

Theorem 2 follows from two direct reductions. AGM-clarity turns the algebraic explanation of a commitment into a polynomial opening, so an accepting evaluation inconsistent with the extracted polynomial breaks WEB. Conversely, a WEB adversary can compute its commitment algebraically from the polynomial opening that it already supplies. Randomizing between the adversarial and the honest evaluation then produces an accepting claim inconsistent with the output of any fixed AGM extractor with at least half the WEB adversary's success probability.

Knowledge soundness against fully algebraic adversaries certifies strictly less. WEB therefore implies the fully algebraic notion for AGM-clear schemes, but the converse fails. We show strictness with an AGM-clear modification of KZG that places the unexplainable group element in the evaluation proof rather than in the commitment, which is forced, since the adversary must explain its commitment in either formulation.

Theorem 7 (Informal). *There is an AGM-clear modification of KZG, with succinct commitments and evaluation proofs, that fails weak evaluation binding against an efficient adversary succeeding with probability one, yet is knowledge-sound against fully algebraic adversaries under the discrete-logarithm and SBDH assumptions.*

The fully algebraic variant therefore lies strictly below WEB, and the equivalence of Theorem 2 holds only for the strongest form of the definition. In Appendix C, we characterize the fully algebraic variant by a restricted form of WEB in which the evaluation proof must itself be produced by the opening algorithm, under additional proof-clarity and algebraicity hypotheses stated there.

The AGM Notion Lies Strictly Below Evaluation Binding. We denote by $\Pi_{\mathsf{KZG}}[\{0,1\}]$ a variant of $\Pi_{\mathsf{KZG}}[\{0\}]$ where the verifier accepts the empty proof for the value one in addition to zero. We show that $\Pi_{\mathsf{KZG}}[\{0,1\}]$ separates knowledge soundness in the AGM from evaluation binding as well. Two empty proofs for values 0 and 1 at the same point certify distinct values, and evaluation binding fails outright, against an efficient adversary succeeding with probability one. On the other hand, evaluation binding is the only property this breaks. The reasons the scheme is AGM-clear and weakly evaluation binding never appeal to which values the empty proof accepts, so both survive, and with them knowledge soundness in the AGM.

Corollary 1 (Informal). *Under the discrete-logarithm assumption, knowledge soundness in the AGM is strictly weaker than evaluation binding for AGM-clear polynomial commitments with an algebraic commitment algorithm.*

The same change has the same effect on the coefficient-wise family (Theorem 4), so neither separation depends on the shape of the commitment. A notion named for knowledge therefore certifies less than binding, and for AGM-clear schemes it sits strictly below the property that KZG was proved to have in the first place.

No Consequence of the AGM Notion Guarantees Extraction. The separations rule out a class of repairs, not only the notion as stated. Let N be any property implied by knowledge soundness in the AGM. For AGM-clear schemes, Theorem 1 makes evaluation binding imply knowledge soundness in the AGM, which in turn implies N, so every evaluation-binding AGM-clear scheme satisfies N. Both $\Pi_{\mathsf{KZG}}[\{0\}]$ and its coefficient-wise counterpart are of that kind, and each admits a prover that is accepted on every challenge and defeats extraction. No consequence of the present notion can therefore guarantee *opening-based* extraction in the standard model, in which the extracted polynomial must open the commitment as well as agree with the accepted evaluation, and this covers every relaxation and equivalent reformulation of it, including variants that merely give the extractor more information, such as the evaluation proof with its algebraic explanation. Restricting attention to succinct schemes is no escape, since $\Pi_{\mathsf{KZG}}[\{0\}]$ has commitments and honest evaluation proofs of a single group element. Appendix B gives a further evaluation-binding, AGM-clear modification of KZG that rules out the weaker target used by Chiesa et al. [CHM+19], which drops commitment consistency but fixes the extracted polynomial before the evaluation challenge, even against unbounded extractors.

Scope and Relation to Prior Work. We now indicate which schemes our results cover, how far they reach beyond the setting in which they are stated, and how they relate to the AGM with oblivious sampling.

AGM-clarity in the literature. With respect to AGM-clarity, the pairing-based constructions fall into the following cases. For most, the commitment key is a basis for exactly the class the scheme commits to, and AGM-clarity holds with no assumption. Some generate the key for a larger class and enforce the smaller one with a degree check inside the evaluation protocol. There, the explanation names a polynomial of the larger class, AGM-clarity does not hold once the extended commitment map is computationally binding, and the characterization of Theorem 2 acquires one further standard-model conjunct, constraining degrees rather than openings. Whether a published analysis takes care of that conjunct depends on the adversary interface it assumes, and the degree-check arguments in the literature establish it against fully algebraic adversaries. See Appendix A for details on specific PCSs.

Randomized and multivariate commitments. Theorem 1 and the WEB characterization (Theorem 2) are stated and proven for randomized polynomial commitments, and they also address the common hiding/zero-knowledge variants of pairing-based schemes. They can be simplified for deterministic commitments by ignoring the randomness in the corresponding definitions. While our formal exposition treats univariate polynomial commitments for clarity, extensions to multivariate commitments are straightforward and require no new ideas beyond routine notation (e.g., component-wise explanations).

Relation to the AGMOS. Everything the AGM certifies about $\Pi_{\mathsf{KZG}}[\{0\}]$ is true since the scheme is evaluation binding against arbitrary standard-model adversaries, and what fails is the reading of the notion as a guarantee of extraction rather than the model. The separation persists in the AGMOS. Belohorec, Hubáček, and Stejskal [BHS26] showed that AGM security lifts to the AGMOS for many games, under the Find Polynomial Representation and Tensor Oracle Find Representation assumptions, and illustrated this on knowledge soundness of polynomial commitments for KZG-like pairing-based schemes. Our results therefore lift to the AGMOS for the schemes their theorem covers, where knowledge soundness in the AGMOS is likewise implied by WEB, and hence by evaluation binding. Our separating schemes satisfy their lifting condition too, so they are knowledge-sound in the AGMOS and the separation from standard-model extraction lifts with them. Theorems 1 and 2 concern the definition alone and do not depend on any of this.

1.2 Open Problems

Establishing standard-model knowledge soundness for practical pairing-based commitments under better-studied falsifiable assumptions remains a major open problem.

Weakening AGM-clarity. AGM-clarity is the hypothesis that makes evaluation binding imply knowledge soundness in the AGM (Theorem 1), and it cannot simply be dropped. The modified-key variant of the coefficient-wise scheme in Remark 5 is evaluation binding with an algebraic commitment algorithm, yet not knowledge-sound in the AGM. Theorem 8 relaxes AGM-clarity in one direction, replacing it by a basis condition on a larger class together with a standard-model obligation on degrees. What remains open is whether a weaker sufficient condition covers schemes of a different shape, and whether that obligation can be established at the unrestricted-proof interface of Definition 4 rather than only against fully algebraic adversaries.

Meaningful AGM notion. Can additional structure support a stronger AGM notion that guarantees standard-model extraction? Such a notion must exclude our evaluation-binding separating schemes, so AGM-clarity and evaluation binding cannot be its only requirements. Identifying such structure is a research problem in its own right, not a definitional refinement. We leave this for future work.

1.3 Additional Related Work

As the AGM of Fuchsbauer, Kiltz, and Loss [FKL18] is a relatively recent framework, its limitations and relationship to other idealized models like the Generic Group Model (GGM) [Sho97, Mau05] remain the subject of ongoing research [KP19, AHK20, ABK$^+$21, ZZK22, AHPP23, JM24, BFHK24]. The AGM is used well beyond polynomial commitments. Fuchsbauer, Kiltz, and Loss [FKL18] proved the knowledge soundness of Groth's zk-SNARK [Gro16] under a q-type discrete-logarithm assumption, where the original analysis was generic. Subsequently, Bünz et al. [BMM$^+$21] and Gailly, Maller, and Nitulescu [GMN22] proved soundness of batch Groth16 proofs. Ghoshal and Tessaro [GT21] study state-restoration witness-extended emulation in the AGM and prove it for variants of Bulletproofs. The Halo2 book claims state-restoration soundness for Halo2 [Com21]. Several works [ZBK$^+$22, PK22, EFG22, ZGK$^+$25] analyze lookup arguments.

Subsequent work. Rothmann [Rot26] formally verifies the security of standard and batched KZG in Isabelle/HOL, including proofs of AGM knowledge soundness by reduction to evaluation binding. His reduction overlaps with Theorem 1 for concrete KZG constructions, whereas we identify AGM-clarity as the condition that makes the implication hold across a broad class, characterize the AGM notion exactly by WEB, and separate it from standard-model extractability. Zhang [Zha26] proves knowledge soundness in the AGM and the random oracle model for two new pairing-based multilinear polynomial commitments, Vela and Carina, with a commitment-time notion that, like Definition 4, gives a single extractor only the commitment and its algebraic explanation. The analysis appeals to our AGM-clarity in interpreting the guarantee it obtains, and both commitments are indeed AGM-clear (Appendix A).

2 Preliminaries

We recall the key notions, starting with polynomial commitment schemes. For conciseness, we consider the group-parameter generation to be a part of the definition of a PCS. The group parameters gp describe the prime-order groups used by the scheme, and $\mathbb{F}$ denotes their scalar field.

Definition 1 (PCS). *A* non-interactive univariate polynomial commitment scheme (PCS) $(\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *consists of the following polynomial-time algorithms:*

Setup: $\mathsf{PGen}(1^\lambda) \to \mathsf{gp}$*; generate group parameters* gp.
Commitment key generation: $\mathsf{KGen}(\mathsf{gp}, d) \to \mathsf{ck}$*; given group parameters* gp *and a degree bound* d*, return a* commitment key ck*, which implicitly contains* gp.
Commitment: $\mathsf{Com}(\mathsf{ck}, f(X); r) \to C$*; for a polynomial* $f(X) \in \mathbb{F}[X]^{\leq d}$*, compute a (possibly randomized)* commitment C.
Opening: $\mathsf{Open}(\mathsf{ck}, C, \alpha, f(X); r) \to (\eta, \pi)$*; to open a commitment* C *to a polynomial* $f(X) \in \mathbb{F}[X]^{\leq d}$ *at a point* $\alpha \in \mathbb{F}$*, return the evaluation* $\eta = f(\alpha)$ *and an* evaluation proof π.
Verification: $\mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) \to \{0, 1\}$*; accept/reject the proof* π *that* η *is the evaluation at* $\alpha \in \mathbb{F}$ *of some* $f(X) \in \mathbb{F}[X]^{\leq d}$ *behind the commitment* C.

We require a PCS to satisfy the following standard correctness requirement:

Completeness: *For every* $d \in \mathsf{poly}(\lambda)$, $\mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda)$, $\mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d)$, *polynomial* $f(X) \in \mathbb{F}[X]^{\leq d}$, *point* $\alpha \in \mathbb{F}$*, and admissible commitment randomness* r*, if* $C = \mathsf{Com}(\mathsf{ck}, f(X); r)$ *and* $(\eta, \pi) \leftarrow \mathsf{Open}(\mathsf{ck}, C, \alpha, f(X); r)$*, then* $\eta = f(\alpha)$ *and* $\mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1$.

Definition 1 covers randomized schemes such as the "blinded" KZG scheme (see, e.g., Section 3.5.3 in Kohrita and Towa [KT24]), where the commitment and opening algorithms take some additional randomness r to achieve hiding.

Throughout, we use the standard typed-game convention: inputs to the PCS algorithms and values in the security experiments must lie in the domains indicated by the syntax above; other adversarial outputs count as losing.

A polynomial commitment should at least prevent a prover from opening one commitment to two different values at the same evaluation point, which is the evaluation binding property.

Definition 2 (Evaluation Binding). *A polynomial commitment scheme* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *satisfies* evaluation binding *if, for any* $d \in \mathsf{poly}(\lambda)$ *and PPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\begin{array}{l|l} \eta \neq \eta' \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \\ \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta', \pi') = 1 & (C, \alpha, \eta, \pi, \eta', \pi') \leftarrow \mathcal{A}(\mathsf{ck}) \end{array}\right] \in \mathsf{negl}(\lambda).$$

We now introduce the Algebraic Group Model and algebraic adversaries.

Definition 3 (AGM). *Let* PGen *be a group-parameter generator. A PPT algorithm* $\mathcal{A}$ *is* algebraic *if, for any group elements it outputs, it also outputs their* algebraic explanations *with respect to its input; an algebraic adversary is also called an* AGM adversary. *Formally, for any* $\mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda)$ *and a vector of input group elements* $\bar{g}$*,*

$$\mathcal{A}(\mathsf{gp}, \bar{g}) \to (\bar{o}, \boldsymbol{\gamma}_{\bar{o}}),$$

where $\boldsymbol{\gamma}_{\bar{o}}$*, the* algebraic explanation *of* $\mathcal{A}$*'s output* $\bar{o}$*, is a matrix over the scalar field* $\mathbb{F}$*, and it holds that*

$$\bar{o} = \boldsymbol{\gamma}_{\bar{o}} \bar{g}. \tag{1}$$

The input or output of $\mathcal{A}$ *may contain some other field elements, which we omit. In the* Algebraic Group Model (AGM)*, all PPT adversaries are algebraic.*

Over bilinear groups, an algebraic adversary explains each group element it outputs relative to the inputs in the same source group (see, e.g., [GPS25]). The property of interest is *knowledge soundness in the AGM*, which requires extraction of a polynomial consistent with the commitment and the accepted evaluation claim. Below, a PPT adversary $\mathcal{A}$ is *algebraic with respect to its commitment* if, whenever $((C, \boldsymbol{\gamma}_C), \alpha, \eta, \pi) \leftarrow \mathcal{A}(\mathsf{ck})$, $\boldsymbol{\gamma}_C$ is a valid algebraic explanation of C with respect to the group elements in ck, as in (1). No algebraicity requirement is imposed on the evaluation proof π.

Definition 4 (Knowledge Soundness in the AGM). *A polynomial commitment scheme* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *is* knowledge-sound in the AGM *if, for any* $d \in \mathsf{poly}(\lambda)$*, there exists a PPT* extractor Ext *returning a polynomial* $f(X) \in \mathbb{F}[X]^{\leq d}$ *together with commitment randomness* r *such that for every PPT adversary* $\mathcal{A}$ *that is algebraic with respect to its commitment,*

$$\Pr\left[\begin{array}{l|l} \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ (C \neq \mathsf{Com}(\mathsf{ck}, f(X); r) & ((C, \boldsymbol{\gamma}_C), \alpha, \eta, \pi) \leftarrow \mathcal{A}(\mathsf{ck}); \\ \vee \;\; f(\alpha) \neq \eta) & (f(X), r) \leftarrow \mathsf{Ext}(\mathsf{ck}, (C, \boldsymbol{\gamma}_C)) \end{array}\right] \in \mathsf{negl}(\lambda).$$

Remark 1 (Algebraicity of the evaluation proof). The extractor interface above is that of the commitment-time formulations in the literature, where a single Ext receives the commitment C with its algebraic explanation $\boldsymbol{\gamma}_C$ and must output f before the evaluation point is chosen (see, e.g., [GWC19, BDFG20, BDFG21]). Previous works quantify over algebraic adversaries in the sense of Definition 3, which explain every group element they output, the evaluation proof included, although their extractor is given no explanation beyond $\boldsymbol{\gamma}_C$. Definition 4 keeps that extractor interface and drops the requirement on π, which enlarges the adversary class and therefore makes it the stronger requirement. Those formulations also differ in the winning condition, which asks only that the accepted evaluation disagree with the extracted polynomial and does not separately require that polynomial to open the commitment (Definition 3.1 in [GWC19], Definition 2.3 in [BDFG20]). Our AGM-clear characterization covers that version too, since the extractor of Lemma 1 returns an opening in any case and the adversary of Lemma 2 wins through evaluation disagreement alone. The implication from evaluation binding in Theorem 1 and the separation in Theorem 3 therefore also hold for fully algebraic adversaries. The exact equivalence with WEB in Theorem 2 does concern Definition 4, whose unrestricted-proof interface lets the reverse reduction forward the evaluation proof supplied by a standard-model WEB adversary, and Theorem 7 shows that this distinction is necessary. In Appendix C, we characterize the fully algebraic variant separately.

3 Evaluation Binding Implies Knowledge Soundness in the AGM

In this section, we show that evaluation binding implies knowledge soundness in the AGM for a large class of polynomial commitments. We first define the polynomial commitments to which our results apply.

Definition 5 (AGM-Clarity). *A PCS* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *is* AGM-clear *if there exists a deterministic polynomial-time (DPT)* revealer Rev *such that, for every* $d \in \mathsf{poly}(\lambda)$*, the following implication holds. If* $\mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda)$*,* $\mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d)$*,* C *is a commitment and* $\boldsymbol{\gamma}_C$ *its algebraic explanation with respect to* ck*, then* $\mathsf{Rev}(\mathsf{ck}, (C, \boldsymbol{\gamma}_C))$ *outputs* $(f(X), r)$ *with* $\deg(f(X)) \leq d$ *and* $\mathsf{Com}(\mathsf{ck}, f(X); r) = C$*.*

Informally, given the algebraic explanation of a commitment, the revealer efficiently extracts a polynomial consistent with it. Randomness is compatible with AGM-clarity. For example, let $\mathbb{G}$ be a prime-order commitment group with scalar field $\mathbb{F}$, written additively, and suppose ck consists of generators $G_0, \ldots, G_d, H$. Let $f(X) = \sum_{i=0}^{d} f_i X^i$. For randomness $r \in \mathbb{F}$, consider the Pedersen-style commitment map $C = \mathsf{Com}(\mathsf{ck}, f(X); r) = \sum_{i=0}^{d} f_i G_i + rH$. An algebraic explanation $(\gamma_0, \ldots, \gamma_d, \rho)$ of C immediately supplies an opening, namely the polynomial $\sum_{i=0}^{d} \gamma_i X^i$ with randomness ρ. The revealer must provide a valid opening, not the original commitment randomness. This does not conflict with hiding. The commitment is hiding for uniformly sampled r, while a revealer recovers an opening from $\boldsymbol{\gamma}_C$, which the adversary that computed C supplies and a party holding only the commitment does not have.

Many polynomial commitments from the literature, such as KZG [KZG10], its multivariate variants [PST13], Gemini [BCHO22], Zeromorph [KT24], Samaritan [GPS25] or Mercury [EG25], are AGM-clear with a revealer that simply outputs the algebraic explanation of the commitment, as are the hiding variants of KZG and Zeromorph, provided the commitment key spans exactly the class of polynomials committed to. The AGM analyses of these and similar schemes argue extraction in exactly this way [CHM+20, GWC19, EG25, Zha26]. An adversary must output coefficients explaining its commitment in the group elements of ck, and the extractor reads them as a polynomial. In each of these constructions, the group elements of ck in the commitment group are exactly those combined by Com.[4] Every element of their span is then a commitment to a polynomial of degree at most d, and its algebraic explanation is an opening. AGM-clarity therefore depends on the entire public input basis rather than the algebraicity of Com alone, and it can fail once ck carries further group elements. Remark 5 gives a scheme with an algebraic commitment algorithm that is not AGM-clear under the discrete-logarithm assumption. Appendix A makes this condition precise and verifies it for the schemes above together with more recent proposals. It also shows that AGM-clarity fails outright once the commitment key spans more polynomials than Com admits and the extended commitment map is computationally binding (Lemma 7), which is the case for schemes enforcing the degree bound inside the evaluation protocol. Theorem 8 gives a separate characterization of knowledge soundness in the AGM for that regime.

For an AGM-clear scheme, the extractor can simply be the revealer. Evaluation binding is then what forces the polynomial it returns to agree with the accepted evaluation.

Theorem 1. *Let $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ be an AGM-clear PCS. If Π is evaluation binding, then it is knowledge-sound in the AGM.*

Proof. We establish the contrapositive. Let Π be an AGM-clear PCS that is not knowledge-sound in the AGM. We show that Π is not evaluation binding.

To break the evaluation binding of Π, we use any adversary $\mathcal{A}$ that wins the knowledge-soundness game against the most basic extractor Ext^* that outputs the polynomial and randomness produced from commitment C and its algebraic explanation $\boldsymbol{\gamma}_C$ by the revealer Rev guaranteed by AGM-clarity. Since Π is not knowledge-sound in the AGM, for any efficient extractor Ext, there exists an adversary $\mathcal{A}$ that is algebraic with respect to its commitment and wins the knowledge-soundness game against Ext with non-negligible probability. In particular, such an adversary exists against the above revealer-based extractor Ext^*. Thus, given the commitment key ck, the adversary $\mathcal{A}$ can, with non-negligible probability, output a commitment C with an algebraic explanation $\boldsymbol{\gamma}_C$ and a point α with a claimed evaluation η and an accepting proof π such that the polynomial returned by $\mathsf{Ext}^*(\mathsf{ck}, (C, \boldsymbol{\gamma}_C))$ evaluates at α to some $\eta' \neq \eta$ and $\mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1$. The other disjunct of the winning condition, $C \neq \mathsf{Com}(\mathsf{ck}, f(X); r)$, cannot occur, since Rev returns an opening of C.

Given a commitment key ck, our adversary $\mathcal{A}'$ for the evaluation binding game proceeds as follows:

1. Run $\mathcal{A}(\mathsf{ck})$ to get a commitment C with an explanation $\boldsymbol{\gamma}_C$, and a point α with a claimed evaluation η and a proof π such that $\mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1$.
2. Run $\mathsf{Rev}(\mathsf{ck}, (C, \boldsymbol{\gamma}_C))$ to obtain a polynomial ψ and randomness r such that $\mathsf{Com}(\mathsf{ck}, \psi(X); r) = C$. Compute the honest opening $(\eta', \pi') \leftarrow \mathsf{Open}(\mathsf{ck}, C, \alpha, \psi(X); r)$, so that $\eta' = \psi(\alpha)$ by completeness.
3. Output $(C, \alpha, \eta, \pi, \eta', \pi')$.

The adversary $\mathcal{A}'$ wins the evaluation binding game with the same probability that $\mathcal{A}$ wins the AGM knowledge-soundness game against Ext^*, which we assume to happen with non-negligible probability. Hence, $\mathcal{A}'$ breaks evaluation binding, establishing the contrapositive and thus the theorem. □

4 A Standard-Model Interpretation of Knowledge Soundness in the AGM

In this section, we characterize the power of knowledge soundness in the AGM (Definition 4) with respect to the standard model. To this end, we use the standard-model notion of *weak evaluation binding*

[4] For a hiding variant, the explanation splits into the coefficients of $f(X)$ and the commitment randomness, and the revealer returns both.

(WEB) [CFT22, CCC+25] and show that it is equivalent to knowledge soundness in the AGM for AGM-clear schemes with an algebraic commitment algorithm.

Definition 6 (Weak Evaluation Binding [CFT22, CCC+25]). *A polynomial commitment scheme* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *satisfies* weak evaluation binding *if, for any* $d \in \mathsf{poly}(\lambda)$ *and PPT adversary* $\mathcal{A}$*, it holds that*

$$\Pr\left[\begin{array}{l|l} \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ \eta \neq f(\alpha) & (f(X), r, \alpha, \eta, \pi) \leftarrow \mathcal{A}(\mathsf{ck}); \\ & C \leftarrow \mathsf{Com}(\mathsf{ck}, f(X); r) \end{array}\right] \in \mathsf{negl}(\lambda).$$

The notion is due to Catalano, Fiore, and Tucker [CFT22], who formulated it for functional commitments, and we use the polynomial-commitment formulation of Campanelli et al. [CCC+25, Definition 20]. In the WEB game, the adversary supplies the polynomial itself and the game commits to it honestly, so the degree bound on f is a typing requirement on the adversary's output rather than a winning condition. The adversary must hand over a genuine opening of the commitment it attacks before it can cheat at a single point. Nothing is therefore required of a prover that produces accepting evaluation claims without knowing any opening, and such provers exist, as the scheme of Section 5 shows.

Accordingly, weak evaluation binding sits below evaluation binding, as observed by Catalano et al. [CFT22] for functional commitments and stated for polynomial commitments by Campanelli et al. [CCC+25, Theorem 1]. The argument mirrors the proof of our Theorem 1: the adversary's own (f, r) opens C, and opening it honestly at α yields a second accepting evaluation contradicting the adversary's claim.

Proposition 1 ([CFT22, CCC+25]). *Let* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *be a PCS. If* Π *is evaluation binding, then it satisfies weak evaluation binding.*

The main result of this section is that this low bar is exactly what knowledge soundness in the AGM clears. To establish the implication from knowledge soundness in the AGM to WEB, we restrict to polynomial commitments with an algebraic commitment algorithm.

Remark 2 (Algebraic commitment algorithm). A polynomial commitment has an *algebraic commitment algorithm* if Com is algebraic in the sense of Definition 3, i.e., running Com on $(\mathsf{ck}, f(X); r)$ yields, together with the commitment C, an algebraic explanation $\boldsymbol{\gamma}_C$ of C with respect to the group elements in ck. This is a natural assumption for the applications considered here, since it makes little sense to restrict adversaries to algebraic behavior while permitting the scheme itself to produce group elements it cannot explain. For the schemes we consider, the requirement is immediate, since their commitment algorithms are linear in the commitment key.

Theorem 2. *Let* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *be an AGM-clear PCS with an algebraic commitment algorithm. Then* Π *satisfies WEB if and only if it is knowledge-sound in the AGM.*

In the next section, we show that the unrestricted-proof interface of Definition 4 is essential to this equivalence (Theorem 7). The proof of Theorem 2 follows from the two reductions in Lemmas 1 and 2 below. The implication from WEB to knowledge soundness in the AGM holds for any AGM-clear PCS, irrespective of whether it has an algebraic commitment algorithm.

Lemma 1. *Let* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *be an AGM-clear PCS. If* Π *satisfies weak evaluation binding, then it is also knowledge-sound in the AGM.*

Proof. We prove the contrapositive, that is, breaking knowledge soundness in the AGM also breaks WEB. As Π is AGM-clear, let Rev be the revealer certifying its AGM-clarity.

As our goal is to create an accepting evaluation proof for a claimed evaluation that is inconsistent with the polynomial, we build an AGM extractor Ext that simply runs the revealer on the commitment and the algebraic explanation. As knowledge soundness in the AGM is broken by assumption, let $\mathcal{A}$ be an

adversary that wins the game against Ext with non-negligible probability. Using this, we build the adversary $\mathcal{A}'$ as follows: $\mathcal{A}'$ runs $\mathcal{A}$ to obtain $((C, \boldsymbol{\gamma}_C), \alpha, \eta, \pi)$. Then it uses Rev to learn the polynomial $f(X)$ and randomness r consistent with the commitment. Finally, it outputs $(f(X), r, \alpha, \eta, \pi)$.

Conditioned on $\mathcal{A}$ winning against Ext, we know that $f(\alpha) \neq \eta$ (since the other condition cannot hold by the properties of the revealer). The evaluation proof π for the claimed evaluation (α, η) also verifies, since $\mathcal{A}$ wins, and the commitment of $f(X)$ with randomness r is C by the properties of the revealer. □

The converse implication from knowledge soundness in the AGM to WEB holds for any PCS with an algebraic commitment algorithm, irrespective of whether it is AGM-clear.

Lemma 2. *Let $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ be a PCS with an algebraic commitment algorithm. If Π is knowledge-sound in the AGM, then it also satisfies weak evaluation binding.*

Proof. We again prove the contrapositive: breaking WEB implies breaking knowledge soundness in the AGM. Let Ext be an extractor for AGM knowledge soundness. Our goal is to provide an adversary $\mathcal{A}'$ breaking Ext. As we assume WEB is broken, let $\mathcal{A}$ be a WEB-winning adversary. The adversary $\mathcal{A}'$ first runs $\mathcal{A}$ to obtain $(f(X), r, \alpha, \eta, \pi)$. If WEB is broken, we know that $C = \mathsf{Com}(\mathsf{ck}, f(X); r)$, and thus, by Remark 2, running Com on $(\mathsf{ck}, f(X); r)$ also yields an algebraic explanation $\boldsymbol{\gamma}_C$ of C.

The adversary $\mathcal{A}'$ must also provide an accepting evaluation proof for a claimed evaluation that is inconsistent with the polynomial extracted by Ext. It thus computes an honest opening $(\eta^f, \pi^f) \leftarrow \mathsf{Open}(\mathsf{ck}, C, \alpha, f(X); r)$ at the same point α, tosses a fair coin, and outputs either (α, η) with its evaluation proof π or (α, η^f) with its evaluation proof π^f, together with the commitment C and algebraic explanation $\boldsymbol{\gamma}_C$. Both branches use the point α, so we write (α, η') for the claimed evaluation and π' for the evaluation proof that $\mathcal{A}'$ outputs, where $(\eta', \pi') \in \{(\eta, \pi), (\eta^f, \pi^f)\}$.

Conditioned on the internal execution winning the WEB game, we have that $\eta \neq \eta^f$, the algebraic explanation $\boldsymbol{\gamma}_C$ is correct, f has degree $\leq d$, $C = \mathsf{Com}(\mathsf{ck}, f(X); r)$, and the evaluation proof π' verifies for the claimed evaluation (α, η') in both branches. Since the input to Ext is independent of the coin toss, let $(\varphi(X), \rho)$ be the output of Ext. Since both candidate proofs verify, it suffices to establish $\varphi(\alpha) \neq \eta'$ with non-negligible probability. This satisfies the evaluation-mismatch disjunct of the knowledge-soundness winning condition; commitment inconsistency would also make the adversary win, so it need not be ruled out. Note that $\varphi(\alpha)$ is a single field element, and therefore at least one of η, η^f must be different from it. Hence, with probability at least $1/2$, we break Ext (conditioned on $\mathcal{A}$ winning the WEB game). Since $\mathcal{A}$ wins the WEB game with non-negligible probability, we also break Ext with non-negligible probability, and the lemma follows. Note that $\mathcal{A}'$ is built from $\mathcal{A}$ alone and never invokes Ext. The same adversary therefore defeats every PPT extractor, which is what refuting the adversary-specific formulation of knowledge soundness in the AGM requires. □

Remark 3 (Universal vs. adversary-specific extraction). Definition 4 asks for a single extractor that succeeds against every adversary, which implies the formulation in which the extractor may depend on the adversary while receiving the same input. Our two reductions establish the equivalence for both variants of Definition 4. Lemma 1 produces one extractor from the revealer, independently of the adversary, and Lemma 2 produces one adversary that defeats every extractor, since its coin is not part of the extractor's input. Weak evaluation binding is therefore equivalent to knowledge soundness in the AGM under either quantification.

Remark 4 (AGM knowledge soundness from evaluation binding through WEB). Theorem 1 also follows by composing Proposition 1 with Lemma 1: evaluation binding implies weak evaluation binding for every PCS, and, for AGM-clear schemes, weak evaluation binding implies knowledge soundness in the AGM. We gave the direct proof in Section 3 because it does not presuppose WEB. The factorization locates the collapse: knowledge soundness in the AGM reaches no further than the standard-model notion of WEB that evaluation binding was already known to imply [CFT22].

5 Separating Knowledge Soundness in the AGM from Other Standard-Model Notions

In this section, we separate knowledge soundness in the AGM from standard-model extraction, from evaluation binding, and, in the fully algebraic formulation, from WEB. The sharpest of the three uses a succinct scheme. A modification of KZG with commitments and honest evaluation proofs of a single group element is evaluation binding, hence knowledge-sound in the AGM by Theorem 1, yet admits a prover that is accepted on every challenge but from which no efficient extractor recovers an opening (Theorem 6). A second modification of the same shape is knowledge-sound in the AGM without being evaluation binding at all (Theorem 5), and a third separates knowledge soundness against fully algebraic adversaries from WEB (Theorem 7).

We present the constructions in a different order. A family of PCSs with a coefficient-wise, injective commitment establishes the first two separations. The commitment's unique opening makes the mechanism visible and rules out even unbounded extraction. Within each family, the two members differ in a single parameter of the verification algorithm, which the AGM notion cannot distinguish. The KZG-based schemes modify the coefficient-wise ones, and establish that the separations hold even for succinct schemes with respect to efficient standard-model extraction (see Remark 6). We note that Külaots et al. [KLPS26] recently separated evaluation binding from extractability in the sense of computational special soundness, and our first separation gives an explicit instance of it. Their own schemes, however, are not AGM-clear.

Since our separations rely on ideas related to oblivious hashing to groups, we first recall some basic properties of admissible encodings as defined in Lipmaa, Parisella, and Siim [LPS23] when introducing the AGM with Oblivious Sampling.

Admissible encodings. Let PGen output group parameters gp describing an additively written cyclic group $\mathbb{G} = \langle [1] \rangle$ of prime order p with scalar field $\mathbb{F} = \mathbb{F}_p$, where $[x] = x[1]$ for $x \in \mathbb{F}$. Let $\mathcal{S}$ be a finite set from which uniform elements can be sampled efficiently, and let $E \colon \mathcal{S} \to \mathbb{G}$ be an admissible encoding as defined by Lipmaa, Parisella, and Siim [LPS23, Section 2.1], such that every group element has at most ρ preimages and $\rho|\mathbb{G}|/|\mathcal{S}|$ is polynomially bounded in λ, uniformly over the generated group parameters. For $s \in \mathcal{S}$, write $x_s \in \mathbb{F}$ for the unique scalar with $E(s) = [x_s]$.

Their Theorem 1, together with the efficient preimage enumeration guaranteed by admissibility, implies that under the discrete-logarithm assumption in $\mathbb{G}$, every PPT algorithm $\mathcal{A}$ satisfies

$$\Pr\left[\mathcal{A}(\mathsf{gp}, s) = x_s \,\middle|\, \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); s \leftarrow \mathcal{S}\right] \in \mathsf{negl}(\lambda).$$

In particular, taking for $\mathcal{A}$ the algorithm that always outputs 0 shows that $x_s \neq 0$ except with negligible probability over setup and the seed.

A Coefficient-Wise Family. The construction carries one parameter, a set $V \subseteq \mathbb{F}$ with $0 \in V$ and efficiently decidable membership. The verification algorithm uses V as the set of values it accepts for a single distinguished commitment, and nothing else depends on it. For every such V, we define a deterministic univariate polynomial commitment scheme $\mathit{\Pi}_{\mathsf{os}}[V] = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$. We assume the group-parameter generator includes descriptions of $\mathcal{S}$ and E in gp. Polynomials of degree at most d are written $f(X) = \sum_{i=0}^{d} f_i X^i$, identifying f with its coefficient vector $(f_0, \ldots, f_d)$. The remaining algorithms are as follows.

Commitment key generation KGen**:** On input (gp, d), sample $s^\star \leftarrow \mathcal{S}$ and output $\mathsf{ck} = (\mathsf{gp}, [1], s^\star, d)$.
Commitment Com**:** On input $(\mathsf{ck}, f(X))$, output $C = ([f_0], \ldots, [f_d])$.
Opening Open**:** On input $(\mathsf{ck}, C, \alpha, f(X))$, output $(\eta, \pi) = \big(f(\alpha), (f_0, \ldots, f_d)\big)$.
Verification Vrf**:** On input $(\mathsf{ck}, C, \alpha, \eta, \pi)$, the verification algorithm distinguishes two types of proofs. A full proof $\pi = (g_0, \ldots, g_d) \in \mathbb{F}^{d+1}$ is accepted if $C = ([g_0], \ldots, [g_d])$ and $\eta = g(\alpha)$. Additionally, the empty proof $\pi = \bot$ is accepted exactly when $C = C^\star$ and $\eta \in V$, where $C^\star = (E(s^\star), [0], \ldots, [0]) \in \mathbb{G}^{d+1}$ is the distinguished commitment.
Finally, the verification algorithm rejects all other inputs.

The scheme is complete because honest openings reveal the coefficient vector used to compute the commitment. Note that, if $E(s^\star) = [x^\star]$, then $C^\star = \mathsf{Com}(\mathsf{ck}, x^\star)$, so $C^\star$ is the valid commitment to the constant polynomial $x^\star$. Importantly, $x^\star$ cannot be recovered efficiently from $s^\star$. Committing coordinate-wise makes Com injective, since $x \mapsto [x]$ is a bijection on a group of prime order, so every commitment has exactly one opening. This is what lets both the mixed-opening case and weak evaluation binding reduce directly to the discrete logarithm problem. Remark 6 explains what changes when the commitment is not injective. We abbreviate $\mathit{\Pi}_{\mathsf{os}} = \mathit{\Pi}_{\mathsf{os}}[\{0\}]$.

Transcript consistency. For the opening-based extraction target considered here, the extractor must output a *transcript-consistent* polynomial, one that opens the commitment and agrees with the accepted evaluation. Specifically, after a prover fixes a commitment C and answers a challenge α with an accepting pair (η, π), the extractor must output a polynomial $f(X) \in \mathbb{F}[X]^{\leq d}$ such that $C = \mathsf{Com}(\mathsf{ck}, f(X))$ and $f(\alpha) = \eta$. For the family $\mathit{\Pi}_{\mathsf{os}}[V]$, the failure of transcript consistency established below is independent of the challenge distribution, the extractor's access to the prover, and the extractor's running time. Consequently, no member of this family admits transcript-consistent extraction in the standard model, even by a non-black-box, non-uniform, or computationally unbounded extractor.

Only evaluation binding depends on V, so we establish the remaining properties, and the failure of extraction, once for the whole family.

Lemma 3. *Let $V \subseteq \mathbb{F}$ satisfy $0 \in V$ and have efficiently decidable membership. Then $\mathit{\Pi}_{\mathsf{os}}[V]$ is AGM-clear and has an algebraic commitment algorithm. Under the discrete-logarithm assumption in $\mathbb{G}$, it holds that*

1. *$\mathit{\Pi}_{\mathsf{os}}[V]$ satisfies weak evaluation binding, and is therefore knowledge-sound in the AGM by Theorem 2, but*
2. *there exists a deterministic standard-model prover whose response to every challenge is accepted and for which, except with negligible probability, no polynomial is transcript-consistent with any such response.*

Proof. AGM-clarity. Let $C = (C_0, \ldots, C_d)$ be a commitment with algebraic explanation $(\gamma_0, \ldots, \gamma_d)$ relative to $[1]$, so that $C_i = [\gamma_i]$ for every $i \in \{0, \ldots, d\}$. The revealer outputs $f_{\boldsymbol{\gamma}}(X) = \sum_{i=0}^{d} \gamma_i X^i$. Then it holds that $\mathsf{Com}(\mathsf{ck}, f_{\boldsymbol{\gamma}}(X)) = ([\gamma_0], \ldots, [\gamma_d]) = C$, which proves the AGM-clarity of the scheme. The commitment algorithm is also algebraic in the sense of Remark 2, as Com computes each coordinate as $[f_i] = f_i[1]$ and thereby produces the explanation $(f_0, \ldots, f_d)$ along with the commitment.

Weak evaluation binding. Let $\mathcal{A}$ be an adversary in the game of Definition 6, and let $(f(X), \alpha, \eta, \pi)$ be its output, omitting the commitment randomness because the scheme is deterministic. The game computes $C = \mathsf{Com}(\mathsf{ck}, f(X))$ itself. Suppose first that π is a full proof $(g_0, \ldots, g_d)$. Acceptance requires $C = ([g_0], \ldots, [g_d])$, while $C = ([f_0], \ldots, [f_d])$ by construction, so $g_i = f_i$ for every i because $\mathbb{G}$ has order p. Hence $g = f$, and acceptance forces $\eta = g(\alpha) = f(\alpha)$, so the winning condition $\eta \neq f(\alpha)$ fails.

A winning adversary therefore uses the empty proof, and acceptance then requires $C = C^\star$. Since $C = \mathsf{Com}(\mathsf{ck}, f(X))$ and the commitment is coordinate-wise injective, this forces $f(X) \equiv x^\star$. The adversary has thus written the discrete logarithm of $E(s^\star)$ into the constant coefficient of its own output. An algorithm that receives (gp, s), runs $\mathcal{A}$ on $\mathsf{ck} = (\mathsf{gp}, [1], s, d)$, and returns f_0 therefore computes x_s whenever $\mathcal{A}$ wins, and the commitment key it supplies has exactly the distribution of KGen. Under the discrete-logarithm assumption in $\mathbb{G}$, the probability that $\mathcal{A}$ wins is negligible. No case distinction on $x^\star$ is needed, since the reduction returns the correct value on every run in which $\mathcal{A}$ wins.

Failure of standard-model extraction. Consider the standard-model prover $\mathsf{P}^\star$ that computes $C^\star$ from the public seed $s^\star$ and fixes this commitment before seeing the challenge. On every challenge $\alpha \in \mathbb{F}$, it returns $\mathsf{P}^\star(\alpha) = (0, \bot)$, which the verifier accepts because $0 \in V$.

Suppose an extractor outputs a polynomial $f(X)$ satisfying $C^\star = \mathsf{Com}(\mathsf{ck}, f(X))$. The injectivity of the coefficient-wise commitment then forces $f(X) \equiv x^\star$. On the other hand, consistency with the accepting response requires $f(\alpha) = 0$. Except for the negligible event $x^\star = 0$, no polynomial satisfies both conditions. Thus, with overwhelming probability over setup, the extractor has no valid output for the accepting transcript. □

The two separations. Only evaluation binding is sensitive to the parameter V, and in the simplest way. With $V = \{0\}$ the empty proof cannot contradict itself, whereas a second value in V makes it contradict itself at every point.

Theorem 3. *The scheme $\Pi_{\mathsf{os}}[\{0\}]$ is AGM-clear. Under the discrete-logarithm assumption in $\mathbb{G}$, it holds that*

1. *$\Pi_{\mathsf{os}}[\{0\}]$ is evaluation binding, and hence knowledge-sound in the AGM, but*
2. *there exists a deterministic standard-model prover whose response to every challenge is accepted and for which, except with negligible probability, no polynomial is transcript-consistent with any such response.*

Proof. AGM-clarity and the second item are given by Lemma 3, so only evaluation binding remains. Two accepting full proofs for the same commitment reveal coefficient vectors $(f_0, \ldots, f_d)$ and $(g_0, \ldots, g_d)$ satisfying $[f_i] = [g_i]$ for every i. Since $\mathbb{G}$ has order p, this implies $f_i = g_i$ for every i, and the two claimed evaluations coincide at every point. Two accepting empty proofs cannot disagree either, because $V = \{0\}$ leaves them no value but zero to claim.

It remains to consider one full proof and one empty proof for the same commitment C. The empty proof forces $C = C^\star$ and claims the value zero. The full proof reveals the unique polynomial that opens $C^\star$, namely $f^\star(X) \equiv x^\star$. Thus, an adversary that produces one accepting proof of each type with different values reveals $x^\star$ in its full proof. Such an adversary computes the discrete logarithm of $E(s^\star)$ from $(\mathsf{gp}, s^\star)$. Under the discrete-logarithm assumption in $\mathbb{G}$, the probability of such a violation is negligible. This argument applies to arbitrary PPT adversaries, not only to algebraic ones. Knowledge soundness in the AGM then follows either from Theorem 1 or from the first item of Lemma 3. □

Why weak evaluation binding survives. The prover $\mathsf{P}^\star$ is accepted on every challenge, and yet Lemma 3 asserts weak evaluation binding for the same scheme. The two coexist because $\mathsf{P}^\star$ never hands over an opening of its commitment. Its accepted claims all take the value zero, whose interpolation is the zero polynomial, but $\mathsf{Com}(\mathsf{ck}, 0) = ([0], \ldots, [0]) \neq C^\star$ with overwhelming probability. This is the asymmetry that Theorem 2 converts into a statement about the AGM, since the only standard-model guarantee the AGM notion carries is one that an always-accepted prover can decline to trigger.

Adding 1 to V destroys evaluation binding without disturbing anything else.

Theorem 4. *Let $V \subseteq \mathbb{F}$ satisfy $\{0, 1\} \subseteq V$ and have efficiently decidable membership. Then $\Pi_{\mathsf{os}}[V]$ is AGM-clear, has an algebraic commitment algorithm, and is not evaluation binding, the violation being exhibited by an efficient adversary that succeeds with probability one. Under the discrete-logarithm assumption in $\mathbb{G}$, it is nonetheless knowledge-sound in the AGM.*

Proof. On input ck, the adversary computes $C^\star$ from the public seed $s^\star$, picks an arbitrary $\alpha \in \mathbb{F}$, and outputs $(C^\star, \alpha, 0, \bot, 1, \bot)$. Both responses are accepted, since the commitment is $C^\star$ and both claimed values lie in V, and the two values differ. The remaining properties are those of Lemma 3. □

Theorems 1 and 4 locate the AGM notion strictly below evaluation binding.

Corollary 1. *Under the discrete-logarithm assumption in $\mathbb{G}$, knowledge soundness in the AGM is strictly weaker than evaluation binding for AGM-clear polynomial commitments with an algebraic commitment algorithm.*

Proof. Every evaluation-binding scheme in the class is knowledge-sound in the AGM by Theorem 1. The scheme $\Pi_{\mathsf{os}}[\{0, 1\}]$ of Theorem 4 belongs to the class and is knowledge-sound in the AGM without being evaluation binding. □

The containment in Proposition 1 is therefore strict. Catalano et al. [CFT22] prove the implication, and Campanelli et al. [CCC+25, Theorem 1] place it inside a hierarchy of binding notions they describe as standard. To the best of our knowledge, neither they nor Balbás et al. [BCFL23], who treat the notions above evaluation binding, observe that the implication cannot be reversed.

Dissecting the parameterization. Knowledge soundness in the AGM holds for every admissible V, as does weak evaluation binding, and transcript-consistent extraction fails for all of them. Evaluation binding, by contrast, holds for $V = \{0\}$ and fails outright once $1 \in V$. A single parameter change thus carries the scheme across the boundary of evaluation binding while the AGM notion is unaffected, and no setting of it recovers extraction. Evaluation binding and extraction are exactly the two properties that the AGM notion does not control.

Computing versus explaining a group element. The construction exploits the distinction between computing a group element and producing its scalar representation relative to the AGM input basis. The only group element in ck is $[1]$, and $s^\star$ is an ordinary seed rather than another basis element, so both KGen and the standard-model prover can compute $E(s^\star)$ without learning $x^\star$. An algebraic adversary that outputs $C^\star$, by contrast, must provide $x^\star$ as the first coordinate of its explanation, so the strategy of $\mathsf{P}^\star$ is efficient in the standard model but cannot satisfy the AGM interface.

Remark 5 (The modified-key variant). Suppose we change the commitment key generation of $\Pi_{\mathsf{os}}[\{0\}]$ so that it outputs $(\mathsf{gp}, [1], E(s^\star), d)$, i.e., ck contains the group element $E(s^\star)$ explicitly instead of the "implicit" seed $s^\star$. The resulting scheme is no longer AGM-clear, as an algebraic adversary could explain $C^\star$ with respect to $E(s^\star)$, and any revealer would recover the constant $x^\star$ from that explanation, contradicting the discrete-logarithm assumption. It is nonetheless still evaluation binding with an algebraic commitment algorithm. An algebraic adversary can now output $C^\star$ by selecting $E(s^\star)$ from the commitment key, and, except with negligible probability, no polynomial is transcript-consistent with the resulting accepting transcript. The variant is therefore not AGM knowledge-sound, so the AGM-clarity hypothesis of Theorem 1 cannot be dropped.

Necessary features of the construction. Two features of the family might invite the objection that it is contrived, but both are unavoidable. That the separating prover must produce a group element it cannot explain was noted above. The second is that the verifier admits $d+1$ accepting evaluations at pairwise-distinct points whose interpolating polynomial does not open the commitment. Suppose a scheme had the opposite property, so that for every commitment and every set of $d+1$ accepting evaluations at pairwise-distinct points the interpolating polynomial opens that commitment. Fix a prover that commits, answers a non-negligible fraction of the challenges, and can be rewound with its commitment and randomness held fixed, and assume as usual that $|\mathbb{F}|$ is superpolynomial and $d \in \mathsf{poly}(\lambda)$. Rewinding until the challenge α of the transcript at hand is joined by d further accepting points, all pairwise distinct, takes expected polynomial time. The interpolant through these $d+1$ points has degree at most d, agrees with the accepted value at α, and opens the commitment by assumption, so it is transcript-consistent. This is the argument by which computational special soundness yields standard-model knowledge soundness [LPS24, BDH$^+$25], read in the other direction: if transcript-consistent extraction fails against such a prover, the scheme must admit accepting evaluations whose interpolant does not open the commitment. Our prover $\mathsf{P}^\star$ meets these conditions with room to spare, as it is deterministic and accepted on every challenge. The empty proofs of $\Pi_{\mathsf{os}}[V]$ supply exactly the required failure, so the shape of the construction is dictated by the question rather than chosen.

Relationship to AGMOS. The standard-model prover $\mathsf{P}^\star$ in our proof of Lemma 3 evaluates E at the public seed $s^\star$, and can therefore compute $C^\star$ without learning $x^\star$. This strategy lies outside the scope of the AGMOS of Lipmaa, Parisella, and Siim [LPS23, Definition 1]. There, an AGMOS adversary calls an oblivious sampling oracle that samples a fresh seed s from an allowed distribution and returns $(E(s), s)$ to the adversary. Their definition is parameterized by the family of allowed sampling distributions, and for the well-spread families they consider, such an oracle cannot target the fixed public seed $s^\star$. Thus, the strategy of $\mathsf{P}^\star$ cannot be trivially simulated by an AGMOS adversary. Under the assumptions of Belohorec et al. [BHS26] together with ours, the separating schemes are themselves knowledge-sound in the AGMOS. Their lifting condition holds for $\Pi_{\mathsf{KZG}}[\{0\}]$, since ordinary proofs use the pairing check of KZG, for which they verify it, and the acceptance rule for the empty proof is linear in the obliviously sampled elements, so its oracle-dependent part vanishes only when their coefficients do. The separation of Theorem 6 therefore carries over unchanged,

because $\mathsf{P}^\star$ and the failure of extraction are statements about the standard model rather than about the model in which knowledge soundness is proved.

A Succinct Variant Based on KZG. The commitments and full proofs of $\Pi_{\mathsf{os}}[V]$ each consist of $d+1$ elements. Both separations also hold when the underlying commitment is KZG, so that honest commitments and honest evaluation proofs are single group elements, although the separation from extraction is now computational. Let gp now describe a bilinear group $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e)$ of prime order p with generators $[1]_1$ and $[1]_2$, where $[x]_i = x[1]_i$ for $x \in \mathbb{F}$, and instantiate the admissible encoding above with $\mathbb{G} = \mathbb{G}_1$, so that $C^\star = E(s^\star) = [x^\star]_1$. The commitment key is $\mathsf{ck} = \big(\mathsf{gp}, [1]_1, [\tau]_1, \ldots, [\tau^d]_1, [1]_2, [\tau]_2, s^\star, d\big)$ where $\tau \leftarrow \mathbb{F}$ and $s^\star \leftarrow \mathcal{S}$ are sampled independently. The algorithms Com, Open and the verification equation are those of KZG [KZG10], so that $\mathsf{Com}(\mathsf{ck}, f(X)) = [f(\tau)]_1$. As before, the seed $s^\star$ enters ck as an ordinary string and not as a group element. The verification algorithm accepts an ordinary KZG proof exactly when the KZG pairing equation holds, and accepts the empty proof $\pi = \bot$ exactly when $\alpha \in \mathbb{F}$, $C = C^\star$ and $\eta \in V$. We write $\Pi_{\mathsf{KZG}}[V]$ for the resulting scheme.

Neither AGM-clarity nor algebraicity of the commitment algorithm depends on V, so we establish both once.

Lemma 4. *Let $V \subseteq \mathbb{F}$ have efficiently decidable membership. Then $\Pi_{\mathsf{KZG}}[V]$ is AGM-clear and has an algebraic commitment algorithm.*

Proof. An algebraic adversary that outputs a commitment $C \in \mathbb{G}_1$ explains it relative to $[1]_1, [\tau]_1, \ldots, [\tau^d]_1$, so its explanation is a vector $(\gamma_0, \ldots, \gamma_d)$ with $C = \sum_{i=0}^{d} \gamma_i [\tau^i]_1 = \big[f_{\boldsymbol{\gamma}}(\tau)\big]_1$ for $f_{\boldsymbol{\gamma}}(X) = \sum_{i=0}^{d} \gamma_i X^i$. The revealer outputs $f_{\boldsymbol{\gamma}}$, which satisfies $\mathsf{Com}(\mathsf{ck}, f_{\boldsymbol{\gamma}}(X)) = C$. The same linearity makes Com algebraic in the sense of Remark 2, with explanation $(f_0, \ldots, f_d)$. □

Theorem 5. *Let $V \subseteq \mathbb{F}$ satisfy $\{0,1\} \subseteq V$ and have efficiently decidable membership. Then $\Pi_{\mathsf{KZG}}[V]$ is AGM-clear, has an algebraic commitment algorithm, and is not evaluation binding, the violation being exhibited by an efficient adversary that succeeds with probability one. Under the discrete-logarithm assumption in $\mathbb{G}_1$ and the SBDH assumption under which KZG is evaluation binding, it is nonetheless knowledge-sound in the AGM.*

Proof. AGM-clarity and the algebraic commitment algorithm follow by Lemma 4.

Weak evaluation binding. Let $(f(X), \alpha, \eta, \pi)$ be the output of an adversary in the game of Definition 6, where $C = \mathsf{Com}(\mathsf{ck}, f(X))$ is computed by the game. If π is an ordinary KZG proof, then opening f honestly at α yields a second accepting proof for the value $f(\alpha) \neq \eta$ at the same point, which contradicts the evaluation binding of KZG. If instead $\pi = \bot$, acceptance requires $C = C^\star$, that is, $[f(\tau)]_1 = [x^\star]_1$, and hence $f(\tau) = x^\star$. An algorithm that receives (gp, s), samples τ itself, assembles ck with the seed s, runs the adversary and returns $f(\tau)$ therefore computes x_s whenever the adversary wins using an empty proof, and the commitment key it supplies has exactly the distribution of KGen. Under the discrete-logarithm assumption in $\mathbb{G}_1$, the probability of such a win is negligible. Knowledge soundness in the AGM now follows from Theorem 2.

Failure of evaluation binding. On input ck, the adversary computes $C^\star$ from the public seed $s^\star$, picks an arbitrary $\alpha \in \mathbb{F}$, and outputs $(C^\star, \alpha, 0, \bot, 1, \bot)$, both of whose responses are accepted because the commitment is $C^\star$ and both claimed values lie in V. □

With Theorem 1 this gives the strict separation of Corollary 1 again, now for a scheme whose honest behavior is ordinary KZG. Taking $V = \{0\}$ instead preserves evaluation binding but rules out efficient extraction.

Theorem 6. *Fix a polynomially bounded degree d. The scheme $\Pi_{\mathsf{KZG}}[\{0\}]$ is AGM-clear and has an algebraic commitment algorithm. Under the discrete-logarithm assumption in $\mathbb{G}_1$ and the SBDH assumption under which KZG is evaluation binding, it holds that*

1. *$\Pi_{\mathsf{KZG}}[\{0\}]$ is evaluation binding, and hence knowledge-sound in the AGM, but*

2. *there exists a deterministic standard-model prover whose response to every challenge is accepted and for which every PPT extractor has negligible probability of producing any polynomial of degree at most d that opens its commitment.*

The second item holds even when the extractor receives the accepting transcript for a uniform challenge sampled after the commitment is fixed, and has access to the prover's code, complete state, and polynomially many executions or rewinds.

Proof. AGM-clarity and the algebraic commitment algorithm follow by Lemma 4.

Evaluation binding. Let $\mathcal{A}$ be a PPT adversary against evaluation binding. Two empty proofs both claim the value zero, so they cannot violate evaluation binding. A successful attack therefore produces either two ordinary KZG proofs, or one ordinary proof for a value $\eta \neq 0$ and one empty proof for the value zero, at the same point and for the same commitment. In the latter case, the commitment must be $C^\star$.

Consider first a successful mixed attack with $\eta = x^\star$. A reduction given (gp, s) samples τ itself, constructs ck with seed s, and runs $\mathcal{A}$. Returning the value in the ordinary proof computes x_s whenever this case occurs. The simulated key has exactly the distribution of KGen, so this case has negligible probability under the discrete-logarithm assumption.

For the remaining cases, a reduction against ordinary KZG evaluation binding receives a KZG key, samples $x \leftarrow \mathbb{F}$, and enumerates the encoding preimages $E^{-1}([x]_1) = \{s_1, \ldots, s_m\}$. Recall that we denote by ρ the maximum number of preimages. The reduction samples j uniformly from $\{1, \ldots, \rho\}$ and aborts if $j > m$. Otherwise, it appends s_j to the key and runs $\mathcal{A}$, so that the distinguished commitment of the simulated key is $C^\star = E(s_j) = [x]_1$ and $x^\star = x$. For fixed group parameters, each seed is selected with probability $1/(\rho p)$, so conditional on not aborting the seed is exactly uniform and independent of the KZG key. The non-abort probability is $|\mathcal{S}|/(\rho p)$, which is bounded below by an inverse polynomial uniformly over the group parameters.

If $\mathcal{A}$ produces two ordinary proofs for different values, the reduction forwards them. If $\mathcal{A}$ instead produces a successful mixed attack with $\eta \neq x$, the reduction pairs its ordinary proof with an honest KZG opening of the constant polynomial x at the same point. Both proofs concern the same commitment, since $\mathsf{Com}(\mathsf{ck}, x) = [x]_1 = E(s_j) = C^\star$, and their claimed values differ. Thus a non-negligible probability of either remaining case contradicts the evaluation binding of ordinary KZG. Knowledge soundness in the AGM follows from Theorem 1.

Failure of efficient extraction. The deterministic prover $\mathsf{P}^\star$ computes $C^\star = E(s^\star)$ from the public seed and fixes this commitment before seeing the challenge. On every challenge $\alpha \in \mathbb{F}$, it returns $(0, \perp)$, which the verifier accepts. Fix any PPT extractor Ext. Sample $\mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda)$ and $\mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d)$, fix $C^\star = E(s^\star)$, and then sample $\alpha \leftarrow \mathbb{F}$. Write $f \leftarrow \mathsf{Ext}(\mathsf{ck}, C^\star, \alpha, 0, \perp)$, allowing the extractor the additional access to $\mathsf{P}^\star$ stated in the theorem. We show that

$$\Pr\left[f \in \mathbb{F}[X]^{\leq d} \wedge \mathsf{Com}(\mathsf{ck}, f(X)) = C^\star\right] \in \mathsf{negl}(\lambda).$$

A reduction given (gp, s) samples $\tau \leftarrow \mathbb{F}$, constructs ck with seed s, computes $C^\star = E(s)$, and samples $\alpha \leftarrow \mathbb{F}$. It runs Ext on the resulting accepting transcript and returns $f(\tau)$ if the output is a polynomial of degree at most d, and zero otherwise. The key and transcript have exactly the distribution of the extraction experiment. The prover also has no secret state, since its commitment is computable from the public key and its response is always $(0, \perp)$. The reduction can therefore simulate all access to its code, state, executions, and rewinds. Whenever the extractor produces a commitment-consistent polynomial,

$$[f(\tau)]_1 = \mathsf{Com}(\mathsf{ck}, f(X)) = C^\star = E(s) = [x_s]_1,$$

and hence $f(\tau) = x_s$. A non-negligible extraction probability would therefore contradict the hardness of recovering the encoded scalar from the seed. Transcript-consistent extraction additionally requires $f(\alpha) = 0$, so its success probability is negligible as well. □

We note that the second item uses only the discrete-logarithm assumption, as its reduction goes through the hardness of recovering the encoded scalar from the seed. SBDH is required for the evaluation binding of KZG, which the first item needs.

Remark 6 (Efficient and unbounded extraction). Theorem 6 also rules out expected polynomial-time extraction. If an extractor has expected running time at most $t(\lambda)$ and non-negligible success probability, there is a polynomial r such that it succeeds with probability at least $1/r(\lambda)$ for infinitely many λ. Truncating it after $2t(\lambda)r(\lambda)$ steps preserves success probability at least $1/(2r(\lambda))$ for those λ, contradicting the theorem. In contrast to Theorem 3, the KZG commitment need not rule out the existence of a transcript-consistent polynomial. For $d \geq 1$ and $\alpha \neq \tau$, the polynomial $f(X) = x^\star(X-\alpha)(\tau-\alpha)^{-1}$ satisfies $\mathsf{Com}(\mathsf{ck}, f(X)) = C^\star$ and $f(\alpha) = 0$. The exceptional event $\alpha = \tau$ has probability $1/p$ for a uniform challenge, and no transcript-consistent polynomial exists on this event when $x^\star \neq 0$. Thus, the succinct construction separates knowledge soundness in the AGM from efficient extraction, whereas the coefficient-wise construction rules out even unbounded extraction.

A Separation for Fully Algebraic Adversaries. We now consider knowledge soundness against fully algebraic adversaries, which must explain all group elements they output, including those in the evaluation proof. Formally, we restrict the adversaries in Definition 4 accordingly, leaving its quantifiers and the extractor's input unchanged. The extractor still receives only the commitment and its algebraic explanation. We show that this restriction admits an AGM-clear scheme that fails even WEB.

Use the commitment key and commitment algorithm of $\Pi_{\mathsf{KZG}}[V]$ above, over asymmetric bilinear groups, with τ and $s^\star$ sampled independently. The group elements of ck in $\mathbb{G}_1$ are exactly $[1]_1, [\tau]_1, \ldots, [\tau^d]_1$, and $s^\star$ is an ordinary seed rather than an additional group element. We change only the evaluation proofs and their verification. An ordinary proof is $(0, \pi_{\mathsf{KZG}})$, where π_{KZG} is checked by the KZG verification equation, and honest openings produce proofs of this form. An exceptional proof is $(1, T)$, with $T \in \mathbb{G}_1$, and is accepted exactly when $C = [0]_1$, $\eta = 1$, and $T = E(s^\star)$. The verifier rejects all other inputs. In particular, T is a group-valued component of the proof, and $E(s^\star)$ is not supplied to the adversary as a group element in the key. Write $\Pi^{\mathsf{proof}}_{\mathsf{KZG}}$ for this scheme. It is complete, since honest proofs use the ordinary verification equation and KZG is complete. Its commitments contain one group element, and its proofs contain one group element and a bit distinguishing the two forms.

Theorem 7. *The scheme $\Pi^{\mathsf{proof}}_{\mathsf{KZG}}$ is AGM-clear, has an algebraic commitment algorithm, and fails WEB against an efficient adversary that succeeds with probability one. Under the discrete-logarithm assumption in $\mathbb{G}_1$ and the SBDH assumption under which KZG is evaluation binding, it is knowledge-sound against fully algebraic adversaries.*

Proof. AGM-clarity. An algebraic explanation $\boldsymbol{\gamma}_C = (\gamma_0, \ldots, \gamma_d)$ of a commitment C gives $C = \sum_{i=0}^{d} \gamma_i [\tau^i]_1 = \left[f_{\boldsymbol{\gamma}}(\tau)\right]_1$, for $f_{\boldsymbol{\gamma}}(X) = \sum_{i=0}^{d} \gamma_i X^i$. Thus the revealer returns $(f_{\boldsymbol{\gamma}}, \perp)$, and the commitment algorithm is algebraic with explanation $(f_0, \ldots, f_d)$, exactly as for KZG.

Failure of WEB. The adversary outputs the zero polynomial, empty commitment randomness, the point $\alpha = 0$, the value $\eta = 1$, and the proof $(1, E(s^\star))$. The WEB game computes $C = \mathsf{Com}(\mathsf{ck}, 0) = [0]_1$, so the proof is accepted although $\eta \neq f(\alpha)$. The adversary therefore wins with probability one.

Knowledge soundness against fully algebraic adversaries. Fix the deterministic extractor that returns $(f_{\boldsymbol{\gamma}}, \perp)$ from the commitment explanation, and let $\mathcal{A}$ be any fully algebraic PPT adversary. The extracted polynomial always has degree at most d and opens C. It remains to show that an accepted evaluation disagrees with $f_{\boldsymbol{\gamma}}(\alpha)$ only with negligible probability.

If $\mathcal{A}$ produces an accepted ordinary proof for $\eta \neq f_{\boldsymbol{\gamma}}(\alpha)$, opening $f_{\boldsymbol{\gamma}}$ honestly at α gives a second ordinary KZG proof for a different value at the same point. An adversary against ordinary KZG evaluation binding can sample $s^\star$ independently, run $\mathcal{A}$ with the augmented key, and output these two KZG proofs. The probability of this event is therefore negligible under the assumption that KZG is evaluation binding.

If $\mathcal{A}$ produces an accepted exceptional proof $(1, T)$, it must also supply an algebraic explanation $\boldsymbol{\delta} = (\delta_0, \ldots, \delta_d)$ of T with $E(s^\star) = T = \sum_{i=0}^{d} \delta_i [\tau^i]_1$. A reduction given (gp, s) samples τ itself, constructs ck with s in place of $s^\star$, and runs $\mathcal{A}$. This gives exactly the distribution of the scheme's commitment key. On every accepted exceptional proof, the reduction returns $\sum_{i=0}^{d} \delta_i \tau^i = x_s$, contradicting the hardness of recovering the encoded scalar from the seed. If an explanation refers to earlier algebraic outputs, substituting

their explanations gives the same representation relative to the commitment key. Hence the probability of an accepted exceptional proof is negligible under the discrete-logarithm assumption. The extractor therefore succeeds except with negligible probability, independently of the choice of $\mathcal{A}$. It also suffices when the extractor receives the evaluation proof and its algebraic explanation, since it can ignore them. □

The WEB adversary in this proof knows the committed polynomial and can explain its commitment. What it cannot explain without breaking the discrete-logarithm assumption is the group element in its exceptional proof. Such an adversary is allowed by Definition 4, but it cannot supply the explanations required of a fully algebraic adversary with non-negligible probability. This is precisely why the reverse reduction in Lemma 2 needs unrestricted evaluation proofs: knowing an opening lets it explain the commitment, but not an arbitrary forged proof.

In Appendix C, we characterize the fully algebraic variant in the standard model as well, under a structural hypothesis on the evaluation proof.

Acknowledgments. We are grateful to Charlotte Hoffmann for helpful discussions during the initial stage of this work. We also gratefully acknowledge the constructive feedback provided by the anonymous reviewers, which substantially contributed to improving the clarity of the presentation of our results. Petr Chmel was partially supported by the Czech Science Foundation GAČR grant No. 24-10306S, Charles Univ. GA UK project No. 70924 and Charles Univ. project UNCE 24/SCI/008. Pavel Hubáček was partially supported by the Academy of Sciences of the Czech Republic (RVO 67985840), Czech Science Foundation GAČR grant No. 25-16311S, and by Zircuit. We have used generative AI for grammar, spell-checking, and/or basic editing.

References

ABK+21. Michel Abdalla, Manuel Barbosa, Jonathan Katz, Julian Loss, and Jiayu Xu. Algebraic adversaries in the universal composability framework. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 311–341. Springer, Cham, December 2021.

AHK20. Thomas Agrikola, Dennis Hofheinz, and Julia Kastner. On instantiating the algebraic group model from falsifiable assumptions. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 96–126. Springer, Cham, May 2020.

AHPP23. Benedikt Auerbach, Charlotte Hoffmann, and Guillermo Pascual-Perez. Generic-group lower bounds via reductions between geometric-search problems: With and without preprocessing. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part III*, volume 14371 of *LNCS*, pages 301–330. Springer, Cham, November / December 2023.

BBB+18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg 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.

BCFL23. David Balbás, Dario Catalano, Dario Fiore, and Russell W. F. Lai. Chainable functional commitments for unbounded-depth circuits. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part III*, volume 14371 of *LNCS*, pages 363–393. Springer, Cham, November / December 2023.

BCHO22. Jonathan Bootle, Alessandro Chiesa, Yuncong Hu, and Michele Orrù. Gemini: Elastic SNARKs for diverse environments. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 427–457. Springer, Cham, May / June 2022.

BDFG20. Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Efficient polynomial commitment schemes for multiple points and polynomials. Cryptology ePrint Archive, Report 2020/081, 2020.

BDFG21. Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Halo infinite: Proof-carrying data from additive polynomial commitments. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 649–680, Virtual Event, August 2021. Springer, Cham.

BDH+25. Juraj Belohorec, Pavel Dvorák, Charlotte Hoffmann, Pavel Hubácek, Kristýna Masková, and Martin Pastyrík. On extractability of the KZG family of polynomial commitment schemes. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VI*, volume 16005 of *LNCS*, pages 584–616. Springer, Cham, August 2025.

BFHK24. Balthazar Bauer, Pooya Farshim, Patrick Harasser, and Markulf Kohlweiss. The uber-knowledge assumption: A bridge to the AGM. *CiC*, 1(3):31, 2024.

BFS20. Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Cham, May 2020.

BHKM26. Juraj Belohorec, Pavel Hubáček, Aleksi Kalsta, and Kristýna Mašková. Chopin: Optimal pairing-based multilinear polynomial commitments from bivariate KZG. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part IX*, volume 16808 of *Lecture Notes in Computer Science*, pages 428–459. Springer, 2026.

BHS26. Juraj Belohorec, Pavel Hubáček, and Dominik Stejskal. On lifting AGM security to AGM with oblivious sampling. *IACR Cryptol. ePrint Arch.*, page 201, 2026.

BMM+21. Benedikt Bünz, Mary Maller, Pratyush Mishra, Nirvan Tyagi, and Psi Vesely. Proofs for inner pairing products and applications. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 65–97. Springer, Cham, December 2021.

CCC+25. Matteo Campanelli, Mario Carrillo, Ignacio Cascudo, Dario Fiore, Danilo Francati, and Rosario Gennaro. On the power of polynomial preprocessing: Proving computations in sublinear time, and more. Cryptology ePrint Archive, Report 2025/238, 2025.

CFT22. Dario Catalano, Dario Fiore, and Ida Tucker. Additive-homomorphic functional commitments and applications to homomorphic signatures. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 159–188. Springer, Cham, December 2022.

CGKY26. Alessandro Chiesa, Ziyi Guan, Christian Knabenhans, and Zihan Yu. On the Fiat-Shamir security of succinct arguments from functional commitments. In Nadia Heninger and Mike Rosulek, editors, *Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part IX*, volume 16808 of *Lecture Notes in Computer Science*, pages 189–213. Springer, 2026.

CHM+19. Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. Cryptology ePrint Archive, Report 2019/1047, 2019.

CHM+20. Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 738–768. Springer, Cham, May 2020.

Com21. Electric Coin Company. Halo2 Protocol. `https://zcash.github.io/halo2/design/protocol.html`, 2021. [Online; accessed 2-October-2025].

EFG22. Liam Eagen, Dario Fiore, and Ariel Gabizon. cq: Cached quotients for fast lookups. *IACR Cryptol. ePrint Arch.*, page 1763, 2022.

EG25. Liam Eagen and Ariel Gabizon. MERCURY: A multilinear polynomial commitment scheme with constant proof size and no prover FFTs. Cryptology ePrint Archive, Report 2025/385, 2025.

FKL18. Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 33–62. Springer, Cham, August 2018.

GMN22. Nicolas Gailly, Mary Maller, and Anca Nitulescu. SnarkPack: Practical SNARK aggregation. In Ittay Eyal and Juan A. Garay, editors, *FC 2022*, volume 13411 of *LNCS*, pages 203–229. Springer, Cham, May 2022.

GPS25. Chaya Ganesh, Sikhar Patranabis, and Nitin Singh. Samaritan: Linear-time prover SNARK from new multilinear polynomial commitments. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part V*, volume 16249 of *LNCS*, pages 456–489. Springer, Singapore, December 2025.

Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Berlin, Heidelberg, May 2016.

GT21. Ashrujit Ghoshal and Stefano Tessaro. Tight state-restoration soundness in the algebraic group model. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part III*, volume 12827 of *LNCS*, pages 64–93, Virtual Event, August 2021. Springer, Cham.

GWC19. Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019.

JM24. Joseph Jaeger and Deep Inder Mohan. Generic and algebraic computation models: When AGM proofs transfer to the GGM. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part V*, volume 14924 of *LNCS*, pages 14–45. Springer, Cham, August 2024.

KLPS26. Erki Külaots, Helger Lipmaa, Roberto Parisella, and Janno Siim. Special soundness and binding properties: A framework for tightly secure zk-SNARKs. In Joan Daemen and Emmanuel Thomé, editors, *EUROCRYPT 2026, Part VII*, volume 16547 of *LNCS*, pages 239–267. Springer, Cham, May 2026.

KP19. Julia Kastner and Jiaxin Pan. Towards instantiating the algebraic group model. Cryptology ePrint Archive, Report 2019/1018, 2019.

KT24. Tohru Kohrita and Patrick Towa. Zeromorph: Zero-knowledge multilinear-evaluation proofs from homomorphic univariate commitments. *Journal of Cryptology*, 37(4):38, October 2024.

KZG10. Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Berlin, Heidelberg, December 2010.

LPS23. Helger Lipmaa, Roberto Parisella, and Janno Siim. Algebraic group model with oblivious sampling. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 363–392. Springer, Cham, November / December 2023.

LPS24. Helger Lipmaa, Roberto Parisella, and Janno Siim. Constant-size zk-SNARKs in ROM from falsifiable assumptions. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 34–64. Springer, Cham, May 2024.

LPS25. Helger Lipmaa, Roberto Parisella, and Janno Siim. On knowledge-soundness of plonk in ROM from falsifiable assumptions. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 362–395. Springer, Cham, August 2025.

Mau05. Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Berlin, Heidelberg, December 2005.

MBKM19. Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge SNARKs from linear-size universal and updatable structured reference strings. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2111–2128. ACM Press, November 2019.

PK22. Jim Posen and Assimakis A. Kattis. Caulk+: Table-independent lookup arguments. *IACR Cryptol. ePrint Arch.*, page 957, 2022.

PST13. Charalampos Papamanthou, Elaine Shi, and Roberto Tamassia. Signatures of correct computation. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 222–242. Springer, Berlin, Heidelberg, March 2013.

Rot26. Tobias Rothmann. On the formal verification of polynomial commitments: two KZG constructions and the algebraic group model. *IACR Cryptol. ePrint Arch.*, 2026:1490, 2026.

Sho97. Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Berlin, Heidelberg, May 1997.

ZBK+22. Arantxa Zapico, Vitalik Buterin, Dmitry Khovratovich, Mary Maller, Anca Nitulescu, and Mark Simkin. Caulk: Lookup arguments in sublinear time. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, pages 3121–3134. ACM, 2022.

ZGK+25. Arantxa Zapico, Ariel Gabizon, Dmitry Khovratovich, Mary Maller, and Carla Ràfols. Baloo: Algebraic lookup arguments. *IACR Commun. Cryptol.*, 2(2):1, 2025.

Zha26. Yuncong Zhang. Vela and Carina: Fast pairing-based multilinear polynomial commitments from reciprocal polynomials. *IACR Cryptol. ePrint Arch.*, 2026:1438, 2026.

ZZK22. Cong Zhang, Hong-Sheng Zhou, and Jonathan Katz. An analysis of the algebraic group model. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part IV*, volume 13794 of *LNCS*, pages 310–322. Springer, Cham, December 2022.

A AGM-Clarity of Pairing-Based PCSs

AGM-clarity (Definition 5) asks for an efficient revealer that turns the algebraic explanation of a commitment into an opening. As observed in Section 3, this works when the group elements of ck in the commitment group are exactly those combined by Com. We make that condition precise in Appendix A.1 and show that it suffices, and Appendix A.2 extends it to hiding variants. A key that spans more polynomials than Com admits, however, rules AGM-clarity out when the extended commitment map is computationally binding, as Appendix A.3 shows. Two schemes in the literature allow such a key, and for them AGM-clarity holds only

Table 1. Commitment keys of the pairing-based polynomial commitments discussed in this appendix, grouped by the route through which each reaches the standard-model characterization. The schemes are KZG [KZG10], PST13 [PST13], Gemini [BCHO22], Mercury [EG25], Vela and Carina [Zha26], Samaritan [GPS25], and Zeromorph [KT24]. In the third block, the key spans more polynomials than the commitment algorithm admits, and the verifier's degree check makes up the difference, so that route is conditional. The index ranges in each row describe the monomial family the key encodes, whose size is the dimension of the class the scheme commits to in the first two blocks and of the larger class the key spans in the third. The trapdoors and the size parameters n, N, b, d, d_i and D are the schemes' own, with $n_{\max}$ Zeromorph's bound on the number of variables and $N_{\max} = 2^{n_{\max}}$.

Scheme	Elements of ck in $\mathbb{G}_1$	Route
KZG	$[\tau^i]_1$, $i \le d$	Read-off alone, Lemma 5
PST13	$[\tau_1^{e_1} \cdots \tau_\mu^{e_\mu}]_1$, $e_i \le d_i$	
Gemini	$[\tau^i]_1$, $i \le D$	
Mercury	$[\tau^i]_1$, $i < n$	
Vela	$[\rho^i]_1$, $i < N$	
Carina	$[\rho^i \sigma^j]_1$, $i, j < b$	
Samaritan, $D = n - 1$	$[\tau^i]_1$, $i < n$	
KZG, hiding	$[\tau^i]_1$, $[\beta\tau^i]_1$, $i \le d$	Read-off with blinders, Lemma 6
Zeromorph, $n = n_{\max}$	$[\tau^i]_1$, $i < N_{\max}$, and $[\varrho]_1$	
Samaritan, $D > n - 1$	$[\tau^i]_1$, $i \le D$	Read-off plus degree check, Theorem 8
Zeromorph, $n < n_{\max}$	$[\tau^i]_1$, $i < N_{\max}$, and $[\varrho]_1$	

when the key stops at the degree bound. For the case when it does not, we give a separate characterization. Table 1 groups the schemes by which of these three routes to a standard-model characterization each takes.

A.1 Commitment Keys That Form a Basis

Each of the schemes below commits by taking a linear combination of the group elements of ck whose coefficients are determined by the committed polynomial. What varies between them is only how a polynomial is turned into a coefficient vector, so we leave that step abstract and state the condition over the class of polynomials that a scheme actually commits to.

Definition 7 (Basis-Linear Commitment). *Let $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ be a PCS with a deterministic commitment algorithm. For $d \in \mathsf{poly}(\lambda)$, $\mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda)$ and $\mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d)$, let $\mathcal{P}$ be the finite-dimensional space of polynomials that Com admits under ck, and let $\bar{g} = (g_1, \ldots, g_m)$ list* all *elements of ck that lie in the group of commitments. The commitment algorithm is* basis-linear *if, for every such ck, there is an $\mathbb{F}$-linear bijection $L \colon \mathcal{P} \to \mathbb{F}^m$, computable and invertible in polynomial time given ck, with*

$$\mathsf{Com}(\mathsf{ck}, f) = \langle L(f), \bar{g} \rangle \qquad \text{for every } f \in \mathcal{P}.$$

In this appendix, $\mathcal{P}$ and $\mathcal{Q}$ are read relative to the commitment key at hand and any further parameters that Com takes, which we abbreviate "Π admits $\mathcal{P}$".

Lemma 5. *Every PCS with a basis-linear commitment algorithm is AGM-clear.*

Proof. Let $\mathsf{Rev}(\mathsf{ck}, (C, \boldsymbol{\gamma}_C))$ return $(L^{-1}(\boldsymbol{\gamma}_C), \bot)$, which is DPT because L^{-1} is computable in polynomial time given ck. The randomness component is empty because Com is deterministic. By the convention for the AGM over bilinear groups from Section 2, an algebraic explanation of a commitment is taken relative to the input group elements lying in the same group, and $\bar{g}$ lists all of them. Thus $\boldsymbol{\gamma}_C$ is a vector in $\mathbb{F}^m$ with $C = \langle \boldsymbol{\gamma}_C, \bar{g} \rangle$, so $f = L^{-1}(\boldsymbol{\gamma}_C)$ lies in $\mathcal{P}$ and satisfies

$$\mathsf{Com}(\mathsf{ck}, f) = \langle L(L^{-1}(\boldsymbol{\gamma}_C)), \bar{g} \rangle = \langle \boldsymbol{\gamma}_C, \bar{g} \rangle = C,$$

which is what Definition 5 requires. □

The condition holds for the schemes in the first block of Table 1, where τ, ρ, and σ denote trapdoors sampled by the key generator and $\bar{g}$ is the vector of group elements of ck lying in the group of commitments. For most of them the situation is the one KZG already illustrates, with $\bar{g}$ the monomial family the table lists and L the corresponding coefficient map. This covers KZG [KZG10], the multivariate variant of Papamanthou, Shi, and Tamassia [PST13], the univariate commitment underlying Gemini [BCHO22], Carina [Zha26], and Samaritan [GPS25] in the parameterization whose key stops at the degree bound. Two cases are worth spelling out.

- Mercury [EG25] has $\bar{g} = ([\tau^i]_1)_{i<n}$ and commits $f \in \mathbb{F}^n$ to $\sum_{i<n} f_i[\tau^i]_1$. Its own extractor is the revealer of Lemma 5, since the proof of knowledge soundness returns the vector that the algebraic adversary supplies with its commitment.
- Vela [Zha26] has $\bar{g} = ([\rho^i]_1)_{i<N}$ and commits a multilinear polynomial by taking its hypercube evaluations as the coefficients of a univariate polynomial of degree below N and committing to that. Its security analysis performs the read-off of Lemma 5 explicitly.

Nothing in Lemma 5 constrains $\bar{g}$ beyond its length. One statement therefore covers the powers of a single trapdoor, Carina's monomials $\rho^i\sigma^j$ in two trapdoors with both exponents below a common bound b, and the monomial sets of the multivariate schemes. Definition 5 is stated for univariate polynomials of degree at most d, whereas several of these schemes commit to multilinear or bivariate ones. This is not a gap, since Definition 7 works over the class $\mathcal{P}$ that Com actually admits, and AGM-clarity is then read over $\mathcal{P}$. Carina and the variant of Papamanthou, Shi, and Tamassia [PST13] commit a multivariate coefficient vector against multivariate monomials, without introducing a univariate polynomial.

A.2 Hiding Variants

Hiding variants extend the commitment key by blinding elements that Com combines along with the rest. The Pedersen-style example following Definition 5 uses a single blinding generator, and the schemes below may need several.

Lemma 6. *Let a PCS Π be as in Definition 7, except that* Com *is randomized and the elements of* ck *in the group of commitments are $\bar{g}$ together with a further vector $\bar{h} = (h_1, \ldots, h_k)$. Suppose the commitment randomness ranges over a space $\mathcal{R}$ with an $\mathbb{F}$-linear bijection $M\colon \mathcal{R} \to \mathbb{F}^k$, computable and invertible in polynomial time given* ck*, such that*

$$\mathsf{Com}(\mathsf{ck}, f; r) = \langle L(f), \bar{g}\rangle + \langle M(r), \bar{h}\rangle.$$

Then Π is AGM-clear.

Proof. An algebraic explanation of C relative to $(\bar{g}, \bar{h})$ splits as $(\boldsymbol{\gamma}, \boldsymbol{\delta}) \in \mathbb{F}^m \times \mathbb{F}^k$ with $C = \langle \boldsymbol{\gamma}, \bar{g}\rangle + \langle \boldsymbol{\delta}, \bar{h}\rangle$. Let $\mathsf{Rev}(\mathsf{ck}, (C, (\boldsymbol{\gamma}, \boldsymbol{\delta})))$ return $(L^{-1}(\boldsymbol{\gamma}), M^{-1}(\boldsymbol{\delta}))$. Then $L^{-1}(\boldsymbol{\gamma}) \in \mathcal{P}$, the randomness $M^{-1}(\boldsymbol{\delta})$ lies in $\mathcal{R}$, and $\mathsf{Com}(\mathsf{ck}, L^{-1}(\boldsymbol{\gamma}); M^{-1}(\boldsymbol{\delta})) = \langle \boldsymbol{\gamma}, \bar{g}\rangle + \langle \boldsymbol{\delta}, \bar{h}\rangle = C$. As noted after Definition 5, the revealer has to produce some valid opening rather than the randomness used by the committer. □

The above lemma applies to two hiding variants from the literature. The Pedersen-style commitment of Kate, Zaverucha, and Goldberg [KZG10] adds a second sequence of powers based on an independent generator $[\beta]_1$, so that $\bar{h} = ([\beta]_1, [\beta\tau]_1, \ldots, [\beta\tau^d]_1)$. Its randomness is a polynomial $\hat{f} \in \mathbb{F}[X]^{\leq d}$, and M is the coefficient map of $\hat{f}$. The hiding variant of Zeromorph [KT24] takes $k = 1$. The elements of its commitment key in $\mathbb{G}_1$ are $[1]_1, [\tau]_1, \ldots, [\tau^{N_{\max}-1}]_1$ together with a single generator $[\varrho]_1$ for an independent ϱ, where $N_{\max} = 2^{n_{\max}}$ for the bound $n_{\max}$ on the number of variables fixed at setup. A commitment to a multilinear f in n variables is $[U_n(f)(\tau)]_1 + r[\varrho]_1$ for a scalar r, where $U_n(f)$ is the univariate representation of f whose coefficients are the values of f on the Boolean hypercube. This covers Zeromorph when n matches $n_{\max}$, and, in Appendix A.3, we treat the remaining case. We make no claim about hiding variants of the other schemes above; Vela and Carina, in particular, are not hiding [Zha26].

A.3 Keys That Span More Than the Admissible Class

Some schemes generate a commitment key large enough for more polynomials than their commitment algorithm admits, and enforce the smaller class inside the evaluation protocol through a degree check. AGM-clarity then fails when the extended commitment map is computationally binding, and no better proof would help. The revealer sees a commitment and its algebraic explanation, with no transcript and no evaluation proof, so a check the verifier performs during an evaluation is not available to it.

Lemma 7. *Let a PCS Π admit the class $\mathcal{P}$, and suppose its commitment map extends to a class $\mathcal{Q}$ having $\mathcal{P}$ as a proper linear subspace, on which it is basis-linear in the sense of Definition 7 and computationally binding. Then Π is not AGM-clear.*

Proof. Suppose Π is AGM-clear with revealer Rev. Sample $\boldsymbol{\gamma} \leftarrow \mathbb{F}^m$ uniformly and set $q = L^{-1}(\boldsymbol{\gamma})$, which is uniform on $\mathcal{Q}$ and so lies outside $\mathcal{P}$ with probability at least $1 - 1/|\mathbb{F}|$, as $\mathcal{P}$ is a proper subspace of $\mathcal{Q}$. Let $C = \mathsf{Com}(\mathsf{ck}, q)$, whose algebraic explanation is $\boldsymbol{\gamma}$. Running $\mathsf{Rev}(\mathsf{ck}, (C, \boldsymbol{\gamma}))$ yields (f, r) with $f \in \mathcal{P}$ and $\mathsf{Com}(\mathsf{ck}, f; r) = C$, so $f \neq q$ whenever $q \notin \mathcal{P}$. The two openings q and f of the same commitment then break binding on $\mathcal{Q}$, with probability at least $1 - 1/|\mathbb{F}|$ over the choice of $\boldsymbol{\gamma}$, and every step runs in polynomial time. □

Samaritan [GPS25] permits the commitment key to carry powers $[\tau^i]_1$ past the degree bound $n-1$ of the committed polynomial, and enforces that bound inside the evaluation protocol through a shifted pairing check. Writing D for the degree of the commitment key, the extended map is KZG on $\mathbb{F}[X]^{\leq D}$, which is binding under D-DLOG, so Lemma 7 applies whenever $D > n-1$. Concretely, the commitment $[\tau^n]_1$ comes with a unit-vector explanation, and a revealer would return some $f \in \mathbb{F}[X]^{<n}$ with $f(\tau) = \tau^n$. Then τ is a root of the nonzero polynomial $X^n - f(X)$, whose roots can be found in randomized polynomial time and tested against $[\tau]_1$. When $D = n-1$, the commitment key spans exactly $\mathbb{F}[X]^{<n}$, so Lemma 5 applies.

Zeromorph [KT24] separates key capacity from the admissible class in the same way. Its setup takes a bound $n_{\max}$ and produces powers through $N_{\max} - 1$ for $N_{\max} = 2^{n_{\max}}$, whereas Com takes the number of variables $n \leq n_{\max}$ as a further input and commits $U_n(f)$, of degree below 2^n. The underlying commitment is binding on $\mathbb{F}[X]^{<N_{\max}}$ under $(N_{\max}-1)$-DLOG, so Lemma 7 applies whenever $n < n_{\max}$, and the blinding of the hiding variant does not change this. Remark 5 exhibits the same phenomenon in isolation, with a commitment key extended by a single group element.

The verifier's degree check nonetheless makes both schemes knowledge-sound in the AGM against fully algebraic adversaries. The read-off still returns a polynomial, only one in the wider class $\mathcal{Q}$, and what the degree check has to supply is that a read-off landing outside $\mathcal{P}$ cannot survive to an accepting transcript. We state this as a property in its own right. It is a standard-model property, and like WEB it has the adversary hand over the polynomial. Remark 7 separates it from the degree binding of the literature, which it resembles in name only.

Definition 8 (Admissible-Opening Soundness). *Let $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ admit the class $\mathcal{P}$ and let $\mathcal{Q} \supseteq \mathcal{P}$ be a class on which Com is also defined. Then Π is* admissible-opening sound *with respect to $\mathcal{Q}$ if, for every $d \in \mathsf{poly}(\lambda)$ and every PPT adversary $\mathcal{A}$,*

$$\Pr\left[\begin{array}{l|l} \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \\ C = \mathsf{Com}(\mathsf{ck}, q; r) \wedge & \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ q \in \mathcal{Q} \setminus \mathcal{P} & (C, q, r, \alpha, \eta, \pi) \leftarrow \mathcal{A}(\mathsf{ck}) \end{array}\right] \in \mathsf{negl}(\lambda).$$

Stating this with an opening the adversary supplies is not a stylistic choice. The alternative reading, that no polynomial of $\mathcal{P}$ opens C, is vacuous. The group of commitments has prime order and $\mathcal{P}$ is closed under scalar multiplication, so the scalar multiples of a single $q \in \mathcal{P}$ with a non-trivial commitment already open every commitment. Remark 6 likewise distinguishes the existence of a transcript-consistent polynomial from the ability to compute one.

Theorem 8. *Let Π admit the class $\mathcal{P}$, with membership in $\mathcal{P}$ decidable in polynomial time given* ck*, let* Com *be basis-linear on a class $\mathcal{Q} \supseteq \mathcal{P}$ spanned by the commitment-group elements of* ck *and computationally binding on $\mathcal{Q}$, and let Π have an algebraic commitment algorithm. Then Π is knowledge-sound in the AGM if and only if it is admissible-opening sound with respect to $\mathcal{Q}$ and weakly evaluation binding.*

Proof. Let Ext compute $q = L^{-1}(\boldsymbol{\gamma}_C)$ and return $(q, \perp)$ if $q \in \mathcal{P}$ and $(0, \perp)$ otherwise. Consider an adversary winning the game of Definition 4, so that $\mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1$. If $q \notin \mathcal{P}$, then q opens C because Com is basis-linear on $\mathcal{Q}$, and the adversary followed by L^{-1} is a PPT algorithm outputting $(C, q, \perp, \alpha, \eta, \pi)$, which contradicts admissible-opening soundness. If $q \in \mathcal{P}$, then $C = \mathsf{Com}(\mathsf{ck}, q)$, so winning forces $q(\alpha) \neq \eta$, and the reduction of Lemma 1 turns the adversary into one against WEB.

Conversely, WEB follows from Lemma 2. For admissible-opening soundness, let $\mathcal{A}$ output $(C, q, r, \alpha, \eta, \pi)$ with $q \in \mathcal{Q} \setminus \mathcal{P}$, $C = \mathsf{Com}(\mathsf{ck}, q; r)$ and $\mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1$. As $\mathcal{A}$ holds q, it also computes the explanation $L(q)$ of C, so it is algebraic with respect to its commitment. Knowledge soundness then supplies (f, r') with $f \in \mathcal{P}$ and $C = \mathsf{Com}(\mathsf{ck}, f; r')$, where $f \neq q$ because $q \notin \mathcal{P}$, so the two openings break binding on $\mathcal{Q}$. □

Composing with Proposition 1 gives the version that applies to these schemes, in which admissible-opening soundness together with evaluation binding implies knowledge soundness in the AGM. The hiding case goes through with the revealer of Lemma 6 in place of L^{-1}, with binding required of the full randomized map $(f, r) \mapsto \langle L(f), \bar{g} \rangle + \langle M(r), \bar{h} \rangle$ on $\mathcal{Q} \times \mathcal{R}$.

Samaritan and Zeromorph meet the hypotheses in the parameterizations of Appendix A.3, with Zeromorph's n fixed by the parameterization rather than chosen per commitment, and with binding on $\mathcal{Q}$ holding under the assumptions their own analyses use. Admissible-opening soundness is what their degree checks are for, and those analyses establish it against adversaries that explain every group element they output. The dependence on those explanations is not incidental. Both checks are shifted pairing equations whose proof elements the verifier never recomputes, and each argument builds from the polynomials behind those elements a nonzero polynomial vanishing at τ, in Zeromorph [KT24, Section 5.1] from the two group elements of the degree proof and in Samaritan [GPS25, Section 4.3] from the aggregate behind cm_t. By Remark 1, Definition 4 leaves exactly those elements unexplained, so at that interface the hypothesis asks for more than the AGM alone supplies. Kohrita and Towa call a standard-model proof of their degree check the main difficulty. Their rewinding argument [KT24, Section 5.4] concerns generic batching and assumes knowledge soundness of the underlying degree-check and evaluation protocols, so it does not on its own establish standard-model security of the concrete shifted pairing check.

Remark 7 (Relation to degree binding). Definition 8 is not the degree binding of Külaots et al. [KLPS26, Definition 8]. Degree binding asks that, for degree bound d, no adversary produce $d + 2$ or more accepting evaluations at pairwise distinct points whose interpolating polynomial has degree above d. The two games differ in what carries the violation and in how many accepting evaluations it takes. No opening of the commitment appears in theirs, and the polynomial carrying the violation need not open the commitment at all, whereas ours turns on a single accepting evaluation together with an opening that the adversary exhibits. We are not aware of an implication in either direction. Definition 8 isolates the class-membership conjunct of the evaluation relation used since Bünz, Fisch, and Szepieniec [BFS20], much as WEB isolates the conjunct tying the opening to the claimed value.

The collapse is therefore not avoided by checking degrees. The characterization simply acquires one further standard-model conjunct, which constrains degrees rather than openings and so leaves the separations of Section 5 untouched.

B Extraction Without Commitment Consistency

The extraction separations in Section 5 require the extracted polynomial to open the commitment and to agree with the accepted evaluation. The definition Chiesa et al. [CHM+19, Definition B.2] use for their compiler drops commitment consistency but requires the extractor to fix a bounded-degree polynomial before

the evaluation challenge. The polynomial must agree with the accepted evaluation, but need not open the commitment. That suffices for their purpose, since the compiler replaces commitments by polynomial oracles and the underlying verifier reads evaluations rather than commitments. The zero-answer provers used above do not defeat this target, so we give a different prover here.

We use a simplified requirement with one commitment and one query, and give the extractor the full state of the commitment stage. The extractor must fix its polynomial before the evaluation point is drawn; Remark 8 explains why this timing matters.

Definition 9 (Commitment-Time Evaluation-Only Extraction). *A PCS* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *admits* commitment-time evaluation-only extraction *if, for every* $d \in \mathsf{poly}(\lambda)$ *and every PPT adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, there is a PPT extractor* Ext *such that for every PPT query sampler* $\mathcal{D}$*,*

$$\Pr\left[\begin{array}{l|l} \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \\ (\deg f > d \vee f(\alpha) \neq \eta) & \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ & (C, \mathsf{st}) \leftarrow \mathcal{A}_1(\mathsf{ck}); \\ & f \leftarrow \mathsf{Ext}(\mathsf{ck}, C, \mathsf{st}); \\ & \alpha \leftarrow \mathcal{D}(\mathsf{ck}, C); \\ & (\eta, \pi) \leftarrow \mathcal{A}_2(\mathsf{st}, \alpha) \end{array}\right] \in \mathsf{negl}(\lambda).$$

The state st *contains the full commitment-stage state, including the random coins used by* $\mathcal{A}_1$*. The query sampler uses fresh randomness after* Ext *outputs* f*.*

The extractor may depend on both stages of $\mathcal{A}$ and may run them before fixing f. The impossibility below holds even with this information and access.

The zero-answer provers do not suffice. For $\Pi_{\mathsf{os}}[\{0\}]$ and $\Pi_{\mathsf{KZG}}[\{0\}]$, the prover $\mathsf{P}^\star$ answers every challenge with the value zero, and the zero polynomial has degree at most d and agrees with every such answer. An extractor that outputs it before seeing α succeeds against this particular prover in Definition 9. Those attacks therefore do not establish a failure of commitment-time evaluation-only extraction.

A scheme breaking commitment-time evaluation-only extraction. Keep the commitment key, the commitment algorithm and the honest opening of $\Pi_{\mathsf{KZG}}[\{0\}]$, defined in Section 5, and change only the rule for the empty proof, which is now accepted exactly when

$$C = C^\star \qquad \text{and} \qquad \eta = \alpha^{d+1}.$$

Ordinary proofs are verified by the KZG equation as before, so the scheme is complete, and its commitments and honest evaluation proofs remain single group elements. We write $\Pi^{\mathsf{pow}}_{\mathsf{KZG}}$ for the result. The exceptional answers follow a polynomial of degree $d+1$, exceeding the permitted bound; Remark 9 relates this to degree binding. The rule $\eta = \alpha^{d+1}$ is due to Külaots et al. [KLPS26, Lemma 4(ii)], who use it to keep evaluation binding while breaking degree binding. What is new here is the commitment it applies to: $C^\star$ is a valid commitment that an algebraic adversary cannot explain without computing $x^\star$, which is what leaves the scheme knowledge-sound in the AGM.

Theorem 9. *Fix a polynomially bounded degree* d *and assume* $(d+1)/|\mathbb{F}|$ *is negligible. The scheme* $\Pi^{\mathsf{pow}}_{\mathsf{KZG}}$ *is AGM-clear and has an algebraic commitment algorithm. Under the discrete-logarithm assumption in* $\mathbb{G}_1$ *and the SBDH assumption under which KZG is evaluation binding [KZG10], it is evaluation binding and hence knowledge-sound in the AGM. There is nonetheless a deterministic standard-model prover that is accepted on every challenge and against which no extractor, of any running time, meets the success requirement of Definition 9 for the uniform query sampler.*

Proof. AGM-clarity and the algebraic commitment algorithm are those of $\Pi_{\mathsf{KZG}}[V]$, established there for every admissible V, and depend only on the commitment key and the commitment algorithm.

Evaluation binding. Consider two accepting proofs for the same commitment at the same point α, claiming $\eta \neq \eta'$. Two empty proofs both claim α^{d+1} and cannot differ. Two ordinary proofs contradict the evaluation binding of KZG. In the remaining case one proof is empty, so the commitment is $C^\star$, and the other is an ordinary proof claiming some value η'. If $\eta' = x^\star$, the reduction of Theorem 6 that samples τ itself and returns the value of the ordinary proof computes x_s. Otherwise the preimage-sampling reduction of the same proof pairs the ordinary proof with an honest KZG opening of the constant polynomial $x^\star$, whose proof is $[0]_1$ and needs no trapdoor, and the two claimed values differ. Neither argument refers to the value the empty proof claims, so replacing zero by α^{d+1} leaves it intact. Knowledge soundness in the AGM follows by Theorem 1.

Failure of evaluation-only extraction. The prover computes $C^\star = E(s^\star)$ from the public seed, fixes it before seeing the challenge, and answers α with $(\alpha^{d+1}, \bot)$, which the verifier accepts. Let Ext be any extractor, even one given the prover's code and full state, and let f be its output. Condition on the key, this state, and all randomness used before f is fixed. The challenge α is then sampled independently and uniformly from $\mathbb{F}$. If $\deg f > d$ the winning condition already holds, so suppose $\deg f \leq d$. Then $X^{d+1} - f(X)$ is a nonzero polynomial of degree $d+1$ and has at most $d+1$ roots, so for uniform α

$$\Pr\big[f(\alpha) = \alpha^{d+1}\big] \leq \frac{d+1}{|\mathbb{F}|}.$$

The accepted value is α^{d+1}, so $f(\alpha) \neq \eta$ except with negligible probability. The bound counts roots and does not limit the running time of Ext. □

Remark 8 (Extraction before the challenge). If the extractor were allowed to see the accepted pair (α, η) before choosing f, it could satisfy the evaluation and degree requirements by returning the constant polynomial η. Definition 9 requires f to be fixed before the evaluation point is drawn, as does the definition of Chiesa et al. The impossibility relies on this timing, not on computational hardness.

Remark 9 (Degree binding). The prover's exceptional answers at $d+2$ pairwise-distinct points interpolate to X^{d+1}. Thus $\Pi^{\mathsf{pow}}_{\mathsf{KZG}}$ fails the degree binding of Külaots et al. [KLPS26, Definition 8].

Marlin does not sidestep Theorem 9. Chiesa et al. prove that their own polynomial commitment satisfies their extraction definition in the AGM [CHM+19, Definition B.2]. Their concrete extractor obtains more than the definition asks: it reads a genuine opening from the algebraic explanation and uses evaluation binding to establish agreement with the accepted evaluation [CHM+19, Lemma B.14].

To apply the attack of Theorem 9 to their game, take one commitment with degree bound d, an empty preliminary challenger message, and a fresh uniform query α. The commitment adversary outputs $C^\star$, and the opening adversary returns $(\alpha^{d+1}, \bot)$, ignoring the opening challenge. This query sampler is admissible in their sense [CHM+19, Definition 6.5]. The extracted polynomial is fixed before α, so the same root bound applies.

Thus AGM knowledge soundness does not imply commitment-time evaluation-only extraction against standard-model adversaries, even for evaluation-binding, AGM-clear schemes. Marlin's compiler can use their extraction theorem within the AGM; our separation rules out reading that theorem as a standard-model extraction guarantee.

C Standard-Model Characterization of Fully Algebraic Knowledge Soundness

For AGM-clear and AGM-proof-clear schemes with algebraic commitment and opening algorithms, we show that knowledge soundness against fully algebraic adversaries is equivalent to restricted WEB. This game restricts the evaluation proof to one generated by the opening algorithm from an adversary-supplied polynomial. The modified KZG construction of Theorem 7 shows that restricted WEB is strictly weaker than WEB (Corollary 2).

Throughout, Π is *knowledge-sound against fully algebraic adversaries* if it satisfies Definition 4 with the adversary additionally returning an algebraic explanation of every group element of π, the quantifiers and the extractor's input being unchanged.

Definition 10 (Restricted Weak Evaluation Binding). *A polynomial commitment scheme* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *satisfies* restricted weak evaluation binding *if, for any* $d \in \mathsf{poly}(\lambda)$ *and PPT adversary* $\mathcal{A}$*,*

$$\Pr\left[\begin{array}{l|l} \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ \eta \neq f(\alpha) & (f(X), r, g(X), r', \alpha) \leftarrow \mathcal{A}(\mathsf{ck}); \\ & C \leftarrow \mathsf{Com}(\mathsf{ck}, f(X); r); \\ & (\eta, \pi) \leftarrow \mathsf{Open}(\mathsf{ck}, C, \alpha, g(X); r') \end{array}\right] \in \mathsf{negl}(\lambda).$$

The adversary supplies two polynomials f and g of degree at most d. The game commits to f and uses g to generate the evaluation proof. There is no requirement that $C = \mathsf{Com}(\mathsf{ck}, g; r')$, and the verifier's acceptance check is retained. As in Definition 6, the degree bound is a typing requirement on the output rather than a winning condition. Nothing in the game refers to a group, a basis or an explanation, and no extractor occurs in it. Running Com and Open and returning the resulting evaluation claim and proof gives an adversary in the game of Definition 6, so WEB implies restricted weak evaluation binding for every PCS.

To connect this game to fully algebraic adversaries, we introduce a condition that allows an accepted explained proof to be replaced by one generated through Open.

Definition 11 (AGM-Proof-Clarity). *Let* $\bar{g}$ *denote the group elements of* ck *and* $\bar{\pi}$ *the group-valued components of an evaluation proof* π*. A PCS* $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ *is* AGM-proof-clear *if there is a DPT* proof revealer Rev_π *such that, for every* $d \in \mathsf{poly}(\lambda)$ *and every PPT adversary* $\mathcal{A}$*,* Rev_π *always returns a polynomial of degree at most* d *together with randomness, and*

$$\Pr\left[\begin{array}{l|l} \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta, \pi) = 1 \wedge & \mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda); \mathsf{ck} \leftarrow \mathsf{KGen}(\mathsf{gp}, d); \\ C = \boldsymbol{\gamma}_C \bar{g} \ \wedge \ \bar{\pi} = \boldsymbol{\gamma}_\pi \bar{g} \wedge & ((C, \boldsymbol{\gamma}_C), \alpha, \eta, (\pi, \boldsymbol{\gamma}_\pi)) \leftarrow \mathcal{A}(\mathsf{ck}); \\ (\eta' \neq \eta \vee & (g(X), r) \leftarrow \mathsf{Rev}_\pi(\mathsf{ck}, (C, \boldsymbol{\gamma}_C), \alpha, \eta, (\pi, \boldsymbol{\gamma}_\pi)); \\ \mathsf{Vrf}(\mathsf{ck}, C, \alpha, \eta', \pi') = 0) & (\eta', \pi') \leftarrow \mathsf{Open}(\mathsf{ck}, C, \alpha, g(X); r) \end{array}\right] \in \mathsf{negl}(\lambda).$$

Over bilinear groups the explanations are taken in each source group separately, as in Definition 3.

AGM-clarity supplies a polynomial f and randomness r with $C = \mathsf{Com}(\mathsf{ck}, f; r)$. AGM-proof-clarity supplies a polynomial g and randomness r' from which the opening algorithm produces an accepted proof of the claimed value. The polynomial g may depend on the entire evaluation transcript and need not open C. Restricted WEB ensures that the resulting accepted value agrees with $f(\alpha)$.

The condition concerns accepted transcripts with valid explanations of both the commitment and the proof. The proof revealer need not reproduce π itself, only lead the opening algorithm to an accepted proof of the same value, allowing fresh opening randomness and different proof encodings. Unlike AGM-clarity, this is a computational reconstruction condition and may rely on a hardness assumption. We verify it for KZG in Remark 11.

Proof-clarity alone does not ensure consistency with the commitment. For example, take coefficient-wise commitments $\mathsf{Com}(\mathsf{ck}, f) = ([f_0], \ldots, [f_d])$ and let $\mathsf{Open}(\mathsf{ck}, C, \alpha, g)$ return $(g(\alpha), g)$. Let the verifier check only that the proof polynomial has degree at most d and evaluates to the claimed value, ignoring C. The proof revealer returns the supplied polynomial if its degree is at most d, and zero otherwise, in both cases with empty randomness. Thus the scheme is AGM-proof-clear. Nevertheless, a fully algebraic adversary can commit to zero and prove the value one using the constant polynomial 1. The zero polynomial is the only polynomial that opens the resulting commitment, so fully algebraic knowledge soundness fails.

Remark 10 (Algebraic opening algorithm). A polynomial commitment has an *algebraic opening algorithm* if Open returns, together with (η, π), an algebraic explanation of $\bar{\pi}$ with respect to $\bar{g}$, on every input and

not only on those where the commitment opens to the polynomial supplied. For the KZG-based schemes above the requirement is immediate, since an honest evaluation proof is a fixed linear combination of the commitment key.

Theorem 10. *Let $\Pi = (\mathsf{PGen}, \mathsf{KGen}, \mathsf{Com}, \mathsf{Open}, \mathsf{Vrf})$ be a PCS. If Π is AGM-clear and AGM-proof-clear and satisfies restricted weak evaluation binding, then it is knowledge-sound against fully algebraic adversaries. If Π has an algebraic commitment algorithm and an algebraic opening algorithm and is knowledge-sound against fully algebraic adversaries, then it satisfies restricted weak evaluation binding.*

The hypotheses are not symmetric, exactly as in Theorem 2, and the converse uses neither revealer.

Proof. For the first implication, let Rev and Rev_π be the revealers certifying AGM-clarity and AGM-proof-clarity, and let Ext return $\mathsf{Rev}(\mathsf{ck}, (C, \boldsymbol{\gamma}_C))$. Let $\mathcal{A}$ be a fully algebraic adversary that wins against Ext with probability $\varepsilon(\lambda)$, and write $((C, \boldsymbol{\gamma}_C), \alpha, \eta, \pi, \boldsymbol{\gamma}_\pi)$ for its output. Construct an adversary $\mathcal{B}$ for restricted WEB that runs $\mathcal{A}$ and applies Ext to obtain (f, r) from the commitment explanation. It then applies $\mathsf{Rev}_\pi(\mathsf{ck}, (C, \boldsymbol{\gamma}_C), \alpha, \eta, (\pi, \boldsymbol{\gamma}_\pi))$ to obtain (g, r') from the evaluation transcript and returns (f, r, g, r', α). It runs in probabilistic polynomial time, since both revealers do.

AGM-clarity supplies $f \in \mathbb{F}[X]^{\leq d}$ with $C = \mathsf{Com}(\mathsf{ck}, f(X); r)$, so whenever $\mathcal{A}$ wins, it does so through $f(\alpha) \neq \eta$. By Definition 11, the probability that $\mathcal{A}$ wins but a fresh call $(\eta', \pi') \leftarrow \mathsf{Open}(\mathsf{ck}, C, \alpha, g(X); r')$ fails to give an accepted proof with $\eta' = \eta$ is negligible. The game of Definition 10 recomputes the same commitment and samples precisely this opening. Thus $\mathcal{B}$ wins with probability at least $\varepsilon(\lambda) - \mathsf{negl}(\lambda)$, so restricted WEB implies that ε is negligible.

For the converse we argue contrapositively. Let $\mathcal{A}$ win the game of Definition 10 with non-negligible probability $\varepsilon(\lambda)$, and let $\mathcal{A}'$ run it to obtain (f, r, g, r', α). Committing to $f(X)$ with randomness r yields C together with its explanation $\boldsymbol{\gamma}_C$ by Remark 2. Opening $g(X)$ at α yields the adversarial pair (η, π) and, by Remark 10, an explanation $\boldsymbol{\gamma}_\pi$ of $\bar{\pi}$, and opening $f(X)$ at α yields (η^f, π^f) with an explanation $\boldsymbol{\gamma}_{\pi^f}$. The adversary $\mathcal{A}'$ tosses a fair coin and outputs $(C, \boldsymbol{\gamma}_C)$ with either $(\alpha, \eta, \pi, \boldsymbol{\gamma}_\pi)$ or $(\alpha, \eta^f, \pi^f, \boldsymbol{\gamma}_{\pi^f})$. Both branches are fully algebraic.

Conditioned on the internal execution winning the restricted-WEB game, both branches verify and $\eta \neq f(\alpha) = \eta^f$, by completeness. Any extractor receives exactly the same input $(\mathsf{ck}, (C, \boldsymbol{\gamma}_C))$ in both branches, independently of the coin. Its polynomial φ therefore differs from at least one of the two claimed values at α, so $\mathcal{A}'$ wins with conditional probability at least $1/2$. If the extractor's output does not open C, both branches win. Hence $\mathcal{A}'$ wins against every extractor with probability at least $\varepsilon(\lambda)/2$. As in Lemma 2, the adversary $\mathcal{A}'$ never invokes the extractor, so by Remark 3 the argument refutes both quantifications. □

Remark 11 (KZG is AGM-proof-clear). We use the $(d, 1)$-power discrete-logarithm assumption [BHS26, Definition 2]: for $\mathsf{gp} \leftarrow \mathsf{PGen}(1^\lambda)$ and $\tau \leftarrow \mathbb{F}$, every PPT algorithm given $(\mathsf{gp}, ([\tau^i]_1)_{i=0}^d, [1]_2, [\tau]_2)$ returns τ only with negligible probability. The commitment key of KZG spans exactly $\mathbb{F}[X]^{\leq d}$, so the map of Definition 7 is invertible and the two explanations read as polynomials f and h with $C = \mathsf{Com}(\mathsf{ck}, f)$ and $\pi = \mathsf{Com}(\mathsf{ck}, h)$. The proof revealer computes $g(X) = (X - \alpha)h(X) + \eta$ and returns $(g, \bot)$ when its degree is at most d, and $(0, \bot)$ otherwise. In the first case $g(\alpha) = \eta$ and the honest opening of g at α commits to $(g(X) - \eta)/(X - \alpha) = h(X)$, so it returns the accepted pair (η, π) itself. The probability of an accepted transcript with valid explanations requiring the fallback is negligible. Indeed, suppose such a transcript has a nonzero coefficient of X^d in h, so that $(X - \alpha)h$ has degree $d + 1$ whereas $f - \eta$ has degree at most d. The verification equation gives $[f(\tau) - \eta]_1 = [(\tau - \alpha)h(\tau)]_1$, so $f(X) - \eta - (X - \alpha)h(X)$ is a nonzero polynomial vanishing at τ. Its roots can be found in randomized polynomial time and tested against $[\tau]_2$, which is present in the key even when $d = 0$. This recovers τ, contradicting the assumption if accepted transcripts requiring the fallback occur with non-negligible probability. Otherwise h has degree at most $d - 1$ or vanishes, so g has degree at most d.

The restriction is not vacuous, and the separating scheme of the body already witnesses this.

Corollary 2. *Under the assumptions of Theorem 7, restricted weak evaluation binding is strictly weaker than WEB.*

Proof. WEB implies restricted weak evaluation binding for every PCS, as noted after Definition 10. For strictness, take the scheme $\Pi^{\mathsf{proof}}_{\mathsf{KZG}}$ of Theorem 7, which fails WEB against an efficient adversary that succeeds with probability one and is knowledge-sound against fully algebraic adversaries. Its commitment algorithm is that of KZG and is therefore algebraic, and Open computes an ordinary KZG proof as a fixed linear combination of the commitment key from the polynomial it is given, so the scheme also has an algebraic opening algorithm in the sense of Remark 10. The second implication of Theorem 10 therefore gives restricted weak evaluation binding. □

Only the two algebraic hypotheses are used here, and neither revealer, so the separation does not depend on $\Pi^{\mathsf{proof}}_{\mathsf{KZG}}$ being AGM-clear or AGM-proof-clear. In this scheme, the opening algorithm produces only ordinary KZG proofs, whereas the attack uses the exceptional verification rule. Restricted WEB therefore holds even though an adversary knowing the committed polynomial can produce an accepting false evaluation.