AICR · OPEN PROBLEM

One-Shot Decryption Without Extractability

Threshold public-key encryption with one-shot decryption from one-shot signatures and plain witness encryption

Status. Statement: AI-written from Alexandru Cojocaru, Aggelos Kiayias, Yu Shen, Petros Wallden, *Proactive Secret Sharing without Erasures*, IACR Cryptology ePrint Archive 2026, not yet checked by a human. Open. Under *extractable* witness encryption the primitive is constructed for every $f < 1/2$ (the paper's Theorem 2, p. 20); the paper reports no progress toward removing extractability and records that post-quantum extractable witness encryption "is only understood from strong knowledge-type assumptions, and obtaining satisfactory instantiations remains an important open problem" (p. 7). The single-receiver (non-threshold) variant is separately argued to be unachievable by a gentle-measurement rewinding attack, so the threshold structure is not negotiable.

Category. *Quantum Cryptography & Unclonability.*

Conjecture (informal). *A one-shot signature is a quantum primitive whose signing key self-destructs: from one key you can produce a signature on one message and never on a second one. Building on this, the paper defines a threshold public-key encryption scheme with "one-shot decryption": a committee of quantum key-holders can jointly decrypt one ciphertext of their choosing, but no matter what the committee does afterwards — even if every single key-holder is subsequently compromised and hands over its entire remaining state — a second ciphertext under the same public keys can never be decrypted. This is the engine that makes proactive secret sharing possible without ever asking anyone to erase anything: the act of decrypting consumes the shares in a physically irreversible way. The paper constructs the primitive, but only by assuming witness encryption with an extractability guarantee, meaning that any adversary who can distinguish two ciphertexts must be accompanied by an efficient algorithm that literally reads a witness out of it. That is a knowledge-type assumption of a kind that is poorly understood post-quantum. The conjecture is that the extractability crutch is unnecessary: one-shot decryption should be achievable with no knowledge assumption at all and, in the reading formalised below, from one-shot signatures together with witness encryption in its ordinary, non-extractable form. Settling it would tell us whether proactive security without erasures rests on standard-model tools or on a knowledge assumption.*

1 The setting

Proactive secret sharing [OY91, HJKY95] keeps a secret alive against a *mobile* adversary that may, over time, corrupt every share-holder, by refreshing shares each epoch. Every scheme in this line, and every application of it — notably evolving-committee secret sharing [BGG+20] and YOSO-style MPC [GHK+21] — assumes that share-holders can *securely erase* their old state. Secure erasure is a strong physical assumption, and the source paper proves that without it proactive secret sharing is impossible for any corruption threshold $f > 0$ in a purely classical model.

The paper's way out is quantum: measurement is irreversible, so a quantum share can be consumed rather than erased. The engine is a new primitive, *threshold public-key encryption with one-shot decryption*

(Definition 3). A single-receiver version of the same idea is shown to be hopeless: because decryption is essentially deterministic, a receiver can run it in superposition, measure the plaintext register and rewind, recovering an undisturbed key — exactly the trick behind single-decryptor encryption [GZ20]. Thresholding defeats the rewinding attack, because a majority of *distinct* one-shot signing keys [AGKZ20, SZ25, HV25] must be genuinely consumed to decrypt, and $2(1-f)n > n$ for $f < 1/2$.

The construction is short: the sender picks a random tag and witness-encrypts [GGSW13] the message under the statement "at least $(1-f)n$ of the committee's one-shot public keys carry a valid signature on this tag"; partial decryption is just signing the tag. Its security proof, however, runs the witness-encryption *extractor* on a successful distinguisher to pull out $(1-f)n$ one-shot signatures, one of which must live under an honest public key, and then forges a second signature under that key to break one-shot-ness. Plain indistinguishability security is useless here, since the statement being encrypted to is *true*: honest parties can sign. So extractability is not a convenience in the proof, it is the entire reduction. The same is true of the proactive secret sharing protocols built on top and of the functional-witness-encryption variant [BCP14] used for evaluating a function on the shared secret without reconstructing it.

This leaves the whole edifice standing on a knowledge-type assumption. Post-quantum extractable witness encryption has no satisfying instantiation; [FHAS24] obtains extractable witness encryption in the algebraic group model but not post-quantum, and [ACL+22] is cited only as possible lattice-based inspiration. (External context, not a claim of the source paper: extractable witness encryption for all of NP with *arbitrary* auxiliary input — which is what Definition 2(ii) demands, with aux an arbitrary mixed state — has been argued to be implausible [GGHW14], a work the source paper does not cite.) The conjecture asks whether the extractability can simply be dropped: is threshold one-shot decryption achievable without any knowledge assumption — concretely, from one-shot signatures plus ordinary witness encryption? A positive answer would put proactive security without erasures on standard-model footing.

Progress to date. The paper proves the positive result under extractable witness encryption (Theorem 2, p. 20) and builds three

proactive secret sharing protocols (Protocols 1–3, Theorems 3–5, pp. 23/26/29), the third of which adds secret usability via functional witness encryption; Theorem 3 tolerates any $f < 1$ against a passive mobile adversary, Theorems 4 and 5 require $f < 1/2$ against an active one. Functional witness encryption is equivalent to extractability obfuscation [BCP14]. The paper identifies where extractability is used — the extractor must output at least $(1-f)n$ one-shot signatures for the reduction to one-shot-signature security to go through — and it points to extractable witness encryption from the algebraic group model, and to lattice-based knowledge assumptions, as the kind of more structured assumption that would already count as progress. It reports no progress toward removing extractability altogether.

Why it matters. It decides whether the first (and so far only) route to proactive security without erasures depends on a knowledge assumption that is unavailable post-quantum, or whether it can be based on standard tools. Stripped of the application, it is a clean question about quantum unclonability: does one-shot decryption — irreversible consumption of a decryption capability, verifiable by nobody and surviving total posterior corruption — follow from one-shot signatures, or is it strictly stronger? Either answer moves a real line: a construction removes a knowledge assumption from a headline feasibility result, and a separation would be the first evidence that thresholded one-shot-ness needs extraction.

Why it is clean. It is a single implication between two primitives, both fully definable in a page, with no reference to the paper's proactive secret sharing protocols, its epoch machinery, or its mobile-adversary model. The hypotheses are the paper's own ingredients minus one adjective. Nothing is sketched or unpublished: one-shot signatures and witness encryption are standard and constructed elsewhere, and the target definition is a self-contained three-property game.

2 Notation and parameters

Throughout, $\lambda \in \mathbb{N}$ is the security parameter, $\mathsf{negl}(\cdot)$ denotes an unspecified negligible function, and QPT abbreviates "quantum polynomial time". For $n \in \mathbb{N}$ we write $[n] = \{1, \ldots, n\}$. All parties and adversaries are QPT and all communication is classical: every

algorithm below outputs classical strings except where a state is written in ket notation, so pk denotes a classical public key and $|\mathsf{sk}\rangle$ a quantum secret key. $\mathcal{M}$ denotes a message space and $|m|$ the length of a message $m \in \mathcal{M}$. For an NP relation R we write $\mathcal{L}_{\mathsf{R}} = \{\mathsf{stmt} : \exists\, \mathsf{w},\ (\mathsf{stmt}, \mathsf{w}) \in \mathsf{R}\}$ for the associated language. The symbol aux denotes an arbitrary mixed quantum state given to an adversary as auxiliary input, possibly entangled with an external register.

Parameters.

- λ — security parameter; all algorithms run in quantum polynomial time in λ and all guarantees are asymptotic in λ.
- n — committee size, i.e. the number of quantum secret-key holders; any polynomial $n = n(\lambda)$.
- f — corruption threshold: the adversary holds at most fn of the n quantum secret keys before decryption, and at least $(1-f)n$ partial decryptions are needed to decrypt; the conjecture is for every constant $0 < f < 1/2$, while Definition 3 is stated for $f \in [0, 1/2)$.
- R — the NP relation the witness encryption scheme is for; the hypothesis quantifies over all NP relations.

3 The conjecture

Definition 1 (One-shot signatures). A *one-shot signature* (OSS) scheme with message space $\mathcal{M} \supseteq \{0,1\}^\lambda$ is a tuple of QPT algorithms $(\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Verify})$:

- $\mathsf{Setup}(1^\lambda) \to \mathsf{CRS}$: outputs a classical common reference string.
- $\mathsf{KeyGen}(\mathsf{CRS}) \to (\mathsf{pk}, |\mathsf{sk}\rangle)$: outputs a classical public key and a quantum signing key.
- $\mathsf{Sign}(\mathsf{CRS}, |\mathsf{sk}\rangle, m) \to \sigma$: outputs a classical signature, consuming the signing key.
- $\mathsf{Verify}(\mathsf{CRS}, \mathsf{pk}, m, \sigma) \to 0/1$.

Correctness: for every $m \in \mathcal{M}$,

$$\Pr[\mathsf{Verify}(\mathsf{CRS}, \mathsf{pk}, m, \sigma) = 1] \geq 1 - \mathsf{negl}(\lambda),$$

where $\mathsf{CRS} \leftarrow \mathsf{Setup}(1^\lambda)$, $(\mathsf{pk}, |\mathsf{sk}\rangle) \leftarrow \mathsf{KeyGen}(\mathsf{CRS})$ and $\sigma \leftarrow \mathsf{Sign}(\mathsf{CRS}, |\mathsf{sk}\rangle, m)$.
Security (one-shot-ness): for every QPT $\mathcal{A}$, sampling $\mathsf{CRS} \leftarrow \mathsf{Setup}(1^\lambda)$ and $(\mathsf{pk}, m_0, m_1, \sigma_0, \sigma_1) \leftarrow \mathcal{A}(\mathsf{CRS})$,

$$\Pr\big[\, m_0 \neq m_1 \ \wedge\ \mathsf{Verify}(\mathsf{CRS}, \mathsf{pk}, m_0, \sigma_0) = 1 \ \wedge\ \mathsf{Verify}(\mathsf{CRS}, \mathsf{pk}, m_1, \sigma_1) = 1 \,\big] \ \leq\ \mathsf{negl}(\lambda).$$

Note that pk may be adversarially generated.

Remark 1 (two harvester-side repairs to Definition 1). The clause $m_0 \neq m_1$ is not printed in the source paper's Definition 5 (p. 12); without it the requirement is unsatisfiable, since honest KeyGen followed by Sign already yields two verifying pairs, so it is added here as clearly intended, matching the definition of [AGKZ20]. The source paper's Definition 5 leaves the message space an abstract $\mathcal{M}$, and its Algorithm 1 (p. 20) signs only tags in $\{0,1\}^\lambda$; the abstract $\mathcal{M} \supseteq \{0,1\}^\lambda$ above follows the paper rather than fixing $\mathcal{M} = \{0,1\}^*$.

Definition 2 (Witness encryption; indistinguishability vs. extractable security). A *witness encryption* scheme for an NP relation R and message space $\mathcal{M}$ is a pair of QPT algorithms $(\mathsf{Enc}, \mathsf{Dec})$ with $\mathsf{Enc}(1^\lambda, \mathsf{stmt}, m) \to \mathsf{ct}$ and $\mathsf{Dec}(\mathsf{ct}, \mathsf{w}) \to m/\perp$, such that
Correctness: for every $m \in \mathcal{M}$ and every $(\mathsf{stmt}, \mathsf{w}) \in \mathsf{R}$, $\Pr[\mathsf{Dec}(\mathsf{Enc}(1^\lambda, \mathsf{stmt}, m), \mathsf{w}) = m] = 1$.
We distinguish two security notions.

(i) *Indistinguishability security* (the standard notion of [GGSW13], supplied here by the harvester; the source paper defines only the extractable variant, its Definition 6, p. 12): for every QPT $\mathcal{A}$, every $\mathsf{stmt} \notin \mathcal{L}_\mathsf{R}$,

every mixed state aux and all $m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|$,

$$\Big|\Pr[\mathcal{A}(\mathsf{aux}, \mathsf{Enc}(1^\lambda, \mathsf{stmt}, m_0)) = 1] - \Pr[\mathcal{A}(\mathsf{aux}, \mathsf{Enc}(1^\lambda, \mathsf{stmt}, m_1)) = 1]\Big| \leq \mathsf{negl}(\lambda).$$

No guarantee whatsoever is required when $\mathsf{stmt} \in \mathcal{L}_\mathsf{R}$.

(ii) *Extractable security:* for every QPT $\mathcal{A}$, every polynomial p and all $m_0, m_1 \in \mathcal{M}$, there exist a QPT extractor $\mathcal{E}$ and a polynomial q such that for every statement stmt and every mixed state aux,

$$\Big|\Pr[\mathcal{A}(\mathsf{aux}, \mathsf{Enc}(1^\lambda, \mathsf{stmt}, m_0)) = 1] - \Pr[\mathcal{A}(\mathsf{aux}, \mathsf{Enc}(1^\lambda, \mathsf{stmt}, m_1)) = 1]\Big| \geq \tfrac{1}{p(\lambda)}$$

implies

$$\Pr[\mathcal{E}(1^\lambda, \mathsf{stmt}, \mathsf{aux}) = \mathsf{w} \text{ with } (\mathsf{stmt}, \mathsf{w}) \in \mathsf{R}] \geq \tfrac{1}{q(\lambda)}.$$

Item (ii) implies item (i), since a false statement has no witness.

Definition 3 (Threshold PKE with one-shot decryption). Fix $n \in \mathbb{N}$ and $f \in [0, 1/2)$. A *threshold public-key encryption scheme with one-shot decryption* for (n, f) is a tuple of QPT algorithms $(\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{PartialDec}, \mathsf{Dec})$:

- $\mathsf{Setup}(1^\lambda) \to \mathsf{CRS}$;
- $\mathsf{KeyGen}(\mathsf{CRS}) \to (\mathsf{pk}, |\mathsf{sk}\rangle)$, a classical public key and a quantum secret key;
- $\mathsf{Enc}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i \in [n]}, m) \to \mathsf{ct}$, encryption to a committee of n public keys;
- $\mathsf{PartialDec}(\mathsf{CRS}, |\mathsf{sk}\rangle, \mathsf{ct}) \to \sigma/\perp$, producing a *classical* partial decryption;
- $\mathsf{Dec}(\{\sigma_j\}_{j \in J}, \mathsf{ct}) \to m/\perp$ for $J \subseteq [n]$ with $|J| > fn$.

Correctness: for every $m \in \mathcal{M}$ and every $J \subseteq [n]$ with $|J| \geq (1-f)n$,

$$\Pr\left[\mathsf{Dec}(\{\sigma_j\}_{j\in J}, \mathsf{ct}) \neq m\right] \leq \mathsf{negl}(\lambda),$$

where $\mathsf{CRS} \leftarrow \mathsf{Setup}(1^\lambda)$, $(\mathsf{pk}_i, |\mathsf{sk}_i\rangle) \leftarrow \mathsf{KeyGen}(\mathsf{CRS})$ for $i \in [n]$, $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i\in[n]}, m)$ and $\sigma_j \leftarrow \mathsf{PartialDec}(\mathsf{CRS}, |\mathsf{sk}_j\rangle, \mathsf{ct})$ for $j \in J$.

Corruption phases. Both remaining properties refer to a *corruption phase*, in which the adversary $\mathcal{A}$ outputs a set $J' \subseteq [n] \setminus J$ of honest indices to corrupt, where J is the current set of indices already controlled by $\mathcal{A}$. The two experiments use different corruption phases, following the source paper's Experiment 1 and Experiment 2 respectively.

- *Security phase* (Experiment 1, p. 18): the experiment aborts if $|J'| + |J| > fn$; otherwise $\mathcal{A}$ receives the current internal state of party j for each $j \in J'$, and $J \leftarrow J \cup J'$. There is no exemption for parties that have already partially decrypted.
- *One-shot-decryption phase* (Experiment 2, p. 19, as printed): the experiment aborts if $|J'| + 1 > fn$; parties whose secret key has already been used to answer a partial-decryption query are exempt; $\mathcal{A}$ receives the current internal state of party j for each non-exempt $j \in J'$; and $J \leftarrow J \cup J'$.

Security: consider $\mathbf{Exp}^{f,n}_{\mathsf{SEC},\mathcal{A}}(\lambda)$:

1. $\mathsf{CRS} \leftarrow \mathsf{Setup}(1^\lambda)$; $\mathcal{A}(\mathsf{CRS})$ outputs maliciously generated public keys $\{\mathsf{pk}_j\}_{j\in J}$ for some $J \subseteq [n]$; abort if $|J| > fn$.
2. $(\mathsf{pk}_i, |\mathsf{sk}_i\rangle) \leftarrow \mathsf{KeyGen}(\mathsf{CRS})$ for $i \in [n] \setminus J$; run a security corruption phase.
3. $\mathcal{A}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i\in[n]})$ outputs m_0, m_1 with $|m_0| = |m_1|$.
4. $b \leftarrow \{0,1\}$, $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i\in[n]}, m_b)$; run a security corruption phase.
5. $b' \leftarrow \mathcal{A}(\mathsf{CRS}, \mathsf{ct})$; output 1 if $b = b'$ and 0 otherwise.

The scheme is *secure* if $\Pr[\mathbf{Exp}^{f,n}_{\mathsf{SEC},\mathcal{A}}(\lambda) = 1] \le \frac{1}{2} + \mathsf{negl}(\lambda)$ for every QPT $\mathcal{A}$.
One-shot decryption: consider $\mathbf{Exp}^{f,n}_{\mathsf{OSD},\mathcal{A}}(\lambda)$:

1. $\mathsf{CRS} \leftarrow \mathsf{Setup}(1^\lambda)$; $\mathcal{A}(\mathsf{CRS})$ outputs $\{\mathsf{pk}_j\}_{j\in J}$, $J \subseteq [n]$; abort if $|J| > fn$.

2. $(\mathsf{pk}_i, |\mathsf{sk}_i\rangle) \leftarrow \mathsf{KeyGen}(\mathsf{CRS})$ for $i \in [n] \setminus J$. Initialise a partial-decryption oracle O^{Dec} holding the honest keys, together with a variable $\mathsf{ct}^* := \varepsilon$. On a query $(\mathsf{pk}_i, \mathsf{ct})$ with $i \in [n] \setminus J$: abort if pk_i has been queried before; if $\mathsf{ct}^* = \varepsilon$ set $\mathsf{ct}^* := \mathsf{ct}$; abort if $\mathsf{ct} \neq \mathsf{ct}^*$; otherwise return $\mathsf{PartialDec}(\mathsf{CRS}, |\mathsf{sk}_i\rangle, \mathsf{ct})$. (So each honest key partially decrypts at most once, and all honest keys that decrypt do so on the same ciphertext.)

3. $\mathcal{A}^{O^{\mathsf{Dec}}}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i\in[n]})$, interleaved with one-shot-decryption corruption phases, outputs $m_{0,0}, m_{0,1}, m_{1,0}, m_{1,1}$ with $|m_{0,0}| = |m_{0,1}|$ and $|m_{1,0}| = |m_{1,1}|$.

4. $b_0, b_1 \leftarrow \{0,1\}$; $\mathsf{ct}_0 \leftarrow \mathsf{Enc}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i\in[n]}, m_{0,b_0})$ and $\mathsf{ct}_1 \leftarrow \mathsf{Enc}(\mathsf{CRS}, \{\mathsf{pk}_i\}_{i\in[n]}, m_{1,b_1})$.

5. $\mathcal{A}^{O^{\mathsf{Dec}}}(\mathsf{CRS}, \mathsf{ct}_0, \mathsf{ct}_1)$, again interleaved with one-shot-decryption corruption phases, outputs (b'_0, b'_1); output 1 if $(b_0, b_1) = (b'_0, b'_1)$ and 0 otherwise.

The scheme has *one-shot decryption* if $\Pr[\mathbf{Exp}^{f,n}_{\mathsf{OSD},\mathcal{A}}(\lambda) = 1] \le \frac{1}{2} + \mathsf{negl}(\lambda)$ for every QPT $\mathcal{A}$.

Remark 2 (on the one-shot-decryption corruption phase and the bound 1/2). Two harvester-side observations on Definition 3. First, Experiment 2 as printed in the source paper (p. 19) hands an exempt party's state over to $\mathcal{A}$ nowhere at all, whereas the paper's prose (p. 17) says the adversary "can corrupt all parties after decryption"; on the prose reading, corrupting a party *after* it has partially decrypted is free of charge and delivers its post-decryption residual state, so $\mathcal{A}$ may end up holding the state of all n parties — which is the notion the paper's proactive secret

sharing application needs. The printed abort test $|J'| + 1 > fn$ also does not accumulate $|J|$ across successive corruption phases, unlike Experiment 1's. Both are typeset above as printed. Second, the threshold $1/2$ rather than $1/4$ in the one-shot-decryption bound is stated by the paper (Definition 10, p. 18; also Definition 9, p. 16) without justification; the reason it is the right bound is that the honest committee does decrypt one of the two ciphertexts, so one bit is always learnable.

Conjecture 1 (one-shot decryption without extractability). *Let f be an arbitrary constant with $0 < f < 1/2$. Then for every polynomial $n = n(\lambda)$ there exists a threshold public-key encryption scheme with one-shot decryption for the parameters (n, f) (Definition 3) — a tuple of QPT algorithms* $(\mathsf{Setup}, \mathsf{KeyGen}, \mathsf{Enc}, \mathsf{PartialDec}, \mathsf{Dec})$ *satisfying correctness, security and one-shot decryption as in Definition 3 — whose security rests on no extractability or other knowledge-type assumption.*
Concretely, in the formalisation adopted here, suppose that

(i) there exists a one-shot signature scheme OSS *(Definition 1) with message space $\mathcal{M} \supseteq \{0,1\}^\lambda$, secure against QPT adversaries; and*

(ii) for every NP *relation* R *and every message space $\mathcal{M}$ there exists a witness encryption scheme for* R *and $\mathcal{M}$ (Definition 2) satisfying correctness and* indistinguishability security *against QPT adversaries (Definition 2, item (i)).*

Then such a scheme exists for every polynomial $n = n(\lambda)$ and every constant $0 < f < 1/2$. In other words, the implication "one-shot signatures + extractable witness encryption $\Rightarrow$ threshold one-shot decryption for every $f < 1/2$", which is Theorem 2 of the source paper, remains true when extractable security (Definition 2, item (ii)) is weakened to indistinguishability security (Definition 2, item (i)); in particular no extractability or other knowledge-type assumption is used.

Remark 3 (scope of the conjecture). The source paper's open question (p. 7) is existential and disjunctive: whether the extractability requirements "can be weakened or eliminated altogether, for example by relying on more structured or restricted knowledge

assumptions, or by developing alternative realizations of one-shot decryption without extractability". The displayed hypotheses (i) and (ii) are the harvester's formalisation of the strongest reading, elimination; plain witness encryption is never named by the paper as the replacement. The paper equally counts weakening extractability to more structured or restricted knowledge assumptions — its examples are the algebraic-group-model construction of [FHAS24] and the lattice-based knowledge assumptions of [ACL+22] — as progress on the same question: "Even weakening extractability to more structured assumptions would already constitute meaningful progress" (p. 7).

Also the harvester's framing rather than the paper's: a negative resolution would exhibit some $f < 1/2$ for which no such scheme exists under these assumptions — in particular, a black-box separation would suffice. The paper contains no oracle or black-box-separation discussion.

4 Bibliography

[ACL+22] M. R. Albrecht, V. Cini, R. W. F. Lai, G. Malavolta, S. A. K. Thyagarajan. *Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable (extended abstract).* CRYPTO 2022, Part II, LNCS 13508, pages 102–132, 2022.

[AGKZ20] R. Amos, M. Georgiou, A. Kiayias, M. Zhandry. *One-shot signatures and applications to hybrid quantum/classical authentication.* 52nd Annual ACM Symposium on Theory of Computing, pages 255–268, 2020.

[BCP14] E. Boyle, K.-M. Chung, R. Pass. *On extractability obfuscation.* TCC 2014, LNCS 8349, pages 52–73, 2014.

[BGG+20] F. Benhamouda, C. Gentry, S. Gorbunov, S. Halevi, H. Krawczyk, C. Lin, T. Rabin, L. Reyzin. *Can a public blockchain keep a secret?* TCC 2020, Part I, LNCS 12550, pages 260–290, 2020.

[FHAS24] N. Fleischhacker, M. Hall-Andersen, M. Simkin. *Extractable witness encryption for KZG commitments and efficient laconic OT.* ASIACRYPT 2024, Part II, LNCS 15485, pages 423–453, 2024.

[GGHW14] S. Garg, C. Gentry, S. Halevi, D. Wichs. *On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input.* CRYPTO 2014 (recalled from memory; not cited by the source paper — verify before use). [UNVERIFIED]

[GGSW13] S. Garg, C. Gentry, A. Sahai, B. Waters. *Witness encryption and its applications.* 45th Annual ACM Symposium on Theory of Computing, pages 467–476, 2013.

[GHK+21] C. Gentry, S. Halevi, H. Krawczyk, B. Magri, J. B. Nielsen, T. Rabin, S. Yakoubov. *YOSO: You only speak once — secure MPC with stateless ephemeral roles.* CRYPTO 2021, Part II, LNCS 12826, pages 64–93, 2021.

[GZ20] M. Georgiou, M. Zhandry. *Unclonable decryption keys.* Cryptology ePrint Archive, Report 2020/877, 2020.

[HJKY95] A. Herzberg, S. Jarecki, H. Krawczyk, M. Yung. *Proactive secret sharing or: How to cope with perpetual leakage.* CRYPTO'95, LNCS 963, pages 339–352, 1995.

[HV25] A. Huang, V. Vaikuntanathan. *A simple and efficient one-shot signature scheme.* Cryptology ePrint Archive, Report 2025/1906, 2025.

[OY91] R. Ostrovsky, M. Yung. *How to withstand mobile virus attacks (extended abstract).* 10th ACM Symposium Annual on Principles of Distributed Computing, pages 51–59, 1991.

[SZ25] O. Shmueli, M. Zhandry. *On one-shot signatures, quantum vs. classical binding, and obfuscating permutations.* CRYPTO 2025, Part II, LNCS 16001, pages 350–383, 2025.