On Removing Interaction from Quantum Proofs

Nicholas Spooner
Cornell University

Max Tromanhauser
Cornell University

Abstract

An important open question in quantum cryptography is the construction of publicly-verifiable NIZKs for QMA. Classically, one can construct NIZKs for NP in the random oracle model (and sometimes in the standard model) by compiling an honest-verifier ZK (HVZK) Σ-protocol for NP using the Fiat–Shamir transformation. Broadbent and Grilo introduced a quantum analog of a Σ-protocol (which they call a Ξ-protocol) in which the prover's first message is quantum, and show that HVZK Ξ-protocols exist for QMA. However, it is not clear how to compile such protocols into NIZKs in the (Q)ROM, because the Fiat–Shamir transformation seems to be incompatible with quantum messages. In this work we give formal evidence that this is indeed the case: we show that if generic "Fiat–Shamir-like" compilers for quantum protocols exist in the QROM (with small completeness and soundness error) then $\mathsf{QMA} = \mathsf{BQP}$.

1 Introduction

This work is motivated by the problem of constructing *nontrivial publicly-verifiable non-interactive (quantum) proof systems* for QMA. A proof system for QMA is a protocol in which a prover aims to convince a verifier that a given input x belongs to a language $\mathfrak{L} \in \mathsf{QMA}$. By *nontrivial* we mean that the system satisfies some useful property, like zero knowledge or succinctness, which is not achieved by simply sending the QMA witness. A proof system is *non-interactive* if the communication consists of a single message from the prover.

One common approach to designing non-interactive proof systems is to start from an *interactive* proof system and then remove the interaction using a generic compilation. This is most commonly achieved using the *Fiat–Shamir (FS) transformation* [FS86], which applies to classical public-coin interactive proofs (i.e., where the verifier only sends uniformly random strings). In a nutshell, the FS transformation has the prover simulate each verifier message by applying a cryptographic hash function h to the protocol transcript so far. Provided the hash function is sufficiently "unpredictable," this forces the prover to commit to all of its prior messages before learning the verifier's next message.

The FS transformation can be proven secure if h is modelled as a (quantum-accessible) random oracle. By applying the transformation to public-coin zero knowledge proofs and succinct arguments, one can obtain non-interactive zero knowledge (NIZK) arguments and succinct non-interactive arguments (SNARGs) in the random oracle model (ROM). While it is known that the Fiat–Shamir transformation cannot be proven secure in the standard model in general [GK03], for certain protocols one can instantiate h with a well-chosen cryptographic hash function to obtain a secure scheme in the standard model [CCH$^+$18].[1]

Unfortunately we do not know how to leverage the FS transformation to obtain publicly-verifiable NIZKs or SNARGs for QMA even in the ROM. The reason is that known interactive proof systems for QMA are either based on protocols for classical verification of quantum computations, which have classical communication but require the verifier to hold private state, or are public-coin but with quantum prover messages [BG20, GJMZ23]. In the former case, one can obtain *designated-verifier* NIZKs/SNARGs for QMA in the ROM by leveraging fully-homomorphic encryption [ACGH20, BKL$^+$22]; it is unclear whether this approach can lead to publicly verifiable NIZKs/SNARGs.

[1] This is a very substantial line of work and so we do not attempt to provide a full list of citations.

For protocols with quantum communication, one possible adaptation of classical Fiat–Shamir might be to coherently compute a classical hash of the first message and measure the output. This is problematic for two reasons. First, the resulting protocol may not be *complete*, because the measurement may substantially disturb the quantum message and cause the verifier to reject. Second, the resulting protocol may not be *sound*, because (e.g.) if the verifier checks the hash in the Z basis, the adversary can apply arbitrary Z operations without affecting that check. (These issues are related to the impossibility of signing quantum states [AGM21].)

In this work we seek to determine whether this difficulty represents a barrier to *any* generic technique for removing interaction from quantum protocols. More precisely, we ask the question:

Is there a generic method for removing interaction from protocols with quantum communication?

1.1 Our Results

This work initiates the study of that question by providing a partial negative answer.

First, we formalize what we mean by a "generic method for removing interaction" by defining an object we call a (Quantum) Non-Interactivity Compiler ((Q)NIC). Informally, a QNIC is a black-box compilation that converts an interactive Σ-protocol with quantum communication (also known as a Ξ-protocol [BG20]) into a publicly-verifiable non-interactive quantum argument. We additionally define a subclass of *straightline* QNICs which capture compilation strategies that require only a single black-box interaction with the interactive proof. The FS transformation is an example of a straightline classical NIC.

Our main contribution is to rule out the existence of straightline QNICs with small completeness and soundness error assuming the existence of quantum commitments[2].

Theorem 1 (Informal)**.** *Let* $\mathbf{O}$ *be an efficiently simulatable distribution over unitary oracles. Assume that*

- *there exist quantum commitments with respect to* $\mathbf{O}$*, and*
- *there exist* $(1-\varepsilon)$*-complete,* δ*-sound, straightline QNICs relative to* $\mathbf{O}$ *where* $\varepsilon(n) + 2\delta(n) \leq \frac{3}{5}$*.*

Then $\mathsf{QMA} = \mathsf{BQP}$*.*

Our full result (stated in Section 5) covers distributions over arbitrary unitary oracles, although the class collapsing implication is less straightforward in the general case. It is important that the result relativizes because Fiat–Shamir is not secure in the plain model in general, so any meaningful impossibility should hold in, for example, the (quantum) random oracle model. In particular, Theorem 1 implies corollaries in the following well-studied settings.

- In the plain model (i.e., $\mathbf{O}$ is an empty distribution) or the common reference string (CRS) model (i.e., where $O_{\mathsf{crs}}|x\rangle = |x \oplus \mathsf{crs}\rangle$ for $O_{\mathsf{crs}} \leftarrow \mathbf{O}$), assuming the existence of both quantum commitments and small-error straightline QNICs implies $\mathsf{QMA} = \mathsf{BQP}$.
- In the QROM or the quantum Haar random oracle model (QHROM), quantum commitments exist unconditionally. In these cases the implication $\mathsf{QMA} = \mathsf{BQP}$ follows from just the existence of small-error straightline QNICs relative to those models.

We require that $\mathbf{O}$ be a distribution over unitary oracles because our decision algorithm requires access to the oracle's inverse. Thankfully this is not an onerous requirement and it is satisfied by most relevant oracle models (in particular, any quantum-accessible classical oracle).

Open Questions. Our result represents a first step towards understanding whether interaction can be generically removed from quantum protocols. There remains plenty of scope for future work:

[2]Note that quantum zero knowledge for any hard-on-average language implies quantum commitments [BCQ23].

- It remains open whether our impossibility result can be extended to a broader class of compilers which either are not restricted to straightline behavior or have larger completeness and soundness error. Note that while the completeness and soundness error of a QNIC can be generically reduced via parallel repetition, the resulting protocol is no longer straightline and so our result does not apply. Note also that there do exist classical NICs which are not straightline, and so could potentially be used in the quantum setting (e.g., the Fischlin transform [Fis05], which requires multiple interactions with the prover to generate a proof).
- Our result applies only to QNICs which yield publicly-verifiable protocols; it remains an intriguing open question to design QNICs that yield designated-verifier protocols.
- Our result rules out QNICs that work for any Ξ-protocol; potentially there exist QNICs that work for a subclass of protocols satisfying some constraint.

There may also be other ways to obtain non-interactive proof systems for QMA that do not look like removing interaction (analogously to e.g. the hidden bits model in the classical setting).

1.2 Related Work

Classically, the FS transformation has been proven secure in the ROM [BR93, PS96], and a long line of work was required to upgrade this to the QROM. Several works [DFG13, ARU14] demonstrated barriers for standard proof techniques, while others demonstrated that particular modifications of the transformation were secure [Unr15, Unr17, KLS18]. The standard Fiat–Shamir was eventually proven secure in the QROM in [DFMS19, LZ19]. These results all concern the transformation applied to classical Σ-protocols which do not support quantum communication. Another line of work has proven that the classical FS transformation is not secure in the plain model in general [GK03]. In contrast, our result is agnostic to the oracle model and so covers (e.g.) the QROM.

Several strategies for developing nontrivial proof systems for QMA have appeared in the literature. It has been shown that there exist public-coin, three-message succinct arguments [GJMZ23] (assuming the quantum PCP conjecture) and zero knowledge proofs [BG20] for QMA if we allow the prover to send quantum messages. If we wish to avoid interaction, the main approaches have been to use a designated-verifier [ACGH20, BKL$^+$22, CVZ20, Shm21, MY22] — i.e., a verifier that receives a secret setup with which it can confirm the validity of the proof — or to rely on heuristic security of post-quantum obfuscation [BM22, BJRT25].

1.3 Technical Overview

Before we attempt to prove our main result, we must formalize the class of transformations which we are considering. At a high level, we'd like to capture strategies which resemble the structure of Figure 1 — i.e., those which, through a straightline interaction with the prover, create a non-interactive quantum proof that can be unfolded into a valid transcript. This model consists of two algorithms: $\mathcal{T}_P$ which, through a single interaction with the prover, generates a quantum proof; and $\mathcal{T}_V$ which confirms the validity of the proof and outputs a transcript consistent with the compiled protocol. The proof is accepted if $\mathcal{T}_V$'s output is accepted by the verifier[3]. Together $\mathcal{T}_P$ and $\mathcal{T}_V$ define a straightline QNIC, and we require them to be "universal" in the sense that, for any secure Ξ-protocol, the resulting non-interactive argument is also secure.

Our strategy towards proving Theorem 1 is to assume that such a $(\mathcal{T}_P, \mathcal{T}_V)$ exist and leverage them to construct a decision algorithm for a QMA-complete language $\mathfrak{L}$. As an intermediate objective, we construct a family of Ξ-protocols, each of which securely proves membership in $\mathfrak{L}$, but if compiled with $(\mathcal{T}_P, \mathcal{T}_V)$, the resulting NIZK will be vulnerable to a forgery attack. That attack is then used directly to construct the decision algorithm for $\mathfrak{L}$.

[3]Our model actually captures strategies which are slightly more general. Rather than having the non-interactive verifier forward $\mathcal{V}$'s output, we allow $\mathcal{T}_V$ to make a final decision after receiving $\mathcal{V}$'s verdict. See Section 3 for details.

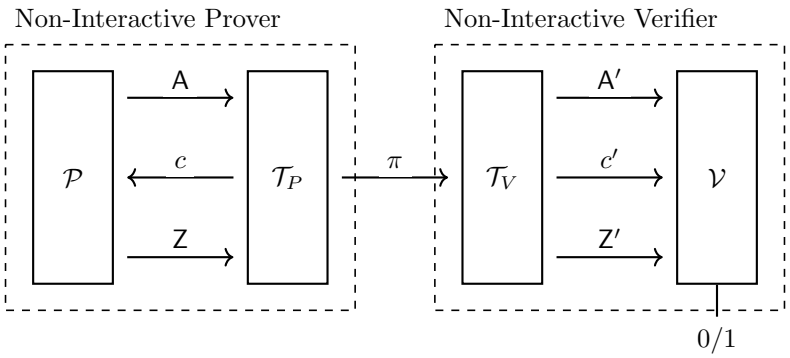


Figure 1: Straightline "Fiat–Shamir-like" compiler (where A, Z, A', and Z' are quantum registers).

The key insight for this counterexample family comes from understanding the security of the classical FS transformation. At its heart, the transform enforces a dependence between the initial message a and the challenge c. That dependence prevents a cheating prover from correlating their message with a favorable challenge. Any prover that succeeds in the non-interactive argument must behave as though it had committed to a before learning c, thereby inheriting the soundness of the interactive protocol.

As we expand our ambitions to include quantum communication, introducing a correlation between a quantum message and a classical challenge introduces a complication: it requires a measurement be performed on the message state. Measurements are necessarily destructive and may irreversibly disturb the message in a way which compromises its completeness. This is the tension which we exploit to build our counterexample.

Our construction starts with an existing Ξ-protocol $(\mathcal{P}, \mathcal{V})$ (e.g., from [BG20]) which we will modify. The QNIC algorithms $(\mathcal{T}_P, \mathcal{T}_V)$ together can be understood as a man-in-the-middle attempting to make some undetectable measurement on the first message. Given this framing, our goal is to modify $\mathcal{P}$ and $\mathcal{V}$ such that any man-in-the-middle adversary must either pass the first message through unaltered or else $\mathcal{V}$ will reject.

Quantum Ciphertext Authentication. Introduced in [AGM18], quantum ciphertext authentication (QCA) provides a natural tool for enforcing this behavior. Any eavesdropper listening in on a QCA-secure quantum encryption scheme will be unable to alter or forge a ciphertext without being caught. The security property is defined through a simulation-based definition where an unbounded adversary is given a single copy of either a 'real' or 'simulated' ciphertext and it attempts to distinguish between the cases.

Formally, the 'real world' is modeled as a quantum channel $\Gamma = \mathbb{E}_k[\mathsf{Dec}_k \circ \mathcal{A} \circ \mathsf{Enc}_k]$ (sometimes referred to as the attacker's 'effective attack map') which composes encryption, the fixed man-in-the-middle attack $\mathcal{A}$, and then decryption. The channel's input message can be adversarially chosen by $\mathcal{A}$ and entangled with an outside register $\mathcal{A}$ controls. In the simulated world, we have a corresponding channel Λ where Enc and Dec are replaced with 'ideal' versions:

- The ciphertext provided by the simulated Enc does not contain the adversary's message but instead encrypts a state which is maximally entangled with a register private to Enc and Dec.
- The simulated Dec then validates the ciphertext and confirms that the entanglement between the message register and the private register has been preserved.

If both conditions are met, Λ outputs the original message. An encryption scheme is QCA-secure if Γ and Λ are negligibly close for all $\mathcal{A}$.

Using a QCA scheme as a building block, our counterexample family is constructed by modifying the Ξ-protocol as follows. The prover's initial message is encrypted under a key k which is part of the protocol description. The verifier sends a uniform challenge as before, and the prover's final message is similarly

$\mathsf{Real}(\rho_{\mathsf{A}} \otimes \rho_{\mathsf{B}})$

1. Sample $k \leftarrow \mathsf{Gen}(1^n)$.
2. Sample encryption randomness r.
3. Prepare
$$\sigma = \mathcal{A} \circ (\mathsf{Enc}_{k;r} \otimes \mathbb{1}_{\mathsf{B}})(\rho_{\mathsf{A}} \otimes \rho_{\mathsf{B}}).$$
4. Measure if A is a valid ciphertext.
5. If yes, output the residual state.
6. Otherwise, output $\perp \otimes \sigma_{\mathsf{B}}$.

$\mathsf{Defer}(\rho_{\mathsf{A}} \otimes \rho_{\mathsf{B}})$

1. Sample $k \leftarrow \mathsf{Gen}(1^n)$.
2. Sample encryption randomness r.
3. Prepare
$$\sigma = \mathcal{A} \circ (\mathsf{Enc}_{k;r} \otimes \mathbb{1}_{\mathsf{B}})(\Phi^+_{\mathsf{ZA}} \otimes \rho_{\mathsf{B}}).$$
4. Measure if A is a valid ciphertext encrypting half of $|\Phi^+\rangle$.
5. If yes, output $\mathsf{Enc}_{k;r}(\rho_{\mathsf{A}}) \otimes \sigma_{\mathsf{B}}$.
6. Otherwise, output $\perp \otimes \sigma_{\mathsf{B}}$.

Figure 2: 'Real' and 'deferred' channels (where $\perp$ is some failure state).

unchanged. The verifier attempts to decrypt the first message register A using k. If this fails, then the verifier rejects; otherwise, it verifies the transcript incorporating the decrypted message.

Note that this construction yields a *family* of protocols: $\mathcal{P}_k$ and $\mathcal{V}_k$ for each key k. Our eventual decision algorithm will compile a randomly-chosen protocol from this family. Importantly, we are not modifying the model to have the prover and verifier share a key: k is part of the *description* of the protocol. The role of k is to exploit the black-box nature of $(\mathcal{T}_P, \mathcal{T}_V)$: the behavior of these algorithms must be independent of k, and so we will be able to obtain security guarantees against them.

Attacking the NIZK. Notice that in the compilation of these protocols the QNIC algorithms $(\mathcal{T}_P, \mathcal{T}_V)$ never see the prover's first message in the clear. Instead, $\mathcal{T}_P$ queries $\mathcal{P}_k$ which provides a message which has already been encrypted. If we switch to a simulated Enc (as in the QCA definition), the message is placed aside and remains untouched until it is decrypted. In this way, we should be able to use the simulated Enc and Dec to 'defer' the choice of message until the moment of decryption, which in our setting is when the unrolled transcript is given to $\mathcal{V}_k$. At that point, we would like to swap the simulated ciphertext for a real ciphertext with some freshly chosen message which depends on the challenge c.

We formalize this idea by introducing two quantum channels (Figure 2): Real, which, when $\mathcal{A} = \mathcal{T}_V \circ \mathcal{T}_P$, corresponds to the honest application of the QNIC; and Defer, which we will use in our attack. We will show that Real and Defer are indistinguishable for all $\mathcal{A}$ when $(\mathsf{Enc}, \mathsf{Dec})$ is QCA-secure. Given that, we construct an attack (depicted in Figure 3) which forges a proof for a randomly sampled counterexample protocol. After switching from Real to Defer we can replace $\mathcal{P}$ with its zero knowledge simulator Sim. Given a statement x, our forgery attack $\mathcal{F}$ proceeds as follows:

1. Sample a random $k \leftarrow \mathsf{Gen}(1^n)$ and run $\mathcal{T}_P$.
2. Once $\mathcal{T}_P$ requests an initial prover message, send a ciphertext encrypting half of a maximally entangled state Φ^+.
3. After $\mathcal{T}_P$ sends a challenge c, run $\mathsf{Sim}(x, c)$ to generate AZ and respond with Z (but hold onto A).
4. After $\mathcal{T}_P$ produces a proof π, provide it to $\mathcal{T}_V$ to get a transcript state on $\mathsf{A'C'Z'}$. Measure if $\mathsf{A'}$ is a valid ciphertext. If not, then abort. Otherwise, replace $\mathsf{A'}$ with an encryption of A output by Sim.
5. Uncompute $\mathcal{T}_V$ to get a forged proof π'.

Assuming Real and Defer are indistinguishable and that Sim outputs a transcript state which is indistinguishable from one which is honestly generated, the forged proof π' should be indistinguishable from an honestly generated proof for x. Because the QNIC is $(1-\varepsilon)$-complete, the proof produced by $\mathcal{F}(x)$ for $x \in \mathfrak{L}$ will be accepted by the NIZK verifier with probability $\geq 1 - \varepsilon(n) - \mathsf{negl}(n)$. In the same way, we know that

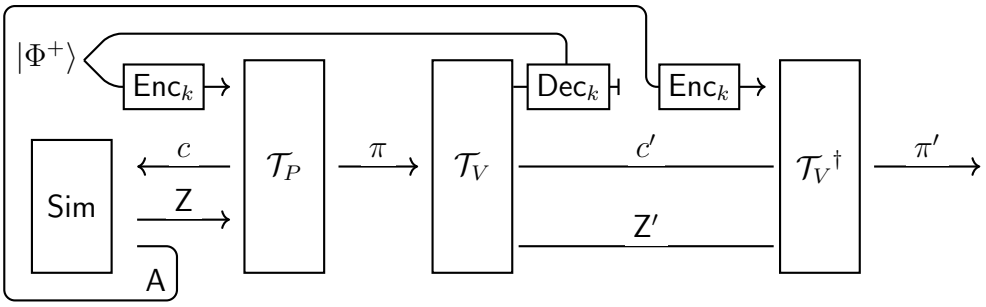


Figure 3: Attack on Straightline QNIC

$\pi' \leftarrow \mathcal{F}(x)$ for $x \notin \mathfrak{L}$ must be accepted with probability $< \delta$ because the QNIC is δ-sound. Thus combining $\mathcal{F}$ with the NIZK verifier yields an efficient decider for the language $\mathfrak{L}$.

Observe that $\mathcal{F}$ relativizes quite straightforwardly because it makes no assumption on the inner workings of $\mathcal{T}_P$ or $\mathcal{T}_V$. The only caveat is that $\mathcal{F}$ must be able to uncompute $\mathcal{T}_V$. Therefore if $\mathcal{T}_V$ has access to some oracle O, we require soundness to hold against adversaries with access to $(O, O^\dagger)$.

Deferring the Message. It remains to be shown that Real and Defer are in fact indistinguishable. At first glance, it might appear that this follows immediately from QCA security: Real and Defer resemble the real and simulated channels Γ and Λ, respectively. However, there is a crucial difference: Γ and Λ output a decrypted plaintext, whereas Real and Defer output a ciphertext. It is therefore not clear that a distance guarantee on Γ and Λ implies a distance guarantee on Real and Defer.

In order to bridge this gap, we introduce a property for QCA schemes which we call *retrospective security*. An encryption scheme is retrospectively secure if an adversary cannot distinguish between the real and simulated channels, even if both k and r are *revealed* upon a successful decryption. This resolves our issue: a distinguisher for Real and Defer can be transformed into a retrospective QCA distinguisher by first reencrypting the plaintext under k, r.

Note that we only disclose the secret key if the adversary's ciphertext is valid. Otherwise, an adversary could simply hold onto its given ciphertext, provide a junk state to Dec, and decrypt the ciphertext itself once given the key. By returning a valid ciphertext to Dec, the adversary is demonstrating that it has 'forgotten' any information it had about the message. Although retrospective security appears to be stronger than plain QCA security, we prove that all QCA-secure schemes satisfy the property.

We prove this implication using a continuity theorem for Stinespring operators from [KSW08]. At a high level, the theorem provides the following guarantee. Let $\mathcal{X}, \mathcal{Y} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{B})$ be channels and let $\mathcal{X}' : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BC})$ be any expansion of $\mathcal{X}$ such that $\mathrm{Tr}_\mathsf{C} \circ \mathcal{X}' = \mathcal{X}$. If $\|\mathcal{X} - \mathcal{Y}\|_\diamond$ is small, then there exists a corresponding expansion $\mathcal{Y}' : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BC})$ such that $\|\mathcal{X}' - \mathcal{Y}'\|_\diamond$ is small. In our setting, we have real and simulated channels Γ and Λ which we know by the QCA guarantee are negligibly close, and we have an expansion of Γ, which we denote Γ_R, which also outputs (k, r) on a new register K. Then by the continuity theorem there exists some expansion Λ_R of Λ to include K such that Γ_R and Λ_R are negligibly close as well. We then show that Λ_R is close to the expansion of Λ which outputs (k, r) on K.

In addition, there is a subtle malleability issue that arises when showing Real and Defer are close. If the attacker in Real is somehow able to map a ciphertext with randomness r to a valid ciphertext with randomness $r' \neq r$, then the output of Real and Defer could be distinguished: Real will produce a ciphertext encrypted with r', whereas Defer always produces a ciphertext encrypted with r. This type of attack does not arise in either the original or retrospective QCA security games, since the encryption randomness of the adversary's ciphertext is discarded. Nonetheless, in Section 4.1 we show via an operator-theoretic argument that ciphertexts in a QCA-secure scheme are not malleable in this way.

Remark 1.1. Readers familiar with the quantum authentication literature may notice that retrospective security closely resembles definitions for key recycling. In particular, [DS18] defines "QCA with key recycling" which is close to our definition but where the simulator outputs a freshly sampled key rather than the key used in the simulation. The authors note that this is a strictly stronger notion than QCA (and therefore our retrospective definition) and provide a simple separating construction where an additional key bit is added to the ciphertext. Another work [GYZ16] defines a security notion "total authentication with key leakage" which is closest to our retrospective definition. In fact, our definition could be understood as a strengthening of their definition from plaintext to ciphertext authentication, if not for our inclusion of the encryption randomness.

Roadmap. The structure for the rest of the paper is as follows. Notation and definitions are laid out in Section 2. In particular, we point readers towards Section 2.3 which describes our unconventional modeling and notational choices for representing Ξ-protocols. In Section 3, we formally define QNICs and discuss the black-box access that they are given to the Ξ-protocols they are compiling. Section 4 proves the indistinguishability of Real and Defer by tackling both gaps described above (in Section 4.2 and Section 4.1). Finally, the formal statement and proof of Theorem 1 appear in Section 5.

Disclosure on Use of AI. The authors used ChatGPT (v5.6 Sol and v6 Astra) as a sounding board, for proofreading, and for some of the technical ideas. Specifically, ChatGPT supplied some of the ideas used in the proof of Lemma 4.2, and the full proof strategy for Lemma 4.4. The paper itself was written and verified in its entirety by the authors, and they take full responsibility for the correctness and content of the work.

2 Preliminaries

By $\mathsf{negl}(n)$ we denote any function which is asymptotically smaller than every inverse polynomial — i.e., that is $o(n^{-c})$ for every $c \in \mathbb{N}$. Given a set S, we write $s \leftarrow S$ to signify that s is uniformly sampled from S. If $\mathbf{D}$ is a distribution, then $x \leftarrow \mathbf{D}$ denotes that x is chosen according to that distribution.

Definition 2.1. Let $(S, \leq)$ be a partially ordered set and $T \subseteq S$. Then T is *downward-closed* if, for all $b \in T$ and $a \leq b$, it holds that $a \in T$.

Definition 2.2. Let X be a Hermitian operator such that $X = \sum_{k=1}^{m} \lambda_k \Pi_k$ for real λ_k and non-zero projections Π_k. Then its *Jordan-Hahn decomposition* is given by

$$X_+ = \sum_{k:\lambda_k>0} \lambda_k \Pi_k, \quad \text{and} \quad X_- = \sum_{k:\lambda_k<0} -\lambda_k \Pi_k$$

where $X_+, X_- \geq 0$, $X = X_+ - X_-$, and $X_+ X_- = 0$.

Fact 2.1. *Let X be Hermitian. Then* $\mathrm{Tr}[X_+] = \frac{1}{2}(\|X\|_1 + \mathrm{Tr}[X])$.

Proof. Because X is Hermitian, $\|X\|_1 = \mathrm{Tr}\sqrt{X^* X} = \mathrm{Tr}[X_+] + \mathrm{Tr}[X_-]$ whereas $\mathrm{Tr}[X] = \mathrm{Tr}[X_+] - \mathrm{Tr}[X_-]$. Adding these equations gives $2\mathrm{Tr}[X_+] = \|X\|_1 + \mathrm{Tr}[X]$. □

2.1 Quantum Information

A pure quantum state is a vector $|\psi\rangle$ in a complex Hilbert space $\mathcal{H}$ such that $\| |\psi\rangle \|_2 = 1$. In this work, $\mathcal{H}$ is finite-dimensional and described as a register — e.g., $\mathcal{H} = \mathsf{A} \otimes \mathsf{B}$ where A and B are subspaces. Let $\mathrm{L}(\mathsf{A})$ denote the space of linear operators on A and $\mathrm{CP}(\mathsf{A}, \mathsf{B})$ denote the cone of completely-positive maps from $\mathrm{L}(\mathsf{A})$ to $\mathrm{L}(\mathsf{B})$. A general (potentially mixed) quantum state is a density matrix $\rho \in \mathrm{L}(\mathsf{A})$ which is positive semidefinite and has unit trace. The density matrix corresponding to the pure state $|\psi\rangle$ is $|\psi\rangle\langle\psi|$ which we occasionally shorten to ψ. Subnormalized quantum states are positive semidefinite operators ρ where

$0 \leq \mathrm{Tr}[\rho] \leq 1$. Given two n-qubit registers A and B, we denote by $|\Phi^+\rangle_{\mathsf{AB}}$ the maximally entangled state $|\Phi^+\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle_{\mathsf{A}} |x\rangle_{\mathsf{B}}$.

Binary measurement of a quantum state can be represented by an operator $0 \leq M \leq \mathbb{1}$, corresponding to the POVM $\{M, \mathbb{1} - M\}$. On an input state ρ, the M outcome occurs with probability $\mathrm{Tr}[M\rho]$ and leaves the residual state $\rho' = \sqrt{M}\rho\sqrt{M}/\mathrm{Tr}[M\rho]$. If $\mathrm{Tr}[M\rho] \geq 1 - \varepsilon$, then $\|\rho - \rho'\|_1 \leq 2\sqrt{\varepsilon}$. We will also use the following bound related to alternating projections.

Lemma 2.1. *Let Π_1, Π_2 be orthogonal projectors, $0 \leq M \leq \mathbb{1}$ commute with Π_1, and ρ be a density matrix with $\Pi_2 \rho \Pi_2 = \rho$. Then*

$$\mathrm{Tr}[M\Pi_1\rho\Pi_1] - \mathrm{Tr}[M\Pi_2\Pi_1\rho\Pi_1\Pi_2] \leq \frac{2}{3\sqrt{3}}.$$

Proof. First, for any operators $0 \leq B \leq A \leq \mathbb{1}$, completing a square gives

$$B - ABA = \frac{3}{2}\left[\left(\mathbb{1} - \frac{A}{\sqrt{3}}\right) B \left(\mathbb{1} - \frac{A}{\sqrt{3}}\right) - \left(A - \frac{\mathbb{1}}{\sqrt{3}}\right) B \left(A - \frac{\mathbb{1}}{\sqrt{3}}\right)\right] \leq \frac{3}{2} A \left(\mathbb{1} - \frac{A}{\sqrt{3}}\right)^2 \leq \frac{2}{3\sqrt{3}}\mathbb{1}.$$

where the final inequality follows by maximizing $\frac{3}{2}t(1 - t/\sqrt{3})^2$ over $0 \leq t \leq 1$, whose maximum occurs at $t = 1/\sqrt{3}$. Apply this bound with $A = \Pi_2\Pi_1\Pi_2$ and $B = \Pi_2\Pi_1 M \Pi_1 \Pi_2$, which satisfy $0 \leq B \leq A \leq \mathbb{1}$. Since M commutes with Π_1,

$$\Pi_2 M \Pi_2 = B + \Pi_2(\mathbb{1} - \Pi_1) M (\mathbb{1} - \Pi_1)\Pi_2 \geq B.$$

Using $\Pi_2\rho\Pi_2 = \rho$ and cyclicity of the trace, we conclude

$$\mathrm{Tr}[M\Pi_1\rho\Pi_1] - \mathrm{Tr}[M\Pi_2\Pi_1\rho\Pi_1\Pi_2] = \mathrm{Tr}[(B - A\Pi_2 M \Pi_2 A)\rho] \leq \mathrm{Tr}[(B - ABA)\rho] \leq \frac{2}{3\sqrt{3}}. \qquad \square$$

Pure quantum states evolve unitarily. General evolution of a quantum state can be represented through completely-positive, trace-preserving maps (also called channels). We describe a channel $\mathcal{N}$ as 'efficient' if there exists a poly-sized quantum circuit (potentially receiving a quantum advice state) which can realize $\mathcal{N}$. If a linear map is completely-positive and trace-decreasing, then we say it is a subchannel. A linear map which acts on maps is called a supermap. In this work we are only concerned with "physically realizable" supermaps which admit a realization $\Theta(\mathcal{N}) = \mathcal{Y} \circ (\mathcal{N} \otimes \mathbb{1}_{\mathsf{E}}) \circ \mathcal{X}$ where $\mathcal{N} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{B})$, $\mathcal{X} : \mathrm{L}(\mathsf{C}) \to \mathrm{L}(\mathsf{AE})$, and $\mathcal{Y} : \mathrm{L}(\mathsf{BE}) \to \mathrm{L}(\mathsf{D})$ are (sub)channels. In this work, we overload $\mathbb{1}$ to denote the identity operator, the channel which acts trivially, and the supermap which acts trivially.

Given a linear map $\mathcal{N} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{B})$, its diamond norm is $\|\mathcal{N}\|_\diamond = \sup_{0 \neq X \in \mathrm{L}(\mathsf{AE})} \|(\mathcal{N} \otimes \mathbb{1}_{\mathsf{E}})(X)\|_1 / \|X\|_1$ where it suffices to take $\mathsf{E} \cong \mathsf{A}$. If $\mathcal{N}$ is Hermiticity-preserving, the supremum may equivalently be taken over density matrices on AE. If $\mathcal{N}$ is completely positive, then $\|\mathcal{N}\|_\diamond = \|\mathcal{N}^\dagger(\mathbb{1})\|_\infty$. The diamond distance induced by that norm obeys the data processing inequality — i.e., it contracts under composition with channels (e.g., $\|\mathcal{N} \circ (\mathcal{X} - \mathcal{Y})\|_\diamond \leq \|\mathcal{X} - \mathcal{Y}\|_\diamond$). More generally, the norm contracts under (physically realizable) supermaps.

Fact 2.2. *Let $\mathcal{N} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BC})$ be completely positive. Then $\|\mathcal{N}\|_\diamond = \|\mathrm{Tr}_{\mathsf{C}} \circ \mathcal{N}\|_\diamond$.*

Proof. Since $\mathcal{N}$ is completely positive, $\|\mathcal{N}\|_\diamond = \|\mathcal{N}^\dagger(\mathbb{1}_{\mathsf{BC}})\|_\infty$. Also, $(\mathrm{Tr}_{\mathsf{C}} \circ \mathcal{N})^\dagger(\mathbb{1}_{\mathsf{B}}) = \mathcal{N}^\dagger(\mathbb{1}_{\mathsf{B}} \otimes \mathbb{1}_{\mathsf{C}})$. $\square$

For all completely-positive maps $\mathcal{N} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{B})$, there exists a linear operator $V : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BE})$ and an additional register E such that $\mathcal{N}(\rho) = \mathrm{Tr}_{\mathsf{E}}[V\rho V^\dagger]$. If $\mathcal{N}$ is a channel, then $V^\dagger V = \mathbb{1}$ and V is an isometry. If $\mathcal{N}$ is a subchannel, then $V^\dagger V \leq \mathbb{1}$ and V is a contraction. The pair (V, E) is called a Stinespring dilation of $\mathcal{N}$. Section 4 relies on the following two results on Stinespring dilations.

Theorem 2 ([KSW08])**.** *Let $\mathcal{N}_1, \mathcal{N}_2 : \mathsf{A} \to \mathsf{B}$ be completely-positive, trace-nonincreasing maps with Stinespring operators $V_1, V_2 : \mathsf{A} \mapsto \mathsf{BE}$ with a common auxiliary register E. Then*

$$\inf_U \|(\mathbb{1}_{\mathsf{B}} \otimes U)V_1 - V_2\|_\infty^2 \leq \|\mathcal{N}_1 - \mathcal{N}_2\|_\diamond \leq 2 \inf_U \|(\mathbb{1}_{\mathsf{B}} \otimes U)V_1 - V_2\|_\infty$$

where the minimization is with respect to all unitaries U on E.

Theorem 3 (Radon-Nikodym Theorem for Completely-Positive Maps [BS86, Rag03])**.** *Let $\mathcal{N}_1, \mathcal{N}_2 : \mathsf{A} \to \mathsf{B}$ be completely-positive maps and let $V : \mathsf{A} \to \mathsf{BE}$ be an operator such that (E, V) is a Stinespring dilation of $\mathcal{N}_2$. Then $\mathcal{N}_1 \leq \mathcal{N}_2$ if and only if there exists an operator $0 \leq Q \leq \mathbb{1}$ on E such that*

$$\mathcal{N}_1(X) = \mathrm{Tr}_{\mathsf{E}}\left[(\mathbb{1}_{\mathsf{B}} \otimes Q^{\frac{1}{2}})VXV^\dagger(\mathbb{1}_{\mathsf{B}} \otimes Q^{\frac{1}{2}})\right]$$

In this work we will be working with both unitary and channel oracles. In both cases, we denote by $\mathcal{A}^O$ the algorithm $\mathcal{A}$ with access to the oracle O. If O is unitary, $\mathcal{A}$ can be understood as having access to O, $O^\dagger$, and O^* as well as their controlled versions. If O is a channel, then $\mathcal{A}$ has access only to the channel. In Theorem 1, by efficiently simulatable we mean the following.

Definition 2.3. An oracle distribution $\mathbf{O}$ is *efficiently simulatable* if there exists a QPT (stateful) algorithm Sim such that, for all QPT $\mathcal{A}$,

$$\left|\Pr_{O \leftarrow \mathbf{O}}[\mathcal{A}^O(1^n) = 1] - \Pr[\mathcal{A}^{\mathsf{Sim}}(1^n) = 1]\right| \leq \mathsf{negl}(n)$$

Theorem 4 ([Zha18])**.** *The QROM distribution is efficiently simulatable.*

Theorem 5 ([GY25, MH25, FLM+26])**.** *The QHROM distribution is efficiently simulatable.*

2.2 Quantum Commitments

In order to allow for the existence of post-quantum zero knowledge, we must assume that quantum commitment schemes exist [BCQ23]. Because our result is concerned with zero knowledge in a relativized model, we define quantum commitments with respect to an oracle distribution.

Definition 2.4. Let $\mathsf{Com} = \{Q_0(\lambda), Q_1(\lambda)\}_{\lambda \in \mathbb{N}}$ be an ensemble of uniformly generated quantum oracle circuits acting on registers CD[4] (where we will drop the security parameter λ to simplify notation). Then Com is a quantum bit commitment scheme with respect to an oracle distribution $\mathbf{O}$ if the following hold.

- **Computational Hiding.** The reduced state on C is computationally indistinguishable between $Q_0^O |0\rangle_{\mathsf{CD}}$ and $Q_1^O |0\rangle_{\mathsf{CD}}$ in expectation over $\mathbf{O}$.
- **(Honest) Statistical Binding.** For all auxiliary states $|\psi\rangle$ on E and unbounded quantum oracle circuits U acting on DE, we have that

$$\mathop{\mathbb{E}}_{O \leftarrow \mathbf{O}} \left\| \left(Q_1^O |0\rangle\langle 0| (Q_1^O)^\dagger \otimes \mathbb{1}_{\mathsf{E}}\right) (\mathbb{1}_{\mathsf{C}} \otimes U_{\mathsf{DE}}^O)(Q_0^O |0\rangle_{\mathsf{CD}} |\psi\rangle_{\mathsf{E}}) \right\| \leq \mathsf{negl}(n)$$

2.3 Proof Systems

As it will be convenient in later sections, we use an unconventional but equivalent model to describe Ξ-protocols (also known as quantum Σ-protocols). Rather than defining the prover as a stateful algorithm which is invoked twice in one interaction, we will split the prover into an initial message ξ^x (which depends on the statement x being proved) and a quantum channel $\mathcal{P} : \mathrm{L}(\mathsf{CS}) \to \mathrm{L}(\mathsf{Z})$ which responds to the verifier's challenge. The state ξ^x includes a message register A and an internal register S which is used by $\mathcal{P}$ to

[4]The size of CD will depend on the security parameter.

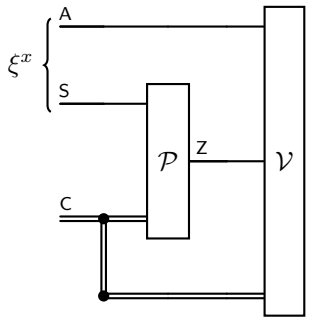


Figure 4: Circuit Representation for an Interactive Ξ-Protocol

formulate its final message on Z. This is equivalent to the stateful prover model because we can simply define ξ^x to be the output of $\mathcal{P}(x)$ along with the prover's residual internal state. Our choice makes clear what is meant by making a single query to the prover.

Although the challenge and verifier output are classical, we can model the verifier as a quantum channel $\mathcal{V} : \mathrm{L}(\mathsf{XACZ}) \to \mathrm{L}(\mathsf{B})$ where B is a single qubit. It can be assumed that $\mathcal{V}$ measures C in the computational basis immediately, and likewise $\mathcal{V}$'s acceptance can be determined by measuring B in the same basis. Given these modeling choices, the protocol interaction can be described by the circuit in Figure 4.

Definition 2.5 (Ξ-Protocols)**.** Let $\mathfrak{L} \subseteq \{0,1\}^n$ be a language, $\{\xi^x\}_{x\in\mathfrak{L}} \subseteq \mathrm{L}(\mathsf{AS})$ be a family of quantum states, $\mathcal{P} : \mathrm{L}(\mathsf{CS}) \to \mathrm{L}(\mathsf{Z})$ be a quantum channel, $\mathcal{V} : \mathrm{L}(\mathsf{XACZ}) \to \mathrm{L}(\mathsf{B})$ be an efficient quantum algorithm and $\pi_{x,c} = (\mathbb{1}_{\mathsf{A}} \otimes \mathcal{P})(\xi^x \otimes |c\rangle\langle c|)$ be a proof state. Then $(\{\xi^x\}, \mathcal{P}, \mathcal{V})$ is a Ξ-*protocol for* $\mathfrak{L}$ if the following holds.

- **Completeness.** $\forall\, x \in \mathfrak{L}$,

$$\mathop{\mathbb{E}}_{c\leftarrow\mathscr{C}} \langle 1|\mathcal{V}(\,|x,c\rangle\langle x,c| \otimes \pi_{x,c})|1\rangle \geq 1 - \mathsf{negl}(n)$$

- **Soundness.** $\forall\ x \notin \mathfrak{L}$, $\rho \in \mathrm{L}(\mathsf{AS})$, and $\mathcal{P}^* : \mathrm{L}(\mathsf{CS}) \to \mathrm{L}(\mathsf{Z})$,

$$\mathop{\mathbb{E}}_{c\leftarrow\mathscr{C}} \langle 1|\mathcal{V}(\,|x,c\rangle\langle x,c| \otimes \pi^*_{x,c})|1\rangle < \mathsf{negl}(n)$$

  where $\pi^*_{x,c} = (\mathbb{1}_{\mathsf{A}} \otimes \mathcal{P}^*)(\rho \otimes |c\rangle\langle c|)$.

- **Special Honest Verifier Zero Knowledge.** $\exists$ an efficient algorithm Sim such that $\forall\, x \in \mathfrak{L}, c \in \mathscr{C}$,

$$\{x, c, \pi \mid \pi \leftarrow \mathsf{Sim}(x, c)\} \overset{c}{\approx} \{x, c, \pi_{x,c}\}$$

Remark 2.1. In general, Ξ-protocols are not required to be overwhelmingly complete nor negligibly sound, but we have made this choice for convenience. As we are working towards an impossibility, this will only strengthen our result. That is to say, if straightline QNICs cannot be universal for this restricted class of Ξ-protocols, then they cannot be universal for a more general definition either.

2.4 Quantum Encryption Schemes

Definition 2.6 (SKQES)**.** Let $\mathsf{QES} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ where $\mathsf{Gen} : \mathbb{N} \to \mathscr{K}$, $\mathsf{Enc} : \mathscr{K} \times \mathrm{L}(\mathsf{M}) \to \mathrm{L}(\mathsf{C})$, and $\mathsf{Dec} : \mathscr{K} \times \mathrm{L}(\mathsf{C}) \to \mathrm{L}(\mathsf{M}) \oplus \perp$ are efficient quantum algorithms. Then QES is a *symmetric-key quantum encryption scheme* if for all $k \in \mathscr{K}$

$$\|\mathsf{Dec}_k \circ \mathsf{Enc}_k - \mathbb{1}_{\mathsf{M}} \oplus \mathbf{0}_\perp\|_\diamond \leq \mathsf{negl}(n)$$

Lemma 2.2 ([AGM18], Lemma 1)**.** *Let* $\mathsf{QES} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a SKQES. Then there exists a family of unitaries* U_k*, a probability distribution* p_k*, a family of quantum states* $|\psi^{(k,r)}\rangle_{\mathsf{T}}$*, and a family of channels* $\mathcal{R}_k : \mathrm{L}(\mathsf{C}) \to \mathrm{L}(\mathsf{M})$ *such that*

$$\mathsf{Enc}_k(\rho_{\mathsf{M}}) = U_k \left(\rho \otimes \sum_r p_k(r) \, |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}| \right) U_k^\dagger$$
$$\mathsf{Dec}_k(\varrho_{\mathsf{C}}) = \mathrm{Tr}_{\mathsf{T}} \left[U_k^\dagger \Pi_k \varrho \Pi_k U_k \right] + \mathcal{R}_k \left((\mathbb{1} - \Pi_k) \varrho (\mathbb{1} - \Pi_k) \right)$$

where $\Pi_k = \sum_r \Pi_{k,r}$ *and* $\Pi_{k,r} = U_k(\mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|)U_k^\dagger$.

As it will be useful later, we will additionally denote by $\mathsf{Enc}_{k;r}$ and $\mathsf{Dec}_{k;r}$ the channels

$$\mathsf{Enc}_{k;r}(\rho_{\mathsf{M}}) = U_k(\rho \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|)U_k^\dagger$$
$$\mathsf{Dec}_{k;r}(\varrho_{\mathsf{C}}) = \mathrm{Tr}_{\mathsf{T}} \left[U_k^\dagger \Pi_{k,r} \varrho \Pi_{k,r} U_k \right] + \mathcal{R}_k \left((\mathbb{1} - \Pi_{k,r}) \varrho (\mathbb{1} - \Pi_{k,r}) \right)$$

Definition 2.7 (Quantum Ciphertext Authenticating (QCA))**.** Let $\mathsf{QES} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a SKQES. Then QES is *QCA-secure* if, for all quantum channels $\mathcal{A} : \mathrm{L}(\mathsf{CR}) \to \mathrm{L}(\mathsf{CR})$, there exists a completely-positive map $\Lambda^{\mathrm{rej}} : \mathrm{L}(\mathsf{R}) \to \mathrm{L}(\mathsf{R})$ such that $\|\Gamma - \Lambda\|_\diamond \leq \mathsf{negl}(n)$ and Λ is trace-preserving where

$$\Gamma = \mathop{\mathbb{E}}_k [\mathsf{Dec}_k \circ \mathcal{A} \circ \mathsf{Enc}_k] \qquad \text{and} \qquad \Lambda = \mathbb{1}_{\mathsf{M}} \otimes \mathop{\mathbb{E}}_{k,r} [\Lambda^{\mathrm{acc}}_{k,r}] + |\perp\rangle\langle\perp|_{\mathsf{M}} \otimes \Lambda^{\mathrm{rej}}$$

are the 'real' and 'simulated' effective attack maps, respectively, given that

$$\Lambda^{\mathrm{acc}}_{k,r}(X_{\mathsf{R}}) = \mathrm{Tr}_{\mathsf{LMT}}[\widetilde{\Pi}_{k,r} \left(\mathcal{A} \circ \mathsf{Enc}_{k;r}(\Phi^+_{\mathsf{LM}} \otimes X) \right)] \quad \text{and} \quad \widetilde{\Pi}_{k,r} = U_k(|\Phi^+\rangle\langle\Phi^+|_{\mathsf{LM}} \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|_{\mathsf{T}})U_k^\dagger.$$

Fact 2.3. *Let* $\mathsf{QES} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a SKQES. Then*

$$\Pi^\perp \mathsf{Dec}_k(\rho)\Pi^\perp = \mathsf{Dec}_k((\mathbb{1} - \Pi_k)\rho(\mathbb{1} - \Pi_k)) \qquad \textit{and} \qquad (\mathbb{1} - \Pi^\perp)\mathsf{Dec}_k(\rho)(\mathbb{1} - \Pi^\perp) = \mathsf{Dec}_k(\Pi_k \rho \Pi_k).$$

3 Modelling Quantum Non-Interactivity Compilers (QNICs)

Our goal in this section is to formalize a broad class of Fiat–Shamir–like transformations which remove interaction from public-coin proof systems. Because we are working towards an impossibility result, we will narrow our focus to compilers for three-message Ξ-protocols. As our main result applies to compilers that are universal for Ξ-protocols, it will also cover any expansion of universality — e.g., to protocols with a polynomial number of rounds.

Recall from Section 2.3 that a Ξ-protocol can be defined by a family of quantum states $\{\xi^x\}_{x \in \mathfrak{L}}$, a quantum channel $\mathcal{P}$, and an efficient quantum algorithm $\mathcal{V}$. In the most general sense, a quantum non-interactivity compiler (QNIC) should be a set of algorithms which, when given black-box access to $(\{\xi^x\}, \mathcal{P}, \mathcal{V})$, can construct a non-interactive argument system. This intuition leads to our first definition.

Definition 3.1 (Quantum Non-Interactivity Compiler (QNIC))**.** Let $\mathcal{T} = (\mathcal{T}_P, \mathcal{T}_V)$ be a pair of quantum poly-time oracle algorithms and $\mathbf{O}$ be a distribution over oracles. Then $\mathcal{T}$ is an (ε, δ)-*QNIC relative to* $\mathbf{O}$ if the following properties hold for all languages $\mathfrak{L}$ and Ξ-protocols $(\{\xi^x\}_{x \in \mathfrak{L}}, \mathcal{P}, \mathcal{V})$.

- $(1 - \varepsilon)$**-Completeness.** $\forall\, x \in \mathfrak{L}$,

$$\Pr_{O \leftarrow \mathbf{O}} \left[\mathcal{T}_V^{O,\mathcal{V}} \circ \mathcal{T}_P^{O,\mathcal{P}}(\xi^x) = 1 \right] \geq 1 - \varepsilon$$

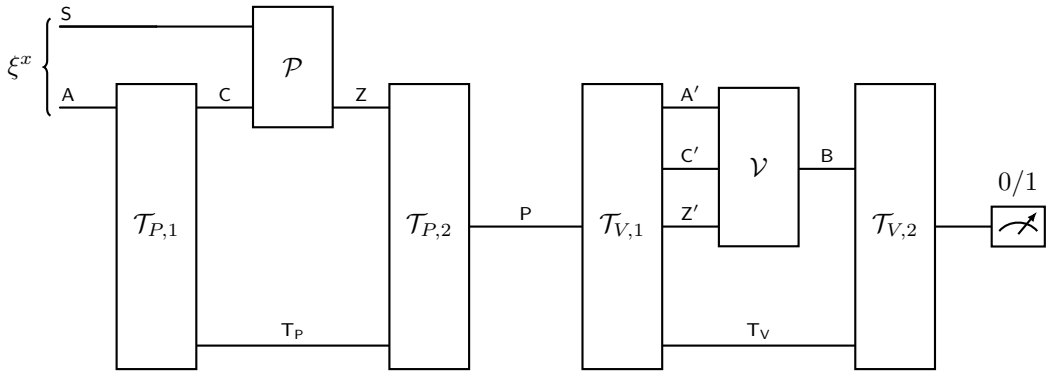


Figure 5: Circuit Representation of a Non-Interactive Argument Compiled via a Straightline QNIC

- δ**-Soundness.** $\forall\, x \notin \mathfrak{L}$ and QPT $\mathcal{A}$,

$$\Pr_{O \leftarrow \mathbf{O}}\left[\mathcal{T}_V^{O,\mathcal{V}}(\rho) = 1 \;\middle|\; \rho \leftarrow \mathcal{A}^O(x)\right] < \delta$$

In this work, we focus on a narrower set of QNICs which are "straightline" in that they execute only a single interaction with the prover and verifier. This does not restrict access to the "global" oracle O.

Definition 3.2 (Straightline QNIC)**.** Let $\mathcal{T} = (\mathcal{T}_P, \mathcal{T}_V)$ be a QNIC. Then $\mathcal{T}$ is *straightline* if $\mathcal{T}_P$ and $\mathcal{T}_V$ make only a single query to their respective $\mathcal{P}$ and $\mathcal{V}$ oracles.

Because $\mathcal{T}_P$ and $\mathcal{T}_V$ make only one query each, they can be decomposed into two channels — $(\mathcal{T}_{P,1}, \mathcal{T}_{P,2})$ and $(\mathcal{T}_{V,1}, \mathcal{T}_{V,2})$, respectively — such that (e.g.) $\mathcal{T}_P^{\mathcal{P}} = \mathcal{T}_{P,2}(\mathbb{1} \otimes \mathcal{P})\mathcal{T}_{P,1}$. In this way, an honestly compiled NIZK argument can be represented as the circuit in Figure 5 with registers labelled.

Remark 3.1. It is worth taking time to consider the quantum oracle access that we allow for our QNIC algorithms. The choice for oracle access in the quantum setting is typically unitary implementations. In our context this would allow (e.g.) $\mathcal{T}_P$ to both compute and uncompute $\mathcal{P}$. In the zero knowledge setting, allowing for inverse access to the prover may be too powerful as it could fundamentally break the zero knowledge of the compiled argument. Indeed a trivial compiler could simply extract a witness from the prover and use it directly as a non-interactive argument.

Because our results are constrained to the single-query setting, we can essentially defer these modeling questions. Under the typical black-box assumption that the oracle's private ancilla are accessible only through queries to the oracle, a single unitary query is exactly equivalent to a single corresponding channel query. Also, with only a single query, inverse access is not meaningful: the unitary implementation can be chosen such that an inverse query on fresh ancilla will act as the identity on the visible registers. Thus for straightline QNICs it suffices to consider channel implementations for $\mathcal{P}$ and $\mathcal{V}$.

Remark 3.2. Note that we do require that the global oracle O is unitary and thus allows for inverse queries. This is because the decision algorithm we construct for Theorem 1 needs to uncompute $\mathcal{T}_{V,1}^O$. Otherwise, we allow arbitrary[5] access to a polynomial number of queries to O. Thankfully this is not a particularly onerous requirement and it is satisfied by most oracle models. In particular, this captures any quantum-accessible classical function.

[5]As suggested in [Zha25], this can include conjugate and/or controlled queries as well.

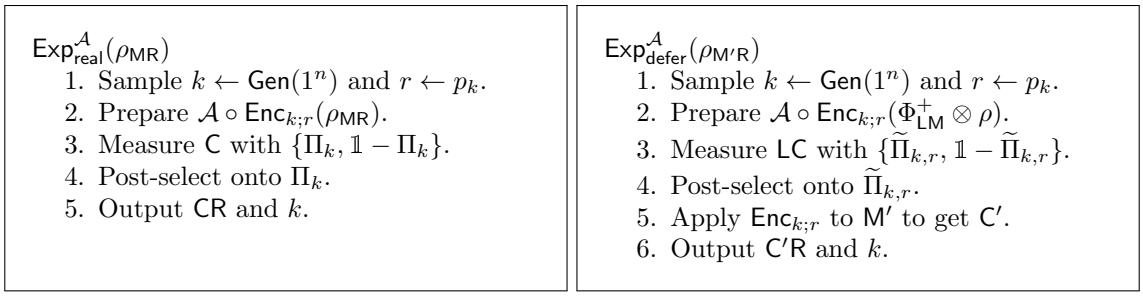


Figure 6: 'Real' and 'Deferring' Experiments

4 Ciphertext Authentication & Retrospective Security

In this section, we show that a quantum authentication scheme can be employed to defer the choice of message being encrypted until after a man-in-the-middle acts. This result is used in Section 5 to reorder the messages in the straightline interaction between the prover and the QNIC algorithms.

Let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a quantum encryption scheme. Recall from Lemma 2.2 that there exists a probability distribution p_k over that scheme's encryption randomness and projectors Π_k and $\widetilde{\Pi}_{k,r}$ which project onto valid real and simulated ciphertexts, respectively. Then we can formalize Real and Defer from Section 1.3 into subchannels $\mathsf{Exp}_{\mathsf{real}}$ and $\mathsf{Exp}_{\mathsf{defer}}$ (see Figure 6) which are parameterized by a man-in-the-middle channel $\mathcal{A} : \mathrm{L}(\mathsf{CR}) \to \mathrm{L}(\mathsf{CR})$. Notice that any action on the message register $\mathsf{M'}$ commutes with all steps prior to Step 5 in $\mathsf{Exp}_{\mathsf{defer}}$. This is exactly the deferral property we want, as it will allow us to change the message in response to $\mathcal{A}$'s actions.

With these formal subchannels defined, the rest of this section is dedicated to proving the following lemma.

Lemma 4.1. *Let* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a SKQES with QCA security and* $\mathcal{A} : \mathrm{L}(\mathsf{CR}) \to \mathrm{L}(\mathsf{CR})$ *be a quantum channel. Then* $\|\mathsf{Exp}^{\mathcal{A}}_{\mathsf{real}} - \mathsf{Exp}^{\mathcal{A}}_{\mathsf{defer}}\|_{\diamond} \leq \mathsf{negl}(n)$.

As mentioned in Section 1.3, this is accomplished through a data processing argument. That is, we demonstrate $\mathsf{Exp}_{\mathsf{real}}$ and $\mathsf{Exp}_{\mathsf{defer}}$ are indistinguishable from the real and simulated effective attack maps from Definition 2.7 composed with a fixed post-processing step. Because the post-processing done in both experiments depends on the choice of classical randomness (k, r), we need retrospective security — i.e., a property of QCA security which allows (k, r) to be revealed to the attacker upon a successful decryption. Its formal definition is given in Section 4.2 along with a proof that all QCA secure schemes satisfy the property.

The second catch in our post-processing argument comes from Step 5 of $\mathsf{Exp}_{\mathsf{defer}}$. Because the re-encryption in that step reuses the original encryption randomness r, the output will be a ciphertext entirely in the image $\Pi_{k,r}$ (i.e., the projector onto valid ciphertexts encrypted with key k and randomness r). In $\mathsf{Exp}_{\mathsf{real}}$, the ciphertext output on C will be in the image of Π_k but could potentially be partly or mostly outside of $\Pi_{k,r}$ for the original r from Step 1. In order to resolve this potential vulnerability, we prove in Section 4.1 that only negligible weight will be placed in the image of $\Pi_k - \Pi_{k,r}$. Put another way, $\mathsf{Exp}_{\mathsf{real}}$'s output is indistinguishable from the output of a modified $\mathsf{Exp}_{\mathsf{real}}$ where $\Pi_{k,r}$ is used instead of Π_k. We prove that this 'randomness preservation lemma' follows from QCA security by analyzing a "Bell projection" supermap which is implicitly used in the QCA simulation.

4.1 Randomness Preservation

Lemma 4.2. *Let* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a SKQES,* $\mathcal{A} : \mathsf{CR} \to \mathsf{CR}$ *be a channel,* Γ *be the corresponding effective attack map, and* Γ^{acc} *be the acceptance branch of that map. Then we define a modified attack map*

$$\widehat{\Gamma} = \mathop{\mathbb{E}}_{k,r} \left[\mathsf{Dec}_{k;r} \circ \mathcal{A} \circ \mathsf{Enc}_{k;r}\right]$$

where $\mathsf{Dec}_{k;r}(X)$ *accepts only ciphertexts which are encrypted with randomness* r*, and a corresponding acceptance branch* $\widehat{\Gamma}^{\mathrm{acc}}$*. If* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is QCA secure, then* $\|\Gamma^{\mathrm{acc}} - \widehat{\Gamma}^{\mathrm{acc}}\|_\diamond \leq \mathsf{negl}(n)$.

In order to prove Lemma 4.2, we start by proving a general technical lemma about subchannels and their behaviour under supermaps. In particular, we are interested in supermaps which are *retractions* in that their output always lies within a particular subspace on which the supermap acts trivially. A retraction is analogous to a projection: it is idempotent, has image M, and fixes every element of M.

Definition 4.1. Let $\mathrm{M} \subseteq \mathrm{CP}(\mathsf{A})$ be a subset of completely positive maps and $\Theta : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{A})$ be a supermap. Then Θ is a *retraction onto* M if $\Theta(\mathcal{N}) \in \mathrm{M}$ for all $\mathcal{N} \in \mathrm{CP}(\mathsf{A})$ and $\Theta(\mathcal{M}) = \mathcal{M}$ for all $\mathcal{M} \in \mathrm{M}$.

Lemma 4.3. *Let* $\mathcal{X}, \mathcal{Z} : \mathrm{L}(\mathsf{AB}) \to \mathrm{L}(\mathsf{CD})$ *be completely positive, trace-nonincreasing maps such that* $\mathcal{Z} \leq \mathcal{X}$ *and* $\Theta : \mathrm{L}(\mathsf{AB}, \mathsf{CD}) \to \mathrm{L}(\mathsf{AB}, \mathsf{CD})$ *be a retraction onto a downward-closed subset* $\mathrm{M} \subseteq \mathrm{CP}(\mathsf{AB}, \mathsf{CD})$ *of completely positive maps. If* $\|\mathcal{X} - \Theta(\mathcal{Z})\|_\diamond \leq \varepsilon$*, then* $\|\mathcal{X} - \mathcal{Z}\|_\diamond \leq \varepsilon + 4\sqrt{\varepsilon}$.

In plain language, we have a subchannel $\mathcal{X}$, a branch of that subchannel $\mathcal{Z}$, and a supermap Θ which is a retraction onto a subspace M. If $\mathcal{X}$ is close to the retraction of $\mathcal{Z}$ onto M, then $\mathcal{Z}$ must be the dominating branch of $\mathcal{X}$. If $\mathcal{X} \in \mathrm{M}$, then $\mathcal{Z} \in \mathrm{M}$ by downward closure and the result becomes trivial.

Before we prove Lemma 4.3, we give motivation by showing that it implies Lemma 4.2. In that setting, Θ becomes the "Bell-projection" described in Fact 4.1 which retracts onto the subspace of maps which act trivially on the first register. This supermap is implicit in the QCA definition as part of the transformation from the real channel Γ to the simulated channel Λ. In fact, we have that $\Theta(\widehat{\Gamma}^{\mathrm{acc}}) = \mathbb{1}_{\mathsf{M}} \otimes \mathbb{E}_{k,r}[\Lambda^{\mathrm{acc}}_{k,r}]$. We start with facts about Θ and its subspace, prove Lemma 4.2, and then prove Lemma 4.3.

Fact 4.1. *The "Bell-projection" supermap* $\Theta : \mathrm{L}(\mathsf{AB}) \to \mathrm{L}(\mathsf{AC})$ *given by*

$$\Theta(\mathcal{N})(\rho) = \mathrm{Tr}_{\mathsf{ZA}'}[(\Phi^+_{\mathsf{ZA}'} \otimes \mathbb{1}_{\mathsf{AC}})(\mathcal{N}_{\mathsf{A}'\mathsf{B}\to\mathsf{A}'\mathsf{C}} \otimes \mathbb{1}_{\mathsf{ZA}})(\Phi^+_{\mathsf{ZA}'} \otimes \rho_{\mathsf{AB}})].$$

is a retraction onto the subspace $\mathrm{M} = \{\mathbb{1}_{\mathsf{A}} \otimes \mathcal{B} : \mathcal{B} \in \mathrm{L}(\mathsf{B}, \mathsf{C})\} \subset \mathrm{L}(\mathsf{AB}, \mathsf{AC})$.

Proof. Clearly $\Theta(\mathcal{N})$ acts as the identity on A by construction for any $\mathcal{N} \in \mathrm{L}(\mathsf{AB}, \mathsf{AC})$. Additionally Θ acts trivially upon any $\mathcal{N} = \mathbb{1}_{\mathsf{A}} \otimes \mathcal{B}$ for $\mathcal{B} \in \mathrm{L}(\mathsf{B}, \mathsf{C})$.

$$\begin{aligned}
\Theta(\mathbb{1}_{\mathsf{A}} \otimes \mathcal{B})(\rho) &= \mathrm{Tr}_{\mathsf{ZA}'}[(\Phi^+_{\mathsf{ZA}'} \otimes \mathbb{1}_{\mathsf{AC}})(\mathbb{1}_{\mathsf{ZAA}'} \otimes \mathcal{B}_{\mathsf{B}\to\mathsf{C}})(\Phi^+_{\mathsf{ZA}'} \otimes \rho_{\mathsf{AB}})] \\
&= \mathrm{Tr}_{\mathsf{ZA}'}[(\mathbb{1}_{\mathsf{ZAA}'} \otimes \mathcal{B}_{\mathsf{B}\to\mathsf{C}})(\Phi^+_{\mathsf{ZA}'} \otimes \rho_{\mathsf{AB}})] \\
&= (\mathbb{1}_{\mathsf{A}} \otimes \mathcal{B}_{\mathsf{B}\to\mathsf{C}})(\rho_{\mathsf{AB}})
\end{aligned}$$

□

Fact 4.2. *The subspace* $\mathrm{M} = \{\mathbb{1}_{\mathsf{A}} \otimes \mathcal{B} : \mathcal{B} \in \mathrm{CP}(\mathsf{B}, \mathsf{C})\} \subset \mathrm{CP}(\mathsf{AB}, \mathsf{AC})$ *is downward-closed.*

Proof. Let $\mathcal{X} = \mathbb{1}_{\mathsf{A}} \otimes \mathcal{B} \in \mathrm{M}$ be a map in the subspace where $\mathcal{B} \in \mathrm{L}(\mathsf{B}, \mathsf{C})$ and $\mathcal{Z}$ be some map such that $\mathcal{Z} \leq \mathcal{X}$. Additionally, let $(V^{\mathcal{B}}, \mathrm{E})$ be a Stinespring dilation of $\mathcal{B}$ and note that $(\mathbb{1} \otimes V^{\mathcal{B}}, \mathsf{E})$ is a Stinespring dilation of $\mathcal{X}$. By Theorem 3, there exists $0 \leq Q \leq \mathbb{1}_{\mathsf{E}}$ such that

$$\mathcal{Z}(X) = \mathrm{Tr}_{\mathsf{E}}[(\mathbb{1}_{\mathsf{AC}} \otimes Q^{\frac{1}{2}})(\mathbb{1}_{\mathsf{A}} \otimes V^{\mathcal{B}})X(\mathbb{1}_{\mathsf{A}} \otimes V^{\mathcal{B}})^\dagger(\mathbb{1}_{\mathsf{AC}} \otimes Q^{\frac{1}{2}})]$$

which acts trivially on the A register. Therefore $\mathcal{Z} \in \mathrm{M}$ and M is downward-closed. □

Proof of Lemma 4.2. This follows from Lemma 4.3 where $\mathcal{X} = \Gamma^{\mathrm{acc}}$, $\mathcal{Z} = \widehat{\Gamma}^{\mathrm{acc}}$, and Θ is the Bell-projection supermap. Note that $\Theta(\widehat{\Gamma}^{\mathrm{acc}}) = \Lambda^{\mathrm{acc}}$ and thus we have the starting assumption that $\|\mathcal{X} - \Theta(\mathcal{Z})\|_\diamond \leq \mathsf{negl}(n)$. Facts 4.1 and 4.2 prove that Θ is a retraction onto a subspace which is downward-closed. It remains then to be shown that $\widehat{\Gamma}^{\mathrm{acc}} \leq \Gamma^{\mathrm{acc}}$.

For brevity, let $\widetilde{X}_{k,r} = \mathcal{A} \circ \mathsf{Enc}_{k;r}(X)$ for a fixed X and a particular (k, r). Then we start by writing Γ^{acc} in terms of its constituent components:

$$\begin{aligned}
\Gamma^{\text{acc}}(X_{\mathsf{MR}}) &= \mathop{\mathbb{E}}_{k,r}[(\mathbb{1} - \Pi^{\perp})\mathsf{Dec}_k(\widetilde{X}_{k,r})(\mathbb{1} - \Pi^{\perp})] \\
&= \mathop{\mathbb{E}}_{k,r}\left[\mathrm{Tr}_{\mathsf{T}}[U_k^{\dagger}\Pi_k\widetilde{X}_{k,r}\Pi_k U_k]\right] \\
&= \mathop{\mathbb{E}}_{k,r}\left[\mathrm{Tr}_{\mathsf{T}}\left[\left(\sum_s \mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}|\right) U_k^{\dagger}\widetilde{X}_{k,r}U_k \left(\sum_s \mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}|\right)\right]\right] \\
&= \mathop{\mathbb{E}}_{k,r}\left[\mathrm{Tr}_{\mathsf{T}}\left[\left(\sum_s \mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}|\right) U_k^{\dagger}\widetilde{X}_{k,r}U_k\right]\right]
\end{aligned}$$

where the final equality uses the cyclic property of the partial trace (and that $\{|\psi^{(k,r)}\rangle\}$ are orthogonal). Through the same process, we find that

$$\widehat{\Gamma}^{\text{acc}}(X_{\mathsf{MR}}) = \mathop{\mathbb{E}}_{k,r}\left[\mathrm{Tr}_{\mathsf{T}}\left[(\mathbb{1} \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|)U_k^{\dagger}\widetilde{X}_{k,r}U_k\right]\right]$$

Therefore

$$(\Gamma^{\text{acc}} - \widehat{\Gamma}^{\text{acc}})(X) = \mathop{\mathbb{E}}_{k,r}\left[\mathrm{Tr}_{\mathsf{T}}\left[\left(\sum_{s\neq r} \mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}|\right) U_k^{\dagger}\widetilde{X}_{k,r}U_k\right]\right] \tag{1}$$

In this way, $\Gamma^{\text{acc}} - \widehat{\Gamma}^{\text{acc}}$ is the composition of three completely positive operations: (1) the channel which encrypts the input, applies $\mathcal{A}$, and reverses the encrypting unitary U_k; (2) the projection onto the subspace given by $\sum_{s\neq r}(\mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}|_{\mathsf{T}})$; and (3) the partial trace. Because complete positivity is preserved through composition, $\Gamma^{\text{acc}} - \widehat{\Gamma}^{\text{acc}}$ is completely positive and $\widehat{\Gamma}^{\text{acc}} \leq \Gamma^{\text{acc}}$. □

Proof of Lemma 4.3. Let $\mathcal{Y} = \mathcal{X} - \mathcal{Z}$. By contractivity of the diamond norm under supermaps and the idempotence of Θ, we can bound the contribution of $\Theta(\mathcal{Y})$ to the diamond norm as

$$\|\Theta(\mathcal{Y})\|_{\diamond} = \|\Theta(\mathcal{X}) - \Theta(\mathcal{Z})\|_{\diamond} = \|\Theta(\mathcal{X}) - \Theta^2(\mathcal{Z})\|_{\diamond} \leq \|\mathcal{X} - \Theta(\mathcal{Z})\|_{\diamond} \leq \varepsilon$$

Our goal now is to find a 'bridge' map $\mathcal{B}$ which is close to both $\mathcal{Y}$ and $\Theta(\mathcal{Y})$ in diamond distance. By Theorem 2, there exist Stinespring contractions $V^{\mathcal{X}}$ and $V^{\Theta(\mathcal{Z})}$ along with an ancilla register E such that $(V^{\mathcal{X}}, \mathsf{E})$ and $(V^{\Theta(\mathcal{Z})}, \mathsf{E})$ are dilations of $\mathcal{X}$ and $\Theta(\mathcal{Z})$ respectively, and

$$\|V^{\mathcal{X}} - V^{\Theta(\mathcal{Z})}\|_{\infty} \leq \sqrt{\|\mathcal{X} - \Theta(\mathcal{Z})\|_{\diamond}} \leq \sqrt{\varepsilon}$$

Because $\mathcal{Y} \leq \mathcal{X}$, we can invoke Theorem 3 to assert the existence of an operator $0 \leq Q \leq \mathbb{1}$ such that

$$\mathcal{Y}(X) = \mathrm{Tr}_{\mathsf{E}}\left[(\mathbb{1}_{\mathsf{CD}} \otimes Q^{\frac{1}{2}})V^{\mathcal{X}}XV^{\mathcal{X}\dagger}(\mathbb{1}_{\mathsf{CD}} \otimes Q^{\frac{1}{2}})\right]$$

Using these ingredients, we construct our bridge map as

$$\mathcal{B}(X) = \mathrm{Tr}_{\mathsf{E}}\left[(\mathbb{1}_{\mathsf{CD}} \otimes Q^{\frac{1}{2}})V^{\Theta(\mathcal{Z})}XV^{\Theta(\mathcal{Z})\dagger}(\mathbb{1}_{\mathsf{CD}} \otimes Q^{\frac{1}{2}})\right]$$

Observe that $\mathcal{Y}$ and $\mathcal{B}$ have Stinespring operators $V^{\mathcal{Y}} = (\mathbb{1} \otimes Q^{\frac{1}{2}})V^{\mathcal{X}}$ and $V^{\mathcal{B}} = (\mathbb{1} \otimes Q^{\frac{1}{2}})V^{\Theta(\mathcal{Z})}$. Using Theorem 2, we can use the operator distance between these to bound the diamond norm of their respective channels.

$$\|\mathcal{Y} - \mathcal{B}\|_{\diamond} \leq 2\|V^{\mathcal{Y}} - V^{\mathcal{B}}\|_{\infty} = 2\|(\mathbb{1}_{\mathsf{CD}} \otimes Q^{\frac{1}{2}})(V^{\mathcal{X}} - V^{\Theta(\mathcal{Z})})\|_{\infty} \leq 2\|V^{\mathcal{X}} - V^{\Theta(\mathcal{Z})}\|_{\infty} \leq 2\sqrt{\varepsilon}$$

where the second inequality relies on Q being a contraction. In this way, $\mathcal{Y}$ and $\mathcal{B}$ are $2\sqrt{\varepsilon}$-close. By contractivity of the diamond norm under supermaps, we know that $\Theta(\mathcal{Y})$ and $\Theta(\mathcal{B})$ are also $2\sqrt{\varepsilon}$-close.

Our final observation is that $\mathcal{B} \leq \Theta(\mathcal{Z})$ by Theorem 3. Thus $\mathcal{B} \in \mathrm{M}$ (because M is downward-closed) and $\Theta(\mathcal{B}) = \mathcal{B}$. Putting all this together,

$$\begin{aligned}
\|\mathcal{X} - \mathcal{Z}\|_\diamond &= \|\mathcal{Y}\|_\diamond \\
&= \|\mathcal{Y} - \mathcal{B} + \mathcal{B}\|_\diamond \\
&\leq \|\mathcal{Y} - \mathcal{B}\|_\diamond + \|\mathcal{B}\|_\diamond \\
&= \|\mathcal{Y} - \mathcal{B}\|_\diamond + \|\Theta(\mathcal{B})\|_\diamond \\
&= \|\mathcal{Y} - \mathcal{B}\|_\diamond + \|\Theta(\mathcal{B}) - \Theta(\mathcal{Y}) + \Theta(\mathcal{Y})\|_\diamond \\
&\leq \|\mathcal{Y} - \mathcal{B}\|_\diamond + \|\Theta(\mathcal{B}) - \Theta(\mathcal{Y})\|_\diamond + \|\Theta(\mathcal{Y})\|_\diamond \\
&\leq 4\sqrt{\varepsilon} + \varepsilon
\end{aligned}$$

□

4.2 Retrospective Security

In order to make a data processing argument where the post-processing channel depends on the encryption randomness, we introduce a property for quantum encryption keys we call "retrospective security."

Definition 4.2 (Retrospective Security)**.** Let $\mathsf{QES} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a SKQES and define the parameterized channel $\mathsf{Reveal}_{k,r} : \mathrm{L}(\mathsf{M} \oplus \perp) \to \mathrm{L}(\mathsf{MK})$ as

$$\mathsf{Reveal}_{k,r}(\rho_{\mathsf{M}}) = (\mathbb{1} - \Pi^\perp)\rho(\mathbb{1} - \Pi^\perp) \otimes |k,r\rangle\langle k,r|_{\mathsf{K}} + \Pi^\perp \rho \Pi^\perp \otimes |\perp\rangle\langle\perp|_{\mathsf{K}}\,.$$

Then QES is QCA with *retrospective security* if, for all channels $\mathcal{A} : \mathrm{L}(\mathsf{CR}) \to \mathrm{L}(\mathsf{CR})$, there exists a completely positive map $\Lambda_{\mathrm{R}}^{\mathrm{rej}} : \mathrm{L}(\mathsf{R}) \to \mathrm{L}(\mathsf{R})$ such that $\|\Gamma_{\mathrm{R}} - \Lambda_{\mathrm{R}}\|_\diamond \leq \mathsf{negl}(n)$ and Λ_{R} is trace-preserving where

$$\Gamma_{\mathrm{R}} = \mathop{\mathbb{E}}_{k,r}\left[\mathsf{Reveal}_{k,r} \circ \mathsf{Dec}_k \circ \mathcal{A} \circ \mathsf{Enc}_{k;r}\right] \qquad \text{and} \qquad \Lambda_{\mathrm{R}} = \mathbb{1}_{\mathsf{M}} \otimes \mathop{\mathbb{E}}_{k,r}\left[\,|k,r\rangle\langle k,r|_{\mathsf{K}} \otimes \Lambda_{k,r}^{\mathrm{acc}}\right] + |\perp\rangle\langle\perp|_{\mathsf{MK}} \otimes \Lambda_{\mathrm{R}}^{\mathrm{rej}}$$

given that $\Lambda_{k,r}^{\mathrm{acc}}$ is as in Definition 2.7.

Before we show that all QCA-secure schemes satisfy this definition, we first show that the randomness preservation of Lemma 4.2 extends to the retrospective setting where (k, r) are included in the output.

Corollary 4.1. *Let* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be a SKQES,* $\mathcal{A} : \mathsf{CR} \to \mathsf{CR}$ *be a channel,* Γ_{R} *be as in Definition 4.2, and* $\Gamma_{\mathrm{R}}^{\mathrm{acc}}$ *be the acceptance branch of that map. Let* $\widehat{\Gamma}_{\mathrm{R}} = \mathbb{E}_{k,r}[\mathsf{Reveal}_{k,r} \circ \mathsf{Dec}_{k;r} \circ \mathcal{A} \circ \mathsf{Enc}_{k;r}]$ *where* $\mathsf{Dec}_{k;r}$ *is as in Lemma 4.2. If* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is QCA secure then* $\|\Gamma_{\mathrm{R}}^{\mathrm{acc}} - \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}}\|_\diamond \leq \mathsf{negl}(n)$.

Proof. By following the same argument as in Lemma 4.2, it can be seen that $\widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}} \leq \Gamma_{\mathrm{R}}^{\mathrm{acc}}$ and therefore $\Gamma_{\mathrm{R}}^{\mathrm{acc}} - \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}}$ is completely positive. By Fact 2.2 and Lemma 4.2,

$$\|\Gamma_{\mathrm{R}}^{\mathrm{acc}} - \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}}\|_\diamond = \|\mathrm{Tr}_{\mathsf{K}} \circ \Gamma_{\mathrm{R}}^{\mathrm{acc}} - \mathrm{Tr}_{\mathsf{K}} \circ \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}}\|_\diamond = \|\Gamma^{\mathrm{acc}} - \widehat{\Gamma}^{\mathrm{acc}}\|_\diamond \leq \mathsf{negl}(n). \qquad \square$$

Lemma 4.4. *If* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is an SKQES with QCA security then it has retrospective security as well.*

We begin by stating an operator-theoretic technical lemma about "classical extensions" of quantum channels. Given a channel $\mathcal{N} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{B})$, we could describe a classical extension of $\mathcal{N}$ as any channel $\mathcal{N}' : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BC})$ such that $\mathrm{Tr}_{\mathsf{C}} \circ \mathcal{N}' = \mathcal{N}$ and the output of $\mathcal{N}'$ on C is always classical. In Lemma 4.5, we show that given any channels $\mathcal{X}$ and $\mathcal{Y}$ which are close (in diamond distance) and any classical extension of $\mathcal{X}$, there exists some corresponding classical extension of $\mathcal{Y}$ such that the extensions are also close.

Lemma 4.5. *Let* $\Theta : \mathrm{L}(\mathsf{A},\mathsf{B}) \to \mathrm{L}(\mathsf{A},\mathsf{B})$ *be a retraction onto a downward-closed subset* $\mathrm{M} \subseteq \mathrm{CP}(\mathsf{A},\mathsf{B})$ *of completely positive maps and let* $\mathcal{X},\mathcal{Y} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{B})$ *be completely positive, trace-nonincreasing maps such that* $\Theta(\mathcal{Y}) = \mathcal{Y}$ *and* $\|\mathcal{X} - \mathcal{Y}\|_\diamond \leq \varepsilon$*. Additionally, let* $\mathcal{X}' : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BC})$ *be a classical extension of* $\mathcal{X}$ *in that* $\mathcal{X}' = \sum_c \mathcal{X}_c \otimes |c\rangle\langle c|_{\mathsf{C}}$ *where* $\sum_c \mathcal{X}_c = \mathcal{X}$*.*

Then there exists $\mathcal{Y}' : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{BC})$ *such that* $(\Theta \otimes \mathbb{1}_{\mathsf{C}})(\mathcal{Y}') = \mathcal{Y}'$ *and* $\|\mathcal{X}' - \mathcal{Y}'\|_\diamond \leq 2\sqrt{\varepsilon}$*.*

Before we prove Lemma 4.5, we give motivation by showing that it implies Lemma 4.4. Given our context, the retrospective effective attack maps can be understood as classical extensions of the original QCA attack maps. In this way, we can use Lemma 4.5 to show that there exists some corresponding extension of Λ which must be close to Γ_{R}. The rest of the proof is showing that extension is close to the constructed Λ_{R}.

Proof of Lemma 4.4. Fix an attack $\mathcal{A}$ and let Γ and Λ be the corresponding real and simulated effective attack maps given by the scheme's QCA security. Observe that Dec performs an explicit accept-or-reject measurement so its output is a classical mixture of the two orthogonal cases. Thus the real map can be decomposed as $\Gamma = \Gamma^{\mathrm{acc}} + \Gamma^{\mathrm{rej}}$ where $\Gamma^{\mathrm{acc}}(\rho) = (\mathbb{1} - \Pi^{\perp})\Gamma(\rho)(\mathbb{1} - \Pi^{\perp})$ and $\Gamma^{\mathrm{rej}}(\rho) = \Pi^{\perp}\Gamma(\rho)\Pi^{\perp}$. By data processing,

$$\left\|\Gamma^{\mathrm{acc}} - \mathbb{1}_{\mathsf{M}} \otimes \mathop{\mathbb{E}}_{k,r}[\Lambda^{\mathrm{acc}}_{k,r}]\right\|_\diamond \leq \mathsf{negl}(n) \qquad \text{and} \qquad \left\|\Gamma^{\mathrm{rej}} - |\perp\rangle\langle\perp|_{\mathsf{M}} \otimes \Lambda^{\mathrm{rej}}\right\|_\diamond \leq \mathsf{negl}(n). \tag{2}$$

In the same way, we can define the real retrospective attack map $\Gamma_{\mathrm{R}} = \Gamma^{\mathrm{acc}}_{\mathrm{R}} + \Gamma^{\mathrm{rej}}_{\mathrm{R}}$ where $\Gamma^{\mathrm{acc}}_{\mathrm{R}}(\rho) = (\mathbb{1} - \Pi^{\perp})\Gamma_{\mathrm{R}}(\rho)(\mathbb{1} - \Pi^{\perp})$ and $\Gamma^{\mathrm{rej}}_{\mathrm{R}}(\rho) = \Pi^{\perp}\Gamma_{\mathrm{R}}(\rho)\Pi^{\perp}$, and the simulated retrospective attack map Λ_{R} where $\Lambda^{\mathrm{rej}}_{\mathrm{R}} = \Lambda^{\mathrm{rej}}$. By the triangle inequality,

$$\|\Gamma_{\mathrm{R}} - \Lambda_{\mathrm{R}}\|_\diamond \leq \|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathbb{1}_{\mathsf{M}} \otimes \mathop{\mathbb{E}}_{k,r}[\Lambda^{\mathrm{acc}}_{k,r} \otimes |k,r\rangle\langle k,r|_{\mathsf{K}}]\|_\diamond + \|\Gamma^{\mathrm{rej}}_{\mathrm{R}} - |\perp\rangle\langle\perp|_{\mathsf{MK}} \otimes \Lambda^{\mathrm{rej}}_{\mathrm{R}}\|_\diamond$$

Because $\Omega^{\mathrm{rej}}_{\mathrm{R}} = (\mathbb{1}_{\mathsf{M}} \otimes |\perp\rangle\langle\perp|_{\mathsf{K}}) \circ \Omega^{\mathrm{rej}}$ for $\Omega \in \{\Gamma, \Lambda\}$ — i.e., the rejection case just composes with a fixed state preparation channel — we can easily bound

$$\|\Gamma^{\mathrm{rej}}_{\mathrm{R}} - |\perp\rangle\langle\perp|_{\mathsf{MK}} \otimes \Lambda^{\mathrm{rej}}_{\mathrm{R}}\|_\diamond \leq \|(\mathbb{1}_{\mathsf{MR}} \otimes |\perp\rangle\langle\perp|_{\mathsf{K}}) \circ (\Gamma^{\mathrm{rej}} - |\perp\rangle\langle\perp|_{\mathsf{M}} \otimes \Lambda^{\mathrm{rej}})\|_\diamond \leq \|\Gamma^{\mathrm{rej}} - |\perp\rangle\langle\perp|_{\mathsf{M}} \otimes \Lambda^{\mathrm{rej}}\|_\diamond \leq \mathsf{negl}(n).$$

The difficulty is in bounding the acceptance term. The main tool for that will be invoking Lemma 4.5 where $\mathcal{X} = \Gamma^{\mathrm{acc}}$, $\mathcal{X}' = \Gamma^{\mathrm{acc}}_{\mathrm{R}}$, $\mathcal{Y} = \mathbb{1}_{\mathsf{M}} \otimes \mathbb{E}_{k,r}[\Lambda^{\mathrm{acc}}_{k,r}]$, and Θ is the Bell-projection supermap. Equation (2) has established that $\|\mathcal{X} - \mathcal{Y}\|_\diamond \leq \mathsf{negl}(n)$, Facts 4.1 and 4.2 prove that Θ is a retraction onto the downward-closed subset $\mathrm{M} = \{\mathbb{1}_{\mathsf{M}} \otimes \mathcal{B} : \mathcal{B} \in \mathrm{L}(\mathsf{R})\}$, and clearly $\mathcal{Y} \in \mathrm{M}$. Thus the lemma yields a $\mathcal{Y}'$ such that $\|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathcal{Y}'\|_\diamond \leq \mathsf{negl}(n)$ and $\Theta(\mathcal{Y}') = \mathcal{Y}'$ which can be used as follows.

$$\begin{aligned}
\|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \Lambda^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond &\leq \|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathcal{Y}' + \mathcal{Y}' - \Theta(\Gamma^{\mathrm{acc}}_{\mathrm{R}}) + \Theta(\Gamma^{\mathrm{acc}}_{\mathrm{R}}) - \Lambda^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond \\
&\leq \|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathcal{Y}'\|_\diamond + \|\mathcal{Y}' - \Theta(\Gamma^{\mathrm{acc}}_{\mathrm{R}})\|_\diamond + \|\Theta(\Gamma^{\mathrm{acc}}_{\mathrm{R}}) - \Lambda^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond \\
&= \|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathcal{Y}'\|_\diamond + \|\Theta(\mathcal{Y}') - \Theta(\Gamma^{\mathrm{acc}}_{\mathrm{R}})\|_\diamond + \|\Theta(\Gamma^{\mathrm{acc}}_{\mathrm{R}}) - \Theta(\widehat{\Gamma}^{\mathrm{acc}}_{\mathrm{R}})\|_\diamond \\
&\leq 2\|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathcal{Y}'\|_\diamond + \|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \widehat{\Gamma}^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond \\
&\leq \mathsf{negl}(n)
\end{aligned}$$

where the equality used the identity $\Lambda^{\mathrm{acc}}_{\mathrm{R}} = \Theta(\widehat{\Gamma}^{\mathrm{acc}}_{\mathrm{R}})$ and the final inequality follows from Corollary 4.1. □

Proof of Lemma 4.5. Let $(V^{\mathcal{X}}, \mathsf{E})$ be a Stinespring dilation of $\mathcal{X}'$ so that $\mathcal{X}'(\rho) = \mathrm{Tr}_{\mathsf{E}}[V^{\mathcal{X}}\rho(V^{\mathcal{X}})^\dagger]$. Note that $\mathcal{X}(\rho) = \mathrm{Tr}_{\mathsf{CE}}[V^{\mathcal{X}}\rho(V^{\mathcal{X}})^\dagger]$ because $\mathrm{Tr}_{\mathsf{C}} \circ \mathcal{X}' = \mathcal{X}$, and therefore $(V^{\mathcal{X}}, \mathsf{CE})$ is a valid Stinespring dilation of $\mathcal{X}$ as well. By Theorem 2, there exists an operator $V^{\mathcal{Y}}$ such that $(V^{\mathcal{Y}}, \mathsf{CE})$ is a Stinespring dilation of $\mathcal{Y}$ (expanding E if necessary) and

$$\|V^{\mathcal{X}} - V^{\mathcal{Y}}\|_\infty \leq \sqrt{\|\mathcal{X} - \mathcal{Y}\|_\diamond} \leq \sqrt{\varepsilon}$$

Let $\mathcal{Y}'_0(\rho) = \mathrm{Tr}_{\mathsf{E}}[V^{\mathcal{Y}}\rho(V^{\mathcal{Y}})^\dagger]$ be the natural extension of $\mathcal{Y}$ to C suggested by its dilation $V^{\mathcal{Y}}$ and $\mathcal{Y}' = \Delta_{\mathsf{C}} \circ \mathcal{Y}'_0$ where we denote by $\Delta_{\mathsf{C}}(\rho) = \sum_c |c\rangle\langle c|\,\rho\,|c\rangle\langle c|$ the dephasing channel on the C register. Because $\mathcal{X}'$ is a classical extension, $\Delta_{\mathsf{C}} \circ \mathcal{X}' = \mathcal{X}'$ and thus

$$\|\mathcal{X}' - \mathcal{Y}'\|_\diamond = \|\Delta_{\mathsf{C}} \circ (\mathcal{X}' - \mathcal{Y}'_0)\|_\diamond \leq \|\mathcal{X}' - \mathcal{Y}'_0\|_\diamond \leq 2\|V^{\mathcal{X}} - V^{\mathcal{Y}}\|_\infty \leq 2\sqrt{\varepsilon}$$

where we've invoked the upper bound from Theorem 2.

It remains to be shown that $(\Theta \otimes \mathbb{1}_{\mathsf{C}})(\mathcal{Y}') = \mathcal{Y}'$. If we define $\mathcal{Y}_c(\rho) = (\mathbb{1}_{\mathsf{B}} \otimes \langle c|_{\mathsf{C}})\mathcal{Y}'(\rho)(\mathbb{1}_{\mathsf{B}} \otimes |c\rangle_{\mathsf{C}})$, then we can decompose $\mathcal{Y}'(\rho) = \sum_c \mathcal{Y}_c(\rho) \otimes |c\rangle\langle c|_{\mathsf{C}}$. Because $\mathrm{Tr}_{\mathsf{C}} \circ \mathcal{Y}' = \mathcal{Y}$, we have that $\mathcal{Y} = \sum_c \mathcal{Y}_c$ and therefore $\mathcal{Y}_c \leq \mathcal{Y}$ for all c. M is downward-closed and $\mathcal{Y} \in \mathrm{M}$ so $\mathcal{Y}_c \in \mathrm{M}$ for all c. Thus we arrive at

$$(\Theta \otimes \mathbb{1}_{\mathsf{C}})(\mathcal{Y}')(\rho) = \sum_c \Theta(\mathcal{Y}_c)(\rho) \otimes |c\rangle\langle c| = \sum_c \mathcal{Y}_c(\rho) \otimes |c\rangle\langle c| = \mathcal{Y}'(\rho)$$

□

4.3 Deferral Lemma

Proof of Lemma 4.1. Fix $\mathcal{A}$. Then we prove our result through a hybrid argument starting with Hyb_1.

$\mathsf{Hyb}_1(\rho)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\mathcal{A} \circ \mathsf{Enc}_{k;r}(\rho_{\mathsf{MR}})$
3. Measure C with $\{\Pi_k, \mathbb{1} - \Pi_k\}$ and post-select onto Π_k.
4. Decrypt C with Dec_k yielding register M'.
5. Encrypt M' with $\mathsf{Enc}_{k;r}$ yielding register C'.
6. Output $\mathsf{C}'\mathsf{R}$ and k.

Claim 4.1. $\|\mathsf{Exp}^{\mathcal{A}}_{\mathsf{defer}} - \mathsf{Hyb}_1\|_\diamond \leq \mathsf{negl}(n)$

Proof. Observe that Hyb_1 and $\mathsf{Exp}^{\mathcal{A}}_{\mathsf{defer}}$ resemble the real and simulated effective attack maps from Definition 4.2. Our goal then is to show that the retrospective security of the SKQES implies our bound. Let $\Gamma^{\mathrm{acc}}_{\mathrm{R}}$ and $\Lambda^{\mathrm{acc}}_{\mathrm{R}}$ be the acceptance branches of the real and simulated channels from Definition 4.2, respectively. By Lemma 4.4, $\|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \Lambda^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond \leq \mathsf{negl}(n)$. Then we define a post-processing channel $\mathcal{N}$ which acts on the output of $\Gamma^{\mathrm{acc}}_{\mathrm{R}}$ or $\Lambda^{\mathrm{acc}}_{\mathrm{R}}$ in the following way:

1. Measure M using $\{\Pi^\perp, \mathbb{1} - \Pi^\perp\}$. Post-select onto $\mathbb{1} - \Pi^\perp$.
2. Measure K in the standard basis to get (k, r).
3. Re-encrypt the M register with $\mathsf{Enc}_{k;r}$ yielding C'.
4. Output $\mathsf{C}'\mathsf{R}$ and k.

By Fact 2.3 (and an analogous fact for the $\widetilde{\Pi}_{k,r}$ projector and the simulated decryption channel), we have that $\mathsf{Exp}^{\mathcal{A}}_{\mathsf{defer}} = \mathcal{N} \circ \Lambda^{\mathrm{acc}}_{\mathrm{R}}$ and $\mathsf{Hyb}_1 = \mathcal{N} \circ \Gamma^{\mathrm{acc}}_{\mathrm{R}}$. Thus we have our desired bound,

$$\|\mathsf{Exp}^{\mathcal{A}}_{\mathsf{defer}} - \mathsf{Hyb}_1\|_\diamond = \|\mathcal{N} \circ \Gamma^{\mathrm{acc}}_{\mathrm{R}} - \mathcal{N} \circ \Lambda^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond \leq \|\Gamma^{\mathrm{acc}}_{\mathrm{R}} - \Lambda^{\mathrm{acc}}_{\mathrm{R}}\|_\diamond \leq \mathsf{negl}(n)$$

□

$\mathsf{Hyb}_2(\rho)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\mathcal{A} \circ \mathsf{Enc}_{k;r}(\rho_{\mathsf{MR}})$
3. Measure C with $\{\Pi_{k,r}, \mathbb{1} - \Pi_{k,r}\}$ and post-select onto $\Pi_{k,r}$.
4. Decrypt C with Dec_k yielding register M'.
5. Encrypt M' with $\mathsf{Enc}_{k;r}$ yielding register C'.
6. Output $\mathsf{C}'\mathsf{R}$ and k.

Claim 4.2. $\|\mathsf{Hyb}_1 - \mathsf{Hyb}_2\|_\diamond \leq \mathsf{negl}(n)$

Proof. Let $\Gamma_{\mathrm{R}}^{\mathrm{acc}}$ and $\widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}}$ be the subchannels as in Corollary 4.1 for the fixed attack $\mathcal{A}$, and let $\mathcal{N}$ be the post-processing map which does the following:

1. Measure K to get outcome (k, r).
2. Apply $\mathsf{Enc}_{k;r}$ to M' to get C'.
3. Output $\mathsf{C}'\mathsf{R}$ and k.

Similar to Fact 2.3, notice that $\mathsf{Dec}_k(\Pi_{k,r}\rho\Pi_{k,r}) = (\mathbb{1} - \Pi^\perp)\mathsf{Dec}_{k;r}(\rho)(\mathbb{1} - \Pi^\perp)$. Therefore $\mathsf{Hyb}_1 = \mathcal{N} \circ \Gamma_{\mathrm{R}}^{\mathrm{acc}}$ and $\mathsf{Hyb}_2 = \mathcal{N} \circ \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}}$. By Corollary 4.1,

$$\|\mathsf{Hyb}_1 - \mathsf{Hyb}_2\|_\diamond = \|\,\mathcal{N} \circ \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}} - \mathcal{N} \circ \Gamma_{\mathrm{R}}^{\mathrm{acc}}\|_\diamond \leq \|\widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}} - \Gamma_{\mathrm{R}}^{\mathrm{acc}}\|_\diamond \leq \mathsf{negl}(n) \qquad \square$$

Claim 4.3. $\|\mathsf{Hyb}_2 - \mathsf{Exp}_{\mathsf{real}}^{\mathcal{A}}\|_\diamond \leq \mathsf{negl}(n)$

Proof. First notice that Steps 4 and 5 act as the identity in Hyb_2 (because of the post-selection in Step 3) and so can be omitted. Then the only remaining difference is in projectors — i.e., $\Pi_{k,r}$ versus Π_k. Therefore we want to prove for all states ρ_{MRE}, where E is an arbitrary reference register on which all maps act trivially,

$$\left\| \mathop{\mathbb{E}}_{k,r} \left[\Pi_k(\mathcal{A} \circ \mathsf{Enc}_{k;r}(\rho))\Pi_k \otimes |k\rangle\langle k|_{\mathsf{K}}\right] - \mathop{\mathbb{E}}_{k,r} \left[\Pi_{k,r}(\mathcal{A} \circ \mathsf{Enc}_{k;r}(\rho))\Pi_{k,r} \otimes |k\rangle\langle k|_{\mathsf{K}}\right] \right\|_1 \leq \mathsf{negl}(n).$$

Fix ρ and let $\sigma_{k,r} = \mathcal{A} \circ \mathsf{Enc}_{k;r}(\rho)$ and $\Pi_{k,\neq r} = \sum_{s \neq r} \Pi_{k,s}$. Then we seek to bound the trace norm of

$$\begin{aligned}
&\mathop{\mathbb{E}}_{k,r} \left[(\Pi_k \sigma_{k,r} \Pi_k - \Pi_{k,r}\sigma_{k,r}\Pi_{k,r}) \otimes |k\rangle\langle k|_{\mathsf{K}}\right] \\
&\quad = \mathop{\mathbb{E}}_{k,r} \left[(\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,\neq r} + \Pi_{k,r}\sigma_{k,r}\Pi_{k,\neq r} + \Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,r}) \otimes |k\rangle\langle k|_{\mathsf{K}}\right]
\end{aligned}$$

Adapting Equation (1) to the retrospective setting and observing $\sum_{s \neq r} \mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}| = U_k^\dagger \Pi_{k,\neq r} U_k$,

$$\begin{aligned}
\|(\Gamma_{\mathrm{R}}^{\mathrm{acc}} - \widehat{\Gamma}_{\mathrm{R}}^{\mathrm{acc}})(\rho)\|_1 &= \left\| \mathop{\mathbb{E}}_{k,r} \mathrm{Tr}_{\mathsf{T}} \left[\left(\sum_{s \neq r} \mathbb{1}_{\mathsf{M}} \otimes |\psi^{(k,s)}\rangle\langle\psi^{(k,s)}| \right) U_k^\dagger \sigma_{k,r} U_k \right] \otimes |k,r\rangle\langle k,r|_{\mathsf{K}} \right\|_1 \\
&= \left\| \mathop{\mathbb{E}}_{k,r} \mathrm{Tr}_{\mathsf{T}} \left[U_k^\dagger \Pi_{k,\neq r} \sigma_{k,r} \Pi_{k,\neq r} U_k \right] \otimes |k,r\rangle\langle k,r|_{\mathsf{K}} \right\|_1 \\
&= \left\| \mathop{\mathbb{E}}_{k,r} \left[\Pi_{k,\neq r} \sigma_{k,r} \Pi_{k,\neq r} \otimes |k\rangle\langle k|_{\mathsf{K}} \right] \right\|_1
\end{aligned}$$

The last equality holds because both operators are positive and have the same trace. By Corollary 4.1, we know that $\| \mathbb{E}_{k,r} [\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,\neq r} \otimes |k\rangle\langle k|_{\mathsf{K}}] \|_1 \leq \mathsf{negl}(n)$, but the cross terms are not accounted for in that setting because decryption traces out the tag register T. By Hölder's inequality,

$$\|\Pi_{k,r}\sigma_{k,r}\Pi_{k,\neq r}\|_1 \leq \|\Pi_{k,r}\sqrt{\sigma_{k,r}}\|_2 \cdot \|\sqrt{\sigma_{k,r}}\,\Pi_{k,\neq r}\|_2 \leq \sqrt{\mathrm{Tr}[\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,\neq r}]}$$

By convexity of the trace norm, $\|X \otimes |k\rangle\langle k|\|_1 = \|X\|_1$, and Jensen's inequality, we have that

$$\begin{aligned}
\left\| \mathop{\mathbb{E}}_{k,r} [\Pi_{k,r}\sigma_{k,r}\Pi_{k,\neq r} \otimes |k\rangle\langle k|_{\mathsf{K}}] \right\|_1 &\leq \mathop{\mathbb{E}}_{k,r} \|\Pi_{k,r}\sigma_{k,r}\Pi_{k,\neq r}\|_1 \\
&\leq \sqrt{\mathop{\mathbb{E}}_{k,r} \mathrm{Tr}[\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,\neq r}]} \\
&= \sqrt{\mathop{\mathbb{E}}_{k,r} \mathrm{Tr}[\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,\neq r} \otimes |k\rangle\langle k|_{\mathsf{K}}]} \\
&= \left\| \mathop{\mathbb{E}}_{k,r} [\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,\neq r} \otimes |k\rangle\langle k|_{\mathsf{K}}] \right\|_1^{\frac{1}{2}} \\
&\leq \mathsf{negl}(n)
\end{aligned}$$

The same argument can be applied to the $\Pi_{k,\neq r}\sigma_{k,r}\Pi_{k,r}$ term. Because all three terms are negligible, our claim follows directly from the triangle inequality. □

Combining Claims 4.1 – 4.3 through the triangle inequality, we have our result. □

5 Impossibility of Straightline QNICs

Now that we have a formal definition settled, we can state a precise version of Theorem 1.

Theorem 1. *Assume there exist quantum commitments relative to the unitary oracle distribution* $\mathbf{O}$ *and let* $(\mathcal{T}_P, \mathcal{T}_V)$ *be an* (ε, δ)*-straightline QNIC with respect to* $\mathbf{O}$ *such that* $\varepsilon(n)+2\delta(n) < \kappa^* - \kappa$ *for* $\kappa^* = 1 - \frac{2}{3\sqrt{3}} \approx \frac{3}{5}$ *and some* $\kappa > 0$*. Then there exists a* QMA*-complete language* $\mathfrak{L}$ *and an efficient algorithm* $\mathcal{A}$ *such that*

$$\begin{aligned}
&\Pr_{O\leftarrow\mathbf{O}}[\mathcal{A}^O(x) = 1] \geq 2\delta + \kappa - \mathsf{negl}(n) && \text{for all } x \in \mathfrak{L} \\
&\Pr_{O\leftarrow\mathbf{O}}[\mathcal{A}^O(x) = 1] < 2\delta + \mathsf{negl}(n) && \text{for all } x \notin \mathfrak{L}
\end{aligned}$$

In particular, if $\kappa \geq 1/\mathsf{poly}(n)$ *and* $\mathbf{O}$ *is efficiently simulatable (as in Definition 2.3), then* $\mathsf{QMA} = \mathsf{BQP}$*.*

In the rest of this section, we work to prove Theorem 1. In Section 5.1, we describe the family of counterexample Ξ-protocols which we use to construct the decision algorithm. Section 5.2 lays out the components for the compiled non-interactive argument and defines some additional notation for use in later sections. In Sections 5.3 and 5.4, we construct two subchannels $\mathcal{A}_1$ and $\mathcal{A}_2$ which capture the contributions from invalid and valid ciphertexts, respectively. Finally, Section 5.5 combines these subchannels into the decision algorithm $\mathcal{A}$ and proves its completeness and soundness.

5.1 The Counterexample Protocol

Let $(\{\xi^x\}_{x\in\mathfrak{L}}, \mathcal{P}, \mathcal{V})$ be a Ξ-protocol for a QMA-complete language $\mathfrak{L}$ with $(1-\mathsf{negl}(n))$-completeness, $\mathsf{negl}(n)$-soundness, and a SHVZK simulator Sim — e.g., a parallel repetition of the Ξ-protocol from [BG20] which exists if we assume quantum commitments. Let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a SKQES with QCA security. In order to align notation between these protocols, let M denote the message register of ξ^x and A denote a ciphertext register so $\mathsf{Enc} : \mathrm{L}(\mathsf{M}) \to \mathrm{L}(\mathsf{A})$ and $\mathsf{Dec} : \mathrm{L}(\mathsf{A}) \to \mathrm{L}(\mathsf{M} \oplus \bot)$. Thus $\xi^x \in \mathrm{L}(\mathsf{MS})$ and the prover and verifier

are channels $\mathcal{P} : \mathrm{L}(\mathsf{CS}) \to \mathrm{L}(\mathsf{Z})$ and $\mathcal{V} : \mathrm{L}(\mathsf{MCZ}) \to \mathrm{L}(\mathsf{B})$. Without loss of generality, we can assume that both $\mathcal{P}$ and $\mathcal{V}$ measure the C register immediately. Although it is technically possible to query them in superposition, any coherence will be lost. Finally, let $\mathcal{S} : \mathrm{L}(\mathsf{C}) \to \mathrm{L}(\mathsf{MZ})$ be a quantum channel corresponding to the Ξ-protocol's SHVZK simulator Sim — i.e., for a fixed x,

$$\mathcal{S}(\rho_{\mathsf{C}}) = \sum_{c \in \mathscr{C}} \langle c|\rho|c\rangle \otimes \mathsf{Sim}(x, c).$$

Given these ingredients, we build a family of Ξ-protocols $\{(\xi_k^x, \mathcal{P}, \mathcal{V}_k)\}$ indexed by a secret key k as follows. The state $\xi_k^x = (\mathsf{Enc}_k \otimes \mathbb{1}_{\mathsf{S}})(\xi^x) \in \mathrm{L}(\mathsf{AS})$ is updated such that its message register is encrypted under the parameterized key. The prover $\mathcal{P}$ is unchanged. On input ρ_{ACZ}, the counterexample verifier $\mathcal{V}_k$ performs the following steps:

1. Apply $\mathsf{Dec}_k \otimes \mathbb{1}_{\mathsf{CZ}}$ to ρ.
2. Measure M with $\{\Pi^{\perp}, \mathbb{1} - \Pi^{\perp}\}$. Reject if the outcome is $\Pi^{\perp}$.
3. Apply $\mathcal{V}$ and forward its output.

Using Fact 2.3, we can write this channel as $\mathcal{V}_k(\rho_{\mathsf{ACZ}}) = \mathcal{V} \circ \mathsf{Dec}_k(\Pi_k \rho \Pi_k) + \mathrm{Tr}\left[(\mathbb{1} - \Pi_k)\rho\right] |0\rangle\langle 0|$.

5.2 The Compiled Counterexample NIZK

Before we describe our decision algorithm, we must first lay down a few ingredients starting with the QNIC algorithms. Recall that $(\mathcal{T}_P, \mathcal{T}_V)$ are single-query oracle algorithms where $\mathcal{T}_P$ prepares the non-interactive argument and $\mathcal{T}_V$ verifies it. Because $\mathcal{T}_V$ makes only one query, it can be decomposed into $\mathcal{T}_{V,1}$ and $\mathcal{T}_{V,2}$ such that $\mathcal{T}_V^{\mathcal{V}} = \mathcal{T}_{V,2}(\mathbb{1} \otimes \mathcal{V})\mathcal{T}_{V,1}$. Additionally, for a fixed oracle O, let $(T_{V,1}, |0\rangle_{\mathsf{E}})$ be a Stinespring dilation of $\mathcal{T}_{V,1}^{O}$ and we denote by $\Pi^{T_V} = T_{V,1}(\mathbb{1} \otimes |0\rangle\langle 0|_{\mathsf{E}})T_{V,1}^{\dagger}$ the projector onto the subspace of $T_{V,1}$'s image. Given a fixed oracle O and dilation $T_{V,1}$, we denote the sequential composition of $\mathcal{T}_P$ with $T_{V,1}$ as

$$\mathcal{T}_{P+V,1}^{(\cdot)}(\rho_{\mathsf{A}}) = T_{V,1}\left(\mathcal{T}_P^{(O,\cdot)}(\rho) \otimes |0\rangle\langle 0|_{\mathsf{E}}\right) T_{V,1}^{\dagger}.$$

Let $\sigma_k = \mathcal{T}_{P+V,1}^{\mathcal{P}}(\xi_k^x)$ be the state right before $\mathcal{T}_V$'s query to $\mathcal{V}_k$ and $\{M_k^{\mathrm{acc}}, M_k^{\mathrm{rej}}\}$ be the measurement for whether that state is both valid (i.e., in the $T_{V,1}$ subspace) and accepted by $\mathcal{T}_V$ — i.e.,

$$M_k^{\mathrm{acc}} = \Pi^{T_V}(\mathcal{T}_{V,2} \circ \mathcal{V}_k)^{\dagger}(|1\rangle\langle 1|)\Pi^{T_V}, \qquad M_k^{\mathrm{rej}} = \mathbb{1} - M_k^{\mathrm{acc}}$$

Remark 5.1. In order to keep notation clear and uncluttered, we will often suppress the parameterization for the oracle O. So that it is not completely forgotten, we will include O in expectations (e.g., $\mathbb{E}_{O,k}[\ldots]$). Keep in mind that all objects downstream from $\mathcal{T}$ will be dependent on O — e.g., Π^{T_V}, M_k^{acc}, σ_k, etc.

Recall the three projectors defined in Lemma 2.2 and Definition 2.7: Π_k which projects onto the space of valid ciphertexts encrypted with the key k, $\Pi_{k,r}$ which narrows this to ciphertexts which were specifically encrypted with randomness r, and $\widetilde{\Pi}_{k,r}$ which additionally restricts onto those ciphertexts which encrypt half of the maximally entangled state $|\Phi^+\rangle$. For each $x \in \mathfrak{L}$ and k, we can define the subnormalized states

$$\Pi_k \sigma_k \Pi_k, \quad \text{and} \quad (\mathbb{1} - \Pi_k)\sigma_k(\mathbb{1} - \Pi_k)$$

which correspond to the projection of x's honest proof onto valid and invalid ciphertexts at the point where $\mathcal{V}_k$ is decrypting.

Because $\mathcal{V}_k(\sigma_k) = \mathcal{V}_k(\Pi_k \sigma_k \Pi_k) + \mathcal{V}_k((\mathbb{1} - \Pi_k)\sigma_k(\mathbb{1} - \Pi_k))$, the expected completeness for a particular

statement x on average over all protocols in the family can be written

$$\begin{aligned}\mathop{\mathbb{E}}_{O,k} \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_V^{\mathcal{V}_k} \circ \mathcal{T}_P^{\mathcal{P}}(\xi_k^x)] &= \mathop{\mathbb{E}}_{O,k} \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2}(\mathcal{V}_k(\Pi_k\sigma_k\Pi_k) + \mathcal{V}_k((\mathbb{1}-\Pi_k)\sigma_k(\mathbb{1}-\Pi_k)))] \\ &= \underbrace{\mathop{\mathbb{E}}_{O,k} \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2} \circ \mathcal{V}_k(\Pi_k\sigma_k\Pi_k)]}_{\mu_1^x} + \underbrace{\mathop{\mathbb{E}}_{O,k} \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2} \circ \mathcal{V}_k((\mathbb{1}-\Pi_k)\sigma_k(\mathbb{1}-\Pi_k)))]}_{\mu_2^x}\end{aligned} \tag{3}$$

where we've labeled two categories of completeness with μ_1^x and μ_2^x. The second category, μ_2^x, captures the event where the non-interactive argument accepts despite the underlying verifier $\mathcal{V}_k$ rejecting the ciphertext. Intuitively, the only way that the compiled argument could be deciding the language is through its interaction with $\mathcal{P}$ and $\mathcal{V}_k$. One might therefore expect this value to be small.

Recall that we are assuming there exists $\kappa > 0$ such that $\varepsilon + 2\delta \leq \kappa^* - \kappa$. In Section 5.3, we will construct a subchannel corresponding to the invalid ciphertext completeness contribution μ_2^x, and in Section 5.4 we do the same for the valid ciphertext contribution μ_1^x. In Section 5.5, we combine the two subchannels as complementary branches of a single algorithm.

These subchannels, $\mathcal{A}_1$ and $\mathcal{A}_2$, and their associated hybrids will all output a single classical bit. As in the previous section, post-selecting onto a projector retains the corresponding measurement outcome without renormalizing: the trace of the output includes the probability of that outcome. The omitted outcome contributes the zero operator $\mathbf{0}$, not the state $|0\rangle\langle 0|$. For these subchannels, we use the convention

$$\Pr_O[\mathcal{A}_i^O(x) = 1] = \mathop{\mathbb{E}}_O \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{A}_i^O(x)],$$

and likewise for the hybrids. Thus all acceptance probabilities below are unconditional weights, not probabilities conditioned on the retained outcome.

5.3 Case 1: Invalid Ciphertexts

In this case, the QNIC algorithms are mauling the ciphertext such that it does not decrypt but they will still accept. Because the decryption fails, $\mathcal{T}$ will never see the first message in the clear. We use an encrypted EPR state and retain the branch that fails its authentication test. QCA security relates this branch to the rejecting branch for an encryption of a fixed state.

$\mathcal{A}_1^O(x)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\sigma_{k,r} = \mathcal{T}_{P+V,1}^{\mathcal{S}} \circ \mathsf{Enc}_{k;r}(\,|\Phi^+\rangle\langle\Phi^+|_{\mathsf{LM}} \otimes |x\rangle\langle x|_{\mathsf{X}})$.
3. Post-select onto $\mathbb{1} - \widetilde{\Pi}_{k,r}$.
4. Discard $\mathsf{ALM'CZ}$, prepare B as $|0\rangle$, apply $\mathcal{T}_{V,2}$, and forward the output.

Lemma 5.1. $\Pr_O[\mathcal{A}_1^O(x) = 1] < \delta(n) + \mathsf{negl}(n)$ *for* $x \notin \mathfrak{L}$.

Proof. Fix an oracle O. Let Γ and Λ be the real and simulated effective attack maps from Definition 2.7 for the attack $\mathcal{T}_{P+V,1}^{\mathcal{S}}$, and let Γ^{rej} be the rejecting branch of Γ. By QCA security,

$$\|\Gamma^{\mathrm{rej}} - |\bot\rangle\langle\bot|_{\mathsf{M}} \otimes \Lambda^{\mathrm{rej}}\|_\diamond \leq \mathsf{negl}(n).$$

Let $\mathcal{T}_{V,2}^{\mathrm{acc}}$ be the accepting branch of $\mathcal{T}_{V,2}$ and let $\mathcal{N}$ be the post-processing subchannel which post-selects onto $\Pi^\bot$, discards MCZ and any unused simulator output $\mathsf{M'}$, prepares B as $|0\rangle$, and applies $\mathcal{T}_{V,2}^{\mathrm{acc}}$. Applying $\mathcal{N}$ to the simulated rejecting branch on input $|0,x\rangle\langle 0,x|$ gives exactly the acceptance weight of $\mathcal{A}_1^O(x)$. Applying it

to the real rejecting branch instead uses an encryption of $|0\rangle\langle 0|_{\mathsf{M}}$, decrypts after the attack, and post-selects onto $\Pi^{\perp}$. By contractivity, these acceptance weights differ by at most $\mathsf{negl}(n)$.

Let $\psi_k = \mathcal{T}_P^{\mathcal{S}} \circ \mathsf{Enc}_k(\,|0\rangle\langle 0|_{\mathsf{M}} \otimes |x\rangle\langle x|_{\mathsf{X}})$. This is an efficiently preparable state because $\mathcal{T}_P$ and $\mathcal{S}$ are efficient, so the compiled argument's soundness applies to ψ_k for every k. Averaging over O, k, we obtain

$$\begin{aligned}\Pr_O[\mathcal{A}_1^O(x) = 1] &\le \mathop{\mathbb{E}}_{O,k} \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2} \circ \mathcal{V}_k^{\mathrm{rej}} \circ \mathcal{T}_{V,1}(\psi_k \otimes |x\rangle\langle x|)] + \mathsf{negl}(n)\\ &\le \mathop{\mathbb{E}}_{O,k} \mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2} \circ \mathcal{V}_k \circ \mathcal{T}_{V,1}(\psi_k \otimes |x\rangle\langle x|)] + \mathsf{negl}(n)\\ &< \delta(n) + \mathsf{negl}(n),\end{aligned}$$

where $\mathcal{V}_k^{\mathrm{rej}}(\rho) = \mathrm{Tr}[(\mathbb{1} - \Pi_k)\rho]\,|0\rangle\langle 0|$ and the second inequality follows from $\mathcal{V}_k^{\mathrm{rej}} \le \mathcal{V}_k$. □

Lemma 5.2. *$|\Pr_O[\mathcal{A}_1^O(x) = 1] - \mu_2^x| \le \mathsf{negl}(n)$ for $x \in \mathfrak{L}$.*

Proof. This is proven through a series of hybrids such that Hyb_2 accepts with probability μ_2^x.

$\mathsf{Hyb}_1^O(x)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\mathcal{T}_{P+V,1}^{\mathcal{P}} \circ \mathsf{Enc}_{k;r}(\,|\Phi^+\rangle\langle\Phi^+|_{\mathsf{LM}} \otimes \xi^x_{\mathsf{XM'S}})$.
3. Post-select onto $\mathbb{1} - \widetilde{\Pi}_{k,r}$.
4. Discard $\mathsf{ALM'CZ}$, prepare B as $|0\rangle$, apply $\mathcal{T}_{V,2}$, and forward the output.

Claim 5.1. *$|\Pr_O[\mathcal{A}_1^O(x) = 1] - \Pr_O[\mathsf{Hyb}_1^O(x) = 1]| \le \mathsf{negl}(n)$ for $x \in \mathfrak{L}$.*

Proof. Fix an oracle O. Observe that the $\mathsf{M'}$ register is only discarded in Steps 3 – 4 and so does not affect the probability that either hybrid outputs 1. Thus the only relevant difference between $\mathcal{A}_1$ and Hyb_1 is in the oracle $\mathcal{S}$ and its internal register S. Towards contradiction, assume there exists a polynomial p such that

$$|\Pr[\mathcal{A}_1(x) = 1] - \Pr[\mathsf{Hyb}_1(x) = 1]| > 1/p(n).$$

Then we can break the SHVZK of the underlying Ξ-protocol. The reduction implements the measurement defining $\mathcal{A}_1$ (or Hyb_1), outputs 0 on the omitted outcome, and otherwise forwards the final output, with $\mathcal{T}_P$'s oracle simulated by the SHVZK challenger. That is to say, when $\mathcal{T}_P$ sends C to its oracle, measure it in the computational basis, send the outcome c to the challenger, receive $\rho_{\mathsf{M'Z}}$ in response, and forward the Z register back to $\mathcal{T}_P$. If the SHVZK challenger sampled ρ from an honest execution, the reduction accepts with exactly the acceptance weight of Hyb_1. Otherwise, its acceptance probability is the acceptance weight of $\mathcal{A}_1$. By our assumption, this is an inverse poly distinguisher and we arrive at a contradiction. □

$\mathsf{Hyb}_2^O(x)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\mathsf{Dec}_k \circ \mathcal{T}_{P+V,1}^{\mathcal{P}} \circ \mathsf{Enc}_{k;r}(\xi^x)$.
3. Post-select onto $\Pi^{\perp}$.
4. Discard MCZ, prepare B as $|0\rangle$, apply $\mathcal{T}_{V,2}$, and forward the output.

Claim 5.2. *$|\Pr_O[\mathsf{Hyb}_1^O(x) = 1] - \Pr_O[\mathsf{Hyb}_2^O(x) = 1]| \le \mathsf{negl}(n)$ for $x \in \mathfrak{L}$.*

Proof. Fix an oracle O. Let Γ (resp. Λ) be the real (resp. simulated) effective attack map from Definition 2.7 for the fixed attack $\mathcal{T}^{\mathcal{P}}_{P+V,1}$. Let $\mathcal{N}$ be the post-processing subchannel from the proof of Lemma 5.1. Then

$$\begin{aligned}
|\Pr[\mathsf{Hyb}_2^O(x)=1]-\Pr[\mathsf{Hyb}_1^O(x)=1]| &= \left|\mathrm{Tr}\left[\mathcal{N}\circ\Gamma^{\mathrm{rej}}(\xi^x)\right]-\mathrm{Tr}\left[\mathcal{N}\circ(\,|\bot\rangle\langle\bot|_{\mathsf{M}}\otimes\Lambda^{\mathrm{rej}})(\xi^x)\right]\right| \\
&= \left\|\mathcal{N}\circ\Gamma^{\mathrm{rej}}(\xi^x)-\mathcal{N}\circ(\,|\bot\rangle\langle\bot|_{\mathsf{M}}\otimes\Lambda^{\mathrm{rej}})(\xi^x)\right\|_1 \\
&\le \|\mathcal{N}\circ\Gamma^{\mathrm{rej}}-\mathcal{N}\circ(\,|\bot\rangle\langle\bot|_{\mathsf{M}}\otimes\Lambda^{\mathrm{rej}})\|_\diamond \\
&\le \|\Gamma^{\mathrm{rej}}-\,|\bot\rangle\langle\bot|_{\mathsf{M}}\otimes\Lambda^{\mathrm{rej}}\|_\diamond \\
&\le \mathsf{negl}(n)
\end{aligned}$$

□

Using Fact 2.3, we get

$$\begin{aligned}
\Pr_O[\mathsf{Hyb}_2^O(x)=1] &= \mathop{\mathbb{E}}_{O,k}\mathrm{Tr}\left[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2}\left(\mathrm{Tr}_{\mathsf{MCZ}}\left[\Pi^{\bot}(\mathsf{Dec}_k\circ\mathcal{T}^{\mathcal{P}}_{P+V,1}(\xi_k^x))\Pi^{\bot}\right]\otimes|0\rangle\langle 0|_{\mathsf{B}}\right)\right] \\
&= \mathop{\mathbb{E}}_{O,k}\mathrm{Tr}\left[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2}\left(\mathrm{Tr}_{\mathsf{MCZ}}\left[\mathsf{Dec}_k((\mathbb{1}-\Pi_k)(\mathcal{T}^{\mathcal{P}}_{P+V,1}(\xi_k^x))(\mathbb{1}-\Pi_k))\right]\otimes|0\rangle\langle 0|_{\mathsf{B}}\right)\right] \\
&= \mathop{\mathbb{E}}_{O,k}\mathrm{Tr}\left[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2}\circ\mathcal{V}_k((\mathbb{1}-\Pi_k)\sigma_k(\mathbb{1}-\Pi_k))\right] \\
&= \mu_2^x
\end{aligned}$$

Combining Claims 5.1 – 5.2, we have our result. □

5.4 Case 2: Valid Ciphertexts

On the valid ciphertext branch, we can use our results from Section 4 to defer the first message until after receiving the challenge. Because of that reordering, we can replace the prover with the simulator and create what should be a convincing proof without knowledge of the witness.

$\mathcal{A}_2^O(x)$

1. Sample $k\leftarrow\mathsf{Gen}(1^n)$ and $r\leftarrow p_k$.
2. Prepare $\sigma_{k,r}=\mathcal{T}^{\mathcal{S}}_{P+V,1}\circ\mathsf{Enc}_{k;r}(\,|\Phi^+\rangle\langle\Phi^+|_{\mathsf{LM}}\otimes|x\rangle\langle x|_{\mathsf{X}})$.
3. Post-select onto $\widetilde{\Pi}_{k,r}$, then discard AL.
4. Encrypt M' with $\mathsf{Enc}_{k;r}$ yielding register A'.
5. Measure $\mathsf{A}'\mathsf{CZE}$ with $\{M_k^{\mathrm{acc}},M_k^{\mathrm{rej}}\}$. If the outcome is M_k^{acc}, output 1. Otherwise output 0.

Lemma 5.3. $\Pr[\mathcal{A}_2(x)=1]<\delta(n)$ *for* $x\notin\mathfrak{L}$.

Proof. By the δ-soundness assumption, for all k, efficiently preparable states $\psi_k\in\mathrm{L}(\mathsf{P})$, and $x\notin\mathfrak{L}$,

$$\mathop{\mathbb{E}}_{O}\mathrm{Tr}[(\mathcal{T}_{V,2}\circ\mathcal{V}_k)^\dagger(|1\rangle\langle 1|)T_{V,1}(\psi_k\otimes|x\rangle\langle x|_{\mathsf{X}}\otimes|0\rangle\langle 0|_{\mathsf{E}})T_{V,1}^\dagger]<\delta$$

Towards contradiction, assume $\Pr_O[\mathcal{A}_2^O(x)=1]\ge\delta$. For fixed k, let ψ_k be the subnormalized state before Step 5 of $\mathcal{A}_2^O$, averaged over r; its trace includes the probability of the retained outcome in Step 3. Then by our assumption, $\mathbb{E}_{O,k}\,\mathrm{Tr}[M_k^{\mathrm{acc}}\psi_k]\ge\delta$. By an averaging argument, there exists some k^* such that $\mathbb{E}_O\,\mathrm{Tr}[M_{k^*}^{\mathrm{acc}}\psi_{k^*}]\ge\delta$. Let φ be the state prepared in the following way:

1. Run Steps 1 – 4 of $\mathcal{A}_2^O$ with $k=k^*$, physically measuring $\{\widetilde{\Pi}_{k^*,r},\mathbb{1}-\widetilde{\Pi}_{k^*,r}\}$ in Step 3. If the outcome is $\mathbb{1}-\widetilde{\Pi}_{k^*,r}$, output any fixed state on P and stop.
2. Measure with $\{\Pi^{T_V},\mathbb{1}-\Pi^{T_V}\}$. If the outcome is Π^{T_V}, apply $T_{V,1}^\dagger$ and output the residual state on P.
3. Otherwise, output any fixed state on P.

This prepares a normalized state efficiently, without conditioning on either measurement outcome. If we apply $T_{V,1}$ to φ and measure with $(\mathcal{T}_{V,2} \circ \mathcal{V}_{k^*})^\dagger(|1\rangle\langle 1|)$, the branch retaining both outcomes contributes exactly $\mathrm{Tr}[M^{\mathrm{acc}}_{k^*}\psi_{k^*}]$ to acceptance. The other branches contribute nonnegative amounts, so the total accepting probability is $\geq \delta$ in expectation over O. This contradicts δ-soundness and the lemma holds. □

Lemma 5.4. $\Pr[\mathcal{A}_2(x) = 1] \geq \mu_1^x - \frac{2}{3\sqrt{3}} - \mathsf{negl}(n)$ *for* $x \in \mathfrak{L}$.

Proof. This is proven through a series of subchannel hybrids where Hyb_3 is the trace-preserving special case corresponding to the honest execution.

$\mathsf{Hyb}_1^O(x)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\sigma_{k,r} = \mathcal{T}^{\mathcal{P}}_{P+V,1} \circ \mathsf{Enc}_{k;r}(\,|\Phi^+\rangle\langle\Phi^+|_{\mathsf{LM}} \otimes \xi^x_{\mathsf{XM'S}})$.
3. Post-select onto $\widetilde{\Pi}_{k,r}$, then discard AL.
4. Encrypt $\mathsf{M'}$ with $\mathsf{Enc}_{k;r}$ yielding register $\mathsf{A'}$.
5. Measure $\mathsf{A'CZE}$ with $\{M^{\mathrm{acc}}_k, M^{\mathrm{rej}}_k\}$. If the outcome is M^{acc}_k, output 1. Otherwise output 0.

Claim 5.3. $|\Pr_O[\mathcal{A}_2^O(x) = 1] - \Pr_O[\mathsf{Hyb}_1^O(x) = 1]| \leq \mathsf{negl}(n)$ *for* $x \in \mathfrak{L}$.

Proof. By the SHVZK argument in Claim 5.1, with the reduction outputting 0 on the omitted measurement outcome. Thus it compares the unconditional acceptance weights of the two subchannels. □

$\mathsf{Hyb}_2^O(x)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\sigma_{k,r} = \mathcal{T}^{\mathcal{P}}_{P+V,1} \circ \mathsf{Enc}_{k;r}(\xi^x_{\mathsf{XMS}})$.
3. Post-select onto Π_k.
4. Measure ACZE with $\{M^{\mathrm{acc}}_k, M^{\mathrm{rej}}_k\}$. If the outcome is M^{acc}_k, output 1. Otherwise output 0.

Claim 5.4. $|\Pr_O[\mathsf{Hyb}_1^O(x) = 1] - \Pr_O[\mathsf{Hyb}_2^O(x) = 1]| \leq \mathsf{negl}(n)$ *for* $x \in \mathfrak{L}$.

Proof. Fix an oracle O. Steps 1 – 4 of Hyb_1^O are exactly the steps in $\mathsf{Exp}^{\mathcal{A}}_{\mathsf{defer}}(\xi^x)$ for $\mathcal{A} = \mathcal{T}^{\mathcal{P}}_{P+V,1}$ and similarly Steps 1 – 3 of Hyb_2^O correspond to $\mathsf{Exp}^{\mathcal{A}}_{\mathsf{real}}(\xi^x)$. By Lemma 4.1, these subchannels are negligibly close in diamond distance[6]. Therefore, the final measurement accepts in the two hybrids with probabilities differing by at most $\mathsf{negl}(n)$. □

$\mathsf{Hyb}_3^O(x)$
1. Sample $k \leftarrow \mathsf{Gen}(1^n)$ and $r \leftarrow p_k$.
2. Prepare $\sigma_{k,r} = \mathcal{T}^{\mathcal{P}}_{P+V,1} \circ \mathsf{Enc}_{k;r}(\xi^x_{\mathsf{XMS}})$.
3. Measure ACZE with $\{M^{\mathrm{acc}}_k, M^{\mathrm{rej}}_k\}$. If the outcome is M^{acc}_k, output 1. Otherwise output 0.

Claim 5.5. $\Pr_O[\mathsf{Hyb}_3^O(x) = 1] - \Pr_O[\mathsf{Hyb}_2^O(x) = 1] \leq \mu_2^x + \frac{2}{3\sqrt{3}}$ *for* $x \in \mathfrak{L}$.

[6]Note that both experiments output k so the measurement can depend on k.

Proof. Fix O, k, r. By construction, $\Pi^{T_V}\sigma_{k,r}\Pi^{T_V} = \sigma_{k,r}$. Moreover, $(\mathcal{T}_{V,2} \circ \mathcal{V}_k)^\dagger(|1\rangle\langle 1|)$ commutes with Π_k, since $\mathcal{V}_k$ measures ciphertext validity before proceeding. Apply Lemma 2.1 with projectors Π_k, Π^{T_V}, measurement operator $(\mathcal{T}_{V,2} \circ \mathcal{V}_k)^\dagger(|1\rangle\langle 1|)$, and state $\sigma_{k,r}$ to obtain

$$\mathrm{Tr}[\,|1\rangle\langle 1|\,\mathcal{T}_{V,2} \circ \mathcal{V}_k(\Pi_k\sigma_{k,r}\Pi_k)] - \mathrm{Tr}[M_k^{\mathrm{acc}}\Pi_k\sigma_{k,r}\Pi_k] \leq \frac{2}{3\sqrt{3}}.$$

Taking the expectation over (O, k, r) and using $\mathbb{E}_r[\sigma_{k,r}] = \sigma_k$ gives

$$\mu_1^x - \Pr_O[\mathsf{Hyb}_2^O(x) = 1] \leq \frac{2}{3\sqrt{3}}.$$

Since $\Pr_O[\mathsf{Hyb}_3^O(x) = 1] = \mu_1^x + \mu_2^x$, the claim follows. □

Hyb_3^O is equivalent to the honest execution of the compiled non-interactive argument for x with a protocol randomly sampled from our family. Combining Claims 5.3 – 5.5, we get

$$\Pr_O[\mathcal{A}_2^O(x) = 1] \geq \Pr_O[\mathsf{Hyb}_3^O(x) = 1] - \mu_2^x - \frac{2}{3\sqrt{3}} - \mathsf{negl}(n) = \mu_1^x - \frac{2}{3\sqrt{3}} - \mathsf{negl}(n).$$ □

5.5 Proof of Theorem 1

The subchannels $\mathcal{A}_1$ and $\mathcal{A}_2$ both begin by preparing $\sigma_{k,r}$. They then retain complementary outcomes of the same projective measurement, and their subsequent steps are trace preserving. Consequently, for every fixed (O, x),

$$\mathrm{Tr}[\mathcal{A}_1^O(x)] + \mathrm{Tr}[\mathcal{A}_2^O(x)] = \mathop{\mathbb{E}}_{k,r} \mathrm{Tr}[(\mathbb{1} - \widetilde{\Pi}_{k,r})\sigma_{k,r}] + \mathop{\mathbb{E}}_{k,r} \mathrm{Tr}[\widetilde{\Pi}_{k,r}\sigma_{k,r}] = 1.$$

Thus their sum defines a channel $\mathcal{A}^O = \mathcal{A}_1^O + \mathcal{A}_2^O$. Operationally, $\mathcal{A}$ performs the common preparation once and measures $\{\widetilde{\Pi}_{k,r}, \mathbb{1} - \widetilde{\Pi}_{k,r}\}$. On outcome $\mathbb{1} - \widetilde{\Pi}_{k,r}$, it continues as in $\mathcal{A}_1$; on outcome $\widetilde{\Pi}_{k,r}$, it continues as in $\mathcal{A}_2$. It always forwards the resulting output bit, without post-selection. In particular, for all x, the acceptance weights add exactly:

$$\Pr_O[\mathcal{A}^O(x) = 1] = \Pr_O[\mathcal{A}_1^O(x) = 1] + \Pr_O[\mathcal{A}_2^O(x) = 1]. \tag{4}$$

All steps are efficient, so $\mathcal{A}$ is itself efficient.

Lemma 5.5. $\Pr_O[\mathcal{A}^O(x) = 1] \leq 2\delta + \mathsf{negl}(n)$ *if* $x \notin \mathfrak{L}$.

Proof. By Lemmas 5.1 and 5.3, each subchannel has acceptance weight $< \delta + \mathsf{negl}(n)$. The result follows from Equation (4). □

Lemma 5.6. $\Pr_O[\mathcal{A}^O(x) = 1] \geq 2\delta + \kappa - \mathsf{negl}(n)$ *if* $x \in \mathfrak{L}$.

Proof. By Equation (4) and Lemmas 5.2 and 5.4,

$$\begin{aligned}
\Pr_O[\mathcal{A}^O(x) = 1] &\geq \mu_2^x + \mu_1^x - \frac{2}{3\sqrt{3}} - \mathsf{negl}(n) \\
&\geq 1 - \varepsilon - \frac{2}{3\sqrt{3}} - \mathsf{negl}(n) \\
&= \kappa^* - \varepsilon - \mathsf{negl}(n) \\
&\geq 2\delta + \kappa - \mathsf{negl}(n).
\end{aligned}$$

where the second inequality uses our completeness bound and the last uses the hypothesis $\varepsilon + 2\delta < \kappa^* - \kappa$. □

References

[ACGH20] Gorjan Alagic, Andrew M. Childs, Alex B. Grilo, and Shih-Han Hung. Non-interactive classical verification of quantum computation. In *TCC (3)*, volume 12552 of *Lecture Notes in Computer Science*, pages 153–180. Springer, 2020.

[AGM18] Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Unforgeable quantum encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, Lecture Notes in Computer Science, pages 489–519. Springer, 2018.

[AGM21] Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Can you sign a quantum state? *Quantum*, 5:603, 2021.

[ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '14, pages 474–483, 2014.

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, MIT, Cambridge, Massachusetts, USA, January 10-13, 2023*, volume 251 of *LIPIcs*, pages 24:1–24:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[BG20] Anne Broadbent and Alex B. Grilo. Qma-hardness of consistency of local density matrices with applications to quantum zero-knowledge. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 196–205. IEEE, 2020.

[BJRT25] James Bartusek, Ruta Jawale, Justin Raizes, and Kabir Tomer. A new approach to arguments of quantum knowledge. *IACR Cryptol. ePrint Arch.*, page 2155, 2025.

[BKL+22] James Bartusek, Yael Tauman Kalai, Alex Lombardi, Fermi Ma, Giulio Malavolta, Vinod Vaikuntanathan, Thomas Vidick, and Lisa Yang. Succinct classical verification of quantum computation. In *Proceedings of the 42nd Annual International Cryptology Conference*, CRYPTO '22, pages 195–211, 2022.

[BM22] James Bartusek and Giulio Malavolta. Indistinguishability obfuscation of null quantum circuits and applications. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, Berkeley, CA, USA, January 31 - February 3, 2022*, volume 215 of *LIPIcs*, pages 15:1–15:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, CCS '93, pages 62–73, 1993.

[BS86] VP Belavkin and P Staszewski. A radon-nikodym theorem for completely positive maps. *Reports on Mathematical Physics*, 24(1):49–55, 1986.

[CCH+18] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, and Ron D. Rothblum. Fiat–Shamir from simpler assumptions. Cryptology ePrint Archive, Report 2018/1004, 2018.

[CVZ20] Andrea Coladangelo, Thomas Vidick, and Tina Zhang. Non-interactive zero-knowledge arguments for qma, with preprocessing. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 799–828. Springer, 2020.

[DFG13] Özgür Dagdelen, Marc Fischlin, and Tommaso Gagliardoni. The Fiat-Shamir transformation in a quantum world. In *Proceedings of the 19th International Conference on the Theory and Application of Cryptology and Information Security*, ASIACRYPT '13, pages 62–81, 2013.

[DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. Cryptology ePrint Archive, Report 2019/190, 2019.

[DS18] Yfke Dulek and Florian Speelman. Quantum ciphertext authentication and key recycling with the trap code. In Stacey Jeffery, editor, *13th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2018, Sydney, Australia, July 16-18, 2018*, volume 111 of *LIPIcs*, pages 1:1–1:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[Fis05] Marc Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In *Proceedings of the 25th Annual International Cryptology Conference*, CRYPTO '05, pages 152–168, 2005.

[FLM+26] Ben Foxman, Alex Lombardi, Fermi Ma, Barak Nehoran, and John Wright. Quantum lazy sampling and path recording for any group. *CoRR*, abs/2606.30281, 2026.

[FS86] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In *Proceedings of the 6th Annual International Cryptology Conference*, CRYPTO '86, pages 186–194, 1986.

[GJMZ23] Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1579–1588. ACM, 2023.

[GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the Fiat-Shamir paradigm. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '03, pages 102–113, 2003.

[GY25] Dmitry Grinko and Satoshi Yoshida. Quantum simulation of random unitaries from clebsch-gordan transforms. *arXiv preprint arXiv:2509.26623*, 2025.

[GYZ16] Sumegha Garg, Henry Yuen, and Mark Zhandry. New security notions and feasibility results for authentication of quantum data. *CoRR*, abs/1607.07759, 2016.

[KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. In *Proceedings of the 37th Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT '17, pages 552–586, 2018.

[KSW08] Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner. The information-disturbance tradeoff and the continuity of stinespring's representation. *IEEE Trans. Inf. Theory*, 54(4):1708–1717, 2008.

[LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. Cryptology ePrint Archive, Report 2019/262, 2019.

[MH25] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 806–809. ACM, 2025.

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Classically verifiable NIZK for QMA with preprocessing. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology - ASIACRYPT*

2022 - 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5-9, 2022, Proceedings, Part IV, volume 13794 of *Lecture Notes in Computer Science*, pages 599–627. Springer, 2022.

[PS96] David Pointcheval and Jacques Stern. Security proofs for signature schemes. In *Proceedings of the 14th Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT '96, pages 387–398, 1996.

[Rag03] Maxim Raginsky. Radon–nikodym derivatives of quantum operations. *Journal of Mathematical Physics*, 44(11):5003–5020, 2003.

[Shm21] Omri Shmueli. Multi-theorem designated-verifier NIZK for QMA. 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 375–405. Springer, 2021.

[Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *Proceedings of the 34th Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT '15, pages 755–784, 2015.

[Unr17] Dominique Unruh. Post-quantum security of Fiat-Shamir. In *Proceedings of the 23rd International Conference on the Theory and Applications of Cryptology and Information Security*, ASIACRYPT '17, pages 65–95, 2017.

[Zha18] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. Cryptology ePrint Archive, Report 2018/276, 2018.

[Zha25] Mark Zhandry. How to model unitary oracles. In *CRYPTO 2025*, 2025.