Unclonable Non-Interactive Zero-Knowledge

Ruta Jawale* Dakshita Khurana*

Abstract

A non-interactive ZK (NIZK) proof enables verification of NP statements without revealing secrets about them. However, an adversary that obtains a NIZK proof may be able to clone this proof and distribute arbitrarily many copies of it to various entities: this is inevitable for any proof that takes the form of a classical string. In this paper, we ask whether it is possible to rely on quantum information in order to build NIZK proof systems that are impossible to clone.

We define and construct *unclonable non-interactive zero-knowledge proofs (of knowledge)* for NP. Besides satisfying the zero-knowledge and proof of knowledge properties, these proofs additionally satisfy unclonability. Very roughly, this ensures that no adversary can split an honestly generated proof of membership of an instance x in an NP language $\mathcal{L}$ and distribute copies to multiple entities that all obtain accepting proofs of membership of x in $\mathcal{L}$. Our result has applications to *unclonable signatures of knowledge*, which we define and construct in this work; these *non-interactively* prevent replay attacks.

*University of Illinois at Urbana-Champaign, USA. Email:{jawale2,dakshita}@illinois.edu

Contents

1 Introduction

Zero-knowledge (ZK) [GMR89] proofs allow a prover to convince a verifier about the truth of an (NP) statement, without revealing secrets about it. These are among the most widely used cryptographic primitives, with a rich history of study.

Enhancing Zero-knowledge. ZK proofs for NP are typically defined via the simulation paradigm. A simulator is a polynomial-time algorithm that mimics the interaction of an adversarial verifier with an honest prover, given only the statement, i.e., $x \in \mathcal{L}$, for an instance x of an NP language $\mathcal{L}$. A protocol satisfies zero-knowledge if it admits a simulator that generates a view for the verifier, which is indistinguishable from the real view generated by an honest prover. This captures the intuition that any information obtained by a verifier upon observing an honestly generated proof, could have been generated by the verifier "on its own" by running the simulator.

Despite being widely useful and popular, there are desirable properties of proof systems that (standard) simulation-based security does not capture. For example, consider (distributions over) instances x of an NP language $\mathcal{L}$ where it is hard to find an NP witness w corresponding to a given instance x. In an "ideal" world, given just the description of one such NP statement $x \in \mathcal{L}$, it is difficult for an adversary to find an NP witness w, and therefore to output *any* proofs of membership of $x \in \mathcal{L}$. And yet, upon obtaining a *single proof* of membership of $x \in \mathcal{L}$, it may suddenly become feasible for an adversary to make many copies of this proof, thereby generating *several* correct proofs of membership of $x \in \mathcal{L}$.

Unfortunately, this attack is inevitable for classical non-interactive proofs: given any proof string, an adversary can always make multiple copies of it. And yet, there is hope to prevent such an attack quantumly, by relying on the *no-cloning* principle.

Indeed, a recent series of exciting works have combined cryptography with the no-cloning principle to develop quantum money [Wie83, AC13, FGH+12, Zha19a, Kan18], quantum tokens for digital signatures [BS16], quantum copy-protection [Aar09, AP21, ALL+21, CLLZ21], unclonable encryption [Got03, BL20, AK21, MST21, AKL+22], unclonable decryption [GZ20], one-out-of-many unclonable security [KN23], and more. In this work, we combine zero-knowledge and unclonability to address a question first posed by Aaronson [Aar09]:

Can we construct unclonable quantum proofs?
How do these proofs relate to quantum money or copy-protection?

1.1 Our Results

We define and construct unclonable non-interactive zero-knowledge proofs of knowledge (NIZKPoK). We obtain a construction in the common reference string (CRS) model, as well as one in the quantum(-accessible) random oracle model (QROM). The CRS model allows a trusted third-party to set up a structured string that is provided to both the prover and verifier. On the other hand, the QROM allows both parties quantum access to a truly random function $\mathcal{O}$.

In what follows, we describe our contributions in more detail.

1.1.1 Definitional Contributions

Before discussing how we formalize the concept of unclonability for NIZKs, it will be helpful to define hard distributions over NP instance-witness pairs.

Hard Distributions over Instance-Witness Pairs. Informally, an efficiently samplable distribution over instance-witness pairs of a language $\mathcal{L}$ is a "hard" distribution if given an instance sampled randomly from this distribution, it is hard to find a witness. Then, unclonable security requires that no adversary given an instance x sampled randomly from the distribution, together with an honestly generated proof, can output *two accepting proofs* of membership of $x \in \mathcal{L}$.

More specifically, a hard distribution $(\mathcal{X}, \mathcal{W})$ over $R_{\mathcal{L}}$ satisfies the following: for any polynomial-sized (quantum) circuit family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$,

$$\Pr_{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)}[C_\lambda(x) \in R_{\mathcal{L}}(x)] \leq \mathsf{negl}(\lambda).$$

For the sake of simplifying our subsequent discussions and definitions, let us fix a NP language $\mathcal{L}$ with corresponding relation $\mathcal{R}$. Let $(\mathcal{X}, \mathcal{W})$ be some hard distribution over $\mathcal{R}$.

A Weaker Definition: Unclonable Security. For NIZKs satisfying standard completeness, soundness and ZK, we define a simple, natural variant of unclonable security as follows. Informally, a proof system satisfies unclonable security if, given an honest proof for an instance and witness pair (x, w) sampled from a hard distribution $(\mathcal{X}, \mathcal{W})$, no adversary can produce two proofs that verify with respect to x except with negligible probability.

Definition 1.1. (Unclonable Security of NIZK). A NIZK proof $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ satisfies unclonable security if for every language $\mathcal{L}$ and every hard distribution $(\mathcal{X}, \mathcal{W})$ over $R_{\mathcal{L}}$, for every poly-sized circuit family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$,

$$\Pr_{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)}\left[\mathsf{Verify}(\mathsf{crs}, x, \pi_1) = 1 \bigwedge \mathsf{Verify}(\mathsf{crs}, x, \pi_2) = 1 \middle| \begin{matrix} \mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda) \\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w) \\ \pi_1,\pi_2\leftarrow C_\lambda(x,\pi) \end{matrix}\right] \leq \mathsf{negl}(\lambda).$$

In the definition above, we aim to capture the intuition that one of the two proofs output by the adversary can be the honest proof they received, but the adversary cannot output any other correct proof for the same statement. Of course, such a proof is easy to generate if the adversary is able to find the witness w for x, which is exactly why we require hardness of the distribution $(\mathcal{X}, \mathcal{W})$ to make the definition non-trivial.

We also remark that unclonable security of proofs *necessitates* that the proof π keep hidden any witnesses w certifying membership of x in $\mathcal{L}$, as otherwise an adversary can always clone the proof π by generating (from scratch) another proof for x given the witness w.

A Stronger Definition: Unclonable Extractability. We can further strengthen the definition above to require that any adversary generating two (or more) accepting proofs of membership of $x \in \mathcal{L}$ given a single proof, must have generated one of the two proofs "from scratch" and must therefore "know" a valid witness w for x. This will remove the need to refer to hard languages.

In more detail, we will say that a proof system satisfies *unclonable extractability* if, from any adversary $\mathcal{A}$ that on input a single proof of membership of $x \in \mathcal{L}$ outputs two proofs for x, then we can extract a valid witness w from $\mathcal{A}$ for at least one of these statements with high probability. Our (still, simplified) definition of unclonable extractability is as follows.

Definition 1.2 (Unclonable Extractability.)**.** A proof $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ satisfies unclonable security there exists a QPT extractor $\mathcal{E}$ which is an oracle-aided circuit such that for every language $\mathcal{L}$ with corresponding relation $\mathcal{R}_{\mathcal{L}}$ and for every non-uniform polynomial-time quantum adversary

$\mathcal{A}$, for every instance-witness pair $(x, w) \in \mathcal{R}_{\mathcal{L}}$ and $\lambda = \lambda(|x|)$, such that there is a polynomial $p(\cdot)$ satisfying:

$$\Pr\left[\mathsf{Verify}(\mathsf{crs}, x, \pi_1) = 1 \bigwedge \mathsf{Verify}(\mathsf{crs}, x, \pi_2) = 1 \middle| \begin{matrix} \mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi \leftarrow \mathsf{Prove}(\mathsf{crs}, x, w) \\ \pi_1, \pi_2 \leftarrow \mathcal{A}_\lambda(\mathsf{crs}, x, \pi, z) \end{matrix}\right] \geq \frac{1}{p(\lambda)},$$

there is also a polynomial $q(\cdot)$ such that

$$\Pr\left[(x, w_{\mathcal{A}}) \in \mathcal{R}_{\mathcal{L}} | w_{\mathcal{A}} \leftarrow \mathcal{E}^{\mathcal{A}}(x)\right] \geq \frac{1}{q(\lambda)}.$$

In fact, in the technical sections, we further generalize this definition to consider a setting where the adversary obtains an even larger number (say $k-1$) input proofs on instances $x_1, \ldots, x_{k-1}$, and outputs k or more proofs. Then we require the extraction of an NP witness corresponding to any proofs that attempt to "clone" honestly generated proofs (i.e. the adversary outputs two or more proofs w.r.t. the same instance $x_i \in \{x_1, \ldots, x_{k-1}\}$). All our theorem statements hold w.r.t. this general definition. Finally, we also consider definitions and constructions in the quantum-accessible random oracle model (QROM); these are natural generalizations of the definitions above, so we do not discuss them here.

We also show that the latter definition of unclonable extractability implies the former, i.e. unclonable security. Informally, this follows because the extractor guaranteed by the definition of extractability is able to obtain a witness w for x from any adversary, which contradicts hardness of the distribution $(\mathcal{X}, \mathcal{W})$. We refer the reader to Appendix A for a formal proof of this claim.

1.1.2 Realizations of Unclonable NIZK, and Relationship with Quantum Money

We obtain realizations of unclonable NIZKs in both the common reference string (CRS) and the quantum random oracle (QRO) models, assuming public-key quantum money mini-scheme and other (post-quantum) standard assumptions. We summarize these results below.

Theorem 1.3 (Informal)**.** *Assuming public-key quantum money mini-scheme, public-key encryption, perfectly binding and computationally hiding commitments, and adaptively sound NIZK proofs for* NP*, there exists an unclonable NIZKPoK scheme in the CRS model.*

Theorem 1.4 (Informal)**.** *Assuming public-key quantum money mini-scheme and honest verifier zero-knowledge proofs of knowledge (HVZKPoKs) for* NP*, there exists an unclonable NIZKPoK scheme in the QROM.*

Theorem 1.5 (Informal)**.** *Assuming public-key quantum money mini-scheme, public-key encryption, post-quantum perfectly binding and computationally hiding commitments, and simulation-sound NIZK proofs for* NP*, there exists an unclonable signature of knowledge in the CRS model.*

Is Quantum Money *necessary* for Unclonable NIZKs? Our work builds unclonable NIZKs for NP by relying on any (public-key) quantum money scheme (mini-scheme), in conjuction with other assumptions such as NIZKs for NP. Since constructions of public-key quantum money mini-scheme are only known based on post-quantum indistinguishability obfuscation [AC13, Zha19b], it is natural to wonder whether the reliance on quantum money is inherent. We show that this is indeed the case, by proving that unclonable NIZKs in fact imply public-key quantum money mini-scheme.

Theorem 1.6 (Informal)**.** *Unclonable NIZKs for NP imply public-key quantum money mini-scheme.*

1.1.3 Applications

Unclonable Signatures of Knowledge. A (classical) signature scheme asserts that a message m has been signed on behalf of a public key pk. However, in order for this signature to be authenticated, the public key pk must be proven trustworthy through a certification chain rooted at a trusted public key PK. However, as [CL06] argue, this reveals too much information; it should be sufficient for the recipient to only know that *there exists* a public key pk with a chain of trust from PK. To solve this problem, [CL06] propose *signatures of knowledge* which allow a signer to sign *on behalf of an instance x of an* NP*-hard language* without revealing its corresponding witness w. Such signatures provide an anonymity guarantee by hiding the pk of the sender.

While this is ideal for many applications, anonymity presents the following downside: a receiver cannot determine whether they were the intended recipient of this signature. In particular, anonymous signatures are more susceptible to *replay attacks*. Replay attacks are a form of passive attack whereby an adversary observes a signature and retains a copy. The adversary then leverages this signature, either at a later point in time or to a different party, to impersonate the original signer. The privacy and financial consequences of replay attacks are steep. They can lead to data breach attacks which cost millions of dollars annually and world-wide [IBM23].

In this work, we construct a signature of knowledge scheme which is the first *non-interactive* signature in the CRS model that is *naturally secure against replay attacks*. Non-interactive, replay attack secure signatures have seen a lot of recent interest including a line of works in the bounded quantum storage model [BS23b] and the quantum random oracle model [BS23a]. Our construction is in the CRS model and relies on the hardness of NP problems, plausible cryptographic assumptions, and the axioms of quantum mechanics. We accomplish this by defining *unclonable signatures of knowledge*: if an adversary, given a signature of a message m with respect to an instance x, can produce two signatures for m which verify with respect to the same instance x, then our extractor is able to extract a witness for x.

Our construction involves showing that an existing compiler can be augmented using unclonable NIZKs to construct unclonable signatures of knowledge. The authors of [CL06] construct signatures of knowledge from CPA secure dense cryptosystems [SP92, SCP00] and simulation-sound NIZKs for NP [Sah99, SCO+01]. Signatures of knowledge are signature schemes in the CRS model for which we associate an instance x in a language $\mathcal{L}$. This signature is simulatable, so there exists a simulator which can create valid signatures without knowledge of a witness for x. Additionally, the signature is extractable which means there is an extractor which is given a trapdoor for the CRS and a signature, and is able to produce a witness for x. We show that, by switching the simulation-sound NIZKs for unclonable simulation-extractable NIZKs (and slightly modifying the compiler), we can construct unclonable signatures of knowledge.

Relationship with Revocation. A recent exciting line of work obtains *certified deletion* for time-lock puzzles [Unr14], non-local games [FM18], information-theoretic proofs of deletion with partial security [CW19], encryption schemes [BI20, BK23], device-independent security of one-time pad encryption with certified deletion [KT20], public-key encryption with certified deletion [HMNY21], commitments and zero-knowledge with certified everlasting hiding [HMNY22], and fully-homomorphic encryption with certified deletion [Por22, BK23, BKP23, BGG+23]. While certified everlasting deletion of secrets has been explored in the context of *interactive* zero-knowledge proofs [HMNY22], there are no existing proposals for *non-interactive* ZK satisfying variants of certified deletion. Our work provides a pathway to building such proofs.

In this work, we construct a quantum *revocable anonymous credentials* protocol by way of the hardness of NP problems, plausible cryptographic assumptions, and the axioms of quantum mechanics. Our work follows a line of work on (classical) revocation for anonymous credentials schemes [BCC+09, CKS10, AN11].

In particular, our construction involves noting that NIZK proof systems that are unclonable can also be viewed as supporting a form of certified deletion/revocation, where in order to delete, an adversary must simply return the entire proof. In other words, the (quantum) certificate of deletion is the proof itself, and this certificate can be verified by running the NIZK verification procedure on the proof. The unclonability guarantee implies that an adversary cannot keep with itself or later have the ability to generate *another proof* for the same instance x. In the other direction, in order to offer certifiable deletion, a NIZK must necessarily be unclonable. To see why, note that if there was an adversary who could clone the NIZK, we could use this adversary to obtain two copies, and provably delete one of them. Even though the challenger for the certifiable deletion game would be convinced that its proof was deleted, we would still be left with another correct proof.

1.2 Related Works

This work was built upon the foundations of and novel concepts introduced by prior literature. We will briefly touch upon some notable such results in this section.

Unclonable Encryptions. *Unclonable encryption* [Got03, BL20, AK21, MST21, AKL+22] imagines an interaction between three parties in which one party receives a quantum ciphertext and splits this ciphertext in some manner between the two remaining parties. At some later point, the key of the encryption scheme is revealed, yet both parties should not be able to simultaneouly recover the underlying message. While our proof systems share the ideology of unclonability, we do not have a similar game-based definition of security. This is mainly due to proof systems offering more structure which can take advantage of to express unclonability in terms of simulators and extractors.

Signature Tokens. Prior work [BS17] defines and constructs *signature tokens* which are signatures which involve a quantum signing token which can only be used once before it becomes inert. The setting they consider is where a client wishes to delegate the signing process to a server, but does not wish the server to be able to sign more than one message. They rely on quantum money [AC13] and the no-cloning principle to ensure the signature can only be computed once. For our unclonable signatures of knowledge result, we focus on the setting where a client wishes to authenticate themselves to a server and wants to prevent an adversary from simultaneously, or later, masquerading as them.

One-shot Signatures. The authors of [AGKZ20] introduce the notion of *one-shot signatures* which extend the concept of signature tokens to a scenario where the client and server only exchange classical information to create a one-use quantum signature token. They show that these signatures can be plausibly constructed in the CRS model from post-quantum indistinguishability obfuscation. Unless additional measures for security, which we discussed in our applications section, are employed, classical communication can be easily copied and replayed at a later point. In contrast, we prevent an adversary from simultaneously, or later, authenticating with the client's identity.

Post-quantum Fiat-Shamir. Our QROM results are heavily inspired by the recent post-quantum Fiat-Shamir result [LZ19] which proves the post-quantum security of NIZKs in the compressed quantum(-accessible) random oracle model (compressed QROM). These classical NIZKs are the result of applying Fiat-Shamir to post-quantum sigma protocols which are HVZKPoKs. We further extend, and crucially rely upon, their novel proof techniques to prove extractability (for PoK) and programmability (for ZK) to achieve extractability and programmability for some protocols which output quantum proofs.

2 Technical Overview

In this section, we give a high-level overview of our construction and the techniques underlying our main results.

2.1 Unclonable Extractable NIZKs in the CRS Model

Our construction assumes the existence of public-key encryption, classical bit commitments where honestly generated commitment strings are perfectly binding, along with

- *Public-key quantum money mini-scheme* (which is known assuming post-quantum $i\mathcal{O}$ [Zha19b]). At a high level, public-key quantum money mini-scheme consists of two algorithms: Gen and Ver. Gen on input a security parameter, outputs a quantum banknote $|\$\rangle$ along with a classical serial number s. Ver is public, takes a quantum money banknote, and outputs either a classical serial number s, or $\perp$ indicating that its input is an invalid banknote. The security guarantee is that no efficient adversary given an honest banknote $|\$\rangle$ can output two notes $|\$_1\rangle$ and $|\$_2\rangle$ that both pass the verification and have serial numbers equal to that of $|\$\rangle$.
- *Post-quantum NIZKs for NP*, which are known assuming the post-quantum hardness of LWE. These satisfy (besides completeness) (1) soundness, i.e., no efficient prover can generate accepting proofs for false NP statements, and (2) zero-knowledge, i.e., the verifier obtains no information from an honestly generated proof beyond what it could have generated on *its own* given the NP statement itself.

Construction. Given these primitives, the algorithms $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ of the unclonable extractable NIZK are as follows.

$\underline{\text{SETUP}}(1^\lambda)$: The setup algorithm samples a public key pk, the common reference string crs of a classical (post-quantum) NIZK for NP, along with a perfectly binding, computationally hiding classical commitment to 0^λ with uniform randomness t, i.e. $c = \mathsf{Com}(0^\lambda; t)$. It outputs $(\mathsf{pk}, \mathsf{crs}, c)$.

$\underline{\text{PROVE}}$: Given the CRS $(\mathsf{pk}, \mathsf{crs}, c)$, instance x and witness w, output $(|\$\rangle, s, ct, \pi)$ where

- The state $|\$\rangle \leftarrow \mathsf{Gen}$ is generated as a quantum banknote with associated serial number s.
- The ciphertext $ct = \mathsf{Enc}_{\mathsf{pk}}(w; u)$ is an encryption of the witness w with randomness u.
- The proof string π is a (post-quantum) NIZK for the following statement using witness (w, u):

$$\mathsf{EITHER}\ (\exists w, u : ct = \mathsf{Enc}_{\mathsf{pk}}(w; u) \wedge R_L(x, w) = 1)\ \mathsf{OR}\ (\exists r : c = \mathsf{Com}(s; r))\,,$$

where we recall that pk and c were a part of the CRS output by the Setup algorithm.

VERIFY: Given CRS $(\mathsf{pk}, \mathsf{crs}, c)$, instance x and proof $(|\$\rangle, s, ct, \pi)$, check that (1) $\mathsf{Ver}(|\$\rangle)$ outputs s and (2) π is an accepting NIZK proof of the statement above.

Analysis. Completeness, soundness/proof of knowledge and ZK for this construction follow relatively easily, so we focus on unclonable extractability in this overview. Recall that unclonable extractability requires that no adversary, given an honestly generated proof for $x \in \mathcal{L}$, can split this into *two accepting proofs* for $x \in \mathcal{L}$ (as long as it is hard to find a witness for x). Towards a contradiction, suppose an adversary splits a proof into 2 accepting proofs $(|\$_1\rangle, s_1, ct_1, \pi_1)$, $(|\$_2\rangle, s_2, ct_2, \pi_2)$. Then,

- If $s_1 = s_2 = s$, the adversary given one bank note with serial number s generated two valid banknotes $|\$\rangle_1$ and $|\$\rangle_2$ that both have the same serial number s. This contradicts the security of quantum money.
- Otherwise, there is a bit b such that $s_b \neq s$. Then, consider an indistinguishable hybrid where the adversary obtains a simulated proof generated *without witness* w as follows: (1) sample quantum banknote $|\$\rangle$ with serial number s, (2) sample public key pk along with secret key sk, (3) generate $c = \mathsf{Com}(s; t)$, $ct = \mathsf{Enc}_{\mathsf{pk}}(0; u)$, (4) generate proof π using witness t (since $c = \mathsf{Com}(s; t)$) instead of using witness w. Send common reference string $(\mathsf{pk}, \mathsf{crs}, c)$ and proof $(|\$\rangle, s, ct, \pi)$ to the adversary. Now, the proof that the adversary generates with $s_b \neq s$ *must* contain $\mathsf{ct}_b = \mathsf{Enc}_{\mathsf{pk}}(w; u)$, since c being generated as a commitment to $s \neq s_b$ along with the perfect binding property implies that $(\nexists\, r : c = \mathsf{Com}(s_b; r))$. That is, given instance x, the adversary can be used to compute a witness w for x by decrypting ciphertext ct_b, thereby contradicting unclonable extractability.

Having constructed unclonable extractable arguments in the CRS model, in the next section, we analyze a construction of unclonable extractable arguments in the QROM.

2.2 Unclonable Extractable NIZK in the QROM

We now turn our attention to the QRO setting in which we demonstrate a protocol which is provably unclonable. Our construction assumes the existence of public-key quantum money mini-scheme and a *post-quantum sigma protocol for NP*. A sigma protocol (P, V) is an interactive three-message honest-verifier protocol: the prover sends a commitment message, the verifier sends a uniformly random challenge, and the prover replies by opening its commitment at the locations specified by the random challenge.

Construction. The algorithms (PROVE, VERIFY) of the unclonable extractable NIZK in the QROM are as follows.

PROVE: Given an instance x and witness w, output $(|\$\rangle, s, \alpha, \beta, \gamma)$ where

- The quantum banknote $|\$\rangle$ is generated alongside associated serial number s.
- P is run to compute the sigma protocol's commitment message as α given (x, w) as input.
- The random oracle is queried on input (α, s, x) in order to obtain a challenge β.
- P is run, given as input (x, w, α, β) and its previous internal state, to compute the sigma protocol's commitment openings as γ.

VERIFY: Given instance x and proof $(|\$\rangle, s, \alpha, \beta, \gamma)$, check that (1) the quantum money verifier accepts $(|\$\rangle, s)$, (2) the random oracle on input (α, s, x) outputs β, and (3) V accepts the transcript (α, β, γ) with respect to x.

Analysis. Since the completeness, proof of knowledge and zero-knowledge properties are easy to show, we focus on unclonable extractability. Suppose an adversary was able to provide two accepting proofs $\pi_1 = (|\$_1\rangle, s_1, \alpha_1, \beta_1, \gamma_1)$ and $\pi_2 = (|\$_2\rangle, s_2, \alpha_2, \beta_2, \gamma_2)$ for an instance x for which it received an honestly generated proof $\pi = (|\$\rangle, s, \alpha, \beta, \gamma)$. Then,

- Suppose $s_1 = s_2 = s$. In this case, the adversary given one bank note with serial number s generated two valid banknotes $|\$\rangle_1$ and $|\$\rangle_2$ that both have the same serial number s. This contradicts the security of quantum money.
- Otherwise, there is a bit $b \in [1, 2]$ such that $s_b \neq s$. By the zero-knowledge property of the underlying HVZK sigma protocol, this event also occurs when the proof π that the adversary is given is replaced with a simulated proof. Specifically, we build a reduction that locally programs the random oracle at location (α, s, x) in order to generate a simulated proof for the adversary. Since the adversary's own proof for $s_b \neq s$ is generated by making a distinct query $(\alpha_b, s_b, x) \neq (\alpha, s, x)$, the programming on (α, s, x) does not affect the knowledge extractor for the adversary's proof, which simply rewinds the (quantum) random oracle to extract a witness for x, following [LZ19]. This allows us to obtain a contradiction, showing that our protocol must be unclonable.

2.3 Unclonable NIZKs imply Quantum Money Mini-Scheme

Finally, we discuss why unclonable NIZKs satisfying even the weaker definition of unclonable security (i.e., w.r.t. hard distributions) imply public-key quantum money mini-scheme. Given an unclonable NIZK, we build a public-key quantum money mini-scheme as follows.

Construction. Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \mathsf{NP}$. Let $\Pi = (\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ be an unclonable NIZK protocol for $\mathcal{L}$.

GEN(1^λ): Sample $(x, w) \leftarrow (\mathcal{X}, \mathcal{W})$, $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, x)$, and an unclonable NIZK proof π as $\mathsf{Prove}(\mathsf{crs}, x, w)$. Output quantum banknote $|\$\rangle = \pi$, and associated serial number $s = (\mathsf{crs}, x)$.

VER$(|\$\rangle, s)$: Given a quantum banknote $|\$\rangle$ and a classical serial number s as input, parse $|\$\rangle = \pi$ and $s = (\mathsf{crs}, x)$, and output the result of $\mathsf{Verify}(\mathsf{crs}, x, \pi)$.

Analysis. The correctness of the quantum money scheme follows from the completeness of the unclonable NIZK Π. We will now argue that this quantum money scheme is unforgeable. Suppose an adversary $\mathcal{A}$ given a quantum banknote and classical serial number $(|\$\rangle, s)$ was able to output two banknotes $(|\$_0\rangle, |\$_1\rangle)$ both of which are accepted with respect to s. We can use $\mathcal{A}$ to define a reduction to the uncloneability of our NIZK Π as follows:

- The NIZK uncloneability challenger outputs a hard instance-witness pair (x, w), a common reference string crs, and an unclonable NIZK π to the reduction.
- The reduction outputs a banknote $(|\$\rangle, s)$ to the adversary, where $|\$\rangle = \pi$ and $s = (\mathsf{crs}, x)$. It receives two quantum banknotes $(|\$_0\rangle, |\$_1\rangle)$ from $\mathcal{A}$, and finally outputs two proofs (π_0, π_1)

where $\pi_0 = |\$_0\rangle$ and $\pi_1 = |\$_1\rangle$.

If $\mathcal{A}$ succeeds in breaking unforgeability, then the quantum money verifier accepts both banknotes ($|\$\rangle_0 = \pi_0, |\$\rangle_1 = \pi_1$), with respect to the same serial number $s = (\mathsf{crs}, x)$. By syntax of the verification algorithm, this essentially means that both *proofs* (π_0, π_1) are accepting proofs for membership of the same instance $x \in \mathcal{L}$, w.r.t. crs, leading to a break in the unclonability of NIZK.

2.4 Unclonable Signatures of Knowledge

Informally, a signature of knowledge has the following property: if an adversary, given a signature of a message m with respect to an instance x, can produce two signatures for m which verify with respect to the same instance x, then the adversary *must know* (and our extractor will be able to extract) a witness for x.

We obtain unclonable signatures of knowledge assuming the existence of an unclonable extractable *simulation-extractable* NIZK for NP. Simulation-extractability states that an adversary which is provided any number of simulated proofs for instance and witness pairs of their choosing, cannot produce an accepting proof π for an instance x which they have not queried before and where extraction fails to find an accepting witness w. Our unclonable extractable NIZK for NP in the CRS model can, with some extra work, be upgraded to simulation-extractable.

We informally describe the construction of signatures of knowledge from such a NIZK below.

Construction. Let $(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge, unclonable-extractable protocol for NP. Let $\mathcal{R}$ be the NP relation corresponding to $\mathcal{L}$.

SETUP: The setup algorithm samples a common reference string crs of an unclonable-extractable simulation-extractable NIZK for NP. It outputs crs.

SIGN: Given the CRS crs, instance x, witness w, and message m, output signature π where

- The proof string π is an unclonable-extractable simulation-extractable NIZK with tag m using witness w of the following statement:

$$(\exists w : (x, w) \in \mathcal{R}) .$$

VERIFY: Given CRS crs, instance x, message m, and signature π, check that π is an accepting NIZK proof with tag m of the statement above.

Analysis. The simulatability (extractability) property follows from the zero-knowledge (resp. simulation-extractability) properties of the NIZK. Suppose an adversary $\mathcal{A}$ given a signature σ was able to forge two signatures $\sigma_1 = \pi_1$ and $\sigma_2 = \pi_2$, and, yet, our extractor was to fail to extract a witness w from $\mathcal{A}$. Then,

- Either both proofs π_1 and π_2 are accepting proofs for membership of the same instance w.r.t. crs. However, this contradicts the unclonability of the NIZK.
- Otherwise there exists a proof π_i (where $i \in \{1, 2\}$) for an instance which $\mathcal{A}$ has not previously seen a proof for. We can switch to a hybrid where our signatures contain simulated proofs for the NIZK. But now, we have that the verifier accepts a proof for an instance which $\mathcal{A}$ has not seen a simulated proof for and, yet, we cannot extract a witness from $\mathcal{A}$. This contradicts the simulation extractability of the NIZK.

Roadmap. In Section 4, we define and construct unclonable NIZKs in the CRS model, and in Section 5, in the QROM. Along the way, we also show that unclonable NIZKs imply quantum money (in the CRS and QRO model respectively). Later, we show how to define and construct unclonable signatures of knowledge from unclonable NIZKs in the CRS model.

3 Preliminaries

3.1 Post-Quantum Commitments and Encryption

Definition 3.1 (Post-Quantum Commitments)**.** Com is a post-quantum commitment scheme if it has the following syntax and properties.
Syntax.

- $c \leftarrow \mathsf{Com}(m; r)$: The polynomial-time algorithm Com on input a message m and randomness $r \in \{0,1\}^{r(\lambda)}$ outputs commitment a c.

Properties.

- **Perfectly Binding**: For every $\lambda \in \mathbb{N}^+$ and every m, m', r, r' such that $m \neq m'$,

$$\mathsf{Com}(m; r) \neq \mathsf{Com}(m'; r').$$

- **Computational Hiding**: There exists a negligible function $\mathsf{negl}(\cdot)$ for every unbounded-size quantum circuit $\mathcal{D}$, every sufficiently large $\lambda \in \mathbb{N}^+$, and every m, m',

$$\left| \Pr_{r \xleftarrow{\$} \{0,1\}^{r(\lambda)},\, c \leftarrow \mathsf{Com}(m;r)} [\mathcal{D}(c) = 1] - \Pr_{r \xleftarrow{\$} \{0,1\}^{r(\lambda)},\, c' \leftarrow \mathsf{Com}(m';r)} [\mathcal{D}(c') = 1] \right| \leq \mathsf{negl}(\lambda).$$

Theorem 3.2 (Post-Quantum Commitment)**.** *[LS19] Assuming the polynomial quantum hardness of LWE, there exists a non-interactive commitment with perfect binding and computational hiding (Definition 3.1).*

Definition 3.3 (Post-Quantum Public-Key Encryption)**.** $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is a post-quantum public-key encryption scheme if it has the following syntax and properties.
Syntax.

- $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$: The polynomial-time algorithm Gen on input security parameter 1^λ outputs a public key pk and a secret key sk.
- $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m; r)$: The polynomial-time algorithm Enc on input a public key pk, message m and randomness $r \in \{0,1\}^{r(\lambda)}$ outputs a ciphertext c.
- $m \leftarrow \mathsf{Dec}(\mathsf{sk}, c)$: The polynomial-time algorithm Dec on input a secret key sk and a ciphertext c outputs a message m.

Properties.

- **Perfect Correctness**: For every $\lambda \in \mathbb{N}^+$ and every m, r,

$$\Pr_{(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)} [\mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{pk}, m; r)) = m] = 1.$$

- **Indistinguishability under Chosen-Plaintext (IND-CPA) Secure**: There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ and every sufficiently large $\lambda \in \mathbb{N}^+$

$$\left| \Pr_{\substack{(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^\lambda)\\(m_0,m_1,\zeta)\leftarrow\mathcal{A}_0(1^\lambda,\mathsf{pk})\\ c\leftarrow\mathsf{Enc}(\mathsf{pk},m_0)}} [\mathcal{A}_1(1^\lambda, c, \zeta) = 1] - \Pr_{\substack{(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^\lambda)\\(m_0,m_1,\zeta)\leftarrow\mathcal{A}_0(1^\lambda,\mathsf{pk})\\ c\leftarrow\mathsf{Enc}(\mathsf{pk},m_1)}} [\mathcal{A}_1(1^\lambda, c, \zeta) = 1] \right| \leq \mathsf{negl}(\lambda).$$

3.2 Sigma protocols

Definition 3.4 (Post-Quantum Sigma Protocol for NP)**.** [LZ19] Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$\Pi = (\mathsf{P} = (\mathsf{P.Com}, \mathsf{P.Prove}), \mathsf{V} = (\mathsf{V.Ch}, \mathsf{V.Ver}))$ is a post-quantum sigma protocol if it has the following syntax and properties.

Syntax. The input 1^λ is left out when it is clear from context.

- $(\alpha, \mathsf{st}) \leftarrow \mathsf{P.Com}(1^\lambda, x, w)$: The probabilistic polynomial-size circuit P.Com on input an instance and witness pair $(x, w) \in \mathcal{L}_\lambda$ outputs a commitment α and an internal prover state st.
- $\beta \leftarrow \mathsf{V.Ch}(1^\lambda, x, \alpha)$: The probabilistic polynomial-size circuit V.Ch on input an instance x outputs a uniformly random challenge β.
- $\gamma \leftarrow \mathsf{P.Prove}(1^\lambda, x, w, \mathsf{st}, \beta)$: The probabilistic polynomial-size circuit P.Prove on input an instance and witness pair $(x, w) \in \mathcal{L}_\lambda$, an internal prover state st, and a challenge β outputs the partial opening (to α as indicated by β) γ.
- $\mathsf{V.Ver}(1^\lambda, x, \alpha, \beta, \gamma) \in \{0, 1\}$: The probabilistic polynomial-size circuit V.Ver on input an instance x, a commitment α, a challenge β, and a partial opening γ outputs 1 iff γ is a valid opening to α at locations indicated by β.

Properties.

- **Perfect Completeness.** For every $\lambda \in \mathbb{N}$ and every $(x, w) \in \mathcal{R}_\lambda$,

$$\Pr_{\substack{(\alpha,\mathsf{st})\leftarrow\mathsf{P.Com}(x,w)\\ \beta\leftarrow\mathsf{V.Ch}(x,\alpha)\\ \gamma\leftarrow\mathsf{P.Prove}(x,w,\mathsf{st},\beta)}} [\mathsf{V.Verify}(x, \alpha, \beta, \gamma) = 1] = 1$$

- **Computational Honest-Verifier Zero-Knowledge with Quantum Simulator.** There exists a quantum polynomial-size circuit Sim and a negligible function $\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{D}$, every sufficiently large $\lambda \in \mathbb{N}$, and every $(x, w) \in \mathcal{R}_\lambda$,

$$\left| \Pr_{\substack{(\alpha,\mathsf{st})\leftarrow\mathsf{P.Com}(x,w)\\ \beta\leftarrow\mathsf{V.Ch}(x,\alpha)\\ \gamma\leftarrow\mathsf{P.Prove}(x,w,\mathsf{st},\beta)}} [\mathcal{D}(x, \alpha, \beta, \gamma) = 1] - \Pr_{(\alpha,\beta,\gamma)\leftarrow\mathsf{Sim}(1^\lambda,x)} [\mathcal{D}(x, \alpha, \beta, \gamma) = 1] \right| \leq \mathsf{negl}(\lambda).$$

- **Proof of Knowledge with Quantum Extractor.** There exists an oracle-aided quantum polynomial-size circuit Ext, a constant c, a polynomial $p(\cdot)$, and negligible functions $\mathsf{negl}_0(\cdot)$, $\mathsf{negl}_1(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ where
  - $\mathcal{A}_0(x)$ is a unitary U_x followed by a measurement and
  - $\mathcal{A}_1(x, |\mathsf{st}\rangle, \beta)$ is a unitary $V_{x,\beta}$ onto the state $|\mathsf{st}\rangle$ followed by a measurement,

  and every x with associated $\lambda \in \mathbb{N}$ satisfying

$$\Pr_{\substack{(\alpha,|\mathsf{st}\rangle)\leftarrow\mathcal{A}_0(x)\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_1(x,|\mathsf{st}\rangle,\beta)}}[\mathsf{V.Ver}(x,\alpha,\beta,\gamma)=1] \geq \mathsf{negl}_0(\lambda)$$

  we have

$$\Pr\left[(x, \mathsf{Ext}^{\mathcal{A}(x)}(x)) \in \mathcal{R}_\lambda\right] \geq \frac{1}{p(\lambda)} \cdot \left(\Pr_{\substack{(\alpha,|\mathsf{st}\rangle)\leftarrow\mathcal{A}_0(x)\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_1(x,|\mathsf{st}\rangle,\beta)}}[\mathsf{V.Ver}(x,\alpha,\beta,\gamma)=1] - \mathsf{negl}_0(\lambda) \right)^c - \mathsf{negl}_1(\lambda).$$

  When we say Ext has oracle access to $\mathcal{A}(x)$, we mean that Ext has oracle access to both unitaries $U_x, V_{x,\beta}$ and their inverses $U_x^\dagger, V_{x,\beta}^\dagger$.

- **Unpredictable Commitment.** There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every sufficiently large $\lambda \in \mathbb{N}$ and every $(x, w) \in \mathcal{R}_\lambda$,

$$\Pr_{\substack{(\alpha,\mathsf{st})\leftarrow\mathsf{P.Com}(x,w)\\ (\alpha',\mathsf{st}')\leftarrow\mathsf{P.Com}(x,w)}}[\alpha=\alpha'] \leq \mathsf{negl}(\lambda).$$

We note that the unpredictable commitment property in the definition above may appear to be an unusual requirement, but this property is w.l.o.g. for post quantum sigma protocols as shown in [LZ19]. In particular, any sigma protocol which does not have unpredictable commitments, can be modified into one that does: the prover can append a random string r to the end of their commitment message α, and the verifier can ignore this appended string r when they perform their checks.

3.3 NIZKs in the CRS model

We consider the common reference string model.

Definition 3.5 (Post-Quantum (Quantum) NIZK for NP in the CRS Model)**.** Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ is a non-interactive post-quantum (quantum) zero-knowledge argument for NP in the CRS model if it has the following syntax and properties.

Syntax. The input 1^λ is left out when it is clear from context.

- $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)$: The probabilistic polynomial-size circuit Setup on input 1^λ outputs a common reference string crs.

- $\pi \leftarrow \mathsf{P}(1^\lambda, \mathsf{crs}, x, w)$: The probabilistic (quantum) polynomial-size circuit P on input a common reference string crs and instance and witness pair $(x, w) \in \mathcal{R}_\lambda$, outputs a proof π.
- $\mathsf{V}(1^\lambda, \mathsf{crs}, x, \pi) \in \{0, 1\}$: The probabilistic (quantum) polynomial-size circuit V on input a common reference string crs, an instance x, and a proof π outputs 1 iff π is a valid proof for x.

Properties.

- **Perfect Completeness.** For every $\lambda \in \mathbb{N}$ and every $(x, w) \in \mathcal{R}_\lambda$,

$$\Pr_{\substack{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ \pi \leftarrow \mathsf{P}(\mathsf{crs}, x, w)}} [\mathsf{V}(\mathsf{crs}, x, \pi) = 1] = 1.$$

- **Adaptive Computational Soundness.** There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A}$ and every sufficiently large $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (x, \pi) \leftarrow \mathcal{A}(\mathsf{crs})}} [\mathsf{V}(x, \mathsf{crs}, \pi) = 1 \ \wedge \ x \notin \mathcal{L}_\lambda] \leq \mathsf{negl}(\lambda).$$

- **Adaptive Computational Zero-Knowledge.** There exists a probabilistic (quantum) polynomial-size circuit $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ and a negligible function $\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A}$, every polynomial-size quantum circuit $\mathcal{D}$, and every sufficiently large $\lambda \in \mathbb{N}$,

$$\left| \Pr_{\substack{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda) \\ (x, w, \zeta) \leftarrow \mathcal{A}(\mathsf{crs}) \\ \pi \leftarrow \mathsf{P}(\mathsf{crs}, x, w)}} [\mathcal{D}(\mathsf{crs}, x, \pi, \zeta) = 1] - \Pr_{\substack{(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x, w, \zeta) \leftarrow \mathcal{A}(\mathsf{crs}) \\ \pi \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x)}} [\mathcal{D}(\mathsf{crs}, x, \pi, \zeta) = 1] \right| \leq \mathsf{negl}(\lambda).$$

Theorem 3.6 (Post-Quantum NIZK argument for NP in the CRS Model)**.** *[PS19] Assuming the polynomial quantum hardness of LWE, there exists a non-interactive adaptively computationally sound, adaptively computationally zero-knowledge argument for* NP *in the common reference string model (Definition 3.5).*

Definition 3.7 (Post-Quantum (Quantum) Simulation-Sound NIZK for NP in CRS Model)**.** Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ is a post-quantum (quantum) non-interactive simulation-sound, adaptive multi-theorem computational zero-knowledge protocol for NP in the CRS model if it has the following syntax and properties.

- Π is a post-quantum (quantum) non-interactive zero-knowledge argument for NP in the CRS model (Definition 3.5).
- **Adaptive Multi-Theorem Computational Zero-Knowledge.** [FLS90] There exists a probabilistic (quantum) polynomial-size circuit $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ [1] and a negligible function

[1] Sim_1 ignores the second term (a witness w) in the queries it receives from $\mathcal{A}$.

$\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A}$, every polynomial-size quantum circuit $\mathcal{D}$, and every sufficiently large $\lambda \in \mathbb{N}$,

$$\left| \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{P}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs}) = 1] \right| \leq \mathsf{negl}(\lambda).$$

- **Simulation Soundness.** [Sah99, SCO+01] Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator given by the adaptive multi-theorem computational zero-knowledge property. There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every oracle-aided polynomial-size quantum circuit $\mathcal{A}$ and every sufficiently large $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x,\pi) \leftarrow \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})}} [\mathsf{V}(\mathsf{crs}, x, \pi) = 1 \wedge x \notin Q \wedge x \notin \mathcal{L}] \leq \mathsf{negl}(\lambda),$$

  where Q is the list of queries from $\mathcal{A}$ to Sim_1.

REMARK 3.1. In Definition 3.7, adaptive multi-theorem computational zero-knowledge implies adaptive computational zero-knowledge.

REMARK 3.2. As defined in Definition 3.7, a simulation-sound zero-knowledge protocol has adaptive computational soundness (Definition 3.5).

Theorem 3.8 (Simulation Sound Compiler)**.** *[SCO+01] Given one-way functions and a single-theorem NIZK proof system for* NP*, then there exists a non-interactive simulation sound, adaptively multi-theorem computationally zero-knowledge proof for* NP *in the common reference string model (Definition 3.7).*

Corollary 3.9 (Post-Quantum Simulation Sound NIZK for NP)**.** Assuming the polynomial quantum hardness of LWE, there exists a post-quantum non-interactive simulation sound, adaptively multi-theorem computationally zero-knowledge proof for NP in the common reference string model (Definition 3.7).

Proof. This follows from Theorem 3.6 and Theorem 3.8. □

3.4 NIZKs in the QRO model

We now consider the quantum random oracle model. For sake of completeness, we briefly outline a definition for a quantum random oracle.

Definition 3.10. A quantum random oracle $\mathcal{O}$ is a random function which support quantum queries and allows for the following accesses:

- **Query Access.** On input a message, $\mathcal{O}$ outputs a uniformly random value. This is the usual access provided. When quantum access may be invoked, we denote the oracle as $|\mathcal{O}\rangle$.
- **Programmability Access.** Given programmability access, $\mathcal{O}$ can be set to output a specified value on a specified input. An arbitrary number of distinct points can be programmed.
- **Extractability Access.** Given extractability access, specific queries to $|\mathcal{O}\rangle$ can be read.

Definition 3.11 ((Quantum) Post-Quantum NIZKPoK for NP in QROM)**.** [LZ19] Let $\mathcal{O}$ be a random oracle. Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$\Pi = (\mathsf{P}, \mathsf{V})$ is a (quantum) non-interactive zero-knowledge proof of knowledge protocol with respect to a random oracle if it has the following syntax and properties.

Syntax. The input 1^λ is left out when it is clear from context.

- $\pi \leftarrow \mathsf{P}^{\mathcal{O}}(1^\lambda, x, w)$: The random oracle-aided (quantum) probabilistic polynomial-size circuit P on input an instance and witness pair $(x, w) \in \mathcal{R}_\lambda$, outputs a proof π.
- $\mathsf{V}^{\mathcal{O}}(1^\lambda, x, \pi) \in \{0, 1\}$: The random oracle-aided (quantum) probabilistic polynomial-size circuit V on input an instance x and a proof π, outputs 1 iff π is a valid proof for x.

Properties.

- **Perfect Completeness.** For every $\lambda \in \mathbb{N}$ and every $(x, w) \in \mathcal{R}_\lambda$,

$$\Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathsf{P}^{\mathcal{O}}(x,w)}}[\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] = 1.$$

- **Zero-Knowledge with Quantum Simulator.** There exists a quantum polynomial-size circuit Sim which ignores its second input and a negligible function $\mathsf{negl}(\cdot)$ such that for every oracle-aided polynomial-size quantum circuit $\mathcal{D}$ which is limited to making queries $(x, \omega) \in \mathcal{R}_\lambda$ on input 1^λ, and every sufficiently large $\lambda \in \mathbb{N}$,

$$\left| \Pr\left[\mathcal{D}^{\mathsf{Sim}, |\mathcal{O}_{\mathsf{Sim}}\rangle}(1^\lambda) = 1\right] - \Pr_{\mathcal{O}}[\mathcal{D}^{\mathsf{P}^{\mathcal{O}}, |\mathcal{O}\rangle}(1^\lambda) = 1] \right| \leq \mathsf{negl}(\lambda)$$

  where Sim simulates the random oracle $|\mathcal{O}_{\mathsf{Sim}}\rangle$.

- **Proof of Knowledge with Quantum Extractor.** There exists an oracle-aided quantum polynomial-size circuit extractor Ext that simulates a random oracle $|\mathcal{O}_{\mathsf{Ext}}\rangle$, a constant c, a polynomial $p(\cdot)$, and negligible functions $\mathsf{negl}_0(\cdot)$, $\mathsf{negl}_1(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A}$ and every x with associated $\lambda \in \mathbb{N}$ satisfying

$$\Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x)}}[\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] \geq \mathsf{negl}_0(\lambda)$$

  we have

$$\Pr\left[(x, \mathsf{Ext}^{\mathcal{A}^{|\mathcal{O}_{\mathsf{Ext}}\rangle}(x)}(x)) \in \mathcal{R}_\lambda\right] \geq \frac{1}{p(\lambda)} \cdot \left(\Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x)}}[\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] - \mathsf{negl}_0(\lambda) \right)^c - \mathsf{negl}_1(\lambda).$$

Theorem 3.12 (NIZKPoK in QROM [Unr17, LZ19])**.** *Let Π be a post-quantum sigma protocol (Definition 3.4). The Fiat-Shamir heuristic applied to Π yields a classical post-quantum NIZKPoK in the QROM (Definition 3.11).*

3.5 Quantum Money

Definition 3.13 (Public Key Quantum Money Mini-Scheme)**.** [AC13, Zha19b] $(\mathsf{Gen}, \mathsf{Ver})$ is a public key quantum money scheme if it has the following syntax and properties.
Syntax.

- $(|\$\rangle, s) \leftarrow \mathsf{Gen}(1^\lambda)$: The quantum polynomial-time algorithm Gen on input security parameter 1^λ outputs a quantum banknote $|\$\rangle$ along with a classical serial number s.
- $\mathsf{Ver}(|\$\rangle, s) \in \{0,1\}$: The quantum polynomial-time algorithm Ver on input a quantum banknote $|\$\rangle$ and a classical serial number s outputs 1 or 0.

Properties.

- **Perfect Correctness**: For every $\lambda \in \mathbb{N}^+$,

$$\Pr_{(|\$\rangle, s) \leftarrow \mathsf{Gen}(1^\lambda)}[\mathsf{Ver}(|\$\rangle, s) = 1] = 1.$$

- **Unforgeable**: There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every sufficiently large $\lambda \in \mathbb{N}^+$ and every polynomial-size quantum circuit $\mathcal{A}$,

$$\Pr_{\substack{(|\$\rangle, s) \leftarrow \mathsf{Gen}(1^\lambda) \\ (|\$_0\rangle, s_0, |\$_1\rangle, s_1) \leftarrow \mathcal{A}(|\$\rangle, s)}}[s_0 = s_1 = s \ \wedge \ \mathsf{Ver}(|\$_0\rangle, s_0) = 1 \ \wedge \ \mathsf{Ver}(|\$_1\rangle, s_1) = 1] \leq \mathsf{negl}(\lambda).$$

- **Unpredictable Serial Numbers**: There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every sufficiently large $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{(|\$\rangle, s) \leftarrow \mathsf{Gen}(1^\lambda) \\ (|\$'\rangle, s') \leftarrow \mathsf{Gen}(1^\lambda)}}[s = s'] \leq \mathsf{negl}(\lambda).$$

REMARK 3.3 (Unpredictable Serial Numbers). The unpredictable serial numbers property follows, w.l.o.g., from unforgeability. We will briefly outline the reduction. Say that Gen produced two quantum banknotes $|\$\rangle$ and $|\$'\rangle$ which had the same serial number s with noticeable probability. Then an adversary $\mathcal{A}$ that receives $(|\$\rangle, s)$ from Gen could run Gen again to produce $(|\$'\rangle, s)$ with noticeable probability. This means that $\mathcal{A}$ would have produced two quantum banknotes $|\$\rangle$ and $|\$'\rangle$ which Verify would accept with respect to the same serial number that $\mathcal{A}$ received, s.

Theorem 3.14 (Quantum Money from Subspace Hiding Obfuscation [AC13, Zha19b])**.** *If injective one-way functions and post-quantum iO exist, then public-key quantum money exists (Definition 3.13).*

Definition 3.15 (Public Key Quantum Money Mini-Scheme in QROM)**.** $(\mathsf{Gen}, \mathsf{Ver})$ is a public key quantum money scheme with respect to a quantum random oracle $\mathcal{O}$ if it has the following syntax and properties.
Syntax.

- $(|\$\rangle, s) \leftarrow \mathsf{Gen}^{\mathcal{O}}(1^\lambda)$: The random oracle-aided quantum polynomial-time algorithm Gen on input a security parameter 1^λ outputs a quantum banknote $|\$\rangle$ along with a classical serial number s.

- $\mathsf{Ver}^{\mathcal{O}}(|\$\rangle, s) \in \{0,1\}$: The random oracle-aided quantum polynomial-time algorithm Ver on input a quantum banknote $|\$\rangle$ and a classical serial number s outputs 1 or 0.

Properties.

- **Perfect Correctness**: For every $\lambda \in \mathbb{N}^+$,

$$\Pr_{(|\$\rangle,s)\leftarrow \mathsf{Gen}^{\mathcal{O}}(1^\lambda)}[\mathsf{Ver}^{\mathcal{O}}(|\$\rangle, s) = 1] = 1.$$

- **Unforgeable**: There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every sufficiently large $\lambda \in \mathbb{N}^+$ and every random oracle-aided polynomial-size quantum circuit $\mathcal{A}$,

$$\Pr_{\substack{(|\$\rangle,s)\leftarrow \mathsf{Gen}^{\mathcal{O}}(1^\lambda)\\ (|\$_0\rangle,s_0,|\$_1\rangle,s_1)\leftarrow \mathcal{A}^{\mathcal{O}}(|\$\rangle,s)}}[s_0 = s_1 = s \ \wedge \ \mathsf{Ver}^{\mathcal{O}}(|\$_0\rangle, s_0) = 1 \ \wedge \ \mathsf{Ver}^{\mathcal{O}}(|\$_1\rangle, s_1) = 1] \leq \mathsf{negl}(\lambda).$$

- **Unpredictable Serial Numbers**: There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every sufficiently large $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{(|\$\rangle,s)\leftarrow \mathsf{Gen}^{\mathcal{O}}(1^\lambda)\\ (|\$'\rangle,s')\leftarrow \mathsf{Gen}^{\mathcal{O}}(1^\lambda)}}[s = s'] \leq \mathsf{negl}(\lambda).$$

The unpredicable serial number property is w.l.o.g., just as above.

3.6 Quantum Signature of Knowledge

Definition 3.16 (Quantum SimExt-secure Signature [CL06])**.** Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$. Let a message space $\mathcal{M}$ be given such that it can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$(\mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify})$ is a SimExt-secure quantum signature of knowledge of a witness with respect to $\mathcal{L}$ and $\mathcal{M}$ if it has the following syntax and properties.

Syntax. The input 1^λ is left out when it is clear from context.

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$: The probabilistic polynomial-time algorithm Setup on input 1^λ outputs a common reference string crs and a trapdoor td.

- $\sigma \leftarrow \mathsf{Sign}(1^\lambda, \mathsf{crs}, x, w, m)$: The polynomial-time quantum algorithm Sign on input a common reference string crs, an instance and witness pair $(x, w) \in \mathcal{R}_\lambda$, and a message $m \in \mathcal{M}_\lambda$, outputs a signature σ.

- $\mathsf{Verify}(1^\lambda, \mathsf{crs}, x, m, \sigma) \in \{0,1\}$: The polynomial-time quantum algorithm Verify on input a common reference string crs, an instance x, a message $m \in \mathcal{M}_\lambda$, and a signature σ, outputs 1 iff σ is a valid signature of m with respect to crs, $\mathcal{R}_\lambda$, and x.

Properties.

- **Correctness**: For every sufficiently large $\lambda \in \mathbb{N}$, every $(x, w) \in \mathcal{R}_\lambda$, and every $m \in \mathcal{M}_\lambda$,

$$\Pr_{\substack{\mathsf{crs}\leftarrow \mathsf{Setup}(1^\lambda)\\ \sigma\leftarrow \mathsf{Sign}(\mathsf{crs},x,w,m)}}[\mathsf{Verify}(\mathsf{crs}, x, m, \sigma) = 1] = 1.$$

- **Simulation**: There exists a quantum polynomial-size circuit simulator $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$, where Sim_1 ignores its second query input (a witness w), and a negligible function $\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{A}$ and every sufficiently large $\lambda \in \mathbb{N}$,

$$\left| \Pr_{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)}[\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)}[\mathcal{A}^{\mathsf{Sign}(\mathsf{crs},\cdot,\cdot,\cdot)}(\mathsf{crs}) = 1] \right| \leq \mathsf{negl}(\lambda).$$

- **Extraction**: Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator given by the simulation property. There exists a quantum polynomial-size circuit Ext and a negligible function $\mathsf{negl}(\cdot)$ such that for every oracle-aided polynomial-size quantum circuit $\mathcal{A}$ and every sufficiently large $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,m,\sigma)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,m,\sigma)}} [\mathsf{Verify}(\mathsf{crs},x,m,\sigma) = 1 \wedge (x,m) \notin Q \wedge (x,w) \notin \mathcal{R}_\lambda] \leq \mathsf{negl}(\lambda)$$

  where Q is the list of queries from $\mathcal{A}$ to Sim_1.

4 Unclonable Non-Interactive Zero-Knowledge in the CRS Model

4.1 Simulation-Extractable Definition

Definition 4.1 (Post-Quantum (Quantum) Simulation-Extractable NIZK for NP in CRS Model)**.** Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ is a post-quantum (quantum) non-interactive simulation-extractable zero-knowledge argument for NP in the CRS model if it has the following syntax and properties.

- Π is a post-quantum (quantum) non-interactive simulation sound, adaptive multi-theorem computational zero-knowledge argument for NP in the CRS model (Definition 3.7).

- **Simulation Extractability.** Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator given by the adaptive multi-theorem computational zero-knowledge property. There exists a (quantum) polynomial-time circuit Ext and a negligible function $\mathsf{negl}(\cdot)$ such that for every oracle-aided polynomial-size quantum circuit $\mathcal{A}$ and every $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}} [\mathsf{V}(\mathsf{crs},x,\pi) = 1 \wedge x \notin Q \wedge (x,w) \notin \mathcal{R}] \leq \mathsf{negl}(\lambda),$$

  where Q is the list of queries from $\mathcal{A}$ to Sim_1.

REMARK 4.1. As defined in Definition 4.1, a simulation-extractable zero-knowledge protocol has simulation soundness [Sah99, SCO+01], is a proof of knowledge, and has adaptive computational soundness (Definition 3.5).

Simulation-Extractable Non-Interactive ZK for $\mathcal{L} \in \mathsf{NP}$

Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be a non-interactive simulation sound, adaptively multi-theorem computationally zero-knowledge protocol for NP, and $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a post-quantum perfectly correct, IND-CPA secure encryption scheme. Let $\mathcal{R}$ be the relation with respect to $\mathcal{L} \in \mathsf{NP}$.

$\underline{\text{SETUP}}(1^\lambda)$: Compute $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$, and $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi) \leftarrow \Pi.\mathsf{Setup}(1^\lambda)$. Output $(\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi), \mathsf{td} = (\mathsf{sk}, \mathsf{td}_\Pi))$.

$\underline{\text{PROVE}}(\mathsf{crs}, x, w)$:

- Compute $\mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, w; r)$ for r sampled uniformly at random.
- Let $x_\Pi = (\mathsf{pk}, x, \mathsf{ct})$ be an instance of the following language $\mathcal{L}_\Pi$:

$$\{(\mathsf{pk}, x, \mathsf{ct}) \ : \ \exists (w, r) \ : \ \mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, w; r) \ \wedge \ (x, w) \in \mathcal{R}\}.$$

- Compute proof $\pi_\Pi \leftarrow \Pi.\mathsf{P}(\mathsf{crs}_\Pi, x_\Pi, (w, r))$ for language $\mathcal{L}_\Pi$.
- Output $\pi = (\mathsf{ct}, \pi_\Pi)$.

$\underline{\text{VERIFY}}(\mathsf{crs}, x, \pi)$:

- Output $\Pi.\mathsf{V}(\mathsf{crs}_\Pi, x_\Pi, \pi_\Pi)$.

Figure 1: Unclonable Non-Interactive Quantum Protocol for $\mathcal{L} \in \mathsf{NP}$

Theorem 4.2 (Post-Quantum Simulation-Extractable NIZK for NP in the CRS Model)**.** *Let* NP *relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given.*

Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be a non-interactive post-quantum simulation sound, adaptively multi-theorem computationally zero-knowledge protocol for NP *(Definition 3.7). Let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a post-quantum perfectly correct, IND-CPA secure encryption scheme (Definition 3.3).*

$(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ as defined in Figure 1 will be a non-interactive post-quantum simulation-extractable, adaptively multi-theorem computationally zero-knowledge argument for $\mathcal{L}$ in the common reference string model (Definition 4.1).

Proof. **Perfect Completeness.** Completeness follows from the perfect completeness of Π.

Adaptively Multi-theorem Computationally Zero-Knowledge. Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. We define Sim_0 with oracle access to $\Pi.\mathsf{Sim}_0$ as follows:

Input: 1^λ.

(1) Compute $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$.

(2) Send 1^λ to $\Pi.\mathsf{Sim}_0$. Receive $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi)$ from $\Pi.\mathsf{Sim}_0$.

(3) Output $(\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi), \mathsf{td} = (\mathsf{sk}, \mathsf{td}_\Pi))$.

We define Sim_1 with oracle access to $\Pi.\mathsf{Sim}_1$ as follows:

Input: $\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi)$, $\mathsf{td} = (\mathsf{sk}, \mathsf{td}_\Pi)$, x.

(1) Compute $\mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, 0; r)$ for r sampled uniformly at random.
(2) Define $x_\Pi = (\mathsf{pk}, x, \mathsf{ct})$.
(3) Send $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi, x_\Pi)$ to $\Pi.\mathsf{Sim}_1$. Receive π_Π.
(4) Output $\pi = (\mathsf{ct}, \pi_\Pi)$.

Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\left| \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{P}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs}) = 1] \right| \geq \frac{1}{p(\lambda)}. \tag{1}$$

We will first switch the honest proofs for simulated proofs, using the adaptive multi-theorem zero-knowledge of Π. Later, we will see how we can switch the encryption of a valid witness to an encryption of 0, by using the security of the encryption scheme.

Towards this end, we define an intermediary circuit $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1)$ which encrypts a valid witness, but provides simulated proofs through $\Pi.\mathsf{Sim}_1$. We define $\mathcal{B}_0$ to be equivalent to Sim_0. We define $\mathcal{B}_1$ with oracle access to $\Pi.\mathsf{Sim}_1$ as follows:

Input: $\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi)$, $\mathsf{td} = (\mathsf{sk}, \mathsf{td}_\Pi)$, x, w.
(1) Compute $\mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, w; r)$ for r sampled uniformly at random.
(2) Define $x_\Pi = (\mathsf{pk}, x, \mathsf{ct})$.
(3) Send $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi, x_\Pi)$ to $\Pi.\mathsf{Sim}_1$. Receive π_Π.
(4) Output $\pi = (\mathsf{ct}, \pi_\Pi)$.

Claim 4.3. There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every oracle-aided polynomial-size quantum circuit $\mathcal{A}$,

$$\left| \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{P}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathcal{B}_0(1^\lambda)} [\mathcal{A}^{\mathcal{B}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) = 1] \right| \leq \mathsf{negl}(\lambda).$$

We will later see a proof of Claim 4.3. For now, assuming that this claim holds, by Equation (2), this claim, and a union bound, there exists a polynomial $p'(\cdot)$ such that

$$\left| \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathcal{B}_0(1^\lambda)} [\mathcal{A}^{\mathcal{B}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs}) = 1] \right| \geq \frac{1}{p'(\lambda)}.$$

We define a series of intermediary hybrids starting from encrypting all real witnesses to encrypting all zeros. The first intermediary hybrid switches the encryption sent in the last query from an encryption of a witness to an encryption of 0. We continue switching the encryption in the second to last query and so on, until we've switched the first proof that the adversary makes.

Let $q(\cdot)$ be a polynomial denoting the maximum number of queries that $\mathcal{A}$ makes. By a union bound and Equation (2), there must exist a hybrid indexed by i (where we switch the ciphertext in the ith proof from encrypting a witness to encrypting 0) where $\mathcal{A}$ first distinguishes between the two ciphertexts with advantage $1/(p'(\lambda)q(\lambda))$. That is,

$$\left| \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1^{(i+1)}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1^{(i)}(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs}) = 1] \right| \geq \frac{1}{p'(\lambda)q(\lambda)}. \tag{2}$$

where $\mathsf{Sim}_1^{(j)}$ is a stateful algorithm which sends real proofs for the first $j-1$ queries and sends simulated proofs for the remaining queries.

We can use $\mathcal{A}$ to define a reduction that breaks the IND-CPA security of the encryption scheme as follows:

Reduction: to IND-CPA of encryption scheme given oracle access to $\mathcal{A}$, Sim_0, and Sim_1.
Hardwired with: i.
(1) Compute $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$.
(2) Compute $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi) \leftarrow \Pi.\mathsf{Sim}_0(1^\lambda)$.
(3) Define $\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi)$ and $\mathsf{td} = (\mathsf{sk}, \mathsf{td}_\Pi)$.
(4) Send crs to $\mathcal{A}$.
(5) On the first $i-1$ queries (x, w) from $\mathcal{A}$: send $\pi \leftarrow \mathcal{B}_0(\mathsf{crs}, x, w)$ to $\mathcal{A}$.
(6) On the ith query (x, w) from $\mathcal{A}$: send $(w, 0)$ to the challenger, receive ct from the challenger, define $x_\Pi = (\mathsf{pk}, x, \mathsf{ct})$, send $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi, x_\Pi)$ to $\Pi.\mathsf{Sim}_1$, receive π_Π from $\Pi.\mathsf{Sim}_1$, and send $\pi = (\mathsf{ct}, \pi_\Pi)$ to $\mathcal{A}$.
(7) On any queries (x, w) after the ith: send $\pi \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x)$ to $\mathcal{A}$.
(8) Output the result of $\mathcal{A}$.

The view of $\mathcal{A}$ matches that of $\mathsf{Sim}_1^{(i+1)}$ or $\mathsf{Sim}_1^{(i)}$. As such, this reduction should have the same advantage at breaking the IND-CPA security of the encryption scheme. We reach a contradiction. Now, all that remains to prove that our earlier claim holds.

Proof of Claim 4.3. Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\left| \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{P}(\mathsf{crs}, \cdot, \cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs}, \mathsf{td}) \leftarrow \mathcal{B}_0(1^\lambda)} [\mathcal{A}^{\mathcal{B}_1(\mathsf{crs}, \mathsf{td}, \cdot, \cdot)}(\mathsf{crs}) = 1] \right| \geq \frac{1}{p(\lambda)}.$$

We define a reduction to the multi-theorem zero-knowledge property of Π as follows:

Reduction: to multi-theorem zero-knowledge of Π given oracle access to $\mathcal{A}$.
(1) Compute $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$.
(2) Receive (real or simulated) crs_Π from the challenger.
(3) Send $\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi)$ to $\mathcal{A}$.
(4) On query (x, w) from $\mathcal{A}$: compute $\mathsf{ct} = \mathsf{Enc}(\mathsf{pk}, w; r)$ for r samples uniformly at random, send $x_\Pi = (\mathsf{pk}, x, \mathsf{ct})$ to the challenger, receive (real or simulated) π_Π from the challenger, send $\pi = (\mathsf{ct}, \pi_\Pi)$ to $\mathcal{A}$.
(5) Output the result of $\mathcal{A}$.

The view of $\mathcal{A}$ matches that of Setup and P or $\mathcal{B}_0$ and $\mathcal{B}_1$. As such, this reduction should have the same advantage at breaking the multi-theorem zero-knowledge property of Π. We reach a contradiction, hence our claim must be true. □

This concludes our proof. Hence our protocol must be multi-theorem zero-knowledge.

Simulation Extractable. Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator, with oracle access to $\Pi.\mathsf{Sim}$, as defined in the proof that Figure 1 is adaptive multi-theorem computational zero-knowledge. We define Ext as follows:

Input: $\mathsf{crs} = (\mathsf{pk}, \mathsf{crs}_\Pi)$, $\mathsf{td} = (\mathsf{sk}, \mathsf{td}_\Pi)$, x, $\pi = (\mathsf{ct}, \pi_\Pi)$.
(1) Output $\mathsf{Dec}(\mathsf{sk}, \mathsf{ct})$ as w.

Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}}[\mathsf{V}(\mathsf{crs},x,\pi)=1\wedge x\notin Q\wedge(x,w)\notin\mathcal{R}]\geq\frac{1}{p(\lambda)},$$

where Q is the list of queries from $\mathcal{A}$ to Sim_1. Since V accepts the output of $\mathcal{A}$, then $\Pi.\mathsf{V}$ must accept $(\mathsf{crs}_\Pi, x_\Pi, \pi_\Pi)$. Since $x\notin Q$, then x_Π which contains x must not have been sent as a query to $\Pi.\mathsf{Sim}_1$. By the definition of Ext and the perfect correctness of the encryption scheme, $x_\Pi\notin\mathcal{L}_\Pi$. Hence, we have that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}}[\Pi.\mathsf{V}(\mathsf{crs}_\Pi,x_\Pi,\pi_\Pi)=1\wedge x_\Pi\notin Q_\Pi\wedge x_\Pi\notin\mathcal{L}_\Pi]\geq\frac{1}{p(\lambda)},$$

where Q_Π is the list of queries, originating from $\mathcal{A}$, that Sim_1 makes to $\Pi.\mathsf{Sim}_1$. We define a reduction to the simulation soundness property of Π as follows:

Reduction: to simulation soundness of Π given oracle access to $\mathcal{A}$.
(1) Compute $(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^\lambda)$.
(2) Receive crs_Π from the challenger.
(3) Send $\mathsf{crs}=(\mathsf{pk},\mathsf{crs}_\Pi)$ to $\mathcal{A}$.
(4) On query x from $\mathcal{A}$: compute $\mathsf{ct}=\mathsf{Enc}(\mathsf{pk},0;r)$ for r samples uniformly at random, send $x_\Pi=(\mathsf{pk},x,\mathsf{ct})$ to the challenger, receives π_Π from the challenger, send $\pi=(\mathsf{ct},\pi_\Pi)$ to $\mathcal{A}$.
(5) Receive $(x,\pi=(\mathsf{ct},\pi_\Pi))$ from $\mathcal{A}$. Define $x_\Pi=(\mathsf{pk},x,\mathsf{ct})$.
(6) Output (x_Π,π_Π).

The view of $\mathcal{A}$ matches that of Sim_0 and Sim_1. As such, this reduction should have the same advantage at breaking the simulation soundness property of Π. We reach a contradiction, hence our protocol must be simulation extractable. □

Corollary 4.4 (Post-Quantum Simulation-Extractable NIZK for NP in the CRS Model)**.** Assuming the polynomial quantum hardness of LWE, there exists a simulation-extractable, adaptively multi-theorem computationally zero-knowledge argument for NP in the common reference string model (Definition 4.1).

Proof. This follows from Corollary 3.9 and Theorem 4.2. □

4.2 Unclonability Definitions

We consider two definitions of unclonability for NIZKs. The first one, motivated by simplicity, informally guarantees that no adversary given honestly proofs for "hard" instances is able to output more than one accepting proof for the same instance.

Definition 4.5 ((Quantum) Hard Distribution)**.** Let an NP relation $\mathcal{R}$ be given. $(\mathcal{X},\mathcal{W})$ is a (quantum) hard distribution over $\mathcal{R}$ if the following properties hold.

- **Syntax.** $(\mathcal{X},\mathcal{W})$ is indexable by a security parameter $\lambda\in\mathbb{N}$. For every choice of $\lambda\in\mathbb{N}$, the support of $(\mathcal{X}_\lambda,\mathcal{W}_\lambda)$ is over instance and witness pairs (x,w) such that $x\in\mathcal{L}$, $|x|=\lambda$, and $(x,w)\in\mathcal{R}$.

- **Hardness.** For every polynomial-sized (quantum) circuit family $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda\in\mathbb{N}}$,

$$\Pr_{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)}[(x,\mathcal{A}_\lambda(x))\in\mathcal{R}] \leq \mathsf{negl}(\lambda).$$

Definition 4.6. (Unclonable Security for Hard Instances). A proof $(\mathsf{Setup},\mathsf{P},\mathsf{V})$ satisfies unclonable security for a language $\mathcal{L}$ with corresponding relation $\mathcal{R}_\mathcal{L}$ if for every polynomial-sized quantum circuit family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$, and for every hard distribution $\{\mathcal{X}_\lambda,\mathcal{W}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\mathcal{R}_\mathcal{L}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\lambda\in\mathbb{N}$,

$$\Pr_{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)}\left[\mathsf{V}(\mathsf{crs},x,\pi_1)=1\bigwedge\mathsf{V}(\mathsf{crs},x,\pi_2)=1\middle|\begin{smallmatrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{P}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow C_\lambda(x,\pi)\end{smallmatrix}\right]\leq\mathsf{negl}(\lambda).$$

We will now strengthen this definition to consider a variant where from any adversary $\mathcal{A}$ that on input a single proof of membership of $x\in\mathcal{L}$ outputs two proofs for x, we can extract a valid witness w for x with high probability. In fact, we can further generalize this definition to a setting where the adversary obtains an even larger number (say $k-1$) input proofs on instances $x_1,\ldots,x_{k-1}$, and outputs k or more proofs. Then we require the extraction of an NP witness corresponding to any proofs that are *duplicated* (i.e. two or more proofs w.r.t. the same instance $x_i\in\{x_1,\ldots,x_{k-1}\}$). We write this definition below.

Definition 4.7 ($(k-1)$-to-k-Unclonable Extractable NIZK)**.** Let security parameter $\lambda\in\mathbb{N}$ and NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given. Let $\Pi=(\mathsf{Setup},\mathsf{P},\mathsf{V})$ be given such that Setup, P and V are $\mathsf{poly}(\lambda)$-size quantum algorithms. We have that for any $(x,w)\in\mathcal{R}$, $(\mathsf{crs},\mathsf{td})$ is the output of Setup on input 1^λ, P receives an instance and witness pair (x,w) along with crs as input and outputs π, and V receives an instance x, crs, and proof π as input and outputs a value in $\{0,1\}$.

Π is a non-interactive $(k-1)$-to-k-unclonable zero-knowledge quantum protocol for language $\mathcal{L}$ if the following holds:

- Π is a quantum non-interactive zero-knowledge protocol for language $\mathcal{L}$ (Definition 3.5).
- $(k-1)$**-to-**k**-Unclonable with Extraction**: There exists an oracle-aided polynomial-size quantum circuit $\mathcal{E}$ such that for every polynomial-size quantum circuit $\mathcal{A}$, for every tuple of $k-1$ instance-witness pairs $(x_1,\omega_1),\ldots,(x_{k-1},\omega_{k-1})\in\mathcal{R}$, for every x where we define
  - $\mathcal{I}\subseteq[k-1]$ such that $|\mathcal{I}|\geq 1$, $x_i=x$ for all $i\in\mathcal{I}$, and $x_\iota\neq x$ for all $\iota\notin\mathcal{I}$, and
  - $\mathcal{J}\subseteq[k]$ such that $|\mathcal{J}|\geq\max\{2,|\mathcal{I}|\}$, $x_j=x$ for all $j\in\mathcal{J}$, and $x_\iota\neq x$ for $\iota\notin\mathcal{J}$,

  such that there is a polynomial $p(\cdot)$ where

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{P}(\mathsf{crs},x_\iota,w_\iota)\\ \{\widetilde{x}_\iota,\widetilde{\pi}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi}_\iota)=1\right]\geq\frac{1}{p(\lambda)},$$

  then there is also a polynomial $q(\cdot)$ such that

$$\Pr_{w\leftarrow\mathcal{E}^{\mathcal{A}}(x_1,\ldots,x_{k-1},x,\mathcal{I},\mathcal{J})}[(x,w)\in\mathcal{R}]\geq\frac{1}{q(\lambda)}.$$

We describe a useful lemma to compare our two definitions.

Lemma 4.8. Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be a 1-to-2-unclonable with extraction, non-interactive zero-knowledge quantum protocol (Definition 4.7). Then, Π satisfies Definition 4.6.

For a proof of Lemma 4.8, we refer to Appendix A.

4.3 Unclonable NIZK Implies Public-Key Quantum Money Mini-Scheme

Public-Key Quantum Money Mini-Scheme

Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \mathsf{NP}$. Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be an unclonable non-interactive zero-knowledge protocol for $\mathcal{L}$.

$\underline{\text{GEN}}(1^\lambda)$: Sample a hard instance-witness pair $(x, w) \leftarrow (\mathcal{X}, \mathcal{Y})$, a common reference string $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, x)$, and a proof $\pi \leftarrow \mathsf{P}(\mathsf{crs}, x, w)$. Output $(|\$\rangle = \pi, s = (\mathsf{crs}, x))$.

$\underline{\text{VERIFY}}(|\$\rangle, s)$: Parse $|\$\rangle = \pi$ and $s = (\mathsf{crs}, x)$. Output $\mathsf{V}(\mathsf{crs}, x, \pi)$.

Figure 2: Public-Key Quantum Money Mini-Scheme from an Unclonable Non-Interactive Quantum Protocol

Theorem 4.9. *Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \mathsf{NP}$. Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ satisfy Definition 4.6. Then $(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ implies a public-key quantum money scheme mini-scheme (Definition 3.13) as described in Figure 2.*

Proof. **Perfect Correctness.** This follows directly from the perfect completeness of Π.
Unforgeability. Let $p(\cdot)$ be a polynomial and $\mathcal{A}$ be a quantum polynomial-time adversary such that for an infinite number of $\lambda \in \mathbb{N}^+$,

$$\Pr_{\substack{(|\$\rangle, s) \leftarrow \mathsf{Gen}(1^\lambda) \\ (|\$_0\rangle, s_0, |\$_1\rangle, s_1) \leftarrow \mathcal{A}(|\$\rangle, s)}} [s_0 = s_1 = s \ \wedge \ \mathsf{Ver}(|\$_0\rangle, s_0) = 1 \ \wedge \ \mathsf{Ver}(|\$_1\rangle, s_1) = 1] \geq \frac{1}{p(\lambda)}.$$

We construct a reduction that breaks the uncloneability definition. The challenger samples a hard instance-witness pair $(x, w) \leftarrow (\mathcal{X}, \mathcal{Y})$, a common reference string $\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda, x)$, and a proof $\pi \leftarrow \mathsf{P}(\mathsf{crs}, x, w)$. The challenger then forwards (crs, x, π) to the reduction. The reduction then sets $|\$\rangle = \pi$ and $s = (\mathsf{crs}, x)$. The reduction sends $(|\$\rangle, s)$ to the adversary $\mathcal{A}$ who returns back $(|\$_0\rangle, s_0, |\$_1\rangle, s_1)$. The reduction then parses and sets $\pi_i = |\$_i\rangle$ for $i \in \{0, 1\}$. The reduction then sends π_0 and π_1 back to the challenger.

When the serial numbers are the same, $s = s_0 = s_1$, we have that the common reference string and instance will be the same for all the proofs π, π_0, π_1. The quantum money state can be parsed as the proof as shown in the construction. When the verification algorithm of the quantum money algorithm accepts both quantum money states $|\$_0\rangle$ and $|\$_1\rangle$ with respect to s, we know that that V would accept both proofs π_0 and π_1 with respect to (crs, x). As such, we will have that the advantage that $\mathcal{A}$ has at breaking the unforgeability of our quantum money scheme directly translates to the advantage of the reduction at breaking the uncloneability of Π. □

4.4 Construction and Analysis of Unclonable NIZK from Public-Key Quantum Money

Unclonable Non-Interactive ZK for $\mathcal{L} \in \mathsf{NP}$

Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be a non-interactive simulation-extractable, adaptively multi-theorem computationally zero-knowledge protocol for NP, Com be a post-quantum perfectly binding, computationally hiding commitment scheme, and $(\mathsf{NoteGen}, \mathsf{Ver})$ be a public-key quantum money scheme. Let $\mathcal{R}$ be the relation with respect to $\mathcal{L} \in \mathsf{NP}$.

$\underline{\text{SETUP}}(1^\lambda)$: Sample the common reference string $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi) \leftarrow \Pi.\mathsf{Setup}(1^\lambda)$, and s^*, r^* uniformly at random. Define $c = \mathsf{Com}(s^*; r^*)$ and output $(\mathsf{crs} = (\mathsf{crs}_\Pi, c), \mathsf{td} = \mathsf{td}_\Pi)$.

$\underline{\text{PROVE}}(\mathsf{crs}, x, w)$:

- Compute a quantum note and associated serial number $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}$.
- Let $x_\Pi = (c, x, s)$ be an instance of the following language $\mathcal{L}_\Pi$:

$$\{(c, x, s) \ : \ \exists z \ : \ (x, z) \in \mathcal{R} \ \vee \ c = \mathsf{Com}(s; z)\}.$$

- Compute proof $\pi_\Pi \leftarrow \Pi.\mathsf{P}(\mathsf{crs}_\Pi, x_\Pi, w)$ for language $\mathcal{L}_\Pi$.
- Output $\pi = (|\$\rangle, s, \pi_\Pi)$.

$\underline{\text{VERIFY}}(\mathsf{crs}, x, \pi)$:

- Check that $\mathsf{Ver}(|\$\rangle, s)$ outputs 1 and that $\Pi.\mathsf{V}(\mathsf{crs}_\Pi, x_\Pi, \pi_\Pi)$ outputs 1.
- If both checks pass, output 1. Otherwise, output 0.

Figure 3: Unclonable Non-Interactive Quantum Protocol for $\mathcal{L} \in \mathsf{NP}$

Theorem 4.10. *Let $k(\cdot)$ be a polynomial. Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given.*

Let $(\mathsf{NoteGen}, \mathsf{Ver})$ be a public-key quantum money mini-scheme (Definition 3.13) and Com be a post-quantum commitment scheme (Definition 3.1). Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be a non-interactive post-quantum simulation-extractable, adaptive multi-theorem computational zero-knowledge protocol for NP (Definition 4.1).

$(\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ as defined in Figure 3 will be a non-interactive quantum simulation-extractable, adaptive multi-theorem computationally zero-knowledge, and $(k-1)$-to-k-unclonable with extraction protocol for $\mathcal{L}$ in the common reference string model (Definition 4.7).

Proof. **Perfect Completeness.** Completeness follows from perfect correctness of the public key quantum money scheme, and perfect completeness of Π.

Adaptive Multi-Theorem Computational Zero-Knowledge. Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. We define Sim_0 with oracle access to $\Pi.\mathsf{Sim}_0$ as follows:

Input: 1^λ.
(1) Send 1^λ to $\Pi.\mathsf{Sim}_0$. Receive $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi)$ from $\Pi.\mathsf{Sim}_0$.
(2) Sample s^*, r^* uniformly at random. Define $c = \mathsf{Com}(s^*; r^*)$.
(3) Output $\mathsf{crs} = (\mathsf{crs}_\Pi, c)$ and $\mathsf{td} = \mathsf{td}_\Pi$.

We define Sim_1 with oracle access to $\Pi.\mathsf{Sim}_1$ as follows:

Input: $\mathsf{crs} = (\mathsf{crs}_\Pi, c)$, $\mathsf{td} = \mathsf{td}_\Pi$, x.
(1) Sample $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$.
(2) Define $x_\Pi = (c, x, s)$. Send $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi, x_\Pi)$ to $\Pi.\mathsf{Sim}_1$. Receive π_Π from $\Pi.\mathsf{Sim}_1$.
(3) Output $\pi = (|\$\rangle, s, \pi_\Pi)$.

Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\left| \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{P}(\mathsf{crs},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs}) = 1] \right| \geq \frac{1}{p(\lambda)}.$$

We define a reduction to the multi-theorem zero-knowledge property of Π as follows:

Reduction: to zero-knowledge of Π given oracle access to $\mathcal{A}$.
(1) Receive (real or simulated) crs_Π from the challenger.
(2) Sample s^*, r^* uniformly at random. Define $c = \mathsf{Com}(s^*; r^*)$ and $\mathsf{crs} = (\mathsf{crs}_\Pi, c)$.
(3) Send crs to $\mathcal{A}$.
(4) On query (x, w) from $\mathcal{A}$: sample $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$, define $x_\Pi = (c, x, s)$ and $w_\Pi = w$, send (x_Π, w_Π) to the challenger, receive (real or simulated) π_Π from the challenger, define $\pi = (|\$\rangle, s, \pi_\Pi)$, send π to $\mathcal{A}$.
(5) Output the result of $\mathcal{A}$.

The view of $\mathcal{A}$ matches that of our protocol in Figure 3 or Sim_0 and Sim_1. As such, this reduction should have the same advantage at breaking the adaptive multi-theorem computational zero-knowledge property of Π. We reach a contradiction, hence our protocol must be multi-theorem zero-knowledge.

Simulation-Extractability. Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. Let $\Pi.\mathsf{Ext}$ be the simulation-extraction extractor of Π with respect to $\Pi.\mathsf{Sim}$. Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator, with oracle access to $\Pi.\mathsf{Sim}$, as defined in the proof that Figure 3 is adaptive multi-theorem computational zero-knowledge. We define Ext with oracle access to $\Pi.\mathsf{Ext}$ as follows:

Input: $\mathsf{crs} = (\mathsf{crs}_\Pi, c)$, $\mathsf{td} = \mathsf{td}_\Pi$, x, $\pi = (|\$\rangle, s, \pi_\Pi)$.
(1) Define $x_\Pi = (c, x, s)$. Send $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi, x_\Pi, \pi_\Pi)$ to $\Pi.\mathsf{Ext}$. Receive w_Π from $\Pi.\mathsf{Ext}$.
(2) Output w_Π as w.

Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x,\pi) \leftarrow \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs}) \\ w \leftarrow \mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}} [\mathsf{V}(\mathsf{crs}, x, \pi) = 1 \wedge x \notin Q \wedge (x, w) \notin \mathcal{R}] \geq \frac{1}{p(\lambda)},$$

where Q is the list of queries from $\mathcal{A}$ to Sim_1. Since Sim_1 forwards oracle queries to $\Pi.\mathsf{Sim}_1$ which contain any query it receives from $\mathcal{A}$, we know that $x_\Pi \notin Q_\Pi$ where Q_Π is the list of queries from Sim_1 to $\Pi.\mathsf{Sim}_1$. Furthermore, since V accepts the output π from $\mathcal{A}$, then $\Pi.\mathsf{V}$ must accept the proof

π_Π. As such, we have that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}}[\Pi.\mathsf{V}(\mathsf{crs}_\Pi,x_\Pi,\pi_\Pi)=1\wedge x_\Pi\notin Q_\Pi\wedge(x,w)\notin\mathcal{R}]\geq\frac{1}{p(\lambda)}. \tag{3}$$

However, we make the following claim which is in direct contradiction with Equation (3).

Claim 4.11. Let Ext be as defined earlier, in the current proof of simulation-extractability. There exists a negligible function $\mathsf{negl}(\cdot)$ such that for every polynomial-size quantum circuit $\mathcal{B}$,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{B}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}}[\Pi.\mathsf{V}(\mathsf{crs}_\Pi,x_\Pi,\pi_\Pi)=1\wedge x_\Pi\notin Q_\Pi\wedge(x,w)\notin\mathcal{R}]\leq\mathsf{negl}(\lambda)$$

where Q is the list of queries from $\mathcal{B}$ to Sim_1.

Proof of Claim 4.11. We proceed by contradiction. Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{B}$ be given such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{B}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}}[\Pi.\mathsf{V}(\mathsf{crs}_\Pi,x_\Pi,\pi_\Pi)=1\wedge x_\Pi\notin Q_\Pi\wedge(x,w)\notin\mathcal{R}]\geq\frac{1}{p(\lambda)} \tag{4}$$

where Q is the list of queries from $\mathcal{B}$ to Sim_1. Given Equation (4), we may be in one of the two following cases: either the extractor $\Pi.\mathsf{Ext}$ extracts w_Π from $\mathcal{B}$ such that $(x_\Pi,w_\Pi)\notin\mathcal{R}_\Pi$, or the extractor $\Pi.\mathsf{Ext}$ extracts w_Π from $\mathcal{B}$ such that $(x_\Pi,w_\Pi)\in\mathcal{R}_\Pi$. We consider these two scenarios separately and show that each reaches a contradiction.

Scenario One

Say that the extractor $\Pi.\mathsf{Ext}$ extracts w_Π from $\mathcal{B}$ such that $(x_\Pi,w_\Pi)\notin\mathcal{R}_\Pi$. By applying a union bound to Equation (5), we have that this event could happen with at least $1/2p(\lambda)$ probability. Symbolically,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{B}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}}[\Pi.\mathsf{V}(\mathsf{crs}_\Pi,x_\Pi,\pi_\Pi)=1\wedge x_\Pi\notin Q_\Pi\wedge(x_\Pi,w_\Pi)\notin\mathcal{R}_\Pi]\geq\frac{1}{2p(\lambda)}. \tag{5}$$

By using the advantage of $\mathcal{B}$ in this game, we can show a reduction that breaks the simulation-extractability of Π. We will now outline this reduction.

Reduction: to simulation-extractability of Π given oracle access to $\mathcal{B}$.

(1) Receive crs_Π from the challenger.

(2) Sample s^*,r^* uniformly at random. Define $c=\mathsf{Com}(s^*;r^*)$.

(3) Define $\mathsf{crs}=(\mathsf{crs}_\Pi,c)$ and $\mathsf{td}=\mathsf{td}_\Pi$. Send crs to $\mathcal{B}$.

(4) On query x from $\mathcal{B}$: sample $(|\$\rangle,s)\leftarrow\mathsf{NoteGen}(1^\lambda)$, define $x_\Pi=(c,x,s)$, send x_Π to the challenger, receive π_Π from the challenger, define $\pi=(|\$\rangle,s,\pi_\Pi)$, and send π to $\mathcal{B}$.

(5) Receive $(x,\pi=(|\$\rangle,s,\pi_\Pi))$ from $\mathcal{B}$. Define $x_\Pi=(c,x,s)$.

(6) Output (x_Π,π_Π).

Given the event in Equation (5) holds, then the reduction will return an accepting proof π_Π for an instance x_Π which it has not previously queried on and, yet, the extraction $\Pi.\mathsf{Ext}$ will fail. With advantage $1/2p(\lambda)$), the reduction will succeed at breaking simulation-extractability of Π, thus reaching a contradiction.

Scenario Two

Alternatively, say that the extractor $\Pi.\mathsf{Ext}$ extracts w_Π from $\mathcal{B}$ such that $(x_\Pi, w_\Pi) \in \mathcal{R}_\Pi$. By applying a union bound to Equation (3), we have that this event could happen with at least $1/2p(\lambda)$ probability. In summary, we have that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\(x,\pi)\leftarrow\mathcal{B}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot)}(\mathsf{crs})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},x,\pi)}} [\Pi.\mathsf{V}(\mathsf{crs}_\Pi, x_\Pi, \pi_\Pi) = 1 \wedge x_\Pi \notin Q_\Pi \wedge (x,w) \notin \mathcal{R} \wedge (x_\Pi, w_\Pi) \in \mathcal{R}_\Pi] \geq \frac{1}{2p(\lambda)}. \quad (6)$$

Since Ext outputs $w = w_\Pi$, by the definition of $\mathcal{L}_\Pi$ and the perfect binding of Com, we must have that $\mathcal{B}$ has found an opening to the commitment c in the crs, that is that $s = s^*$ and $w_\Pi = r^*$. We can use $\mathcal{B}$ to break the hiding of the commitment. We will now outline this reduction.

Reduction: to hiding of Com given oracle access to $\mathcal{B}$.

(1) Compute $(\mathsf{crs}_\Pi, \mathsf{td}_\Pi) \leftarrow \Pi.\mathsf{Sim}_0(1^\lambda)$ from the challenger.

(2) Sample s_0, s_1 uniformly at random. Send (s_0, s_1) to the challenger. Receive c.

(3) Define $\mathsf{crs} = (\mathsf{crs}_\Pi, c)$ and $\mathsf{td} = \mathsf{td}_\Pi$. Send crs to $\mathcal{B}$.

(4) On query x from $\mathcal{B}$: compute $\pi \leftarrow \mathsf{Sim}_1(\mathsf{crs}, \mathsf{td}, x)$, and send π to $\mathcal{B}$.

(5) Receive $(x, \pi = (|\$\rangle, s, \pi_\Pi))$ from $\mathcal{B}$.

(6) Compute $w \leftarrow \mathsf{Ext}(\mathsf{crs}, \mathsf{td}, x, \pi)$.

(7) If $s = s_b$ for $b \in \{0,1\}$, then output b. Else, output s_b for b chosen uniformly at random.

Given the event in Equation (5) holds, then the reduction will, with advantage $1/q(\lambda)$ for some polynomial $q(\cdot)$, succeed at breaking the hiding of Com, thus reaching a contradiction. □

Since Equation (3) directly contradicts Claim 4.11 which we have proven, then we have reached a contradiction. Therefore, the protocol must be simulation extractable.

Unclonable Extractability. Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. Let $\Pi.\mathsf{Ext}$ be the simulation-extraction extractor of Π with respect to $\Pi.\mathsf{Sim}$. Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator, with oracle access to $\Pi.\mathsf{Sim}$, as defined in the proof that Figure 3 is adaptive multi-theorem computational zero-knowledge. Let Ext be the extractor, based on Sim, as defined in the proof that Figure 3 is simulation-extractable. We define $\mathcal{E}$ with oracle access to Sim, Ext, and some $\mathcal{A}$ as follows:

Hardwired: $x_1, \ldots, x_{k-1}$, x, $\mathcal{I}$, $\mathcal{J}$

(1) Send 1^λ to Sim_0. Receive $(\mathsf{crs}, \mathsf{td})$ from Sim_0.

(2) For $\iota \in [k-1]$: send $(\mathsf{crs}, \mathsf{td}, x_\iota)$ to Sim_1, and receive π_ι from Sim_1.

(3) Send $(\mathsf{crs}, \{x_\iota, \pi_\iota\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$.

(4) Define j' uniformly at random from $\mathcal{J}$.

(5) Output $\mathsf{Ext}(\mathsf{crs}, \mathsf{td}, \widetilde{x_{j'}}, \widetilde{\pi_{j'}})$ and w.

Let $\mathcal{A}$, $(x_1, w_1), \ldots, (x_{k-1}, w_{k-1}) \in \mathcal{R}$, x, $\mathcal{I} \subseteq [k-1]$, $\mathcal{J} \subseteq [k]$, polynomial $p(\cdot)$, and negligible function $\mathsf{negl}_1(\cdot)$ be given such that the verifier V accepts all proofs which $\mathcal{A}$ outputs indexed by $\mathcal{I}$,

and the extractor $\mathcal{E}$ is unable to extract a valid witness. Restated more formally, that is that both

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{P}(\mathsf{crs},x_\iota,w_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi_\iota})=1\right]\geq\frac{1}{p(\lambda)},\text{ and} \tag{7}$$

$$\Pr_{w\leftarrow\mathcal{E}^{\mathcal{A}}(x_1,\ldots,x_{k-1},x,\mathcal{I},\mathcal{J})}[(x,w)\in\mathcal{R}]\leq\mathsf{negl}_1(\lambda). \tag{8}$$

Given Equation (7), we may be in one of the two following cases: either $\mathcal{A}$ generate two accepting proofs which have the same serial number as an honestly generated proof, or $\mathcal{A}$ does not. We consider these two scenarios separately and show that each reaches a contradiction.

Scenario One

Say that $\mathcal{A}$ generates two accepting proofs which have the same serial number as an honestly generated proof. By applying a union bound to Equation (7), we have that this event could happen with at least $1/2p(\lambda)$ probability. Symbolically,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{P}(\mathsf{crs},x_\iota,w_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi_\iota})=1\bigwedge\exists i\in\mathcal{I}\ \exists j,\ell\in\mathcal{J}\text{ s.t. }s_i=\widetilde{s_j}=\widetilde{s_\ell}\right]\geq\frac{1}{2p(\lambda)}. \tag{9}$$

Through a hybrid argument, we can fix indices $i\in\mathcal{I}$ and $j,\ell\in\mathcal{J}$ which gives us the same event with an advantage of $1/(2k^3p(\lambda))$. By using the advantage of $\mathcal{A}$ in this game, we can show a reduction that breaks the unforgeability of the quantum money scheme. We will now outline this reduction.

Reduction: to unforgeability of quantum money scheme given oracle access to $\mathcal{A}$ and $\mathcal{O}$.

Hardwired with: $(x_1,w_1),\ldots,(x_{k-1},w_{k-1})$, x, $\mathcal{I}$, $\mathcal{J}$, i,j,ℓ.

(1) Compute $(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)$ where $\mathsf{crs}=(\mathsf{crs}_\Pi,c)$.

(2) Receive $(|\$\rangle,s)\leftarrow\mathsf{NoteGen}$ from the challenger.

(3) Define $|\$\rangle_\ell=|\$\rangle$, $s_\ell=s$, and $x_\Pi=(c,x_\ell,s_\ell)$. Compute $\pi_{\Pi,\ell}\leftarrow\Pi.\mathsf{P}(\mathsf{crs}_\Pi,x_\Pi,w_\ell)$. Define $\pi_\ell=(|\$\rangle_\ell,s_\ell,\pi_{\Pi,\ell})$.

(4) Define $\pi_\iota\leftarrow\mathsf{P}(\mathsf{crs},x_\iota,w_\iota)$ for $\iota\in[k-1]\setminus\{\ell\}$.

(5) Send $\{x_\iota,\pi_\iota\}_{\iota\in[k-1]}$ to $\mathcal{A}$.

(6) Receive $\{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$. Parse $\widetilde{\pi_i}=(\widetilde{|\$_i\rangle},\widetilde{s_i},\widetilde{\pi_{\Pi,i}})$ and $\widetilde{\pi_j}=(\widetilde{|\$_j\rangle},\widetilde{s_j},\widetilde{\pi_{\Pi,j}})$.

(7) Send $(\widetilde{|\$_i\rangle},\widetilde{|\$_j\rangle})$ to the challenger.

Given the event in Equation (9) holds (for the afore mentioned fixed indices), then the reduction will return two quantum money states with the same serial number as the challenger sent. With advantage $1/(2k^3p(\lambda))$, the reduction will succeed at breaking unforgeability of the quantum money scheme, thus reaching a contradiction.

Scenario Two

Alternatively, say that $\mathcal{A}$ does not generate two accepting proofs which have the same serial number as an honestly generated proof. By the pigeon-hole principle, this means that $\mathcal{A}$ generates an accepting proof with a serial number which is not amongst the ones it received. By applying a union bound to Equation (7), we have that this event could happen with at least $1/2p(\lambda)$ probability.

In summary, we have that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{P}(\mathsf{crs},x_\iota,w_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi_\iota})=1\ \bigwedge\ \exists j\in\mathcal{J}\text{ s.t. }\widetilde{s_j}\notin\{s_\iota\}_{\iota\in[k-1]}\right]\geq\frac{1}{2p(\lambda)}. \quad (10)$$

Through an averaging argument, we can fix index $j\in\mathcal{J}$ which gives us the same event with an advantage of $1/(2kp(\lambda))$. We will now switch to a hybrid where we provide $\mathcal{A}$ with simulated proofs at indices $\mathcal{I}$.

Claim 4.12. There exists a polynomial $q(\cdot)$ such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\left(\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi_\iota})=1\right)\ \bigwedge\ \widetilde{s_j}\notin\{s_\iota\}_{\iota\in[k-1]}\right]\geq\frac{1}{q(\lambda)}. \quad (11)$$

We will later see a proof of Claim 4.12. For now, assuming that this claim holds, by the definition of $\mathcal{E}$, Equation (8), and Equation (11), there exists a polynomial $q'(\cdot)$ such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})\\ b'\xleftarrow{\$}\mathcal{J}\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},\widetilde{x_{b'}},\widetilde{\pi_{b'}})}}\left[\left(\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi_\iota})=1\right)\ \wedge\ \widetilde{s_j}\notin\{s_\iota\}_{\iota\in[k-1]}\ \wedge\ (x,w)\notin\mathcal{R}\right]\geq\frac{1}{q'(\lambda)}.$$

We will additionally have that $j'=j$ with advantage at least $1/(kq'(\lambda))$. Since V accepts $\widetilde{\pi_j}$ with respect to x, $\Pi.\mathsf{V}$ must accept $\widetilde{\pi_{\Pi,j}}$ with respect to $\widetilde{x_{\Pi,j}}=(c,x,\widetilde{s_j})$. Since $\widetilde{s_j}\notin\{s_\iota\}_{\iota\in[k-1]}$, we have that $\Pi.\mathsf{Sim}_1$, through Sim_1, has not previously received $\widetilde{x_{\Pi,j}}$ as a query. As such, we have that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})\\ w\leftarrow\mathsf{Ext}(\mathsf{crs},\mathsf{td},\widetilde{x_j},\widetilde{\pi_j})}}[\Pi.\mathsf{V}(\mathsf{crs}_\Pi,\widetilde{x_{\Pi,j}},\widetilde{\pi_{\Pi,j}})=1\wedge\widetilde{x_{\Pi,}}\notin Q_\Pi\wedge(x,w)\notin\mathcal{R}]\geq\frac{1}{kq'(\lambda)}. \quad (12)$$

We now define $\mathcal{B}$ with oracle access to $\mathcal{A}$ and Sim_1 [2]:

Hardwired: $x_1,\ldots,x_{k-1}$, j

Input: crs

(1) For $\iota\in[k-1]$: send x_ι to Sim_1, and receive π_ι from Sim_1.

(2) Send $(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$.

(3) Output $(\widetilde{x_j},\widetilde{\pi_j})$.

Given that the event in Equation (12) holds, then $\mathcal{B}$ contradicts Claim 4.11. Thus, all that remains to be proven is Claim 4.12.

[2] Here, $\mathcal{B}$ is given oracle access to Sim_1 which has the terms $(\mathsf{crs},\mathsf{td})$ fixed by the output of Sim_0.

Proof of Claim 4.12. We proceed by contradiction. Let $\mathsf{negl}(\cdot)$ be a negligible function such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_\iota)\\ \{\widetilde{x}_\iota,\widetilde{\pi}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\left(\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi}_\iota)=1\right)\bigwedge\widetilde{s}_j\notin\{s_\iota\}_{\iota\in[k-1]}\right]\leq\mathsf{negl}(\lambda). \tag{13}$$

By Equation (10) and Equation (13), there exists a polynomial $q^*(\cdot)$ such that

$$\left|\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{P}(\mathsf{crs},x_\iota,w_\iota)\\ \{\widetilde{x}_\iota,\widetilde{\pi}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\left(\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi}_\iota)=1\right)\bigwedge\widetilde{s}_j\notin\{s_\iota\}_{\iota\in[k-1]}\right]\right.$$

$$\left.-\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_\iota)\\ \{\widetilde{x}_\iota,\widetilde{\pi}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})}}\left[\left(\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi}_\iota)=1\right)\bigwedge\widetilde{s}_j\notin\{s_\iota\}_{\iota\in[k-1]}\right]\right|\geq\frac{1}{q^*(\lambda)}. \tag{14}$$

By using the advantage of $\mathcal{A}$ in this game, we can show a reduction that breaks the multi-theorem zero-knowledge of Figure 3. We will now outline this reduction.

Reduction: to multi-theorem zero-knowledge of our protocol given oracle access to $\mathcal{A}$.
Hardwired: $(x_1,w_1),\ldots,(x_{k-1},w_{k-1})$, j
(1) Receive (real or simulated) crs from the challenger.
(2) For $\iota\in[k-1]$: send (x_ι,w_ι) to the challenger, and receive (real or simulated) π_ι from the challenger.
(3) Send $(\mathsf{crs},\{x_\iota,\pi_\iota\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x}_\iota,\widetilde{\pi}_\iota\}_{\iota\in[k]}$ from $\mathcal{A}$.
(4) Parse $\widetilde{\pi}_b=(\widetilde{|\$_b\rangle},\widetilde{s}_b,\widetilde{\pi_{\Pi,b}})$.
(5) Output $\left(\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}(\mathsf{crs},x,\widetilde{\pi}_\iota)=1\right)\bigwedge\widetilde{s}_j\notin\{s_\iota\}_{\iota\in[k-1]}$.

Given that $\mathcal{A}$ is able to change its output dependent on which of the two worlds in Equation (14) that it is in, then the reduction will be able to distinguish between receiving honest proofs or simulated proofs. With advantage $1/q^*(\lambda)$, the reduction will succeed at breaking the adaptive multi-theorem computational zero-knowledge of our protocol, thus reaching a contradiction. □

By completing the proofs of our claim, we have concluding the proof of our theorem statement. □

Corollary 4.13. Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, and post-quantum iO exists, there exists a non-interactive adaptive knowledge sound, adaptive computationally zero-knowledge, and $(k-1)$-to-k-unclonable with extraction protocol for NP in the common reference string model (Definition 4.7).

Proof. This follows from Theorem 3.2, Corollary 4.4, Theorem 3.14, and Theorem 4.10. □

We have thus shown that Figure 3 is an unclonable NIZK PoK in the CRS model as defined according to our proposed unclonability definition, Definition 4.7.

In the upcoming sections, we will consider unclonable proof systems in the QROM.

5 Unclonable NIZK in the Quantum Random Oracle Model

5.1 A Modified Sigma Protocol

We will begin by introducing a slightly modified sigma protocol. In the coming sections, our construction will involve applying Fiat-Shamir to this modified protocol.

Theorem 5.1. *Let a post-quantum sigma protocol with unpredictable commitments Π be given (see Definition 3.4). Let $\mathcal{R}_\Pi$ be an NP relation. Let $\mathcal{R} = \{((x,\mathcal{S}),w) \ : \ (x,w) \in \mathcal{R}_\Pi \wedge \mathcal{S} \neq \emptyset\}$. We argue that the following protocol will be a post-quantum sigma protocol with unpredictable commitments (see Definition 3.4):*

- $\mathsf{P.Com}(1^\lambda,(x,\mathcal{S}),w)$*: Sends* (x,α,s) *to* V *where* $(\alpha,\mathsf{st}) \leftarrow \Pi.\mathsf{P.Com}(1^\lambda,x,w)$ *and* s *is sampled from* $\mathcal{S}$*.*
- $\mathsf{V.Ch}(1^\lambda,(x,\mathcal{S}),(x,\alpha,s))$*: Sends* β *to* P *where* $\beta \leftarrow \Pi.\mathsf{V.Ch}(1^\lambda,x,\alpha)$*.*
- $\mathsf{P.Com}(1^\lambda,(x,\mathcal{S}),w,\mathsf{st},\beta)$*: Sends* γ *to* V *where* $\gamma \leftarrow \Pi.\mathsf{P.Prove}(1^\lambda,x,w,\mathsf{st},\beta)$*.*
- $\mathsf{V.Ver}(1^\lambda,(x,\mathcal{S}),(x,\alpha,s),\beta,\gamma)$*: Outputs* 1 *iff* $s \in \mathsf{Support}(\mathcal{S})$ *and* $\Pi.\mathsf{V.Ver}(1^\lambda,x,\alpha,\beta,\gamma) = 1$*.*

Proof. **Perfect completeness** This follows directly from the perfect completeness of Π.

Proof of Knowledge with Quantum Extractor. Let $\Pi.\mathsf{Ext}$ be the proof of knowledge quantum extractor for Π. Let constant c_Π, polynomial $p_\Pi(\cdot)$, and negligible functions $\mathsf{negl}_{0,\Pi}(\cdot),\mathsf{negl}_{1,\Pi}(\cdot)$ be given such that for any quantum $\mathcal{A}_\Pi = (\mathcal{A}_{0,\Pi},\mathcal{A}_{1,\Pi})$ where

- $\mathcal{A}_{0,\Pi}(x)$ is a unitary U_x followed by a measurement and
- $\mathcal{A}_{1,\Pi}(x),|\mathsf{st}\rangle,\beta)$ is a unitary $V_{x,\beta}$ onto the state $|\mathsf{st}\rangle$ followed by a measurement,

and any x with associated $\lambda \in \mathbb{N}$ satisfying

$$\Pr_{\substack{(\alpha,|\mathsf{st}\rangle)\leftarrow\mathcal{A}_{0,\Pi}(x)\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_{1,\Pi}(x,|\mathsf{st}\rangle,\beta)}}[\Pi.\mathsf{V.Ver}(x,\alpha,\beta,\gamma)=1] \geq \mathsf{negl}_{0,\Pi}(\lambda) \tag{15}$$

we have

$$\Pr\left[(x,\Pi.\mathsf{Ext}^{\mathcal{A}_\Pi(x)}(x)) \in \mathcal{R}_\Pi\right]$$
$$\geq \frac{1}{p(\lambda)} \cdot \left(\Pr_{\substack{(\alpha,|\mathsf{st}\rangle)\leftarrow\mathcal{A}_{0,\Pi}(x)\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_{1,\Pi}(x,|\mathsf{st}\rangle,\beta)}}[\Pi.\mathsf{V.Ver}(x,\alpha,\beta,\gamma)=1] - \mathsf{negl}_{0,\Pi}(\lambda) \right)^{c_\Pi} - \mathsf{negl}_{1,\Pi}(\lambda).$$

We define Ext [3] with oracle-access to $\Pi.\mathsf{Ext}$ and some $\mathcal{A}$ as follows:

[3] An extractor whose local code is implementable as a simple unitary which allows for straightforward rewinding.

Input: x, $\mathcal{S}$.
(1) Given (x, α, s) from $\mathcal{A}_\Pi$: send α to Π.Ext, receive β from Π.Ext, and send β to $\mathcal{A}_\Pi$.
(2) Upon receiving γ from $\mathcal{A}_\Pi$: send γ to Π.Ext.
(3) Output the result of Π.Ext as w.

We define the following set of parameters: $c = c_\Pi$, $p(\cdot) = p_\Pi(\cdot)$, $\mathsf{negl}_0(\cdot) = \mathsf{negl}_{0,\Pi}(\cdot)$ and $\mathsf{negl}_1(\cdot) = \mathsf{negl}_{1,\Pi}(\cdot)$.

Let polynomial-size quantum circuit $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ and $(x, \mathcal{S})$ be given such that

$$\Pr_{\substack{((x,\alpha,s),|\mathsf{st}\rangle)\leftarrow\mathcal{A}_0(x,\mathcal{S})\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_1((x,\mathcal{S}),|\mathsf{st}\rangle,\beta)}}[\mathsf{V.Ver}((x,\mathcal{S}),(x,\alpha,s),\beta,\gamma)=1] \geq \mathsf{negl}_0(\lambda).$$

We now define $\mathcal{A}_\Pi = (\mathcal{A}_{0,\Pi}, \mathcal{A}_{1,\Pi})$ with oracle-access to $\mathcal{A}$. $\mathcal{A}_{0,\Pi}$ is hardwired with $\mathcal{S}$, takes input x, sends $(x, \mathcal{S})$ to $\mathcal{A}_0$, receives $((x, \alpha, s), |\mathsf{st}\rangle)$ from $\mathcal{A}_0$, and outputs $(\alpha, |\mathsf{st}\rangle)$. $\mathcal{A}_{1,\Pi}$ is hardwired with $\mathcal{S}$, takes input $(x, |\mathsf{st}\rangle, \beta)$, sends $((x, \mathcal{S}), |\mathsf{st}\rangle, \beta)$ to $\mathcal{A}_1$, receives γ from $\mathcal{A}_1$, outputs γ. By the structure of our proof and definition of our verifier, this means that

$$\begin{aligned}
&\Pr_{\substack{(\alpha,|\mathsf{st}\rangle)\leftarrow\mathcal{A}_{0,\Pi}^{\mathcal{A}_0}(x,\mathcal{S})\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_{1,\Pi}^{\mathcal{A}_1}((x,\mathcal{S}),|\mathsf{st}\rangle,\beta)}}[\Pi.\mathsf{V.Ver}(x,\alpha,\beta,\gamma)=1]\\
&\geq \Pr_{\substack{((x,\alpha,s),|\mathsf{st}\rangle)\leftarrow\mathcal{A}_0(x,\mathcal{S})\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_1((x,\mathcal{S}),|\mathsf{st}\rangle,\beta)}}[\mathsf{V.Ver}((x,\mathcal{S}),(x,\alpha,s),\beta,\gamma)=1] \geq \mathsf{negl}_0(\lambda)
\end{aligned}$$

which satisfies the constraint in Equation (15). This means we have, when combined with our definition of Ext, that

$$\begin{aligned}
&\Pr\left[((x,\mathcal{S}),\mathsf{Ext}^{\mathcal{A}(x,\mathcal{S})}(x,\mathcal{S}))\in\mathcal{R}\right] = \Pr\left[(x,\Pi.\mathsf{Ext}^{\mathcal{A}_\Pi(x,\mathcal{S})}(x))\in\mathcal{R}_\Pi\right]\\
&\geq \frac{1}{p_\Pi(\lambda)}\cdot\left(\Pr_{\substack{((x,\alpha,s),|\mathsf{st}\rangle)\leftarrow\mathcal{A}_{0,\Pi}^{\mathcal{A}_0}(x,\mathcal{S})\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_{1,\Pi}^{\mathcal{A}_1}((x,\mathcal{S}),|\mathsf{st}\rangle,\beta)}}[\Pi.\mathsf{V.Ver}(x,\alpha,\beta,\gamma)=1]-\mathsf{negl}_{0,\Pi}(\lambda)\right)^{c_\Pi}-\mathsf{negl}_{1,\Pi}(\lambda)\\
&\geq \frac{1}{p_\Pi(\lambda)}\cdot\left(\Pr_{\substack{((x,\alpha,s),|\mathsf{st}\rangle)\leftarrow\mathcal{A}_0(x,\mathcal{S})\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_1((x,\mathcal{S}),|\mathsf{st}\rangle,\beta)}}[\mathsf{V.Ver}((x,\mathcal{S}),(x,\alpha,s),\beta,\gamma)=1]-\mathsf{negl}_{0,\Pi}(\lambda)\right)^{c_\Pi}-\mathsf{negl}_{1,\Pi}(\lambda)\\
&\geq \frac{1}{p(\lambda)}\cdot\left(\Pr_{\substack{((x,\alpha,s),|\mathsf{st}\rangle)\leftarrow\mathcal{A}_0(x,s)\\ \beta\leftarrow\{0,1\}^\lambda\\ \gamma\leftarrow\mathcal{A}_1((x,s),|\mathsf{st}\rangle,\beta)}}[\mathsf{V.Ver}((x,s),(x,\alpha,s),\beta,\gamma)=1]-\mathsf{negl}_0(\lambda)\right)^{c}-\mathsf{negl}_1(\lambda).
\end{aligned}$$

Thus showing that our protocol is a proof of knowledge protocol.

Computational Honest-Verifier Zero-Knowledge with Quantum Simulator. Let $\Pi.\mathsf{Sim}$ be the computational honest-verifier zero-knowledge quantum simulator for Π. We define Sim with oracle access to $\Pi.\mathsf{Sim}$ as follows:

Input: x, $\mathcal{S}$.
(1) Compute $(\alpha, \beta, \gamma) \leftarrow \Pi.\mathsf{Sim}(1^\lambda, x)$.
(2) Sample s from $\mathcal{S}$.
(3) Output $((x, \alpha, s), \beta, \gamma)$.

Let a polynomial $p(\cdot)$, a polynomial-size quantum circuit $\mathcal{D}$, $\lambda \in \mathbb{N}$, and $((x, \mathcal{S}), w) \in \mathcal{R}$ be given such that

$$\left| \begin{array}{c} \\ \\ \\ \Pr \\ {\scriptstyle ((x,\alpha,s),\mathsf{st}) \leftarrow \mathsf{P.Com}((x,\mathcal{S}),w)} \\ {\scriptstyle \beta \leftarrow \mathsf{V.Ch}((x,\mathcal{S}),(x,\alpha,s))} \\ {\scriptstyle \gamma \leftarrow \mathsf{P.Prove}((x,\mathcal{S}),w,\mathsf{st},\beta)} \end{array} [\mathcal{D}((x, \mathcal{S}), (x, \alpha, s), \beta, \gamma) = 1] \right.$$

$$\left. - \Pr_{((x,\alpha,s),\beta,\gamma) \leftarrow \mathsf{Sim}(1^\lambda,(x,\mathcal{S}))} [\mathcal{D}((x, \mathcal{S}), (x, \alpha, s), \beta, \gamma) = 1] \right| \geq \frac{1}{p(\lambda)}.$$

We define a reduction to the zero-knowledge property of Π as follows:

Reduction: to zero-knowledge of Π given oracle access to $\mathcal{D}$.
Hardwired with: x, $\mathcal{S}$.
(1) Receive (real or simulated) (α, β, γ) from the challenger.
(2) Sample s from $\mathcal{S}$.
(3) Send $((x, \alpha, s), \beta, \gamma)$ to $\mathcal{D}$. Receive b from $\mathcal{D}$.
(4) Output b.

When the challenger sends a real (or simulated) proof for Π, the reduction generates a proof that is identical to the real (resp. simulated) proof. As such, this reduction preserves the distinguishing advantage of $\mathcal{D}$. This reaches a contradiction against the zero-knowledge property of Π. Hence, our protocol must be zero-knowledge.

Unpredictable Commitment. Let $\mathsf{negl}_\Pi(\cdot)$ be a negligible function for the unpredictable commitment property of Π.

Let a polynomial function $p(\cdot)$, $\lambda \in \mathbb{N}$, and $((x, \mathcal{S}), w) \in \mathcal{R}$ be given such that

$$\Pr_{\substack{((x,\alpha,s),\mathsf{st}) \leftarrow \mathsf{P.Com}((x,\mathcal{S}),w) \\ ((x,\alpha',s'),\mathsf{st}') \leftarrow \mathsf{P.Com}((x,\mathcal{S}),w)}} [(\alpha, s) = (\alpha', s')] \geq \frac{1}{p(\lambda)}.$$

By the definition of the honest prover $\mathsf{P.Com}$,

$$\Pr_{\substack{(\alpha,\mathsf{st}) \leftarrow \Pi.\mathsf{P.Com}(x,w) \\ (\alpha',\mathsf{st}') \leftarrow \Pi.\mathsf{P.Com}(x,w)}} [\alpha = \alpha'] \geq \Pr_{\substack{((x,\alpha,s),\mathsf{st}) \leftarrow \mathsf{P.Com}((x,\mathcal{S}),w) \\ ((x,\alpha',s'),\mathsf{st}') \leftarrow \mathsf{P.Com}((x,\mathcal{S}),w)}} [(\alpha, s) = (\alpha', s')] \geq \frac{1}{p(\lambda)}$$

which is a contradiction. Hence our protocol must have unpredictable commitments. □

Corollary 5.2. The Fiat-Shamir heuristic applied to the post-quantum sigma protocol defined in Theorem 5.1 yields a classical post-quantum NIZKPoK Π' in the QROM (Definition 3.11).

Proof. This follows by Theorem 5.1 and Theorem 3.12. □

5.2 Unclonability Definitions

Unclonable NIZKs in the quantum random oracle model are defined analogously to the CRS model – we repeat these definitions in the QRO model for completeness below.

Definition 5.3. (Unclonable Security for Hard Instances). A proof (P, V) satisfies unclonable security with respect to a quantum random oracle $\mathcal{O}$ if for every language $\mathcal{L}$ with corresponding relation $\mathcal{R}_{\mathcal{L}}$, for every polynomial-sized quantum circuit family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$, and for every hard distribution $\{\mathcal{X}_\lambda, \mathcal{W}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\mathcal{R}_{\mathcal{L}}$, there exists a negligible function $\mathsf{negl}_1(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr_{\substack{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)\\ \pi\leftarrow\mathsf{P}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2\leftarrow C_\lambda(x,\pi)}}\left[\mathsf{V}^{\mathcal{O}}(x,\pi_1)=1 \bigwedge \mathsf{V}^{\mathcal{O}}(x,\pi_2)=1\right] \leq \mathsf{negl}_1(\lambda).$$

Definition 5.4 ($(k-1)$-to-k-Unclonable Extractable NIZK in QROM)**.** Let security parameter $\lambda \in \mathbb{N}$ and NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given. Let $\Pi = (\mathsf{P}, \mathsf{V})$ be given such that P and V are $\mathsf{poly}(\lambda)$-size quantum algorithms. We have that for any $(x, \omega) \in \mathcal{R}$, P receives an instance and witness pair (x, ω) as input and outputs π, and V receives an instance x and proof π as input and outputs a value in $\{0, 1\}$.

Π is a non-interactive $(k-1)$-to-k-unclonable NIZKPoK protocol for language $\mathcal{L}$ with respect to a random oracle $\mathcal{O}$ if the following holds:

- Π is a NIZKPoK protocol for language $\mathcal{L}$ in the quantum random oracle model (Definition 3.11).
- $(k-1)$**-to-**k**-Unclonable with Extraction**: There exists an oracle-aided polynomial-size quantum circuit $\mathcal{E}$ such that for every polynomial-size quantum circuit $\mathcal{A}$, for every tuple of $k-1$ instance-witness pairs $(x_1, \omega_1), \ldots, (x_{k-1}, \omega_{k-1}) \in \mathcal{R}$, for every x where we define
  - $\mathcal{I} \subseteq [k-1]$ such that $|\mathcal{I}| \geq 1$, $x_i = x$ for all $i \in \mathcal{I}$, and $x_\iota \neq x$ for all $\iota \notin \mathcal{I}$, and
  - $\mathcal{J} \subseteq [k]$ such that $|\mathcal{J}| \geq \max\{2, |\mathcal{I}|\}$, $x_j = x$ for all $j \in \mathcal{J}$, and $x_\iota \neq x$ for $\iota \notin \mathcal{J}$,

  such that there is a polynomial $p(\cdot)$ where

$$\Pr_{\substack{\mathcal{O}\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\mathsf{P}^{\mathcal{O}}(x_\iota,\omega_\iota)\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}^{\mathcal{O}}(\{\pi_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}^{\mathcal{O}}(x,\widetilde{\pi_\iota})=1\right] \geq \frac{1}{p(\lambda)},$$

  then there is also a polynomial $q(\cdot)$ such that

$$\Pr_{w\leftarrow\mathcal{E}^{\mathcal{A}}(x_1,\ldots,x_{k-1},x,\mathcal{I},\mathcal{J})}[(x,w)\in\mathcal{R}] \geq \frac{1}{q(\lambda)}.$$

Similar to the previous section, we have the following lemma.

Lemma 5.5. Let $\Pi = (\mathsf{Setup}, \mathsf{P}, \mathsf{V})$ be a a non-interactive 1-to-2-unclonable zero-knowledge quantum protocol (Definition 5.4). Then, Π satisfies Definition 5.3.

For a proof of Lemma 5.5, we refer to Appendix A.

5.3 Unclonable NIZK Implies Public-Key Quantum Money in QROM

Public-Key Quantum Money Mini-Scheme

Let $\mathcal{O}$ be a quantum random oracle. Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \mathsf{NP}$. Let $\Pi = (\mathsf{P}, \mathsf{V})$ be an unclonable non-interactive zero-knowledge protocol for $\mathcal{L}$ in the QROM.

$\underline{\text{GEN}}^{\mathcal{O}}(1^\lambda)$: Sample a hard instance-witness pair $(x, w) \leftarrow (\mathcal{X}, \mathcal{Y})$ and a proof $\pi \leftarrow \mathsf{P}^{\mathcal{O}}(x, w)$. Output $(|\$\rangle = \pi, s = x)$.

$\underline{\text{VERIFY}}^{\mathcal{O}}(|\$\rangle, s)$: Parse $|\$\rangle = \pi$ and $s = x$. Output $\mathsf{V}^{\mathcal{O}}(x, \pi)$.

Figure 4: Public-Key Quantum Money Mini-Scheme from an Unclonable Non-Interactive Quantum Protocol

Theorem 5.6. *Let $\mathcal{O}$ be a quantum random oracle. Let $(\mathcal{X}, \mathcal{W})$ be a hard distribution over a language $\mathcal{L} \in \mathsf{NP}$. Let $\Pi = (\mathsf{P}, \mathsf{V})$ be a 1-to-2 unclonable non-interactive perfectly complete, computationally zero-knowledge protocol for $\mathcal{L}$ in the QRO model (Definition 5.4).*

Then (P, V) implies a public-key quantum money mini-scheme in the QRO model (Definition 3.15) as described in Figure 4.

Proof. **Perfect Correctness.** This follows directly from the perfect completeness of Π.
Unforgeability. Let $p(\cdot)$ be a polynomial and $\mathcal{A}$ be a quantum polynomial-time adversary such that for an infinite number of $\lambda \in \mathbb{N}^+$,

$$\Pr_{\substack{(|\$\rangle, s) \leftarrow \mathsf{Gen}^{\mathcal{O}}(1^\lambda) \\ (|\$_0\rangle, s_0, |\$_1\rangle, s_1) \leftarrow \mathcal{A}^{\mathcal{O}}(|\$\rangle, s)}} [s_0 = s_1 = s \ \wedge \ \mathsf{Ver}^{\mathcal{O}}(|\$_0\rangle, s_0) = 1 \ \wedge \ \mathsf{Ver}^{\mathcal{O}}(|\$_1\rangle, s_1) = 1] \geq \frac{1}{p(\lambda)}.$$

We construct a reduction that breaks the uncloneability definition (Definition 5.3) which we show (in Appendix A) is implied by our definition (Definition 5.4). The challenger, with access to random oracle $\mathcal{O}$, samples a hard instance-witness pair $(x, w) \leftarrow (\mathcal{X}, \mathcal{Y})$ and a proof $\pi \leftarrow \mathsf{P}^{\mathcal{O}}(x, w)$. The challenger then forwards (x, π) to the reduction, which also has oracle access to $\mathcal{O}$. The reduction then sets $|\$\rangle = \pi$ and $s = x$. The reduction sends $(|\$\rangle, s)$ to the adversary $\mathcal{A}$ who returns back $(|\$_0\rangle, s_0, |\$_1\rangle, s_1)$. The reduction then parses and sets $\pi_i = |\$_i\rangle$ for $i \in \{0, 1\}$. The reduction then sends π_0 and π_1 back to the challenger.

When the serial numbers are the same, $s = s_0 = s_1$, we have that the instance will be the same for all the proofs π, π_0, π_1. The quantum money state can be parsed as the proof as shown in the construction. When the verification algorithm of the quantum money algorithm accepts

both quantum money states $|\$_0\rangle$ and $|\$_1\rangle$ with respect to s, we know that that $\mathsf{V}^{\mathcal{O}}$ would accept both proofs π_0 and π_1 with respect to x. As such, we will have that the advantage that $\mathcal{A}$ has at breaking the unforgeability of our quantum money scheme directly translates to the advantage of the reduction at breaking the uncloneability of Π. □

5.4 Construction and Analysis

Lemma 5.7. Let $\lambda, k \in \mathbb{N}$ and a public-key quantum money mini-scheme $(\mathsf{NoteGen}, \mathsf{Ver})$ be given. Let points $q_1, \ldots, q_k$ with the following structure be given: a point q contains a serial number s sampled according to $\mathsf{NoteGen}(1^\lambda)$.

The points $q_1, \ldots, q_k$ must be distinct with overwhelming probability.

Proof. Each point contains a serial number sampled according to the quantum money generation algorithm, $\mathsf{NoteGen}(1^\lambda)$. By the unpredictability of the serial numbers of quantum money (Definition 3.13), all k honestly generated serial numbers must be distinct with overwhelming probability. Hence, these k points will be distinct with overwhelming probability. □

Unclonable NIZK for NP in the QROM

Let $\mathcal{O}$ be a random oracle. Let $\Pi = (\mathsf{P} = (\mathsf{P.Com}, \mathsf{P.Prove}), \mathsf{V} = (\mathsf{V.Ch}, \mathsf{V.Ver}))$ be a post-quantum sigma protocol with unpredictable commitments (see Definition 3.4), and $(\mathsf{NoteGen}, \mathsf{Ver})$ be a public-key quantum money mini-scheme (see Definition 3.13). Let $\mathcal{R}$ be the relation with respect to $\mathcal{L} \in \mathsf{NP}$.

$\underline{\textsc{Prove}}^{\mathcal{O}}(x, \omega)$:

- Compute a quantum note and associated serial number $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$.
- Compute $(\alpha, \zeta) \leftarrow \mathsf{P.Com}(x, \omega)$.
- Query $\mathcal{O}$ at (x, α, s) to get β.
- Compute $\gamma \leftarrow \mathsf{P.Prove}(x, \omega, \beta, \zeta)$.
- Output $\pi = (|\$\rangle, s, \alpha, \beta, \gamma)$.

$\underline{\textsc{Verify}}^{\mathcal{O}}(x, \pi)$:

- Check that $\mathsf{Ver}(|\$\rangle, s)$ outputs 1.
- Check that $\mathcal{O}$ outputs β when queried at (x, α, s).
- Output the result of $\mathsf{V.Ver}(x, \alpha, \beta, \gamma)$.

Figure 5: Unclonable Non-Interactive Quantum Protocol for $\mathcal{L} \in \mathsf{NP}$ in the Quantum Random Oracle Model

We now introduce our construction in Figure 5 and prove the main theorem of this section.

Theorem 5.8. *Let $k(\cdot)$ be a polynomial. Let* NP *relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given.*

Let $(\mathsf{NoteGen}, \mathsf{Ver})$ *be a public-key quantum money mini-scheme (Definition 3.13) and* $\Pi = (\mathsf{P}, \mathsf{V})$ *be a post-quantum sigma protocol (Definition 3.4).*

(P, V) *as defined in Figure 5 will be a non-interactive knowledge sound, computationally zero-knowledge, and* $(k-1)$*-to-*k*-unclonable with extraction protocol for $\mathcal{L}$ in the quantum random oracle model (Definition 3.11).*

Proof. Let the parameters and primitives be as given in the theorem statement. We argue that completeness follows from the protocol construction in Figure 5, and we prove the remaining properties below.

Proof of Knowledge. Let Ext_{FS} be the extractor for Π' in Corollary 5.2 (where Π instantiates Theorem 5.1). Let $\mathcal{R}_{FS}$ be the relation for Π' with respect to $\mathcal{R}$. Let constant c_{FS}, polynomial $p_{FS}(\cdot)$, and negligible functions $\mathsf{negl}_{0,FS}(\cdot), \mathsf{negl}_{1,FS}(\cdot)$ be given such that for any quantum $\mathcal{A}_{FS}$ and any $(x, \mathcal{S})$ with associated $\lambda \in \mathbb{N}$ satisfying

$$\Pr_{\substack{\mathcal{O} \\ \pi_{FS} \leftarrow \mathcal{A}_{FS}^{|\mathcal{O}\rangle}(x,\mathcal{S})}} [\mathsf{V}_{FS}^{\mathcal{O}}((x, \mathcal{S}), \pi_{FS}) = 1] \geq \mathsf{negl}_{0,FS}(\lambda) \tag{16}$$

we have

$$\Pr\left[(x, \mathsf{Ext}_{FS}^{\mathcal{A}_{FS}^{|\mathcal{O}\rangle}(x,\mathcal{S})}(x, \mathcal{S})) \in \mathcal{R}_{FS}\right]$$
$$\geq \frac{1}{p_{FS}(\lambda)} \cdot \left(\Pr_{\substack{\mathcal{O} \\ \pi_{FS} \leftarrow \mathcal{A}_{FS}^{|\mathcal{O}\rangle}(x,\mathcal{S})}} [\mathsf{V}_{FS}^{\mathcal{O}}((x, \mathcal{S}), \pi_{FS}) = 1] - \mathsf{negl}_{0,FS}(\lambda) \right)^{c_{FS}} - \mathsf{negl}_{1,FS}(\lambda).$$

Let $\mathcal{S}$ be the distribution of serial numbers as output by $\mathsf{NoteGen}(1^\lambda)$. We define Ext [4] with oracle-access to Ext_{FS}, $\mathcal{O}$, and some $\mathcal{A}$ as follows:

Hardwired with: $\mathcal{S}$.

Input: x.

(1) Given an oracle-query (x, α, s) from $\mathcal{A}$: send (x, α, s) to $\mathcal{O}$, receive β from $\mathcal{O}$, and send β to $\mathcal{A}$.

(2) Upon receiving $\pi = (|\$\rangle, s, \alpha, \beta, \gamma)$ from $\mathcal{A}$: send $\pi_{FS} = ((x, \alpha, s), \beta, \gamma)$ to Ext_{FS}.

(3) Output the result of Ext_{FS} as w.

We define the following set of parameters: $c = c_{FS}$, $p(\cdot) = p_{FS}(\cdot)$, $\mathsf{negl}_0(\cdot) = \mathsf{negl}_{0,FS}(\cdot)$ and $\mathsf{negl}_1(\cdot) = \mathsf{negl}_{1,FS}(\cdot)$.

Let polynomial-size quantum circuit $\mathcal{A}$ and x be given such that

$$\Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x)}} [\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] \geq \mathsf{negl}_0(\lambda).$$

Let $\mathcal{A}_{FS}$ be defined with oracle-access to some $\mathcal{A}$ and $\mathcal{O}$ as follows:

[4] An extractor whose local code is implementable as a simple unitary which allows for straightforward rewinding.

Input: x, $\mathcal{S}$.

(1) Given a query (x, α, s) from $\mathcal{A}$: send (x, α, s) to $\mathcal{O}$, receive β from $\mathcal{O}$, and send β to $\mathcal{A}$.

(2) Upon receiving $\pi = (|\$\rangle, s, \alpha, \beta, \gamma)$ from $\mathcal{A}$: output $\pi_{FS} = ((x, \alpha, s), \beta, \gamma)$.

By the structure of our proof and definition of our verifier, this means that

$$\Pr_{\substack{\mathcal{O} \\ \pi_{FS} \leftarrow \mathcal{A}_{FS}^{\mathcal{A}(x),|\mathcal{O}\rangle}(x,\mathcal{S})}} [\mathsf{V}_{FS}^{\mathcal{O}}((x,\mathcal{S}), \pi_{FS}) = 1] \geq \Pr_{\substack{\mathcal{O} \\ (|\$\rangle, \pi_{FS}) \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x,\mathcal{S})}} [\mathsf{V}_{FS}^{\mathcal{O}}((x,\mathcal{S}), \pi_{FS}) = 1 \ \wedge \ \mathsf{Ver}(|\$\rangle, s) = 1]$$
$$= \Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x)}} [\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] \geq \mathsf{negl}_0(\lambda) = \mathsf{negl}_{0,FS}(\lambda)$$

which satisfies the constraint in Equation (16). This means we have, when combined with our definition of Ext and $\mathcal{S}$, that

$$\Pr\left[(x, \mathsf{Ext}^{\mathsf{Ext}_{FS}(x),|\mathcal{O}\rangle,\mathcal{A}(x)}(x)) \in \mathcal{R}\right] = \Pr\left[((x,\mathcal{S}), \mathsf{Ext}_{FS}^{\mathcal{A}_{FS}^{\mathcal{A}(x),|\mathcal{O}\rangle}(x,\mathcal{S})}(x,\mathcal{S})) \in \mathcal{R}_{FS}\right]$$
$$\geq \frac{1}{p_{FS}(\lambda)} \cdot \left(\Pr_{\substack{\mathcal{O} \\ \pi_{FS} \leftarrow \mathcal{A}_{FS}^{\mathcal{A}(x),|\mathcal{O}\rangle}(x,\mathcal{S})}} [\mathsf{V}_{FS}^{\mathcal{O}}((x,\mathcal{S}), \pi_{FS}) = 1] - \mathsf{negl}_{0,FS}(\lambda) \right)^{c_{FS}} - \mathsf{negl}_{1,FS}(\lambda)$$
$$\geq \frac{1}{p_{FS}(\lambda)} \cdot \left(\Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x)}} [\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] - \mathsf{negl}_{0,FS}(\lambda) \right)^{c_{FS}} - \mathsf{negl}_{1,FS}(\lambda)$$
$$= \frac{1}{p(\lambda)} \cdot \left(\Pr_{\substack{\mathcal{O} \\ \pi \leftarrow \mathcal{A}^{|\mathcal{O}\rangle}(x)}} [\mathsf{V}^{\mathcal{O}}(x, \pi) = 1] - \mathsf{negl}_0(\lambda) \right)^{c} - \mathsf{negl}_1(\lambda).$$

Thus showing that our protocol is a proof of knowledge protocol.

Zero-Knowledge. Let Sim_{FS} be the simulator for Π' in Corollary 5.2 (where Π instantiates Theorem 5.1). Let $\mathcal{R}_{FS}$ be the relation for Π' with respect to $\mathcal{R}$. We define Sim with oracle-access to Sim_{FS} and program access to some random oracle $\mathcal{O}$ as follows:

Input: x (ignores any witnesses it may receive).

(1) Sample $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$.

(2) Let $\mathcal{S}$ be the distribution where all probability mass is on s.

(3) Compute $((x, \alpha, s), \beta, \gamma) \leftarrow \Pi.\mathsf{Sim}(x, \mathcal{S})$. Allow $\Pi.\mathsf{Sim}$ to program $\mathcal{O}$ at (x, α, s) to return β.

(5) Output $\pi = (|\$\rangle, s, \alpha, \beta, \gamma)$.

Let an oracle-aided distinguisher $\mathcal{D}$ which can only make queries $(x, w) \in \mathcal{R}$, and a polynomial $p(\cdot)$ be given such that

$$\left| \Pr\left[\mathcal{D}^{\mathsf{Sim},|\mathcal{O}\rangle}(1^\lambda) = 1\right] - \Pr_{\mathcal{O}}\left[\mathcal{D}^{\mathsf{P}^{\mathcal{O}},|\mathcal{O}\rangle}(1^\lambda) = 1\right] \right| \geq \frac{1}{p(\lambda)}. \tag{17}$$

We define a reduction to the zero-knowledge property of Π' as follows:

Reduction: to zero-knowledge of Π' given oracle access to $\mathcal{D}$ and program access to $\mathcal{O}$.

For every (x, w) from $\mathcal{D}$:

(1) Sample $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}(1^\lambda)$.
(2) Let $\mathcal{S}$ be the distribution where all probability mass is on s.
(3) Send $((x, \mathcal{S}), w)$ to the challenger. Receive $((x, \alpha, s), \beta, \gamma)$ from the challenger. The challenger will have already programmed $\mathcal{O}$ at (x, α, s) to return β.
(4) Output $\pi = (|\$\rangle, s, \alpha, \beta, \gamma)$.
Output the result of $\mathcal{D}$.

The view of $\mathcal{D}$ matches that of our protocol in Figure 5 or Sim. As such, our reduction should have the same advantage at breaking the zero-knowledge property of Π'. We reach a contradiction, hence our protocol must be zero-knowledge.

Unclonable Extractability. Let Ext be the quantum circuit of the extractor we defined earlier (in our proof that Figure 5 is a proof of knowledge). Let Sim be the quantum circuit of the simulator that we defined earlier (in our proof that Figure 5 is a zero-knowledge protocol). We define an extractor $\mathcal{E}$ with oracle-access to some $\mathcal{A}$ as follows:

Hardwired with: some choice of $\mathcal{I} \subseteq [k-1]$, $\mathcal{J} \subseteq [k]$, $x_1, \ldots, x_{k-1} \in \mathcal{R}$, x.
(1) Samples $\ell \in \mathcal{J}$ uniformly at random.
(2) Instantiates a simulatable and extractable random oracle $\mathcal{O}$. Runs Ext on $\mathcal{O}$ throughout the interaction with $\mathcal{A}$ (which may involve rewinding, in which case we would rewind $\mathcal{A}$ and repeat the following steps). Require that Ext extracts with respect to the ℓth proof output by $\mathcal{A}$.
(3) Compute $\pi_\iota \leftarrow \mathsf{Sim}(x_\iota)$ for $\iota \in [k-1]$ where we store all points Sim would program into a list $\mathcal{P}$.
(4) Send $\{\pi_\iota\}_{\iota \in [k-1]}$ to $\mathcal{A}$.
(5) For every query from $\mathcal{A}$, if the query is in $\mathcal{P}$, then reply with the answer from $\mathcal{P}$. Else, forward the query to $\mathcal{O}$ and send the answer back to $\mathcal{A}$. Let $\widehat{\mathcal{O}}$ denote this modified random oracle.
(6) Receive $\{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]}$ from $\mathcal{A}$. Send $\widetilde{\pi_\ell}$ to Ext.
(7) Outputs the result of Ext as w.

Let $\mathcal{A}$, $(x_1, w_1), \ldots, (x_{k-1}, w_{k-1}) \in \mathcal{R}$, x, $\mathcal{I} \subseteq [k-1]$, $\mathcal{J} \subseteq [k]$, polynomial $p(\cdot)$, and negligible function $\mathsf{negl}_1(\cdot)$ be given such that the verifier V accepts all proofs indexed by $\mathcal{J}$ which $\mathcal{A}^{\mathcal{O}}$ outputs and the extractor $\mathcal{E}$ is unable to extract a valid witness. Restated more formally, that is that both

$$\Pr_{\substack{\mathcal{O} \\ \forall \iota \in [k-1],\, \pi_\iota \leftarrow \mathsf{P}^{\mathcal{O}}(x_\iota, \omega_\iota) \\ \{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]} \leftarrow \mathcal{A}^{\mathcal{O}}(\{\pi_\iota\}_{\iota \in [k-1]})}} \left[\bigwedge_{\iota \in \mathcal{J}} \mathsf{V}^{\mathcal{O}}(x, \widetilde{\pi_\iota}) = 1 \right] \geq \frac{1}{p(\lambda)}, \text{ and} \tag{18}$$

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(\{x_\iota, \omega_\iota\}_{\iota \in [k-1] \setminus \mathcal{I}}, x, \mathcal{I}, \mathcal{J})} [(x, w) \in \mathcal{R}] \leq \mathsf{negl}_1(\lambda). \tag{19}$$

Given Equation (18), we may be in one of the two following cases: either $\mathcal{A}$ generates two accepting proofs which have the same serial number as a honestly generated proof, or $\mathcal{A}$ does not. We consider these two scenarios separately and show that each reaches a contradiction.

Scenario One

Say that $\mathcal{A}$ generates two accepting proofs which have the same serial number as an honestly generated proof. By applying a union bound to Equation (18), we have that this event could happen

with at least $1/2p(\lambda)$ probability. Symbolically,

$$\Pr_{\substack{\mathcal{O} \\ \forall \iota \in [k-1],\, \pi_\iota \leftarrow \mathsf{P}^{\mathcal{O}}(x_\iota, \omega_\iota) \\ \{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]} \leftarrow \mathcal{A}^{\mathcal{O}}(\{\pi_\iota\}_{\iota \in [k-1]})}} \left[\bigwedge_{\iota \in \mathcal{J}} \mathsf{V}^{\mathcal{O}}(x, \widetilde{\pi_\iota}) = 1 \bigwedge \exists i \in \mathcal{I}\ \exists j, \ell \in \mathcal{J} \text{ s.t. } s_i = \widetilde{s_j} = \widetilde{s_\ell} \right] \geq \frac{1}{2p(\lambda)}. \quad (20)$$

Through a hybrid argument, we can fix indices $i \in \mathcal{I}$ and $j, \ell \in \mathcal{J}$ which gives us the same event with an advantage of $1/(2k^3 p(\lambda))$. By using the advantage of $\mathcal{A}$ in this game, we can show a reduction that breaks the unforgeability of the quantum money scheme. We will now outline this reduction.

Reduction: to unforgeability of quantum money scheme given oracle access to $\mathcal{A}$ and $\mathcal{O}$.
Hardwired with: $(x_1, w_1), \ldots, (x_{k-1}, w_{k-1})$, x, $\mathcal{I}$, $\mathcal{J}$, i, j, ℓ.
(1) Receive $(|\$\rangle, s) \leftarrow \mathsf{NoteGen}$ from the challenger.
(2) Define $|\$\rangle_i = |\$\rangle$ and $s_i = s$. Sample $(|\$\rangle_\iota, s_\iota) \leftarrow \mathsf{NoteGen}(1^\lambda)$ for $\iota \in [k-1] \setminus \{i\}$. Compute $(\alpha_\iota, \zeta_\iota) \leftarrow \Pi.\mathsf{P}.\mathsf{Com}(x_\iota, w_\iota)$, query $\mathcal{O}$ at $(x_\iota, \alpha_\iota, s_\iota)$ to get β_ι, compute $\gamma_\iota \leftarrow \Pi.\mathsf{P}.\mathsf{Prove}(x_\iota, w_\iota, \beta_\iota, \zeta_\iota)$, and define $\pi_\iota = (|\$\rangle_\iota, s_\iota, \alpha_\iota, \beta_\iota, \gamma_\iota)$ for $\iota \in [k-1]$.
(3) Send $\{\pi_\iota\}_{\iota \in [k-1]}$ to $\mathcal{A}$.
(4) Receive $\{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]}$ from $\mathcal{A}$.
(5) Send $(\widetilde{|\$\rangle_j}, \widetilde{|\$\rangle_\ell})$ to the challenger.

Given the event in Equation (20) holds, then the reduction will return two quantum money states with the same serial number as the challenger sent. With advantage $1/(2k^3 p(\lambda))$, the reduction will succeed at breaking unforgeability of the quantum money scheme, thus reaching a contradiction.

Scenario Two

Alternatively, say that $\mathcal{A}$ does not generate two accepting proofs which have the same serial number as an honestly generated proof. By the pigeon-hole principle, this means that $\mathcal{A}$ generates an accepting proof with a serial number which is not amongst the ones it received. By applying a union bound to Equation (18), we have that this event could happen with at least $1/2p(\lambda)$ probability. In summary, we have that

$$\Pr_{\substack{\mathcal{O} \\ \forall \iota \in [k-1],\, \pi_\iota \leftarrow \mathsf{P}^{\mathcal{O}}(x_\iota, \omega_\iota) \\ \{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]} \leftarrow \mathcal{A}^{\mathcal{O}}(\{\pi_\iota\}_{\iota \in [k-1]})}} \left[\bigwedge_{\iota \in \mathcal{J}} \mathsf{V}^{\mathcal{O}}(x, \widetilde{\pi_\iota}) = 1 \bigwedge \exists j \in \mathcal{J} \text{ s.t. } \widetilde{s_j} \notin \{s_\iota\}_{\iota \in [k-1]} \right] \geq \frac{1}{2p(\lambda)}. \quad (21)$$

Through an averaging argument, we can fix index $j \in \mathcal{J}$ which gives us the same event with an advantage of $1/(2kp(\lambda))$. We will now switch to a hybrid where we provide $\mathcal{A}$ with simulated proofs at indices $\mathcal{I}$.

Claim 5.9. There exists a polynomial $q(\cdot)$ such that

$$\Pr_{\substack{\{\pi_\iota\}_{\iota \in [k-1]} \leftarrow \mathsf{Sim}(x_1, \ldots, x_{k-1}) \\ \{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]} \leftarrow \mathcal{A}^{\mathsf{Sim}}(\{\pi_\iota\}_{\iota \in [k-1]})}} \left[\bigwedge_{\iota \in \mathcal{J}} \mathsf{V}^{\mathsf{Sim}}(x, \widetilde{\pi_\iota}) = 1 \bigwedge \widetilde{s_j} \notin \{s_\iota\}_{\iota \in [k-1]} \right] \geq \frac{1}{q(\lambda)}. \quad (22)$$

We will later see a proof of Claim 5.9. For now, assuming that this claim holds, we can define an adversary from which Ext can extract a valid witness for x.

Claim 5.10. There exists a polynomial $q'(\cdot)$ such that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(\{x_\iota, \omega_\iota\}_{\iota \in [k-1] \setminus \mathcal{I}}, x, \mathcal{I}, \mathcal{J})} [(x, w) \in \mathcal{R}] \geq \frac{1}{q'(\lambda)}. \tag{23}$$

We will soon see a proof for Claim 5.10. Meanwhile, if this claim is true, then we will have a direct contradiction with Equation (19). Thus, all that remains to be proven are the two claims: Claim 5.9 and Claim 5.10. We start by proving the former claim.

Proof of Claim 5.9. We first need to argue that our strategy is well-defined, that we will be able to independently program these k points. Then we can argue the indistinguishability of switching one-by-one to simulated proofs. We will argue that our simulator will run in expected polynomial time. By Lemma 5.7, the k points which our simulator will program will be distinct with overwhelming probability. Furthermore, since we assumed that our quantum random oracle can be programmed at multiple distinct points Definition 3.10, our simulator is well-defined.

We now argue indistinguishability of the simulated proofs from the honestly generated proofs via a hybrid argument. Suppose for sake of contradiction that the probability difference between Equation (21) and Equation (22) was $1/p'(\lambda)$ for some polynomial $p'(\cdot)$. We construct a series of consecutive hybrids for each $i \in [k-1]$ where we switch the i^{th} proof from prover generated to simulated. By this hybrid argument, there must be some position $\ell \in [k-1]$ where switching the ℓ^{th} proof has a probability difference of at least $1/(kp'(\lambda))$. We now formalize a reduction which can distinguish between these two settings:

Reduction: to zero-knowledge of our protocol given oracle access to $\mathcal{A}$ and some $\mathcal{O}$.
Hardwired with: $(x_1, w_1), \ldots, (x_{\ell-1}, w_{\ell-1})$, x, $\mathcal{I}$, $\mathcal{J}$, j, ℓ.
(1) Receive (real or simulated) π from the challenger and mediated query access to a (real or simulated, respectively) random oracle $\mathcal{O}$.
(2) Define $\pi_\ell = \pi$. Compute $\pi_\iota \leftarrow \mathsf{P}^{\mathcal{O}}(x_\iota, w_\iota)$ for $\iota \in [\ell - 1]$. Compute $\pi_\iota \leftarrow \mathsf{Sim}(x_\iota)$ for $\iota \in \{\ell + 1, \ldots, k-1\}$ where we store all points Sim would program into a list $\mathcal{P}$.
(3) Send $\{\pi_\iota\}_{\iota \in [k-1]}$ to $\mathcal{A}$.
(4) For every query from $\mathcal{A}$, if the query is in $\mathcal{P}$, then reply with the answer from $\mathcal{P}$. Else, forward the query to $\mathcal{O}$ and send the answer back to $\mathcal{A}$. Let $\widehat{\mathcal{O}}$ denote this modified random oracle.
(5) Receive $\{\widetilde{x_\iota}, \widetilde{\pi_\iota}\}_{\iota \in [k]}$ from $\mathcal{A}$.
(6) If $\mathsf{V}^{\widehat{\mathcal{O}}}(x, \widetilde{\pi_\iota}) = 1$ for all $\iota \in \mathcal{J}$ and $\widetilde{s_j} \notin \{s_\iota\}_{\iota \in [k-1]}$, then output 1. Else, output 0.

We first argue that the view that the reduction provides to $\mathcal{A}$ matches one of the games: where all proofs up to the ℓ^{th} are simulated or where all proofs up to and including the ℓ^{th} are simulated. By Lemma 5.7, the point computed or programmed by the challenger will be distinct from the points which the reduction programs. As such, the reduction is allowed to modify [5] the oracle which $\mathcal{A}$ interfaces with (see step **(4)**). In summary, $\mathcal{A}$ will be provided access to an oracle that is consistent with all of the proofs it receives.

Given that $\mathcal{A}$ has a view which directly matches its expected view in either game, then the reduction's advantage is the same as $\mathcal{A}$'s advantage which is at least $1/(kp'(\lambda))$. This is a contradiction with the zero-knowledge property of our protocol. Thus, our original claim must be true. □

[5] In more detail, the reduction can construct a unitary which runs the classical code in step **(4)**. This can then be applied in superposition to a query sent by $\mathcal{A}$.

Now, we continue on to proving the latter claim.

Proof of Claim 5.10. Given that Claim 5.9 holds, this implies that

$$\Pr_{\substack{\{\pi_\iota\}_{\iota\in[k-1]}\leftarrow\mathsf{Sim}(x_1,\ldots,x_{k-1})\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}^{\mathsf{Sim}}(\{\pi_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{V}^{\mathsf{Sim}}(x,\widetilde{\pi_\iota})=1\right]\geq\frac{1}{q(\lambda)},\text{ and} \tag{24}$$

$$\Pr_{\substack{\{\pi_\iota\}_{\iota\in[k-1]}\leftarrow\mathsf{Sim}(x_1,\ldots,x_{k-1})\\ \{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}^{\mathsf{Sim}}(\{\pi_\iota\}_{\iota\in[k-1]})}}\left[\widetilde{s_j}\notin\{s_\iota\}_{\iota\in[k-1]}\right]\geq\frac{1}{q(\lambda)}. \tag{25}$$

We define a reduction to the proof of knowledge property of our protocol, which we subsequently will refer to as $\widetilde{\mathsf{P}}$ (borrowing notation from the definition of PoK in Definition 3.11), as follows:

Reduction: to proof of knowledge of our protocol given oracle access to $\mathcal{A}$ and some $\mathcal{O}$.
Hardwired with: $\mathcal{I}$, $\mathcal{J}$, $x_1,\ldots,x_{k-1}$, x, j.
(1) Sample ℓ from $\mathcal{J}$ uniformly at random.
(2) Receive query access to a (real or extractable) random oracle $\mathcal{O}$.
(3) Compute $\pi_\iota\leftarrow\mathsf{Sim}(x_\iota)$ for $\iota\in[k-1]$ where we store all points Sim would program into a list $\mathcal{P}$.
(4) Send $\{\pi_\iota\}_{\iota\in[k-1]}$ to $\mathcal{A}$.
(5) For every query from $\mathcal{A}$, if the query is in $\mathcal{P}$, then reply with the answer from $\mathcal{P}$. Else, forward the query to $\mathcal{O}$ and send the answer back to $\mathcal{A}$. Let $\widehat{\mathcal{O}}$ denote this modified random oracle.
(6) Receive $\{\widetilde{x_\iota},\widetilde{\pi_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$.
(7) Output $\widetilde{\pi_\ell}$.

First we must argue that $\mathcal{A}$'s view remains identical to the game which appears in both Equation (24) and Equation (25). The oracle which $\mathcal{A}$ interfaces with (see step **(4)**) will be consistent with all of the proofs it receives.

By Equation (24), we have that our adversary $\widetilde{\mathsf{P}}$ will be able to produce an accepting proof with noticeable probability. That is that,

$$\Pr_{\widetilde{\pi_\ell}\leftarrow\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\mathsf{V}^{\widetilde{\mathsf{P}}}(x,\widetilde{\pi_\ell})=1\right]\geq\frac{1}{q(\lambda)}.$$

Say that the output proof (denoted by $\widetilde{\pi_\ell}$) has a serial number which differs from the serial numbers in the proofs provided to $\mathcal{A}$ (denoted by $\{\pi_\iota\}_{\iota\in[k-1]}$). In this case, the verification algorithm V should succeed with the same probability even given oracle access to the unmodified random oracle (denoted by $\mathcal{O}$). We have that

$$\Pr_{\widetilde{\pi_\ell}\leftarrow\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\mathsf{V}^{\mathcal{O}}(x,\widetilde{\pi_\ell})=1\;\middle|\;\widetilde{s_\ell}\notin\{s_\iota\}_{\iota\in[k-1]}\right]=\Pr_{\widetilde{\pi_\ell}\leftarrow\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\mathsf{V}^{\widetilde{\mathsf{P}}}(x,\widetilde{\pi_\ell})=1\right]\geq\frac{1}{q(\lambda)}.$$

Through Equation (25), we reach the conclusion that

$$\begin{aligned}&\Pr_{\widetilde{\pi_\ell}\leftarrow\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\mathsf{V}^{\mathcal{O}}(x,\widetilde{\pi_\ell})=1\right]\\ &\geq\Pr_{\widetilde{\pi_\ell}\leftarrow\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\mathsf{V}^{\mathcal{O}}(x,\widetilde{\pi_\ell})=1\;\middle|\;\widetilde{s_\ell}\notin\{s_\iota\}_{\iota\in[k-1]}\right]\cdot\Pr_{\widetilde{\pi_\ell}\leftarrow\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\widetilde{s_j}\notin\{s_\iota\}_{\iota\in[k-1]}\right]\cdot\Pr\left[\ell=j\right]\\ &\geq\frac{1}{kq(\lambda)^2}.\end{aligned}$$

By the definition of a proof of knowledge (Definition 3.11) which have some parameters polynomial $p^*(\cdot)$ and negligible functions $\mathsf{negl}_0(\cdot)$ and $\mathsf{negl}_1(\cdot)$, we have that there exists some polynomial $q'(\cdot)$ such that

$$\Pr\left[(x, \mathsf{Ext}^{\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(x)}(x)) \in \mathcal{R}_\lambda\right] \geq \frac{1}{p^*(\lambda)} \cdot \left(\Pr_{\widetilde{\pi}_\ell \leftarrow \widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(1^\lambda)}\left[\mathsf{V}^{\mathcal{O}}(x, \widetilde{\pi}_\ell) = 1\right] - \mathsf{negl}_0(\lambda)\right)^c - \mathsf{negl}_1(\lambda) \geq \frac{1}{q'(\lambda)}.$$

We now compare the reduction $\widetilde{\mathsf{P}}$ to the extractor $\mathcal{E}$. Only steps **(1)** and **(6)** differ. If we consider the extractor Ext with oracle access to $\widetilde{\mathsf{P}}$, then this whole extraction strategy is identical to that in the definition of $\mathcal{E}$ (when hardwired with the same ℓ). Hence, by this modularity, we will have that

$$\Pr_{w \leftarrow \mathcal{E}^{\mathcal{A}}(x_1,\ldots,x_{k-1},x,\mathcal{I},\mathcal{J})}\left[(x,w) \in \mathcal{R}\right] = \Pr\left[(x, \mathsf{Ext}^{\widetilde{\mathsf{P}}^{\mathcal{A},|\mathcal{O}\rangle}(x)}(x)) \in \mathcal{R}_\lambda\right] \geq \frac{1}{q'(\lambda)}$$

which completes the proof of our claim. □

By completing the proofs of our claims, we have concluding the proof of our theorem statement. □

Corollary 5.11. Assuming the injective one-way functions exist, and post-quantum iO exists, there exists a non-interactive knowledge sound, computationally zero-knowledge, and $(k-1)$-to-k-unclonable with extraction protocol for NP in the quantum random oracle model (Definition 5.4).

Proof. This follows from Theorem 3.14 and Theorem 5.8. □

We have thus shown that Figure 5 is an unclonable NIZK PoK in the ROM model as defined according to our unclonability definition, Definition 5.4.

6 Unclonable Signatures of Knowledge

6.1 Definition

Definition 6.1 (Unclonable Extractable SimExt-secure Signatures of Knowledge)**.** Let NP relation $\mathcal{R}$ with corresponding language $\mathcal{L}$ be given such that they can be indexed by a security parameter $\lambda \in \mathbb{N}$. Let a message space $\mathcal{M}$ be given such that it can be indexed by a security parameter $\lambda \in \mathbb{N}$.

$(\mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify})$ is an unclonable signature of knowledge of a witness with respect to $\mathcal{L}$ and $\mathcal{M}$ if it has the following properties:

- $(\mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify})$ is a quantum Sim-Ext signature of knowledge (Definition 3.16).
- $(k-1)$**-to-**k**-Unclonable with Extraction**: There exists an oracle-aided polynomial-size quantum circuit $\mathcal{E}$ such that for every polynomial-size quantum circuit $\mathcal{A}$, for every tuple of $k-1$ instance-witness pairs $(x_1, \omega_1), \ldots, (x_{k-1}, \omega_{k-1}) \in \mathcal{R}$, every $\{m_\iota \in \mathcal{M}_\lambda\}_{\iota \in [k-1]}$, for every (x, m) where we define
  - $\mathcal{I} \subseteq [k-1]$ such that $|\mathcal{I}| \geq 1$, $(x_i, m_i) = (x, m)$ for all $i \in \mathcal{I}$, and $(x_\iota, m_\iota) \neq (x, m)$ for all $\iota \notin \mathcal{I}$, and
  - $\mathcal{J} \subseteq [k]$ such that $|\mathcal{J}| \geq \max\{2, |\mathcal{I}|\}$, $(x_j, m_j) = (x, m)$ for all $j \in \mathcal{J}$, and $(x_\iota, m_\iota) \neq (x, m)$ for $\iota \notin \mathcal{J}$,

such that there is a polynomial $p(\cdot)$ where

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\sigma_\iota\leftarrow\mathsf{Sign}(\mathsf{crs},x_\iota,\omega_\iota,m_\iota)\\ \{\widetilde{x}_\iota,\widetilde{m}_\iota,\widetilde{\sigma}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{Verify}(\mathsf{crs},x,m,\widetilde{\sigma}_\iota)=1\right]\geq\frac{1}{p(\lambda)},$$

then there is also a polynomial $q(\cdot)$ such that

$$\Pr_{w\leftarrow\mathcal{E}^{\mathcal{A}}(\{x_\iota,m_\iota\}_{\iota\in[k-1]},x,\mathcal{I},\mathcal{J})}[(x,w)\in\mathcal{R}]\geq\frac{1}{q(\lambda)}.$$

6.2 Construction

Unclonable Signature of Knowledge with CRS

Let $(\mathsf{Setup},\mathsf{P},\mathsf{V})$ be non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge, unclonable-extractable protocol for NP. Let $\mathcal{R}$ be the relation with respect to $\mathcal{L}\in\mathsf{NP}$.

SETUP(1^λ): $(\mathsf{crs},\mathsf{td})\leftarrow\Pi.\mathsf{Setup}(1^\lambda)$.

SIGN(crs,x,w,m):

- Let $x_\Pi=(x,m)$ be an instance and $w_\Pi=w$ be its corresponding witness for the following language $\mathcal{L}_\Pi$:
$$\{(x,m)\ :\ \exists w\ \ :\ \ (x,w)\in\mathcal{R}\}.$$
- Compute $\pi\leftarrow\Pi.\mathsf{P}(\mathsf{crs},x_\Pi,w_\Pi)$.
- Output $\sigma=\pi$.

VERIFY$(\mathsf{crs},x,m,\sigma)$: Output $\Pi.\mathsf{V}(\mathsf{crs},(x,m),\pi)$.

Figure 6: Unclonable Signature of Knowledge in CRS model

Theorem 6.2. *Let $\Pi=(\mathsf{Setup},\mathsf{P},\mathsf{V})$ be a non-interactive simulation-extractable, adaptive multi-theorem computational zero-knowledge, unclonable-extractable protocol for* NP *(Definition 4.7).*

$(\mathsf{Setup},\mathsf{Sign},\mathsf{Verify})$ in Figure 6 is an unclonable-extractable SimExt-secure signature of knowledge (Definition 6.1).

Proof of Theorem 6.2. Correctness follows naturally. It remains to prove simulateability, extractability, and unclonable extractability.

Simulateable. Let $\Pi.\mathsf{Sim}=(\Pi.\mathsf{Sim}_0,\Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. We define Sim_0 with oracle access to $\Pi.\mathsf{Sim}_0$ as follows:

Input: 1^λ.

(1) Send 1^λ to $\Pi.\mathsf{Sim}_0$. Receive $(\mathsf{crs},\mathsf{td})$ from $\Pi.\mathsf{Sim}_0$.

(2) Output crs and td.

We define Sim_1 with oracle access to $\Pi.\mathsf{Sim}_1$ as follows:

Input: crs, td, x, m.

(1) Define $x_\Pi = (x, m)$. Send $(\mathsf{crs}, \mathsf{td}, x_\Pi)$ to $\Pi.\mathsf{Sim}_1$. Receive π from $\Pi.\mathsf{Sim}_1$.

(2) Output $\sigma = \pi$.

Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\left| \Pr_{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda)} [\mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) = 1] - \Pr_{\mathsf{crs} \leftarrow \mathsf{Setup}(1^\lambda)} [\mathcal{A}^{\mathsf{Sign}(\mathsf{crs},\cdot,\cdot,\cdot)}(\mathsf{crs}) = 1] \right| \geq \frac{1}{p(\lambda)}.$$

We define a reduction to the multi-theorem zero-knowledge property of Π as follows:

Reduction: to zero-knowledge of Π given oracle access to $\mathcal{A}$.

(1) Receive (real or simulated) crs from the challenger.

(2) Send crs to $\mathcal{A}$.

(3) On query (x, w, m) from $\mathcal{A}$: send $x_\Pi = (x, m)$ to the challenger, receives (real or simulated) π from the challenger, send $\sigma = \pi$ to $\mathcal{A}$.

(4) Output the result of $\mathcal{A}$.

The view of $\mathcal{A}$ matches that of our protocol in Figure 6 or Sim_0 and Sim_1. As such, this reduction should have the same advantage at breaking the adaptive multi-theorem computational zero-knowledge property of Π. We reach a contradiction, hence our protocol must be simulateable.

Extractable. Let $\Pi.\mathsf{Sim} = (\Pi.\mathsf{Sim}_0, \Pi.\mathsf{Sim}_1)$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. Let $\Pi.\mathsf{Ext}$ be the simulation extractable extractor of Π defined relative to $\Pi.\mathsf{Sim}$. Let $\mathsf{Sim} = (\mathsf{Sim}_0, \mathsf{Sim}_1)$ be the simulator given by the simulation property which uses $\Pi.\mathsf{Sim}$. We define Ext with oracle access to $\Pi.\mathsf{Ext}$ as follows:

Input: crs, td, x, m, $\sigma = \pi$.

(1) Define $x_\Pi = (x, m)$.

(2) Send $(\mathsf{crs}, \mathsf{td}, x_\Pi, \pi)$ to $\Pi.\mathsf{Ext}$. Receive $w_\Pi = w$ from $\Pi.\mathsf{Ext}$.

(3) Output w.

Let a polynomial $p(\cdot)$ and an oracle-aided polynomial-size quantum circuit $\mathcal{A}$ be given such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x,m,\sigma) \leftarrow \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) \\ w \leftarrow \mathsf{Ext}(\mathsf{crs},\mathsf{td},x,m,\sigma)}} [\mathsf{Verify}(\mathsf{crs}, x, m, \sigma) = 1 \wedge (x, m) \notin Q \wedge (x, w) \notin \mathcal{R}_\lambda] \geq \frac{1}{p(\lambda)}$$

where Q is the list of queries from $\mathcal{A}$ to Sim_1. If Verify accepts the output of $\mathcal{A}$, then $\Pi.\mathsf{V}$ must accept $(\mathsf{crs}, x_\Pi, \pi)$. If $(x, m) \notin Q$, then since x_Π contains x, m, x_Π must not be in the queries asked to $\Pi.\mathsf{Sim}_1$. Since $(x, w) \notin \mathcal{R}$, then $x_\Pi \notin \mathcal{L}_\Pi$ by the definition of $\mathcal{L}_\Pi$. As such, it must necessarily be the case that $(x_\Pi, w_\Pi) \notin \mathcal{R}_\Pi$. Hence, we have that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td}) \leftarrow \mathsf{Sim}_0(1^\lambda) \\ (x,m,\sigma) \leftarrow \mathcal{A}^{\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},\cdot,\cdot)}(\mathsf{crs}) \\ w \leftarrow \mathsf{Ext}(\mathsf{crs},\mathsf{td},x,m,\sigma)}} [\Pi.\mathsf{V}(\mathsf{crs}, x_\Pi, \pi) = 1 \wedge x_\Pi \notin Q_\Pi \wedge (x_\Pi, w_\Pi) \notin \mathcal{R}_\Pi] \geq \frac{1}{p(\lambda)}$$

where Q_Π is the list of queries, originating from $\mathcal{A}$, that Sim_1 makes to $\Pi.\mathsf{Sim}_1$. We define a reduction to the simulation extraction property of Π as follows:

Reduction: to simulation extraction of Π given oracle access to $\mathcal{A}$.
(1) Receive crs from the challenger.
(2) Send crs to $\mathcal{A}$.
(3) On query (x, w, m) from $\mathcal{A}$: send $x_\Pi = (x, m)$ to the challenger, receives π from the challenger, send $\sigma = \pi$ to $\mathcal{A}$.
(4) Receive $(x, m, \sigma = \pi)$ from $\mathcal{A}$. Define $x_\Pi = (x, m)$.
(5) Output (x_Π, π).

The view of $\mathcal{A}$ matches that of Sim_0 and Sim_1. As such, this reduction should have the same advantage at breaking the extraction property of Π. We reach a contradiction, hence our protocol must be extractable.

Unclonable Extractability. Let $\Pi.\mathsf{Sim}$ be the adaptive multi-theorem computationally zero-knowledge simulator of Π. Let $\Pi.\mathsf{Ext}$ be the simulation extractable extractor of Π defined relative to $\Pi.\mathsf{Sim}$. Let $\Pi.\mathcal{E}$ be the unclonable extractor of Π. We define $\mathcal{E}$ with oracle-access to $\Pi.\mathcal{E}$, $\Pi.\mathsf{Sim}$, $\Pi.\mathsf{Ext}$, and some $\mathcal{A}$ as follows:

Input: $\{x_\iota, m_\iota\}_{\iota\in[k-1]}$, x, m, $\mathcal{I}$, $\mathcal{J}$
(1) Define $x_{\Pi,\iota} = (x_\iota, m_\iota)$ for $\iota \in [k-1]$.
Sample b uniformly at random from $\{0, 1\}$.
If $b = 0$, then execute the following code:
(2) Sample $\ell \in [k-1]$, $i, j \in [k]$ uniformly at random.
(3) Send $(\{x_{\Pi,\iota}\}_{\iota\in[k-1]}, x_{\Pi,\ell}, \{\ell\}, \{i, j\})$ to $\Pi.\mathcal{E}$. Receive $(\mathsf{crs}, \{\pi_\iota\}_{\iota\in[k-1]})$ from $\Pi.\mathcal{E}$.
(4) Send $(\mathsf{crs}, \{\sigma_\iota = \pi_{\Pi,\iota}\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x_\iota}, \widetilde{m_\iota}, \widetilde{\sigma_\iota} = \widetilde{\pi_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$. Define $\widetilde{x_{\Pi,\iota}} = (\widetilde{x_\iota}, \widetilde{m_\iota})$ for $\iota \in [k]$.
(5) Send $\{x_{\Pi,\ell}, \widetilde{x_{\Pi,\iota}}, \widetilde{\pi_\iota}\}_{\iota\in[k]}$ to $\Pi.\mathcal{E}$. Receive $w_\Pi = w$ from $\Pi.\mathcal{E}$.
(6) Output w.
Otherwise, if $b = 1$, then execute the following code:
(2) Receive $(\mathsf{crs}, \mathsf{td})$ from $\Pi.\mathsf{Sim}_0$.
(3) For $\iota \in [k-1]$: send $(\mathsf{crs}, \mathsf{td}, x_{\Pi,\iota} = (x_\iota, m_\iota))$ to $\Pi.\mathsf{Sim}_1$, receive π_ι from $\Pi.\mathsf{Sim}_1$, and define $\sigma_\iota = \pi_\iota$.
(4) Sample $j \in [k]$ uniformly at random.
(5) Send $(\mathsf{crs}, \mathsf{td}, \widetilde{x_{\Pi,j}}, \widetilde{\pi_j})$ to $\Pi.\mathsf{Ext}$. Receive w_Π from $\Pi.\mathsf{Ext}$.
(6) Output w_Π.

Let $\mathcal{A}$, $(x_1, w_1), \ldots, (x_{k-1}, w_{k-1}) \in \mathcal{R}$, $\{m_\iota \in \mathcal{M}_\lambda\}_{\iota\in[k-1]}$, x, m, $\mathcal{I} \subseteq [k-1]$, $\mathcal{J} \subseteq [k]$, polynomial $p(\cdot)$, and negligible function $\mathsf{negl}(\cdot)$ be given such that the verification algorithm Verify accepts the signatures output by $\mathcal{A}$ which are indexed by $\mathcal{I}$ which $\mathcal{A}$ outputs, and the extractor $\mathcal{E}$ is unable to extract a valid witness. Formally, that is that both

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\sigma_\iota\leftarrow\mathsf{Sign}(\mathsf{crs},x_\iota,\omega_\iota,m_\iota)\\ \{\widetilde{x_\iota},\widetilde{m_\iota},\widetilde{\sigma_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\mathsf{Verify}(\mathsf{crs},x,m,\widetilde{\sigma_\iota})=1\right]\geq\frac{1}{p(\lambda)},\text{ and} \tag{26}$$

$$\Pr_{w\leftarrow\mathcal{E}^{\mathcal{A}}(\{x_\iota,m_\iota\}_{\iota\in[k-1]},x,m,\mathcal{I},\mathcal{J})}[(x,w)\in\mathcal{R}_\lambda]\leq\mathsf{negl}(\lambda). \tag{27}$$

If Verify accepts the signatures $\{\widetilde{\sigma_\iota}\}_{\iota\in\mathcal{J}}$ with respect to (x, m) from $\mathcal{A}$, then $\Pi.\mathsf{V}$ must accept the

proof $\{\widetilde{\pi}_\iota\}_{\iota\in\mathcal{J}}$ with respect to $\{\widetilde{x_{\Pi,\iota}} = (x, m)\}_{\iota\in\mathcal{J}}$. This means that Equation (26) implies that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\sigma_\iota\leftarrow\mathsf{Sign}(\mathsf{crs},x_\iota,\omega_\iota,m_\iota)\\ \{\widetilde{x}_\iota,\widetilde{m}_\iota,\widetilde{\sigma}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\mathcal{J}}\Pi.\mathsf{V}(\mathsf{crs},\widetilde{x_{\Pi,\iota}} = (x,m),\widetilde{\pi}_\iota) = 1\right] \geq \frac{1}{p(\lambda)}. \tag{28}$$

There are two scenarios which could arise: either $\mathcal{A}$ sends two accepting proofs at some indices i and j where the instance for the unclonable NIZK matches one that it received at index ℓ ($\widetilde{x_{\Pi,i}} = \widetilde{x_{\Pi,j}} = x_{\Pi,\ell}$), or there exists an accepting proof at some index j which has an instance for which it has not seen an unclonable NIZK proof for ($\widetilde{x_{\Pi,j}} \notin \{x_{\Pi,\iota}\}_{\iota\in\mathcal{I}}$). We consider these two scenarios separately.

Scenario One

Say that $\mathcal{A}$ sends two accepting proofs at some indices i and j where the instance for the unclonable NIZK matches one that it received at index ℓ ($\widetilde{x_{\Pi,i}} = \widetilde{x_{\Pi,j}} = x_{\Pi,\ell}$). We could fix some indices $i, j \in [k]$ and $\ell \in [k-1]$ by a hybrid argument to get the following from Equation (28) and a union bound,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\sigma_\iota\leftarrow\mathsf{Sign}(\mathsf{crs},x_\iota,\omega_\iota,m_\iota)\\ \{\widetilde{x}_\iota,\widetilde{m}_\iota,\widetilde{\sigma}_\iota\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}\left[\bigwedge_{\iota\in\{i,j\}}\Pi.\mathsf{V}(x_{\Pi,\ell},\widetilde{\pi}_\iota) = 1\right] \geq \frac{1}{2k^3\cdot p(\lambda)}. \tag{29}$$

With probability $1/2$, $\mathcal{E}$ makes use of $\Pi.\mathcal{E}$ algorithm. Given that $\mathcal{E}$ is using $\Pi.\mathcal{E}$, if $\mathcal{E}$'s extracted witness has the property that $(x, w) \notin \mathcal{R}$, then no matter which x_Π (containing x) is given to $\Pi.\mathcal{E}$, they cannot have extracted any valid witness w_Π (namely because $\mathcal{E}$ outputs w contained in w_Π). With probability $1/k^3$, $\mathcal{E}$ will sample i, j and ℓ correctly. As such, we have that

$$\Pr_{w\leftarrow\mathcal{E}^{\mathcal{A}}(\{x_\iota,m_\iota\}_{\iota\in[k-1]},x,m,\mathcal{I},\mathcal{J})}[(x_{\Pi,\ell},w_\Pi)\in\mathcal{R}_\lambda] \leq \mathsf{negl}(\lambda).$$

We will show that we will directly contradict the unclonability of the NIZK protocol Π. We define a reduction to the unclonability of Π as follows:

Reduction: to unclonability of Π given oracle access to $\mathcal{A}$.

Hardwired with: $\{x_\iota, m_\iota\}_{\iota\in[k-1]}$, $\mathcal{I}$, $\mathcal{J}$, i, j, ℓ

(1) Define $x_{\Pi,\iota} = (x_\iota, m_\iota)$ for $\iota \in [k-1]$.

(2) Send $(\{x_{\Pi,\iota}\}_{\iota\in[k-1]}, x_{\Pi,\ell}, \{\ell\}, \{i, j\})$ to the challenger.

We note that the following code is re-windable by the challenger, as necessary:

(3) Receive $(\mathsf{crs}, \{\pi_\iota\}_{\iota\in[k-1]})$ from the challenger. Define and $\sigma_\iota = \pi_{\Pi,\iota}$ for $\iota \in [k-1]$.

(3) Send $(\mathsf{crs}, \{\sigma_\iota\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x}_\iota, \widetilde{m}_\iota, \widetilde{\sigma}_\iota = \widetilde{\pi}_\iota\}_{\iota\in[k]}$ from $\mathcal{A}$.

(4) Output $\{\widetilde{x_{\Pi,\iota}}, \widetilde{\pi}_\iota\}_{\iota\in[k]}$ where $\widetilde{x_{\Pi,\iota}} = (\widetilde{x}_\iota, \widetilde{m}_\iota)$ for $\iota \in [k]$.

The view of $\mathcal{A}$ matches either the real or the simulated game. Additionally, the challenger may run the honest extractor. As such, this reduction should have the same advantage at breaking the unclonability property of Π. This reaches a contradiction.

Scenario Two

Say that $\mathcal{A}$ sends an accepting proof at some index j which has an instance for which it has not seen an unclonable NIZK proof for ($\widetilde{x_{\Pi,j}} \notin \{x_{\Pi,\iota}\}_{\iota\in\mathcal{I}}$). We can fix this index $j \in \mathcal{J}$, by a union

bound and hybrid argument, get that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\mathsf{Setup}(1^\lambda)\\ \forall\iota\in[k-1],\,\sigma_\iota\leftarrow\mathsf{Sign}(\mathsf{crs},x_\iota,\omega_\iota,m_\iota)\\ \{\widetilde{x_\iota},\widetilde{m_\iota},\widetilde{\sigma_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}[\Pi.\mathsf{V}(\mathsf{crs},\widetilde{x_{\Pi,j}},\widetilde{\pi_\iota})=1\ \wedge\ \widetilde{x_{\Pi,j}}\notin\{x_{\Pi,\iota}\}_{\iota\in\mathcal{I}}]\geq\frac{1}{2k\cdot p(\lambda)}. \tag{30}$$

We now switch to a hybrid where the proofs in the signatures are simulated.

Claim 6.3. There exists a polynomial $p'(\cdot)$ such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\Pi.\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\Pi.\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_{\Pi,\iota}=(x_\iota,m_\iota))\\ \{\widetilde{x_\iota},\widetilde{m_\iota},\widetilde{\sigma_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}[\Pi.\mathsf{V}(\mathsf{crs},\widetilde{x_{\Pi,j}},\widetilde{\pi_\iota})=1\ \wedge\ \widetilde{x_{\Pi,j}}\notin\{x_{\Pi,\iota}\}_{\iota\in\mathcal{I}}]\geq\frac{1}{p'(\lambda)}.$$

We will soon see a proof of Claim 6.3. Meanwhile, if this claim is true, then we can make the following reduction to the simulation extractability of Π. If $\widetilde{x_{\Pi,j}}\notin\{x_{\Pi,\iota}\}_{\iota\in\mathcal{I}}$, then $\widetilde{x_{\Pi,j}}$ must not be in the queries $\mathcal{E}$ asks to $\Pi.\mathsf{Sim}_1$. Hence,

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\Pi.\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\Pi.\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_{\Pi,\iota}=(x_\iota,m_\iota))\\ \{\widetilde{x_\iota},\widetilde{m_\iota},\widetilde{\sigma_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})}}[\Pi.\mathsf{V}(\mathsf{crs},\widetilde{x_{\Pi,j}},\widetilde{\pi_\iota})=1\ \wedge\ \widetilde{x_{\Pi,j}}\notin Q_\Pi]\geq\frac{1}{p'(\lambda)}$$

where Q_Π is the list of queries that are made to $\Pi.\mathsf{Sim}_1$.

With probability $1/2$, $\mathcal{E}$ makes use of $\Pi.\mathsf{Sim}$ algorithm. Given that $\mathcal{E}$ is using $\Pi.\mathsf{Sim}$, we can make the following statement from Equation (27): since $(x,w)\notin\mathcal{R}$, then if Ext extracts from $\widetilde{\pi_{\Pi,j}}$ (which happens with $1/k$ probability), $(\widetilde{x_{\Pi,j}},w_\Pi)\notin\mathcal{R}_\Pi$. Hence, we have that there exists a polynomial $q(\cdot)$ such that

$$\Pr_{\substack{(\mathsf{crs},\mathsf{td})\leftarrow\Pi.\mathsf{Sim}_0(1^\lambda)\\ \forall\iota\in[k-1],\,\pi_\iota\leftarrow\Pi.\mathsf{Sim}_1(\mathsf{crs},\mathsf{td},x_{\Pi,\iota})\\ \{\widetilde{x_\iota},\widetilde{m_\iota},\widetilde{\sigma_\iota}\}_{\iota\in[k]}\leftarrow\mathcal{A}(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})\\ w_\Pi\leftarrow\Pi.\mathsf{Ext}(\mathsf{crs},\mathsf{td},\widetilde{x_{\Pi,j}},\widetilde{\pi_j})}}[\Pi.\mathsf{V}(\mathsf{crs},\widetilde{x_{\Pi,j}},\widetilde{\pi_j})=1\wedge\widetilde{x_{\Pi,j}}\notin Q_\Pi\wedge(\widetilde{x_{\Pi,j}},w_\Pi)\notin\mathcal{R}_\Pi]\geq\frac{1}{q(\lambda)}. \tag{31}$$

We define a reduction to the simulation extraction property of Π as follows:

Reduction: to simulation extraction of Π given oracle access to $\mathcal{A}$.

Hardwired with $x_1,m_1,\ldots,x_{k-1},m_{k-1},j$

(1) Receive crs from the challenger.

(2) For $\iota\in[k-1]$: send $x_{\Pi,\iota}=(x_\iota,m_\iota)$ to the challenger, receives π_ι from the challenger, and define $\sigma_\iota=\pi_\iota$.

(3) Send $(\mathsf{crs},\{\sigma_\iota\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x_\iota},\widetilde{m_\iota},\widetilde{\sigma_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$.

(4) Output $(\widetilde{x_{\Pi,j}},\pi_j)$.

The view of $\mathcal{A}$ matches that in Equation (31). As such, this reduction should have the same advantage at breaking the extraction property of Π. We reach a contradiction. As such, it only remains to prove Claim 6.3.

Proof of Claim 6.3. Assume for the sake of contradiction that the claim is false. This means there is an inverse polynomial gap between the game in Claim 6.3 and Equation (30). We define a reduction to the multi-theorem zero-knowledge property of Π as follows:

Reduction: to zero-knowledge of Π given oracle access to $\mathcal{A}$.
(1) Receive (real or simulated) crs from the challenger.
(2) For $\iota \in [k-1]$: send $x_{\Pi,\iota} = (x_\iota, m_\iota)$ to the challenger, receives (real or simulated) π_ι from the challenger, and define $\sigma_\iota = \pi_\iota$.
(3) Send $(\mathsf{crs}, \{\sigma_\iota\}_{\iota\in[k-1]})$ to $\mathcal{A}$. Receive $\{\widetilde{x_\iota}, \widetilde{m_\iota}, \widetilde{\sigma_\iota}\}_{\iota\in[k]}$ from $\mathcal{A}$.
(4) Output $\Pi.\mathsf{V}(\mathsf{crs}, \widetilde{x_{\Pi,j}}, \widetilde{\pi_\iota}) = 1 \wedge \widetilde{x_{\Pi,j}} \notin \{x_{\Pi,\iota}\}_{\iota\in\mathcal{I}}$.

The view of $\mathcal{A}$ matches that of our protocol in Figure 6 or Sim_0 and Sim_1. As such, this reduction should have the same advantage at breaking the adaptive multi-theorem computational zero-knowledge property of Π. We reach a contradiction. □

We reach a contradiction in both scenarios, hence our protocol must be unclonable. □

Corollary 6.4. Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, post-quantum iO exists, there exists an unclonable SimExt-secure signature of knowledge (Definition 6.1).

Proof. This follows from Corollary 4.13 and Theorem 6.2. □

6.3 Revocable Anonymous Credentials

In this section, we will see how to use unclonable signatures of knowledge to construct an anonymous credentials scheme which has a natural revocation property.

Definition 6.5 (Revocable Anonymous Credentials)**.** $(\mathsf{IssuerKeyGen}, \mathsf{Issue}, \mathsf{VerifyCred}, \mathsf{Revoke}, \mathsf{Prove}, \mathsf{VerRevoke})$ is a revocable anonymous credentials scheme with respect to some set of accesses $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$ if it has the following syntax and properties.

Syntax. The input 1^λ is left out when it is clear from context.

- $(\mathsf{nym}, \mathsf{sk}) \leftarrow \mathsf{IssuerKeyGen}(1^\lambda)$: The probabilistic polynomial-time algorithm IssuerKeyGen is run by the issuer of the credentials. It takes input 1^λ; and outputs a pseudonym nym with a secret key sk.
- $\mathsf{cred} \leftarrow \mathsf{Issue}(1^\lambda, \mathsf{nym}, \mathsf{sk}, \mathsf{access})$: The polynomial-time quantum algorithm Issue is run by the issuer of the credentials. It takes input the issuer's keys nym and sk as well as the requested access access $\in \mathcal{S}_\lambda$; and outputs a quantum credential cred along with a classical identifier id.
- $\mathsf{VerifyCred}(1^\lambda, \mathsf{nym}, \mathsf{access}, \mathsf{cred}) \in \{0,1\}$: The polynomial-time quantum algorithm VerifyCred is run by a verifier of the user's credentials. It takes input the issuer's pseudonym nym, the requested access access $\in \mathcal{S}_\lambda$, and a credential cred; and outputs 1 iff cred is a valid credential for access access with respect to nym.
- $\mathsf{revnotice} \leftarrow \mathsf{Revoke}(1^\lambda, \mathsf{nym}, \mathsf{sk}, \mathsf{access})$: The polynomial-time quantum algorithm Revoke is run by the issuer of the credentials. It takes input the issuer's keys nym and sk, and the access access being revoked; and outputs a notice of revocation revnotice.
- $\pi \leftarrow \mathsf{Prove}(1^\lambda, \mathsf{nym}, \mathsf{revnotice}, \mathsf{cred})$: The polynomial-time quantum algorithm Prove is run by the user of the credentials. It takes input the issuer's pseudonym nym, and a revocation notice revnotice, and the credential to be revoked cred which is identified by revnotice; and outputs a proof of revocation π.

- $\mathsf{VerRevoke}(1^\lambda, \mathsf{nym}, \mathsf{sk}, \mathsf{access}, \mathsf{revnotice}, \pi) \in \{0,1\}$: The polynomial-time quantum algorithm VerRevoke is run by the issuer of the credentials. It takes input the issuer's keys nym and sk, the access access being revoked, the revocation notice revnotice, and a proof of revocation π; and outputs 1 iff π is a valid proof that the user's access to the credential identified by id has been revoked.

Properties.

- **Correctness**: For every sufficiently large $\lambda \in \mathbb{N}$, and every $\mathsf{access} \in \mathcal{S}_\lambda$,

$$\Pr_{\substack{(\mathsf{nym},\mathsf{sk})\leftarrow\mathsf{IssuerKeyGen}(1^\lambda)\\ \mathsf{cred}\leftarrow\mathsf{Issue}(1^\lambda,\mathsf{nym},\mathsf{sk},\mathsf{access})}}[\mathsf{VerifyCred}(1^\lambda,\mathsf{nym},\mathsf{access},\mathsf{cred})=1]=1.$$

- **Revocation**: For every polynomial-size quantum circuit $\mathcal{A}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for sufficiently large $\lambda \in \mathbb{N}$, and every $\mathsf{access} \in \mathcal{M}_\lambda$

$$\Pr_{\substack{(\mathsf{nym},\mathsf{sk})\leftarrow\mathsf{IssuerKeyGen}(1^\lambda)\\ \mathsf{cred}\leftarrow\mathsf{Issue}(1^\lambda,\mathsf{nym},\mathsf{sk},\mathsf{access})\\ \mathsf{revnotice}\leftarrow\mathsf{Revoke}(1^\lambda,\mathsf{nym},\mathsf{sk},\mathsf{access})\\ \pi,\mathsf{cred}'\leftarrow\mathcal{A}(1^\lambda,\mathsf{nym},\mathsf{revnotice},\mathsf{cred})}}\left[\begin{array}{c}\mathsf{VerRevoke}(1^\lambda,\mathsf{nym},\mathsf{sk},\mathsf{access},\mathsf{revnotice},\pi)=1\\ \bigwedge \mathsf{VerifyCred}(1^\lambda,\mathsf{nym},\mathsf{access},\mathsf{cred}')=1\end{array}\right]\leq \mathsf{negl}(\lambda).$$

REMARK 6.1. Unlike previous literature, the users that get issued credentials do not have their own identity. We also define algorithms for a three-message revocation process as opposed to the polynomial-message revocation process defined in the literature.

We now introduce a construction based on unclonable signatures of knowledge.

Theorem 6.6. *Let $(\mathcal{X},\mathcal{W})$ be a hard-distribution of instance and witness pairs for some* NP *relation. Let $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$ be some set of accesses. Let* $(\mathsf{Setup},\mathsf{Sign},\mathsf{Verify})$ *be an unclonable-extractable SimExt-secure signature of knowledge for message space $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$ (Definition 6.1).*

$(\mathsf{IssuerKeyGen},\mathsf{Issue},\mathsf{VerifyCred},\mathsf{Revoke},\mathsf{Prove},\mathsf{VerRevoke})$ *defined in Figure 7 is a revocable anonymous credentials scheme (Definition 6.5).*

Proof Sketch. The correctness of this revocable anonymous credentials scheme follows from the correctness of the unclonable signature of knowledge scheme.

We will now sketch the proof of revocation. Say that there exists an adversary $\mathcal{A}$, access access, and polynomial $p(\cdot)$ such that, with probability at least $1/p(\lambda)$: (1) π passes the revocation check, and (2) cred' passes the credential check. This means that both π and cred' are valid signatures with respect to the same crs, x, and access that the signature cred was issued under. This satisfies the "if" condition of the unclonability property of the unclonable signature of knowledge. As such, there exists a polynomial $q(\cdot)$ such that the unclonable signature of knowledge's extractor can produce a witness w for x with probability at least $1/q(\lambda)$. However, this contradicts the hardness of the distribution $(\mathcal{X},\mathcal{W})$. Hence, our protocol must have the revocation property. □

Corollary 6.7. Assuming the polynomial quantum hardness of LWE, injective one-way functions exist, post-quantum iO exists, and the hardness of NP, there exists a revocable anonymous credentials scheme (Definition 6.5).

Proof. This follows from Corollary 6.4 and Theorem 6.6. □

Revocable Anonymous Credentials

Let $(\mathcal{X}, \mathcal{W})$ be a hard-distribution of instance and witness pairs for some NP relation. Let $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$ be some set of accesses. Let $(\mathsf{Setup}, \mathsf{Sign}, \mathsf{Verify})$ be an unclonable-extractable SimExt-secure signature of knowledge for message space $\{\mathcal{S}_\lambda\}_{\lambda\in\mathbb{N}}$ (Definition 6.1).

ISSUERKEYGEN(1^λ):

- $(\mathsf{crs}, \mathsf{td}) \leftarrow \mathsf{Setup}(1^\lambda)$.
- $(x, w) \leftarrow (\mathcal{X}, \mathcal{W})$.
- Output $\mathsf{nym} = (\mathsf{crs}, x)$ and $\mathsf{sk} = (\mathsf{td}, w)$.

ISSUE$(\mathsf{nym}, \mathsf{sk}, \mathsf{access})$:

- $\sigma \leftarrow \mathsf{Sign}(\mathsf{crs}, x, w, \mathsf{access})$.
- Output $\mathsf{cred} = \sigma$.

VERIFYCRED$(\mathsf{nym}, \mathsf{access}, \mathsf{cred})$:

- Output $\mathsf{Verify}(\mathsf{crs}, x, \mathsf{access}, \sigma)$.

REVOKE$(\mathsf{nym}, \mathsf{sk}, \mathsf{access})$:

- Output $\mathsf{revnotice} = \mathsf{access}$.

PROVE$(\mathsf{nym}, \mathsf{revnotice}, \mathsf{cred})$:

- Output $\pi = \mathsf{cred}$.

VERIFYREVOKE$(\mathsf{nym}, \mathsf{sk}, \mathsf{access}, \mathsf{revnotice}, \pi)$:

- Output $\mathsf{VerifyCred}(\mathsf{nym}, \mathsf{access}, \mathsf{cred})$.

Figure 7: Revocable Anonymous Credentials

References

[Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC 2009, Paris, France, 15-18 July 2009*, pages 229–242. IEEE Computer Society, 2009.

[AC13] Scott Aaronson and Paul F. Christiano. Quantum money from hidden subspaces. *Theory Comput.*, 9:349–401, 2013.

[AGKZ20] Ryan Amos, Marios Georgiou, Aggelos Kiayias, and Mark Zhandry. One-shot signatures and applications to hybrid quantum/classical authentication. In Konstantin

Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 255–268. ACM, 2020.

[AK21] Prabhanjan Ananth and Fatih Kaleoglu. Unclonable encryption, revisited. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part I*, volume 13042 of *Lecture Notes in Computer Science*, pages 299–329. Springer, 2021.

[AKL+22] Prabhanjan Ananth, Fatih Kaleoglu, Xingjian Li, Qipeng Liu, and Mark Zhandry. On the feasibility of unclonable encryption, and more. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 212–241. Springer, 2022.

[ALL+21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 526–555. Springer, 2021.

[AN11] Tolga Acar and Lan Nguyen. Revocation for delegatable anonymous credentials. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings*, volume 6571 of *Lecture Notes in Computer Science*, pages 423–440. Springer, 2011.

[AP21] Prabhanjan Ananth and Rolando L. La Placa. Secure software leasing. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 501–530. Springer, 2021.

[BCC+09] Mira Belenkiy, Jan Camenisch, Melissa Chase, Markulf Kohlweiss, Anna Lysyanskaya, and Hovav Shacham. Randomizable proofs and delegatable anonymous credentials. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 108–125. Springer, 2009.

[BGG+23] James Bartusek, Sanjam Garg, Vipul Goyal, Dakshita Khurana, Giulio Malavolta, Justin Raizes, and Bhaskar Roberts. Obfuscation and outsourced computation with certified deletion. Cryptology ePrint Archive, Paper 2023/265, 2023.

[BI20] Anne Broadbent and Rabib Islam. Quantum encryption with certified deletion. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography*, pages 92–122, Cham, 2020. Springer International Publishing.

[BK23] James Bartusek and Dakshita Khurana. Cryptography with certified deletion. In *Crypto 2023 (to appear)*, 2023.

[BKP23] James Bartusek, Dakshita Khurana, and Alexander Poremba. Publicly-verifiable deletion via target-collapsing functions. In *Crypto 2023 (to appear)*, 2023.

[BL20] Anne Broadbent and Sébastien Lord. Uncloneable quantum encryption via oracles. In Steven T. Flammia, editor, *15th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2020, June 9-12, 2020, Riga, Latvia*, volume 158 of *LIPIcs*, pages 4:1–4:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[BS16] Shalev Ben-David and Or Sattath. Quantum tokens for digital signatures. *CoRR*, abs/1609.09047, 2016.

[BS17] Shalev Ben-David and Or Sattath. Quantum tokens for digital signatures. *IACR Cryptol. ePrint Arch.*, page 94, 2017.

[BS23a] Mohammed Barhoush and Louis Salvail. How to sign quantum messages, 2023.

[BS23b] Mohammed Barhoush and Louis Salvail. Powerful primitives in the bounded quantum storage model, 2023.

[CKS10] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. Solving revocation with efficient update of anonymous credentials. In Juan A. Garay and Roberto De Prisco, editors, *Security and Cryptography for Networks, 7th International Conference, SCN 2010, Amalfi, Italy, September 13-15, 2010. Proceedings*, volume 6280 of *Lecture Notes in Computer Science*, pages 454–471. Springer, 2010.

[CL06] Melissa Chase and Anna Lysyanskaya. On signatures of knowledge. In Cynthia Dwork, editor, *Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings*, volume 4117 of *Lecture Notes in Computer Science*, pages 78–96. Springer, 2006.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 556–584. Springer, 2021.

[CW19] Xavier Coiteux-Roy and Stefan Wolf. Proving erasure. In *IEEE International Symposium on Information Theory, ISIT 2019, Paris, France, July 7-12, 2019*, pages 832–836, 2019.

[FGH+12] Edward Farhi, David Gosset, Avinatan Hassidim, Andrew Lutomirski, and Peter W. Shor. Quantum money from knots. In Shafi Goldwasser, editor, *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 276–289. ACM, 2012.

[FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume I*, pages 308–317. IEEE Computer Society, 1990.

[FM18] Honghao Fu and Carl A. Miller. Local randomness: Examples and application. *Phys. Rev. A*, 97:032324, Mar 2018.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.

[Got03] Daniel Gottesman. Uncloneable encryption. *Quantum Inf. Comput.*, 3(6):581–602, 2003.

[GZ20] Marios Georgiou and Mark Zhandry. Unclonable decryption keys. *IACR Cryptol. ePrint Arch.*, page 877, 2020.

[HMNY21] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum encryption with certified deletion, revisited: Public key, attribute-based, and classical communication. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021*, pages 606–636, Cham, 2021. Springer International Publishing.

[HMNY22] Taiga Hiroka, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Certified everlasting zero-knowledge proof for QMA. CRYPTO, 2022. https://ia.cr/2021/1315.

[IBM23] IBM. Cost of a data breach report 2023. Technical report, IBM, 2023.

[Kan18] Daniel M. Kane. Quantum money from modular forms. *CoRR*, abs/1809.05925, 2018.

[KN23] Fuyuki Kitagawa and Ryo Nishimaki. One-out-of-many unclonable cryptography: Definitions, constructions, and more. *IACR Cryptol. ePrint Arch.*, page 229, 2023.

[KT20] Srijita Kundu and Ernest Y. Z. Tan. Composably secure device-independent encryption with certified deletion, 2020.

[LS19] Alex Lombardi and Luke Schaeffer. A note on key agreement and non-interactive commitments. Cryptology ePrint Archive, Paper 2019/279, 2019. https://eprint.iacr.org/2019/279.

[LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum fiat-shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 326–355. Springer, 2019.

[MST21] Christian Majenz, Christian Schaffner, and Mehrdad Tahmasbi. Limitations on uncloneable encryption and simultaneous one-way-to-hiding. *IACR Cryptol. ePrint Arch.*, page 408, 2021.

[Por22] Alexander Poremba. Quantum proofs of deletion for learning with errors. Cryptology ePrint Archive, Report 2022/295, 2022. https://ia.cr/2022/295.

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 89–114. Springer, 2019.

[Sah99] Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th Annual Symposium on Foundations of Computer Science, FOCS '99, 17-18 October, 1999, New York, NY, USA*, pages 543–553. IEEE Computer Society, 1999.

[SCO+01] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 566–598. Springer, 2001.

[SCP00] Alfredo De Santis, Giovanni Di Crescenzo, and Giuseppe Persiano. Necessary and sufficient assumptions for non-iterative zero-knowledge proofs of knowledge for all NP relations. In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *Automata, Languages and Programming, 27th International Colloquium, ICALP 2000, Geneva, Switzerland, July 9-15, 2000, Proceedings*, volume 1853 of *Lecture Notes in Computer Science*, pages 451–462. Springer, 2000.

[SP92] Alfredo De Santis and Giuseppe Persiano. Zero-knowledge proofs of knowledge without interaction (extended abstract). In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*, pages 427–436. IEEE Computer Society, 1992.

[Unr14] Dominique Unruh. Revocable quantum timed-release encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 129–146. Springer, 2014.

[Unr17] Dominique Unruh. Post-quantum security of fiat-shamir. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, volume 10624 of *Lecture Notes in Computer Science*, pages 65–95. Springer, 2017.

[Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983.

[Zha19a] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 239–268. Springer, 2019.

[Zha19b] Mark Zhandry. Quantum lightning never strikes the same state twice. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 408–438. Springer, 2019.

A A Reduction Between Unclonability Definitions

A.1 In the CRS model

For completeness, here we repeat the definitions of unclonability.

Definition A.1. (Unclonable Security for Hard Instances). A proof $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ satisfies unclonable security if for every language $\mathcal{L}$ with corresponding relation $\mathcal{R}_{\mathcal{L}}$, for every polynomial-sized quantum circuit family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$, and for every hard distribution $\{\mathcal{X}_\lambda, \mathcal{W}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\mathcal{R}_{\mathcal{L}}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr_{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)}\left[\mathsf{Verify}(\mathsf{crs},x,\pi_1)=1 \bigwedge \mathsf{Verify}(\mathsf{crs},x,\pi_2)=1 \middle| \begin{smallmatrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow C_\lambda(x,\pi)\end{smallmatrix}\right] \leq \mathsf{negl}(\lambda).$$

Definition A.2. (1-to-2 Unclonable Extractability) A proof $(\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ satisfies unclonable security there exists a QPT extractor $\mathcal{E}$ which is an oracle-aided circuit such that for every language $\mathcal{L}$ with corresponding relation $\mathcal{R}_{\mathcal{L}}$ and for every non-uniform polynomial-time quantum adversary $\mathcal{A}$, for every instance-witness pair $(x, w) \in \mathcal{R}_{\mathcal{L}}$ and $\lambda = \lambda(|x|)$, such that there is a polynomial $p(\cdot)$ satisfying:

$$\Pr\left[\mathsf{Verify}(\mathsf{crs},x,\pi_1)=1 \bigwedge \mathsf{Verify}(\mathsf{crs},x,\pi_2)=1 \middle| \begin{smallmatrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow\mathcal{A}_\lambda(\mathsf{crs},x,\pi,z)\end{smallmatrix}\right] \geq \frac{1}{p(\lambda)}, \tag{32}$$

there is also a polynomial $q(\cdot)$ such that

$$\Pr\left[(x,w_{\mathcal{A}})\in\mathcal{R}_{\mathcal{L}} | w_{\mathcal{A}}\leftarrow\mathcal{E}^{\mathcal{A}}(x)\right] \geq \frac{1}{q(\lambda)}. \tag{33}$$

Claim A.3. Any protocol satisfying Definition A.2 also satisfies Definition A.1.

Proof. Suppose there exists a protocol $\Pi = (\mathsf{Setup}, \mathsf{Prove}, \mathsf{Verify})$ satisfying Definition A.2.

Suppose towards a contradiction that Π does not satisfy Definition A.1. This implies that there is a QPT adversary $\widehat{\mathcal{A}}$, auxiliary input $\widehat{z} = \{\widehat{z}_\lambda\}_{\lambda\in\mathbb{N}}$, a hard distribution $(\mathcal{X}, \mathcal{W})$ over $\mathcal{R}_{\mathcal{L}}$, and a polynomial $p(\cdot)$ such that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}\left[\mathsf{Verify}(\mathsf{crs},x,\pi_1)=1 \bigwedge \mathsf{Verify}(\mathsf{crs},x,\pi_2)=1 \middle| \begin{smallmatrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow\widehat{\mathcal{A}}_\lambda(\mathsf{crs},x,\pi,\widehat{z})\end{smallmatrix}\right] \geq \frac{1}{p(\lambda)}. \tag{34}$$

Let S denote the set of instance-witness pairs $\{(x, w) \in (\mathcal{X}, \mathcal{W})\}$ that satisfy

$$\Pr\left[\mathsf{Verify}(\mathsf{crs},x,\pi_1)=1 \bigwedge \mathsf{Verify}(\mathsf{crs},x,\pi_2)=1 \middle| \begin{smallmatrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow\widehat{\mathcal{A}}_\lambda(\mathsf{crs},x,\pi,\widehat{z})\end{smallmatrix}\right] \geq \frac{1}{2p(\lambda)}. \tag{35}$$

First, we claim that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}[(x,w)\in S] \geq \frac{1}{2p(\lambda)} \tag{36}$$

Suppose not, then by Equation (35),

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}\left[\mathsf{Verify}(\mathsf{crs},x,\pi_1)=1\bigwedge\mathsf{Verify}(\mathsf{crs},x,\pi_2)=1\middle|\begin{matrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow\widehat{\mathcal{A}}_\lambda(\mathsf{crs},x,\pi,\widehat{z})\end{matrix}\right]<\frac{1}{2p(\lambda)}+\frac{1}{2p(\lambda)}.$$

contradicting Equation (34). Thus, Equation (36) must be true.

Consider the extractor $\mathcal{E}$ guaranteed by Definition A.2. Given a sample $(x,w)\leftarrow(\mathcal{X},\mathcal{W})$, we will show that there is a polynomial $p'(\cdot)$ such that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}[\mathcal{E}^{\widehat{\mathcal{A}}}(x,\widehat{z})\in\mathcal{R}_{\mathcal{L}}(x)]\geq\frac{1}{p'(\lambda)}\tag{37}$$

which suffices to contradict hardness of the distribution $(\mathcal{X},\mathcal{W})$, as desired.

Towards showing that Equation (37) holds, recall by Definition A.2 that for every NP instance-witness pair (x,w) such that there is a polynomial $p(\cdot)$ satisfying:

$$\Pr\left[\mathsf{Verify}(\mathsf{crs},x,\pi_1)=1\bigwedge\mathsf{Verify}(\mathsf{crs},x,\pi_2)=1\middle|\begin{matrix}\mathsf{crs}\leftarrow\mathsf{Setup}(1^\lambda)\\ \pi\leftarrow\mathsf{Prove}(\mathsf{crs},x,w)\\ \pi_1,\pi_2\leftarrow\widehat{A}_\lambda(\mathsf{crs},x,\pi,\widehat{z})\end{matrix}\right]\geq\frac{1}{p(\lambda)},$$

there is also a polynomial $q(\cdot)$ such that

$$\Pr\Big[R_L(x,w)=1|w\leftarrow\mathcal{E}^{\widehat{A}}(x,\widehat{z})\Big]\geq\frac{1}{q(\lambda)}$$

This implies that there is a polynomial $q(\cdot)$ such that for every $(x,w)\in S$,

$$\Pr\Big[R_L(x,w)=1|w\leftarrow\mathcal{E}^{\widehat{A}}(x,\widehat{z})\Big]\geq\frac{1}{q(\lambda)}$$

This, combined with Equation (36) implies that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}[R_L(x,w)=1|w\leftarrow\mathcal{E}^{\widehat{A}}(x,\widehat{z})]\geq\frac{1}{2p(\lambda)q(\lambda)}$$

which proves Equation (37) as desired. □

A.2 In the QRO model

For completeness, here we repeat the definitions of unclonability.

Definition A.4. (Unclonable Security for Hard Instances). A proof $(\mathsf{Prove},\mathsf{Verify})$ satisfies unclonable security with respect to a quantum random oracle $\mathcal{O}$ if for every language $\mathcal{L}$ with corresponding relation $\mathcal{R}_{\mathcal{L}}$, for every polynomial-sized quantum oracle-aided circuit family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$, and for every hard distribution $\{\mathcal{X}_\lambda,\mathcal{W}_\lambda\}_{\lambda\in\mathbb{N}}$ over $\mathcal{R}_{\mathcal{L}}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\lambda\in\mathbb{N}$,

$$\Pr_{(x,w)\leftarrow(\mathcal{X}_\lambda,\mathcal{W}_\lambda)}\left[\mathsf{Verify}^{\mathcal{O}}(x,\pi_1)=1\bigwedge\mathsf{Verify}^{\mathcal{O}}(x,\pi_2)=1\middle|\begin{matrix}\pi\leftarrow\mathsf{Prove}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2\leftarrow C_\lambda(x,\pi)\end{matrix}\right]\leq\mathsf{negl}(\lambda).$$

Definition A.5. (1-to-2 Unclonable Extractability) A proof $(\mathsf{Prove}, \mathsf{Verify})$ satisfies unclonable security with respect to a quantum random oracle $\mathcal{O}$ there exists a QPT extractor $\mathcal{E}$ which is an oracle-aided circuit such that for every language $\mathcal{L}$ with corresponding relation $\mathcal{R}_{\mathcal{L}}$ and for every non-uniform polynomial-time quantum adversary $\mathcal{A}$, for every instance-witness pair $(x, w) \in \mathcal{R}_{\mathcal{L}}$ and $\lambda = \lambda(|x|)$, such that there is a polynomial $p(\cdot)$ satisfying:

$$\Pr\left[\mathsf{Verify}^{\mathcal{O}}(x,\pi_1)=1 \bigwedge \mathsf{Verify}^{\mathcal{O}}(x,\pi_2)=1 \middle| \begin{matrix}\pi \leftarrow \mathsf{Prove}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2 \leftarrow \mathcal{A}_\lambda^{\mathcal{O}}(x,\pi,z)\end{matrix}\right] \geq \frac{1}{p(\lambda)}, \tag{38}$$

there is also a polynomial $q(\cdot)$ such that

$$\Pr\Big[(x, w_{\mathcal{A}}) \in \mathcal{R}_{\mathcal{L}} | w_{\mathcal{A}} \leftarrow \mathcal{E}^{\mathcal{A}^{|\mathcal{O}\rangle}}(x)\Big] \geq \frac{1}{q(\lambda)}. \tag{39}$$

Claim A.6. Any protocol satisfying Definition A.5 also satisfies Definition A.4.

Proof. Suppose there exists a protocol $\Pi = (\mathsf{Prove}, \mathsf{Verify})$ satisfying Definition A.5.

Suppose towards a contradiction that Π does not satisfy Definition A.4. This implies that there is a QPT adversary $\widehat{\mathcal{A}}$ with oracle access to some quantum random oracle $\mathcal{O}$, auxiliary input $\widehat{z} = \{\widehat{z}_\lambda\}_{\lambda\in\mathbb{N}}$, a hard distribution $(\mathcal{X}, \mathcal{W})$ over $\mathcal{R}_{\mathcal{L}}$, and a polynomial $p(\cdot)$ such that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}\left[\mathsf{Verify}^{\mathcal{O}}(x,\pi_1)=1 \bigwedge \mathsf{Verify}^{\mathcal{O}}(x,\pi_2)=1 \middle| \begin{matrix}\pi \leftarrow \mathsf{Prove}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2 \leftarrow \widehat{\mathcal{A}}_\lambda^{\mathcal{O}}(x,\pi,\widehat{z})\end{matrix}\right] \geq \frac{1}{p(\lambda)}. \tag{40}$$

Let S denote the set of instance-witness pairs $\{(x, w) \in (\mathcal{X}, \mathcal{W})\}$ that satisfy

$$\Pr\left[\mathsf{Verify}^{\mathcal{O}}(x,\pi_1)=1 \bigwedge \mathsf{Verify}^{\mathcal{O}}(x,\pi_2)=1 \middle| \begin{matrix}\pi \leftarrow \mathsf{Prove}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2 \leftarrow \widehat{\mathcal{A}}_\lambda^{\mathcal{O}}(x,\pi,\widehat{z})\end{matrix}\right] \geq \frac{1}{2p(\lambda)}. \tag{41}$$

First, we claim that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}[(x, w) \in S] \geq \frac{1}{2p(\lambda)} \tag{42}$$

Suppose not, then by Equation (41),

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}\left[\mathsf{Verify}^{\mathcal{O}}(x,\pi_1)=1 \bigwedge \mathsf{Verify}^{\mathcal{O}}(x,\pi_2)=1 \middle| \begin{matrix}\pi \leftarrow \mathsf{Prove}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2 \leftarrow \widehat{\mathcal{A}}_\lambda^{\mathcal{O}}(x,\pi,\widehat{z})\end{matrix}\right] < \frac{1}{2p(\lambda)} + \frac{1}{2p(\lambda)}.$$

contradicting Equation (40). Thus, Equation (42) must be true.

Consider the extractor $\mathcal{E}$ guaranteed by Definition A.5. Given a sample $(x, w) \leftarrow (\mathcal{X}, \mathcal{W})$, we will show that there is a polynomial $p'(\cdot)$ such that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}[\mathcal{E}^{\widehat{\mathcal{A}}}(x, \widehat{z}) \in \mathcal{R}_{\mathcal{L}}(x)] \geq \frac{1}{p'(\lambda)} \tag{43}$$

which suffices to contradict hardness of the distribution $(\mathcal{X}, \mathcal{W})$, as desired.

Towards showing that Equation (43) holds, recall by Definition A.5 that for every NP instance-witness pair (x, w) such that there is a polynomial $p(\cdot)$ satisfying:

$$\Pr\left[\mathsf{Verify}^{\mathcal{O}}(x,\pi_1)=1 \bigwedge \mathsf{Verify}^{\mathcal{O}}(x,\pi_2)=1 \middle| \begin{smallmatrix}\pi\leftarrow\mathsf{Prove}^{\mathcal{O}}(x,w)\\ \pi_1,\pi_2\leftarrow\widehat{A}_\lambda^{|\mathcal{O}\rangle}(\mathsf{crs},x,\pi,\widehat{z})\end{smallmatrix}\right] \geq \frac{1}{p(\lambda)},$$

there is also a polynomial $q(\cdot)$ such that

$$\Pr\Big[R_L(x,w)=1 | w\leftarrow \mathcal{E}^{\widehat{A}}(x,\widehat{z})\Big] \geq \frac{1}{q(\lambda)}$$

This along with Equation (40) implies that there is a polynomial $q(\cdot)$ such that for every $(x, w) \in S$,

$$\Pr\Big[R_L(x,w)=1 | w\leftarrow \mathcal{E}^{\widehat{A}}(x,\widehat{z})\Big] \geq \frac{1}{q(\lambda)}$$

This, combined with Equation (42) implies that

$$\Pr_{(x,w)\leftarrow(\mathcal{X},\mathcal{W})}[R_L(x,w)=1 | w\leftarrow \mathcal{E}^{\widehat{A}}(x,\widehat{z})] \geq \frac{1}{2p(\lambda)q(\lambda)}$$

which proves Equation (43) as desired. □