# Security of the Fischlin Transform in Quantum Random Oracle Model

Christian Majenz and Jaya Sharma

Department of Applied Mathematics and Computer Science, Technical University of Denmark
{chmaj, jaysh}@dtu.dk

**Abstract.** The Fischlin transform yields non-interactive zero-knowledge proofs with straight-line extractability in the classical random oracle model. This is done by forcing a prover to generate multiple accepting transcripts through a proof-of-work mechanism. Whether the Fischlin transform is straight-line extractable against quantum adversaries has remained open due to the difficulty of reasoning about the *likelihood* of query transcripts in the quantum-accessible random oracle model (QROM), even when using the compressed oracle methodology. In this work, we prove that the Fischlin transform remains straight-line extractable in the QROM, via an extractor based on the compressed oracle. This establishes the post-quantum security of the Fischlin transform, providing a post-quantum straight-line extractable NIZK alternative to Pass' transform with smaller proof size. Our techniques are built on different combinations of symmetrization, query amplitude and quantum union bound arguments, as well as tail bounds for sums of independent random variables and for martingales.

**Keywords:** Fischlin transform, QROM, NIZK, post-quantum

# Table of Contents

Security of the Fischlin Transform in Quantum Random Oracle Model . . . 1
*Christian Majenz and Jaya Sharma*
1 Introduction . . . 2
2 Preliminaries . . . 5
2.1 Quantum Random Oracle (QROM) and Compressed Oracle Technique . . . 8
3 Proof of Post Quantum Security of the Fischlin Transform . . . 12
3.1 Extractability – Deterministic Commitments . . . 14
3.2 Extractability – lifting to general prover. . . . 25
3.3 Main Result . . . 29
4 Acknowledgements . . . 29
A Deferred proof of Lemma 9 . . . 32
B A gentle measurement lemma . . . 33
C Bounding Lemmas . . . 34
D Proof of Zero Knowledge . . . 36

## 1 Introduction

Non-interactive zero-knowledge proofs (NIZKs) obtained from interactive protocols via the Fiat–Shamir paradigm form a central tool in modern cryptography. Many applications require *straight-line extractability* – the ability to efficiently extract a witness from any succesful prover, without rewinding. The Fischlin transform [Fis05] provides a prominent approach for achieving straight-line extractable NIZKs from $\Sigma$-protocols in the random oracle model. It constitutes an alternative to Pass' transform [Pas03], using what would now be called a proof-of-work-like mechanism instead of commitments to enable smaller proofs.

With the advent of quantum computing, it has become necessary to analyze the security of random-oracle-based constructions in the presence of quantum adversaries. This has led to the development of the quantum random oracle model (QROM) [BDF+11], in which adversaries may query the random oracle in superposition. Over the past decade, a substantial body of work has established a versatile technical toolbox for proving the post-quantum security of classical constructions in the QROM. These techniques have been successfully applied to a wide range of Fiat–Shamir-type transforms, including Pass' transform [CMS19a,DFMS22a,RT25] and related primitives. However, despite its importance, the Fischlin transform has so far resisted security analysis in the QROM. Existing proof techniques do not readily extend to this setting, and no QROM extractability proof for the Fischlin transform was previously known.

The main difficulty in analyzing the Fischlin transform in the QROM stems from the structure of its extraction mechanism. Informally, a prover succeeds

only if it finds an oracle query whose output satisfies a somewhat sparse acceptance condition, which in turn requires many attempts, except with tiny probability. A classical extractor can thus read off two accepting transcripts from the adversarial prover's oracle query inputs, allowing extraction of a witness via special soundness. In the quantum setting, however, this reasoning becomes problematic. The adversary's output depends only on the single accepting query, the unsuccessful ones are in some sense *counterfactual*. This stands in sharp contrast to Pass' transform, where accepting transcripts contain explicit commitments to responses for all challenges that can be leveraged for extraction even in the quantum setting. As a result, proving that a quantum adversary must invest substantial query complexity in order to succeed requires new techniques.

**Our results.** In this work, we prove the straight-line extractability of the Fischlin transform in the QROM. Our result provides concrete bounds for a wide range of parameter choices that include an infinite family of parameter sets for asymptotic security. In more detail, we have the following

**Theorem 1 (Informal version of Corollary 1).** *Let $\Sigma$ be a* $\Sigma$-protocol *for a witness relation $R$ with special soundness and unique responses. Under suitable conditions on the parameters of the Fischlin transform that allow unbounded values for the parallel repetition parameter of the Fischlin transform, $k$, (and thus for the security paramter,) the Fischlin transform $\mathsf{Fis}[\Sigma]$ is a proof of knowledge with straight-line extractability in QROM. The simulation can be made perfect and the extractor succeeds whenever a $q$-query malicious prover does, except with probability no more than*

$$\varepsilon_{\mathrm{ex}} \leq q^2 \cdot \mathrm{negl}(k).$$

**Technical Overview.** Defining the straight-line extractor for the Fischlin transform is straightforward. We simulate the random oracle using a compressed oracle [Zha19]. Once the prover has finished and its proof is verified, the extractor measures the compressed oracle database register. If the measurement result contains two accepting transcripts for the same repetition with different challenges, special soundness of the $\Sigma$-protocol is used to produce a witness.

Proving that this extractor works turns out to be a formidable challenge. Known techniques for straightline extraction in the QROM are not sufficient. There seem to be two such techniques in the literature. Firstly, in [Unr15], straightline-extraction for commitments is achieved by simulating a length- preserving quantum-accessible random oracle (QRO) in a way that allows finding pre-images. And secondly, [CMS19a] showed how to use the compressed oracle for straightline extraction. This mechanism was further developed [DFMS22b,DFMS22a,RT25] and applied [AHJ+23,HJMN24,BBB+25]. None of these techniques can be used for the Fischlin transform though – the absence of commitments prevents straightforward generalization, and the fact that extractor failure is characterized by the absence of certain patterns in the compressed oracle database, rather than their existence, prevents transition capacity reasoning [CFHL21].

The main work goes into proving that no prover can succeed with significant probability at outputting a proof that involves a fixed commitment vector chosen ahead of time ("deterministic prover") so that the extractor still fails. This result is then lifted to a general prover using an intricate argument based on quantum query amplitudes and a quantum union bound.

To show the result for deterministic provers, our strategy is reminiscent of the Monty Hall paradox. Here, a game master presents three doors. Behind a random one is a car, while there are goats behind the other two. The participant can now choose a door, which is, however, not opened. The game master now opens a *different* door which reveals a goat. The participant can now decide: stick to their choice of door, or choose the only remaining door instead. The optimal strategy is to switch, which reveals the car with probability 2/3. What makes this situation appear paradoxical is the fallacy of regarding the remaining door as a fixed object instead of a random variable. The game master chooses the door they open depending on the randomness of the placement of the car and the participant's choice, which correlates the latter with the car's placement.

When a prover outputs a valid Fischlin proof, a number of hash values are now known (to fulfill the PoW condition, in the following discussion simplified to "output 0"). This can only be achieved by correlating the choice of challenges with the random oracle. Clearly this, in turn, necessitates that some hash outputs not involved in proof verification are now *less* likely to be 0 than suggested by the uniform distribution, otherwise the oracle output distribution would be biased. In the compressed oracle, analogously, the oracle registers corresponding to valid transcripts not involved in the verification of the proof output by the adversarial prover cannot contain the uniform superposition of outputs anymore. This deviation is then exploited by the extractor.

From a technical perspective, the task of making this argument rigorous is formidable. The output length of the random oracle is logarithmic, preventing the usual tight relationship between oracle register content and adversarial view up to negligible error. We therefore need to explicitly analyze every switch between the "compressed" and "uncompressed" bases of the superposition oracle using Chernoff bounds. The high-level strategy is as follows. Pick a number $\eta$ such that (i) it is unlikely that there are $\eta$ oracle inputs in a certain sub-domain that map to zero (by a standard Chernoff bound), but (ii) it is likely to find roughly $\eta - k$ zeroes, where $k$ is the number of parallel repetitions in the Fischlin transform and thus the number of hash values that need to be zero for a Fischlin proof to be accepted. We then use (a somewhat cumbersome compressed oracle version of) the fact that it is unlikely to find $\eta$ zeroes, but likely to find $\eta$ zeroes given that the prover has succeeded and the extractor has failed, to conclude that the probability that the prover succeeds and the extractor fails is small. Proving that it is likely to find $\eta$ zeroes given that the prover has succeeded and the extractor has failed is non-trivial. Conditioning (or rather projecting) on these events renders the oracle register non-uniform and concentration inequalities for independent Bernoulli random variables no longer apply. Fortunately, we can simplify the analysis by symmetrizing the state with respect to permutations

of the oracle database registers not associated to a hash input queried during verification. We can then define a martingale based on the outcomes of sequentially measuring these database registers. As a result of the symmetrization, the conditional expectations of these measurements can be bounded. This bound translates into a bound on the expected number of zeroes, and an application of the Azuma-Hoeffding concentration inequality supplies the necessary tail bound for deviations from that expected value.

Overall, our analysis establishes that, even in the presence of quantum superposition queries, success in the Fischlin transform implicitly certifies the existence of some "hidden" accepting transcripts. While these transcripts may not be classically observable, they manifest themselves in the compressed oracle database, enabling straight-line extraction in the QROM.

**Organization.** Section 2 reviews the QROM and the compressed oracle framework. Section 3 proves straight-line extractability of the Fischlin transform in the QROM. Section 4 establishes zero-knowledge.

## 2 Preliminaries

For asymptotic security analysis, protocol parameters are expressed as a function of the security parameter $\lambda \in \mathbb{N}$. A (possibly negative) function $f(\lambda)$ is called negligible if $|f(\lambda)| \leq 1/p(\lambda)$ for any polynomial $p(\lambda)$ and all sufficiently large $\lambda$. We use negl to denote a negligible function. A function which is not negligible is called noticeable. For any witness relation $R \subset \mathcal{X} \times \mathcal{W}$ we call the elements of $\mathcal{X}$ instances, the elements of $\mathcal{W}$ witnesses, and define the associated language $\mathcal{L}_R = \{x \in \mathcal{X} | \exists w \in \mathcal{W} : (x, w) \in R\}$. We denote by $\delta(\cdot, \cdot)$ the maximal 1-sample computational distinguishing advantage between two probability distributions. We allow algorithms as arguments here, which indicate computational distinguishability of their outputs.

A common idealization is the random oracle model where parties have access to a random function $H$ with some domain and range depending on $\lambda$. Our random oracles have output length at most polynomial in the security parameter.

**Definition 1 (Σ-Protocol).** *A* $\Sigma$-protocol $\Sigma = (\mathcal{P}, \mathcal{V})$ *for a relation* $R \subseteq \mathcal{X} \times \mathcal{W}$ *is a three-round two-party interactive protocol between a Prover* $\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2)$ *and a verifier* $\mathcal{V}$ *where the prover produces* $(a, \mathsf{st}) \leftarrow \mathcal{P}_1(x, w)$ *and sends the commitment* $a \in \mathcal{M}$ *to the verifier, the verifier sends a random challenge* $c \leftarrow \mathcal{C}$ *(*public-coin*) and the prover sends a response* $z \leftarrow \mathcal{P}_2(\mathsf{st}, c)$*. Finally, the verifier outputs the verdict* $b \leftarrow \mathcal{V}(x, a, c, z)$*, with* $b = 1$ *meaning "valid". We denote the part of the interaction that produces the communication transcript as* $(a, c, z) \leftarrow \langle \mathcal{P}(x, w), \mathcal{V} \rangle$*. It has to satisfy the following properties:*

1. *Completeness. For any* $(x, w) \in R$*, any* $(a, c, z) \leftarrow \langle \mathcal{P}(x, w), \mathcal{V} \rangle$ *it holds* $\mathcal{V}(x, a, c, z) = 1$*.*
2. *Special Soundness. There exists a probabilistic polynomial-time (PPT) extractor algorithm* $\mathsf{Ext}_{\mathrm{ss}}$ *such that for all* $x \in \mathcal{X}$ *and for all pairs of transcripts*

*$(a, c, z)$ and $(a, c', z')$ with $c \neq c'$ such that $\mathcal{V}(x, a, c, z) = \mathcal{V}(x, a, c', z') = 1$, the extractor produces a witness $w \leftarrow \mathsf{Ext}_{\mathrm{ss}}(x, a, c, z, c', z')$ with $(x, w) \in R$.*

**Definition 2 (Min-Entropy).** *The **min-entropy** of a discrete random variable $X$ over a finite set $\mathcal{X}$ is defined as:*

$$H_\infty(X) = -\log_2 \left( \max_{x \in \mathcal{X}} \Pr[X = x] \right)$$

*where $\Pr[X = x]$ denotes the probability that $X$ takes the value $x$.*

We now list some additional properties required by our $\Sigma$-protocol the make the transformation from interactive to non-interactive, work:

**Definition 3 (Commitment Entropy).** *Let $\Sigma$ be a $\Sigma$-protocol. We say $\Sigma$ has* commitment entropy $\alpha \in \mathbb{R}_{\geq 0}$ *if , for any $(x, w) \in R$ and $a \leftarrow \mathcal{P}_1(x, w)$ it holds that $H_\infty(a) \geq \alpha$.*

**Definition 4 (Unique Responses).** *Let $\Sigma$ be a $\Sigma$-protocol. We say $\Sigma$ has* unique responses*, if for all tuples $(x, a, c, z, z') \leftarrow A(\lambda)$ with $z \neq z'$, either $\mathcal{V}(x, a, c, z) = 0$ or $\mathcal{V}(x, a, c, z') = 0$.*

**Definition 5 (Honest-Verifier Zero-Knowledge (HVZK)).** *Let $\Sigma$ be a $\Sigma$-protocol for relation $R$. We say $\Sigma$ has* honest-verifier zero-knowledge (HVZK) *if there exists a PPT algorithm $\mathsf{Sim}_{\mathrm{HVZK}}$, the zero-knowledge simulator, such that for any quantum polynomial-time (QPT) algorithm $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$, it holds that*

$$\Bigg| \Pr_{\substack{(x,w,\mathsf{st}) \leftarrow \mathcal{D}_0(1^\lambda) \\ (a,c,z) \leftarrow \langle \mathcal{P}(x,w), \mathcal{V} \rangle}} [1 \leftarrow \mathcal{D}_1(a, c, z, \mathsf{st}) \wedge (x, w) \in R] - \Pr_{\substack{(x,w,\mathsf{st}) \leftarrow \mathcal{D}_0(1^\lambda) \\ (a,c,z) \leftarrow \mathsf{Sim}_{\mathrm{HVZK}}(x)}} [1 \leftarrow \mathcal{D}_1(a, c, z, \mathsf{st}) \wedge (x, w) \in R] \Bigg| \leq \mathrm{negl}(\lambda).$$

Below we sometimes use a stronger zero-knowledge property where the simulator is able to generate proofs for a specific challenge.

**Definition 6 (Special Honest Verifier Zero-Knowledge).** *A $\Sigma$-protocol $\Sigma$ is special honest-verifier zero-knowledge (SHVZK) if it is HVZK and there exists a simulator $\mathsf{Sim}_{\mathrm{SHVZK}}$ such that $\mathsf{Sim}_{\mathrm{HVZK}}(x)$ simulates transcripts $(a, c, z)$ by sampling $c \leftarrow \mathcal{C}$ and running $(a, z) \leftarrow \mathsf{Sim}_{\mathrm{SHVZK}}(x, c)$.*

We now formally describe a Non-Interactive Zero-Knowledge(NIZK) proof with online extractors in the Random Oracle Model.

**Definition 7 (Non-Interactive Zero-Knowledge proof (NIZK)).** *A pair $(\mathcal{P}, \mathcal{V})$ of probabilistic polynomial-time algorithms is called a non-interactive zero-knowledge proof (NIZK) for relation $R$ if the following holds:*
***Completeness.*** *For any $(x, w) \in R$ and any $\pi \leftarrow \mathcal{P}^H(x, w)$ we have* $\Pr\left[\mathcal{V}^H(x, \pi) = 1\right] = 1 - \mathrm{negl}(\lambda)$.

***Zero-Knowledge.*** *There exists a probabilistic polynomial-time simulator* $\mathsf{Sim}$ *such that for any quantum polynomial-time distinguisher* $\mathcal{D} = (\mathcal{D}_0, \mathcal{D}_1)$*, we have*

$$\Big| \Pr\Big[\mathcal{D}_1^H(\pi, \mathsf{st}) = 1 : (x, w, \mathsf{st}) \leftarrow \mathcal{D}_0^H(1^\lambda),\ (x, w) \in R,\ \pi \leftarrow \mathcal{P}^H(x, w)\Big] - \Pr\Big[\mathcal{D}_1^{H'}(\pi, \mathsf{st}) = 1 : (x, w, \mathsf{st}) \leftarrow \mathcal{D}_0^H(1^\lambda), (x, w) \in R, (H', \pi) \leftarrow \mathsf{Sim}(x)\Big]\Big| \leq \mathrm{negl}(\lambda).$$

*Here, the simulator outputting* $H'$ *is to be understood as outputting (reprogramming) instructions how to instantiate* $H'$ *relative to* $H$*.*

In the QROM, straight-line extractability can be defined as follows, following [DFMS22a]. Let $\mathcal{P}^*$ be a dishonest prover that potentially outputs some additional auxiliary (possibly quantum) output $Z$ next to $\pi$. We then consider an interactive algorithm $\mathsf{Ext}$, called *online extractor*, which takes the security parameter $\lambda$ as input and simulates the answers to the oracle queries in the execution of $\mathcal{V}^H \circ \mathcal{P}^{*H}(x)$, which we define to run $(\pi, Z) \leftarrow \mathcal{P}^{*H}(x)$ followed by $v \leftarrow \mathcal{V}^H(x, \pi)$; furthermore, at the end, $\mathsf{Ext}$ outputs $w \in \mathcal{W}$. We denote the execution of $\mathcal{V}^H \circ \mathcal{P}^{*H}(x)$ with the calls to $H$ simulated by $\mathsf{Ext}$, and considering $\mathsf{Ext}$'s final output $w$ as well, as $(\pi, Z; v; w) \leftarrow \mathcal{V}^{\mathsf{Ext}} \circ \mathcal{P}^{*\mathsf{Ext}}(x)$.

**Definition 8 (Definition 3.1 in [DFMS22a] specialized to non-adaptive adversaries).** *A NIZK in the (QROM) for relation* $R \subset \mathcal{X} \times \mathcal{W}$ *is a* proof of knowledge with straight-line extractability *against non-adaptive adversaries if there exists an online extractor* $\mathsf{Ext}$*, and functions* $\varepsilon_{\mathrm{sim}}$ *(the* simulation error*)* and $\varepsilon_{\mathrm{ex}}$ *(the* extraction error*), with the following properties. For any* $\lambda \in \mathbb{N}$*, for any* $x \in \mathcal{X}$ *and for any dishonest prover* $\mathcal{P}^*$ *making no more than* $q$ *queries,*

$$\delta\big([(\pi, Z, v)]_{\mathcal{V}^H \circ \mathcal{P}^{*H}(x)}, [(\pi, Z, v)]_{\mathcal{V}^{\mathsf{Ext}} \circ \mathcal{P}^{*\mathsf{Ext}}(x)}\big) \leq \varepsilon_{\mathrm{sim}}(\lambda, q)$$

*and*

$$\Pr\big[v = \mathtt{accept} \wedge (x, w) \notin R : (\pi, Z; v; w) \leftarrow \mathcal{V}^{\mathsf{Ext}} \circ \mathcal{P}^{*\mathsf{Ext}}(x)\big] \leq \varepsilon_{\mathrm{ex}}(\lambda, q)\,.$$

*Furthermore, the runtime of* $\mathsf{Ext}$ *is polynomial in* $\lambda + q$*, and* $\varepsilon_{\mathrm{sim}}(\lambda, q)$ *and* $\varepsilon_{\mathrm{ex}}(\lambda, q)$ *are negligible in* $\lambda$ *whenever* $q$ *is polynomial in* $\lambda$*.*

**The Fischlin transform.**

We use the version of the Fischlin transform described in [Ks22].

**Definition 9 (The Fischlin transform).** *The security parameter* $\lambda$ *defines the integers* $k, \ell, t$*, related as* $k \cdot \ell = \lambda$ *and* $t = \lceil \log \lambda \rceil \cdot \ell$*. The protocol uses a hash function* $H : \mathbf{X} = \{(\mathbf{a}, i, c_i, z_i)\} = X_{\mathbf{a}} \times X_i \times X_{c_i} \times X_{z_i} \to \mathcal{Y} = \{0,1\}^\ell$*, modeled as a random oracle, and a* $\Sigma$-protocol $\Sigma = ((\mathcal{P}_{\Sigma,1}, \mathcal{P}_{\Sigma,2}), \mathcal{V}_\Sigma)$*. The prover and verifier of the Fischlin transform* $\mathsf{Fis}[\Sigma]$ *are defined as follows.*

$\mathcal{P}^{\mathsf{Fis}[\Sigma]}(x, w)$ :

*1. For each* $i \in [k]$*, compute* $(a_i\,,\ \mathrm{st}_i) \leftarrow \mathcal{P}_{\Sigma,1}(x, w)$

*2. Set* $\boldsymbol{a} = (a_i)_{i \in [k]}$*, and initialize* $c_i = -1$ *for each* $i \in [k]$

*3. For each* $i \in [k]$*, do the following:*

*(a) If $c_i > t$, abort. Otherwise increment $c_i$ and compute $z_i = \mathcal{P}_{\Sigma,2}(\mathrm{st}_i, c_i)$*
*(b) If $H(\boldsymbol{a}, i, c_i, z_i) \neq 0^\ell$, repeat Step 3a*
*4. Output $\pi = (\mathbf{a}, \mathbf{c}, \mathbf{z})$, where $\mathbf{c} = (c_i)_{i\in[k]}$ and $\mathbf{z}$ are defined analogously.*
$\mathcal{V}^{\mathsf{Fis}[\Sigma]}(x, \pi)$ :
*1. Parse $(\mathbf{a}, \mathbf{c}, \mathbf{z}) = \pi$*
*2. For each $i \in [k]$, verify that $H(\boldsymbol{a}, i, c_i, z_i) = 0^\ell$ and $\mathcal{V}_\Sigma(x, (a_i, c_i, z_i)) = 1$, aborting with output 0 if not*
*3. Accept by outputting 1*

For simplicity, we set the parameter $t$ equal to the size of the challenge space of the underlying $\Sigma$-protocol for the remainder of this article.

### 2.1 Quantum Random Oracle (QROM) and Compressed Oracle Technique

In this section we summarize the mathematical tools and notational conventions used throughout the paper. We review basic facts about Hilbert spaces and operators, introduce the computational and Fourier bases for finite Abelian groups, and recall the standard encoding of functions into quantum oracles. These foundations are needed to formalize the compressed oracle technique and to analyze the behavior of quantum adversaries making oracle queries.

**Hilbert Spaces and Operators.** We work over finite-dimensional complex Hilbert spaces. Unless stated otherwise, $\mathcal{H} = \mathbb{C}^d$ for some $d$, with the usual braâĂŞket notation. Linear operators between spaces $\mathcal{H}$ and $\mathcal{H}'$ are denoted $\mathcal{L}(\mathcal{H}, \mathcal{H}')$, with $\mathcal{L}(\mathcal{H})$ for endomorphisms. A (pure) quantum state is a unit vector $|\psi\rangle \in \mathcal{H}$.

**Norms.** For $A \in \mathcal{L}(\mathcal{H}, \mathcal{H}')$, $\|A\|$ denotes the operator norm. If $\mathcal{H} = \bigoplus_{i=1}^m \mathcal{H}_i$ and $A$ acts as $B_i$ on each $\mathcal{H}_i$, then $\|A\| = \max_i \|B_i\|$. When $A$ is defined only on a subspace, it can be extended by zero outside without ambiguity in its norm.

Any finite set $\mathcal{Y}$ defines an associated quantum register with state space $\mathbb{C}\mathcal{Y}$ computational basis $\{|y\rangle\}_{y\in\mathcal{Y}}$. We will consider the extension $\overline{\mathcal{Y}} = \mathcal{Y} \cup \{\perp\}$ and the superspace $\mathbb{C}\overline{\mathcal{Y}} \supset \mathbb{C}\mathcal{Y}$.

We will denote the function table stored in a quantum register by $|H\rangle = \bigotimes_x |H(x)\rangle \in (\mathbb{C}\mathcal{Y})^{\otimes\mathcal{X}}$.We implicitly consider the different registers to be labeled by $x \in \mathcal{X}$ in the obvious way.

**Compressed Oracle Technique.** The compressed oracle formalism introduced by Zhandry [Zha19] is a way to simulate a quantum-accessible random oracle in a way that provides additional features for reductions. Instead of storing the full random function, we maintain a *compressed database $D$* that records only the points queried so far and their outputs. This suffices for simulating any adversary's interaction while simplifying the analysis.

Instead of a uniformly random choice of function, consider a superposition $\frac{1}{\sqrt{|\mathfrak{H}|}} \sum_H |H\rangle$, where $\mathfrak{H}$ is the set of functions $H$ is sampled uniformly from. This is the purified oracle which is indistinguishable from the original random oracle for any (quantum) query algorithm since the queries commute with measuring the superposition. The initial state of the oracle registers is given by

$$|\Pi_0\rangle = \frac{1}{\sqrt{|\mathfrak{H}|}} \sum_H |H\rangle = \frac{1}{\sqrt{|\mathfrak{H}|}} \bigotimes_x \left( \sum_y |y\rangle_{D_x} \right) =: \bigotimes_x |+_{\mathcal{Y}}\rangle_{D_x}. \tag{1}$$

**Compressed Oracle.** The Compressed Oracle[1] is now obtained by redefining each local register $D_x$ to have state space $\mathbb{C}\overline{\mathcal{Y}}$, embedding the superposition oracle into it, and applying the basis change unitary

$$\mathrm{Comp} = |\perp\rangle\langle +^{\ell}| + |+^{\ell}\rangle\langle \perp| + (\mathbb{1} - |+^{\ell}\rangle\langle +^{\ell}| - |\perp\rangle\langle \perp|). \tag{2}$$

This compression operator acts on a single subregister of the compressed oracle database register. As for any operator, we denote by $\mathrm{Comp}_{D_x}$ the compression operator applied to a particular register $D_x$. For the compression operator, we additionally abuse notation slightly by writing

$$\mathrm{Comp}_{D_S} = \bigotimes_{x \in S} \mathrm{Comp}_{D_x}$$

and thus also $\mathrm{Comp}_D = \left(\mathrm{Comp}^{\otimes |\mathcal{X}|}\right)_D$.

The compression operator $\mathrm{Comp}_D := \bigotimes_x \mathrm{Comp}_{D_x}$ maps $|\Pi_0\rangle$ to

$$\begin{aligned} |\Delta_0\rangle := \mathrm{Comp}\,|\Pi_0\rangle &= \left( \bigotimes_x \mathrm{Comp}_{D_x} \right) |+_{\mathcal{Y}}^{|\mathcal{X}|}\rangle \\ &= \bigotimes_x \mathrm{Comp}_{D_x} |+_{\mathcal{Y}}\rangle = \bigotimes_x |\perp\rangle = |\perp^{|\mathcal{X}|}\rangle, \end{aligned}$$

which is the quantum representation of the trivial database that maps any $x \in \mathcal{X}$ to $\perp$.

As a consequence, the internal state of the compressed oracle after $q$ queries is supported by computational basis states $|D\rangle$ for which $D(x) = \perp$ (respectively $\hat{D}(x) = \perp$ ) for all but (at most) $q$ choices of $x$. We call the number of registers that are in a state other than $|\perp\rangle$ the *size* of the database.

An oracle query to the superposition oracle is implemented using the unitary operator $O$ given by

$$O : |\mathbf{x}\rangle_X |y\rangle_Y \otimes |H\rangle_D \mapsto |\mathbf{x}\rangle_X |y \oplus H(x)\rangle_Y \otimes |H\rangle_D. \tag{3}$$

When the compressed oracle is queried, a unitary $O_{XYD}$, acting on the query registers $X$ and $Y$ and the oracle register $D$, is applied, given by

$$O_{XYD} = \sum_x |x\rangle \langle x|_X \otimes O^x_{YD_x},$$

[1] More precisely we are using the *pre-compressed* oracle, the compressed oracle is obtained from the pre-compressed oracle via the standard sparse representation

with

$$O^x_{YD_x} = \mathrm{Comp}_{D_x}\mathrm{CNOT}_{YD_x}\mathrm{Comp}_{D_x}$$

where $\mathrm{CNOT}_{YD_x}|y\rangle|y_x\rangle = |y \oplus y_x\rangle|y_x\rangle$ for $y, y_x \in \{0,1\}^\ell$ and acts as identity on $|y\rangle|\perp\rangle$

**Query algorithms.** A quantum algorithm interacting with a compressed oracle works as follows. It has registers $XYZO$, where $X$ is the query input, $Y$ the query output, $Z$ is the local workspace of the algorithm and $O$ is the register that is measured to produce the final output. $D$ is the compressed oracle database register that is inaccessible to the algorithm except via oracle queries.

Let $|\phi_i\rangle$ denote the adversary's state before the $(i+1)$-th query, and $|\phi'_i\rangle$ the state immediately after that query,

$$|\phi'_i\rangle = O_{XYD}|\phi_i\rangle.$$

The adversary now applies a unitary $U$[2] on $XYZ$ to obtain

$$|\phi_{i+1}\rangle = U_{XYZO}|\phi'_i\rangle.$$

Thus, after $q$ queries, the overall computation evolves as:

1. Start in some initial state $|\phi_0\rangle$.
2. For each $1 \le i \le q$:
   (a) Apply the oracle $O_{XYD}$ to obtain $|\phi'_i\rangle$.
   (b) Apply $U_{XYZO}$ to obtain $|\phi_i\rangle$.
3. At the end, perform a computational basis measurement of register $O$.

*Reprogrammming in the QROM.* We need an adaptive reprogramming lemma from [GHHM21].

Let $\mathcal{X}, \mathcal{X}', \mathcal{Y}$ be finite sets. Consider the game $\mathtt{Repro}_b$, where $b \in \{0,1\}$:

1. Sample an initial random oracle $O_0 : \mathcal{X} \to \mathcal{Y}$.
2. Define $O_1$ to be $O_0$, but with access to a Reprogram oracle: on input a probability distribution $p$ on $\mathcal{X} \times \mathcal{X}'$, it samples $(x, x') \leftarrow p$ and $y \leftarrow \mathcal{Y}$, sets $O_1(x) := y$, and returns $(x, x')$.
3. An adversary $A$ is given oracle access to $O_b$ (quantum) and to Reprogram, and outputs a bit $b'$.

The distinguishing advantage is

$$\left|\Pr[\mathtt{Repro}_1^A \implies 1] - \Pr[\mathtt{Repro}_0^A \implies 1]\right|.$$

Let $p_X$ be the marginal of $p$ on $X$. Define

$$p^{(r)}_{\max} := \mathbb{E}\left[\max_x p^{(r)}_X(x)\right],$$

where $p^{(r)}$ is the adversary's $r$-th input to the Reprogram oracle, and the expectation is over the adversary's randomness up to round $r$.

[2] WLOG, this unitary does not depend on $i$: We can include a counter register $Z_N$ in $Z$ which keeps track of the number of queries that have been applied thus far, and set $U$ as a controlled unitary controlled on $Z_N$ and increment $Z_N$.

**Lemma 1.** *[Simplified from Theorem 1 in [GHHM21]] For any adversary $D$ making $R$ reprogramming calls and $q$ quantum queries in total,*

$$\left|\Pr[\texttt{Repro}_1^D \implies 1] - \Pr[\texttt{Repro}_0^D \implies 1]\right| \leq \sum_{r=1}^{R} \left( \sqrt{q \cdot p_{\max}^{(r)}} + \tfrac{1}{2} q \cdot p_{\max}^{(r)} \right).$$

*Quantum union bound* We need a lemma referred to as the quantum union bound.

**Lemma 2 (Theorem 1.1 in [OV22]).** *Let $|\psi\rangle \in \mathcal{H}$ be a quantum state and let $P_i$, $i = 1, ..., N$ be projectors on $\mathcal{H}$ such that $\langle\psi| P_i |\psi\rangle \geq 1 - \varepsilon_i$. Then*

$$\left\| \left( \prod_{i=1}^{N} P_i \right) |\psi\rangle \right\|^2 \geq 1 - 4 \sum_{i=1}^{N} \varepsilon_i.$$

*Chernoff Bounds.* The following lemma is taken from combining Theorem 4.4 and 4.5 from [MU05].

**Lemma 3 (Chernoff Bounds).** *Let $X_1, \ldots, X_n$ be independent Bernoulli variables such that $\Pr(X_i = 1) = p_i$. Let*

$$X = \sum_{i=1}^{n} X_i \quad \text{and} \quad \mu = \mathbb{E}[X].$$

*Then the following Chernoff bounds hold:*

1. *(**Upper tail** ) For $0 < \delta \leq 1$,*

$$\Pr(X \geq (1+\delta)\mu) \leq e^{-\mu\delta^2/3}. \tag{4.2}$$

2. *(**Lower tail**) For $0 < \delta \leq 1$,*

$$\Pr(X \leq (1-\delta)\mu) \leq e^{-\mu\delta^2/2}. \tag{4.4}$$

The random variables $X_i$ defined by sequential measurements form a martingale difference sequence rather than independent trials. We therefore need to apply the following martingale Chernoff bound derived from Azuma-Hoeffding Inequality below, from Theorem 12.6 in [MU05]

**Lemma 4 (Azuma–Hoeffding Inequality).** *Let $(Y_0, Y_1, \ldots, Y_n)$ be a martingale such that for all $i = 1, \ldots, n$,*

$$|Y_i - Y_{i-1}| \leq c_i \quad \text{almost surely.}$$

*Then for any $t > 0$,*

$$\Pr[\, Y_n - Y_0 \geq t \,] \;\leq\; \exp\left( -\frac{t^2}{2\sum_{i=1}^{n} c_i^2} \right).$$

## 3 Proof of Post Quantum Security of the Fischlin Transform

We are now set to prove that the Fischlin transform is a NIZK (Definition 7) with straight-line extractability (Definition 8) and zero knowledge in Quantum Random Oracle Model.

**Proof Intuition.** We begin by describing the proof strategy for the deterministic prover case (Section 3.1), which contains the key ideas.

*The counting argument and why independence breaks.* The high-level strategy is a counting argument by contradiction. A valid Fischlin proof requires $k$ hash outputs to equal $0^\ell$, an event that occurs independently with probability $2^{-\ell}$ per query. By a standard Chernoff bound (Lemma 3), the total number of $0^\ell$ outcomes in the oracle database concentrates tightly around its mean $\mu = 2^{-\ell}kN$, so seeing significantly more than $\mu$ registers in state $|0^\ell\rangle$ is exponentially unlikely (Lemma 5). The plan is therefore to show that if the prover succeeds *and* the extractor fails, the database must contain *many more* than $\mu$ registers in state $|0^\ell\rangle$, contradicting this tail bound.

The difficulty is that conditioning on both events simultaneously destroys independence. The prover-success projector $P_s$ forces the $k$ oracle registers corresponding to the verified transcripts into the state $|0^\ell\rangle$. The extractor-failure projector $E_f$ then further constrains the *remaining* registers, correlating them with the verified ones. Once both projectors have been applied, the oracle outputs stored in the database registers are no longer independent, and the Chernoff bound used in Lemma 5 does not apply. A new strategy is needed to lower-bound the number of $0^\ell$ outcomes in this correlated state.

*Decomposition, symmetrization, and the martingale.* The resolution proceeds in three steps. First, we *decompose* the post-success state. When the success projector forces $k$ registers to $|0^\ell\rangle$ and we switch to the compressed basis, the term $(\mathrm{Comp}|0^\ell\rangle)^{\otimes k}$ is not simply $|0^\ell\rangle^{\otimes k}$; it is a superposition that Lemma 6 decomposes by the *subset $S$ of registers that carry* $|0^\ell\rangle$:

$$(\mathrm{Comp}|0^\ell\rangle)^{\otimes k} \;=\; \sum_{\substack{S\subset[k]\\ |S|\geq(1-\gamma)k}} |0^\ell\rangle_S^{\otimes|S|}\,|I^{(|S|)}\rangle_{S^c} \;+\; |\delta_{\gamma,k}\rangle,$$

where the tail $|\delta_{\gamma,k}\rangle$ is small. This isolates the "bulk" of the state, in which at least $(1-\gamma)k$ of the $k$ registers carry $|0^\ell\rangle$ in the compressed picture, from a negligible remainder. Crucially, applying $E_f$ to this decomposed state forces most of the *remaining* $m = k(N-1)$ database registers (those outside the verified transcript) into the compressed initial state $|\perp\rangle$; and since $\mathrm{Comp}|\perp\rangle = |+^\ell\rangle$, each such register contributes a $0^\ell$ outcome with probability $2^{-\ell}$ when measured.

Second, we *symmetrize* the resulting state over permutations of these $m$ registers (Equation (13)). The post-conditioning state is not independent across

registers, so we cannot apply Chernoff directly. Symmetrization is the substitute: it ensures the state lies in the subspace $W_n^m$ (where at least $n$ of the $m$ registers are in state $|+^\ell\rangle$), which is closed under permutations. For any such symmetric state, Lemma 7 gives a lower bound on the probability of obtaining $0^\ell$ when measuring any single register; and, crucially, the post-measurement state remains in the same class of subspace $W_{n-1}^{m-1}$, so the bound applies recursively to each subsequent register.

Third, this recursive structure defines a *martingale*. Let $X_i$ indicate whether measuring the $i$-th register yields $0^\ell$, and let $Z_i = \mathbb{E}[X_i \mid X_{<i}]$ be the conditional expectation, which Lemma 7 lower-bounds at each step. The sum $\hat{Z}_i = \sum_{j \le i}(Z_j - X_j)$ is a martingale with bounded differences $|\hat{Z}_i - \hat{Z}_{i-1}| \le 1$, so the Azuma–Hoeffding inequality (Lemma 4) applies. Combined with the lower bound on the cumulative mean $\mu' = \sum_i Z_i$ from Lemma 9, this gives: conditioned on prover success and extractor failure, the database contains at least $\eta \approx \mu'$ registers in state $|0^\ell\rangle$ with high probability. Since $\eta > \mu$ under our parameter choices, this contradicts the upper tail bound of Lemma 5, completing the argument.

The most involved part of the proof is now done. We lift the result to general malicious provers in Section 3.2. This can be done by decomposing a general prover's final state according to when it has first queried an input starting with a fixed commitment vector. As the relevant part of the compressed oracle database is untouched before that query, we can *post-select* on that query with input starting with a fixed commitment vector actually being made, without query cost. This allows the application of the result for deterministic provers to bound the probability of prover success and extractor failure *conditioned* on a fixed commitment vector.

**Prover and Extractor Definitions.** We describe how the prover and extractor work in our analysis for QROM-extractability of the Fischlin Transform described in Definition 9. While it is obviously important for the security of the NIZK that not all transcripts are valid, we will only explicitly keep track of the part of the compressed oracle register that corresponds to transcripts that are accepted by the $\Sigma$-protocol-verifier, as this turns out to be sufficient for extractability.

**Definition 10 (Extractor $\mathsf{Ext}_{\mathsf{Fis}}$).** *Let $\Sigma$ be a $\Sigma$-protocol. The quantum straight-line extractor for the Fischlin transform is given the statement $x$ and simulates the quantum random oracle using a compressed oracle. After the adversary has finished, it receives the proof $\pi$ output by the adversary.*

*The extractor works as follows.*

$\mathsf{Ext}_{Fis}(x)$ *:*

*1. The prover receives oracle access to a compressed oracle with database $D$ and produces a proof $\pi$.*

*2. Verify the proof. If verification fails, abort.*

*3. Measure the compressed oracle database register. Let $((\boldsymbol{a}, i, c, z), (\boldsymbol{a}, i, c', z'))$ be the lexicographically first pair of inputs in the measured database such that*

$(c,z) \neq (c',z')$, *and* $\mathcal{V}_\Sigma(x,a_i,c,z) = \mathcal{V}_\Sigma(x,a_i,c',z') = 1$. *If no such query is found, abort.*
*4. Output* $\mathrm{Ext}_{\mathrm{ss}}(x,a_i,c,z,c',z')$.

### 3.1 Extractability – Deterministic Commitments

We first prove that the extractor will succeed in extracting a witness whenever a prover outputs a valid proof *that starts with a fixed vector of commitments*. We will later lift that result to arbitrary provers (see Section 3.2).

In this section, we consider therefore a modified NIZK $\Pi_{\mathbf{Fis}}^{(\mathbf{a}_0)}$ parameterized by a commitment vector $\mathbf{a}_0$ where the verifier only accepts if the proof is valid *and* the proof starts with the commitment vector $\mathbf{a}_0$. We can thus further restrict our attention to only the compressed oracle registers for inputs of the form $(\mathbf{a},i,c,z)$ such that $\mathbf{a} = \mathbf{a}_0$, and $(a_i,c,z)$ is accepted by $\mathcal{V}_\Sigma$. As we assume the $\Sigma$-protocol to have unique responses, any input is specified among the remaining set of inputs by the pair $(i,c) \in [k] \times \mathcal{C}$.

An arbitrary dishonest prover $\mathcal{P}^*$ for the Fischlin transform is a quantum query algorithm as described in Section 2.1, with the output register $O$ having the right structure to hold a proof $\pi = (\mathbf{a},\mathbf{c},\mathbf{z})$. The joint state of the prover and the oracle in the uncompressed basis (i.e., in the superposition oracle picture, not the pre-compressed oracle) after the unitary part of a malicious prover has finished but $O$ has not been measured to produce the prover's output is denoted as

$$|\phi\rangle = \sum_{\mathbf{x},H} 2^{-\frac{lkN}{2}} |\mathbf{x}\rangle_O |\phi^{\mathbf{x},H}\rangle_Z |H\rangle_D \tag{4}$$

Here, $D = (D_{(i,c)})_{i\in[k],c\in\mathcal{C}}$ and we subsume the malicious prover's query input and output registers as well as the remainder of the compressed oracle database in the register $Z$. The marginal of the register $D$ remains uniform (as the query operator is a controlled unitary with control register $D$).

To bound the probability for the event where the prover succeeds and the extractor fails, the analysis in the following sections is used.

To facilitate our statistical arguments we define the following projector.

$$P_\eta = \sum_{\bar{\omega}(H)\geq\eta} |H\rangle\langle H|_D$$

where $\bar{\omega}(H)$ is defined as $\bar{\omega}(H) = \left|\left\{\mathbf{x} \mid H(\mathbf{x}) = 0^\ell\right\}\right|$.

Applying the projector $(P_\eta)_D$ to the state $|\phi\rangle_{OZD}$ yields a vector of small norm for $\eta$ significantly bigger than $2^{-l}kN$ (where $N = \Omega(2^\ell \log k)$ is the challenge space size) by a standard Chernoff bound.

**Lemma 5.** *Let* $\mu = 2^{-l}kN$, $\delta > 0$ *and* $\eta = (1+\delta)\mu$. *We have the following probability bound*

$$||P_\eta|\phi\rangle||^2 \leq \exp\left(-\frac{\delta^2\mu}{3}\right) =: \varepsilon''.$$

*Proof.* Let $X_{i,c}$, $i \in [k], c \in \mathcal{C}$ be random variables defined as measurement results of the state $|\phi\rangle$ by

$$X_{i,c} = \begin{cases} 1 & \text{if measuring register } D_{i,c} \text{ yields outcome } 0^\ell \\ 0 & \text{else} \end{cases},$$

Clearly $||P_\eta|\phi\rangle||^2 = \sum_{i\in[k],c\in\mathcal{C}} X_{i,c} =: \mathfrak{S}$. Also, measuring the subregisters of $D$ in the computational basis yields independent uniformly random $l$-bit strings,[3] so $\Pr[X_{i,c} = 1] = 2^{-l}$ independently for all $(i, c)$. By a standard Chernoff bound, Lemma 3, we thus have

$$\mathbb{P}\left[(\mathfrak{S} \geq \eta = (1+\delta)\mu)\right] \leq \mathbb{P}\left[|\mathfrak{S} - \mu| \geq \delta\mu\right] \leq \exp\left(-\frac{\delta^2\mu}{3}\right).$$

□

Thinking classically for a moment, the proof strategy is via contradiction: a malicious prover succeeding *and* the extractor failing with significant probability would imply that $||P_\eta|\phi\rangle||^2$ must in fact be large, contradicting the above tail bound. As we are dealing with a malicious quantum prover and with the compressed oracle that is directly manipulated (measured) by the extractor, the quantum analogue of this argument is significantly more involved.

Denote $|\phi_{\text{succ}}\rangle = P_s|\phi\rangle$, where

$$P_s = \sum_{\mathbf{x}} |\mathbf{x}\rangle\langle\mathbf{x}|_O \otimes |0^\ell\rangle\langle 0^\ell|^{\otimes k}_{D_{\mathsf{hi}(\mathbf{x})}} \tag{5}$$

Is the success projector, i.e., the probability that the malicious prover succeeds is $\| \, |\phi_{\text{succ}}\rangle \, \|^2$. Here, $D_{\mathsf{hi}(\mathbf{x})}$ is the subregister of $D$ containing the hash inputs needed to verify the proof $\mathbf{x}$, i.e., $\mathsf{hi}(\mathbf{x}) = \{(\mathbf{a}, i, c_i, z_i) \mid i \in [k], V_\Sigma(a_i, c_i, z_i) = 1\}$.

We now see the state of the oracle in the compressed picture. Writing $\mathbf{x} = (\mathbf{a}, \mathbf{c}, \mathbf{z})$, we apply Comp,

$$\begin{aligned} |\phi'_{\text{succ}}\rangle = \text{Comp}P_s|\phi\rangle &= \sum_{\mathbf{x}} |\mathbf{x}\rangle_O |\phi^{\mathbf{x}}\rangle_{Z D_{\mathsf{hi}^c}} (\text{Comp}|0^\ell\rangle^{\otimes k})_{D_{\mathsf{hi}(x)}} \\ &= \sum_{\mathbf{x}} |\mathbf{x}\rangle_O |\phi^{\mathbf{x}}\rangle_{Z D_{\mathsf{hi}^c}} \left( \sum_{\substack{S\subset \mathsf{hi}(\mathbf{x}) \\ s=|S|\geq(1-\gamma)k}} |0^\ell\rangle^{\otimes s}_{D_S} |\Gamma\rangle_{D_{\mathsf{hi}(\mathbf{x})\setminus S}} + |\delta_{\gamma,k}\rangle_{D_{\mathsf{hi}(\mathbf{x})}} \right), \end{aligned} \tag{6}$$

where we have used Lemma 6, and $0 < \gamma = \gamma(l, k) \leq \frac{1}{2}$ is chosen later in Lemma 13..

**Lemma 6.** *There exist sub-normalized states $|\Gamma^{(s)}\rangle$, $0 \leq s \leq k$ such that*

$$(\text{Comp}|0^\ell\rangle)^{\otimes k}_D = \sum_{\substack{S\subset[k] \\ s=|S|\geq(1-\gamma)k}} |0^\ell\rangle^{\otimes s}_{D_S} |\Gamma^{(|S|)}\rangle_{D_{S^c}} + |\delta_{\gamma,k}\rangle,$$

[3] essentially by the correctness of the superposition oracle methodology

*such that* $\langle 0^\ell|_{D_i} |\Gamma^{(s)}\rangle_{D_{S^c}} = 0$ *for all* $S$ *and all* $i \in S^c$*, where* $D = D_1 \dots D_k$*, and the tail bound is given by* $|| \, |\delta_{\gamma,k}\rangle \, ||^2 \leq e^{-\left(\gamma - 2\cdot 2^{-l}\right)k/2} =: \varepsilon_{\gamma,k}$.

*Proof.* We bound

$$\begin{aligned}
(\mathrm{Comp}|0^\ell\rangle)^{\otimes k} &= \left((|\perp\rangle\langle +^\ell| + |+^\ell\rangle\langle\perp| + (\mathbb{1} - |+^\ell\rangle\langle +^\ell| - |\perp\rangle\langle\perp|)\,|0^\ell\rangle\right)^{\otimes k} \\
&= \left(|0^\ell\rangle + 2^{-l/2}|\perp\rangle - 2^{-l/2}\,|+^\ell\rangle\right)^{\otimes k} \\
&= \left(|0^\ell\rangle + 2^{-l/2}|\perp\rangle - 2^{-l} \sum_{z\in\{0,1\}^\ell} |z\rangle\right)^{\otimes k} \\
&= \left((1-2^{-l})|0^\ell\rangle + 2^{-l/2}|\perp\rangle - 2^{-l} \sum_{\substack{z\in\{0,1\}^\ell \\ z\neq 0^\ell}} |z\rangle\right)^{\otimes k} \\
&= \sum_{\substack{S\subset[k] \\ s=|S|\geq(1-\gamma)k}} |0^\ell\rangle_{D_S}^{\otimes s} |\Gamma^{(s)}\rangle_{D_{S^c}} + \sum_{\substack{S\subset[k] \\ s=|S|<(1-\gamma)k}} |0^\ell\rangle_{D_S}^{\otimes s} |\Gamma^{(s)}\rangle_{D_{S^c}}, \\
&= \sum_{\substack{S\subset[k] \\ s=|S|\geq(1-\gamma)k}} |0^\ell\rangle_{D_S}^{\otimes s} |\Gamma^{(s)}\rangle_{D_{S^c}} + |\delta_{\gamma,k}\rangle
\end{aligned} \tag{7}$$

where $\langle 0|_{D_i} |\Gamma^{(|S|)}\rangle_{D_{S^C}} = 0$ for all $S \subset [k]$ and $i \in S^C$.

We now bound $|| \, |\delta_{\gamma,k}\rangle \, ||^2$. We can use a standard Chernoff bound for the number of results $0^\ell$ we would obtain were we to measure the entire state in the computational basis. The mean is

$$\mu = k\left(1 - 2^{-l}\right)^2$$

Towards using a Chernoff bound we define $\delta$ via $(1-\delta)\mu = (1-\gamma)k$, i.e.,

$$\delta = 1 - \frac{(1-\gamma)k}{\mu} = 1 - \frac{(1-\gamma)k}{kp} = 1 - \frac{(1-\gamma)}{p}$$

As $|| \, |\delta_{\gamma,k}\rangle \, ||^2$ is the probability of obtaining less then $(1-\gamma)k$ outcomes $0^\ell$, we then obtain

$$\begin{aligned}
&|| \, |\delta_{\gamma,k}\rangle \, ||^2 \leq e^{-\delta^2\mu/2} = e^{-[1-(1-\gamma)/p]^2 \times kp/2} = e^{-[p-(1-\gamma)]^2 k/2p} \\
&\leq e^{-\left[\left(1-2^{-l}\right)^2 - (1-\gamma)\right]^2 k/2p} \leq e^{-\left[2^{-2l} - 2\cdot 2^{-l} + \gamma\right]k/2p} \leq e^{-\left(\gamma - 2\cdot 2^{-l}\right)k/2}
\end{aligned}$$

□

**Analyzing the extractor.** Denote by $\mathcal{E}_f$, and $\mathcal{E}_s = \mathbb{1} - \mathcal{E}_f$ the projections onto the subspace where the extractor $\mathsf{Ext}_{\mathsf{Fis}}$ (definition 10 ) of the Fischlin Transform

fails, and succeeds, respectively. The operator $\mathcal{E}_f$ projects onto the states with only one register $D_{\mathbf{a},i,c,z}$ for each $i$ that is not in the state $|\perp\rangle$ :

$$(\mathcal{E}_f)_D = \left( \sum_{H:\mathsf{Ext}_{\mathsf{Fis}} \text{ fails}} |H\rangle\langle H| \right)_D = \bigotimes_{i=1}^{k} \left(P^N_{\perp,1}\right)_{D_{i,*}},$$

where $P^N_{\perp,1}$ is the projector onto the subspace

$$W^N_{\perp,1} = \mathrm{span}\{|\psi\rangle_{D_{i,c}} |\perp^{N-1}\rangle_{D_{i,c^c}} \,|\, |\psi\rangle \in \mathbb{C}^{2^\ell}\},$$

and for each repetition index $i \in [k]$, we write

$$D_{i,*} \coloneqq \{D_{(\mathbf{a}_0,i,c,z)} : c \in \mathcal{C}\}$$

for the collection of compressed oracle registers corresponding to all valid transcripts with fixed commitment vector $\mathbf{a}_0$ and repetition index $i$, ranging over all challenges $c \in \mathcal{C}$. Above, we have also used $D_{i,c^c} = (D_{i,c'})_{c' \in \mathcal{C}\setminus\{c\}}$.

We need to bound the probability that the extractor fails and the prover succeeds, so we study the norm of the state Here $\mathrm{Comp}_*$ indicates compression operator Comp applied to each register of set $*$.

$$\begin{aligned}
&\mathrm{Comp}_D \mathcal{E}_f \mathrm{Comp}_D P_s |\phi\rangle \\
&= \mathrm{Comp}_D \mathcal{E}_f \sum_{\mathbf{x}} |\mathbf{x}\rangle_O |\phi^{\mathbf{x}}\rangle_{ZD_{\mathsf{hi}(x)C}} \left( \sum_{\substack{S \subset \mathsf{hi}(\mathbf{x}) \\ s=|S|\geq(1-\gamma)k}} |0^\ell\rangle^{\otimes s}_{D_S} |\Gamma^{(s)}\rangle_{D_{\mathsf{hi}(\mathbf{x})\setminus S}} + |\delta_{\gamma,k}\rangle_{D_{\mathsf{hi}(\mathbf{x})}} \right) \\
&= \sum_{\mathbf{x}} |\mathbf{x}\rangle_O \sum_{\substack{S \subset \mathsf{hi}(\mathbf{x}) \\ s=|S|\geq(1-\gamma)k}} |\psi^{\mathbf{x}}_S\rangle_{ZD_{S^C}} \left(\mathrm{Comp}|0^\ell\rangle\right)^{\otimes s}_{D_S} + |\delta'_{\gamma,k}\rangle_{OZD} \qquad (8)
\end{aligned}$$

where we implicitly define $|\psi^{\mathbf{x}}_S\rangle_{ZD_{S^C}} = \mathrm{Comp}_{D_{S^C}} \left|\hat{\psi}^{\mathbf{x}}_S\right\rangle_{ZD_{S^C}}$ with

$$\left|\hat{\psi}^{\mathbf{x}}_S\right\rangle_{ZD_{S^C}} |0^\ell\rangle^{\otimes s}_{D_S} = \mathcal{E}_f |\phi^{\mathbf{x}}\rangle_{ZD_{\mathsf{hi}(x)C}} |\Gamma^{(s)}\rangle_{D_{\mathsf{hi}(x)\setminus S}} |0^\ell\rangle^{\otimes s}_{D_S}$$

and have set and

$$|\delta'_{\gamma,k}\rangle_{OZD} = \mathrm{Comp}_D \mathcal{E}_f \sum_{\mathbf{x}} |\mathbf{x}\rangle_O |\phi^{\mathbf{x}}\rangle_{ZD_{\mathsf{hi}(x)C}} \left( |\delta_{\gamma,k}\rangle_{D_{\mathsf{hi}(\mathbf{x})}} \right). \qquad (9)$$

Note that

$$\| \, |\delta'_{\gamma,k}\rangle \, \|^2 \leq \varepsilon_{\gamma,k} \qquad (10)$$

since both $\mathcal{E}_f$ and $\mathrm{Comp}_D$ are contractions, so the bound follows directly from Lemma 6.

We will now show that measuring a normalized version of

$$|\psi^{\mathbf{x}}\rangle_{ZD} = \sum_{\substack{S\subset \mathsf{hi}(\mathbf{x}) \\ s=|S|\geq(1-\gamma)k}} |\psi^{\mathbf{x}}_S\rangle_{ZD_{S^C}} \left(\mathrm{Comp}|0^\ell\rangle\right)^{\otimes s}_{D_S} \tag{11}$$

in the computational basis yields many 0s. Comparing to the final prover-oracle state without conditioning on success and extractor failure yields a bound on $\|\mathcal{E}_f \mathrm{Comp} P_s|\phi\rangle\|^2$, i.e., the probability that the prover succeeds, but the extractor still fails.

For $m, n \in \mathbb{N}$, $m \geq n$, we define the subspaces

$$W^m_n = \mathrm{span}\Big\{ |+^\ell\rangle^{\otimes s}_{D_T} |\psi\rangle_{D_{T^c}} \\ \Big| T\subset[m], s \coloneqq |T| \geq n, |\psi\rangle \in \left(\mathbb{C}^{2^\ell}\right)^{\otimes m-s} \Big\} \subseteq \left(\mathbb{C}^{2^\ell}\right)^{\otimes m}.$$

The projector $\mathcal{E}_f$ projects onto a subspace where for each $i$, only one register $D_{\mathbf{a},i,c,z}$ is not in the state $|\perp\rangle$. As $\mathrm{Comp}\,|\perp\rangle = |+^\ell\rangle$ and $1 \leq i \leq k$, we have

$$|\psi^{\mathbf{x}}\rangle \in W^m_n \otimes \mathcal{H}_{D_{\mathsf{hi}(x)}} \otimes \mathcal{H}_Z, \tag{12}$$

where $m = k(N-1)$, $n = m - \gamma k$, $W^m_n$ is viewed as a subspace of the state space of $D_{\mathsf{hi}(\mathbf{x})^c}$, and $\mathcal{H}_{D_{\mathsf{hi}(x)}}$ and $\mathcal{H}_Z$ are the state spaces of registers $D_{\mathsf{hi}(x)}$ and $Z$, respectively.

In the following, we denote $\tilde{D} = D_{\mathsf{hi}(x)^c}$ and its subregisters by $\tilde{D}_i$, $1 \leq i \leq m$. For simplicity, we sometimes write $|\alpha\rangle_I \coloneqq |\alpha\rangle_{\tilde{D}_I}$ to indicate that for some $I \subset [m]$, the register $\tilde{D}_I$ is in state $|\alpha\rangle$.

We want to bound $\left\|(P_\eta)_{\tilde{D}}|\psi^{\mathbf{x}}\rangle\right\|$ for some integer $\eta$. This quantity is invariant under permutation of the $m$ subregisters of $\tilde{D}$. We therefore analyze a symmetrized version of the state.

Consider the permutation action of any permutation $\pi \in S_m$, where $S_m$ is the group of permutations of $m$ elements, on the subregisters of $\tilde{D}$,

$$\pi\,|y_1\rangle\,|y_2\rangle\ldots|y_m\rangle = \left|y_{\pi^{-1}(1)}\right\rangle \left|y_{\pi^{-1}(2)}\right\rangle \cdots \left|y_{\pi^{-1}(m)}\right\rangle.$$

We can now define the state

$$|\psi^{\mathbf{x}}_{\mathrm{sym}}\rangle_{ZDG} = \frac{1}{\sqrt{m!}} \sum_{\pi\in S_m} \pi_{\tilde{D}}\,|\psi^{\mathbf{x}}\rangle_{ZD}\,|\pi\rangle_G\,, \tag{13}$$

where $G$ is a register keeping track of which permutation was applied to $\tilde{D}$, and observe that $\left\|(P_\eta)_{\tilde{D}}|\psi^{\mathbf{x}}\rangle\right\| = \left\|(P_\eta)_{\tilde{D}}|\psi^{\mathbf{x}}_{\mathrm{sym}}\rangle_{ZDG}\right\|$ since $P_\eta$ commutes with permutations of the $m$ subregisters. We will now lower-bound $\left\|(P_\eta)_{\tilde{D}}|\psi^{\mathbf{x}}_{\mathrm{sym}}\rangle_{ZDG}\right\|$ using the following strategy.

1. We define a sequence of binary random variables indicating whether measuring the $\tilde{D}_i$ resulted in $0^\ell$ or not

2. We show a lower bound on the conditional expectations of this sequence of random variables
3. We use the Azuma-Hoeffding martingale concentration inequality to show a lower bound on the number of $0^\ell$ outcomes that holds with high probability

We begin with a lemma that will facilitate the lower bound on the conditional expectation. Informally, this lemma shows that, for a symmetric state $|\psi\rangle$ supported on subspace $W_n^m$ on register $\tilde{D}$ which is spanned by states with most sub-registers in state $|+^\ell\rangle$, the probability of obtaining $0^\ell$ when measuring the first sub-register is not too small. This is because due to the symmetry assumption, the amplitude of the first sub-register being in state $|+^\ell\rangle$ is at least $\sqrt{m/n}$.

**Lemma 7.** *Let $|\psi\rangle_{\tilde{D}E} \in W_n^m \otimes \mathcal{H}_E$ be a normalized quantum state on registers $\tilde{D} = \tilde{D}_1 \dots \tilde{D}_m$ and an additional register $E$ such that for all operators $A$ acting on $\tilde{D}$ and for all $\pi \in S_m$ it holds that $\|A_{\tilde{D}}\pi_{\tilde{D}} |\psi\rangle_{\tilde{D}E}\| = \|A_{\tilde{D}} |\psi\rangle_{\tilde{D}E}\|$. Then*

$$\left\| \langle 0^\ell|_{\tilde{D}_1} |\psi\rangle_{\tilde{D}E} \right\|^2 \geq \frac{2^{-l}n}{m} - 2^{1-l/2}\sqrt{\frac{n(m-n)}{m^2}}$$

*Proof.* We can decompose

$$|\psi\rangle_{\tilde{D}E} =: \sum_{\substack{I\subset[m]\\|I|\geq n}} \left|+^\ell\right\rangle_I |\psi_I\rangle_{I^C E}$$

with the convention that $\langle +^\ell|_i |\psi_I\rangle_{I^c E} = 0$ for all $i \notin I$. We can use this decomposition to derive expressions for the parts of the state where the first subregister is in state $|+^\ell\rangle$, and orthogonal to it, respectively,

$$\begin{aligned}
|\psi\rangle_{\tilde{D}E} &= \sum_{\substack{I\subset[m]\\|I|\geq n}} |+^\ell\rangle_I^{\otimes|I|} |\psi_I\rangle_{I^C E} = \sum_{\substack{I\subset[m]\\|I|\geq n\wedge 1\in I}} |+^\ell\rangle_I^{\otimes|I|} |\psi_I\rangle_{I^C E} + \sum_{\substack{I\subset[m]\\|I|\geq n\wedge 1\notin I}} |+^\ell\rangle_I^{\otimes|I|} |\psi_I\rangle_{I^C E} \\
&= |+^\ell\rangle_1 \left( \sum_{\substack{I'\subset[m]\setminus 1\\|I'|\geq n-1}} |+^\ell\rangle_{I'}^{\otimes|I'|} |\psi_{I'\cup\{1\}}\rangle_{(I'\cup\{1\})^C E} \right) + \sum_{\substack{I\subset[m]\\|I|\geq n\wedge 1\notin I}} |+^\ell\rangle_I^{\otimes|I|} |\psi_I\rangle_{I^C E} \\
&=: |+^\ell\rangle_1 |\psi^+\rangle_{1^c E} + |\psi^-\rangle_{\tilde{D}E},
\end{aligned} \tag{14}$$

where $\langle +^\ell|_1 |\psi^-\rangle_{\tilde{D}E} = 0$. By the assumed permutation symmetry of $|\psi\rangle$, the norm of $|\psi_I\rangle$ only depends on $|I|$. We thus define $\zeta_i = \| |\psi_{[i]}\rangle \|^2$ and note that by the normalization of $|\psi\rangle$ we have

$$\sum_{i=n}^m \binom{m}{i} \zeta_i = 1.$$

We can now bound

$$\| |\psi_+\rangle \|^2 = \sum_{\substack{I' \subset [m]\setminus 1 \\ |I'| \geq n-1}} \| |\psi_{I' \cup \{1\}}\rangle \|^2 = \sum_{i=n}^{m} \binom{m-1}{i-1} \zeta_i = \sum_{i=n}^{m} \frac{i}{m} \binom{m}{i} \zeta_i$$
$$\geq \sum_{i=n}^{m} \frac{n}{m} \binom{m}{i} \zeta_i = \frac{n}{m} =: p$$

Therefore we get, for some $\alpha \in [0,1]$,

$$\left\| \langle 0^\ell |_1 |\psi\rangle_{\tilde{D}E} \right\|^2 = \| \langle 0|+\rangle |\psi^+\rangle_{1^c E} + \langle 0|_1 |\psi^-\rangle_{\tilde{D}E} \|^2$$
$$\geq \left( 2^{-l/2} \| |\psi^+\rangle \| - \alpha \| |\psi^-\rangle \| \right)^2$$
$$\geq \left( 2^{-l/2} \sqrt{p} - \alpha \sqrt{1-p} \right)^2 = 2^{-l} p + \alpha^2 (1-p) - \alpha 2^{1-l/2} \sqrt{p(1-p)}$$
$$\geq \frac{2^{-l} n}{m} - 2^{1-l/2} \sqrt{\frac{n(m-n)}{m^2}}.$$

□

**Bounding $\left\| (\boldsymbol{P_\eta})_{\tilde{\boldsymbol{D}}} |\psi^{\mathbf{x}}\rangle \right\|$ using martingale tail bound.** We will now follow the same strategy as in Lemma 5 to show that after applying the extractor failure projector, there are, in fact, more than $\eta$ registers in state $|0^\ell\rangle$ with high probability. There is, however, one crucial difference: the random variables obtained by measuring the sub-registers of $\tilde{D}$ in the state $|\psi^{\mathbf{x}}\rangle$ are not independent. This renders the standard Chernoff bound we used in Lemma 5 inapplicable. We therefore resort to a martingale tail bound for our analysis.

In the following, for ease of notation, let

$$|\psi\rangle = \frac{1}{\| |\psi^{\mathbf{x}}_{\text{sym}}\rangle \|} |\psi^{\mathbf{x}}_{\text{sym}}\rangle .$$

Let $X_i$, $0 \leq i \leq k(N-1)$ be the random variables[4] defined analogously to the proof of Lemma 5, i.e.,

$$X_i = \begin{cases} 1 & \text{if measuring register } \tilde{D}_i \text{ of state } |\psi\rangle \text{ yields } 0^\ell \\ 0 & \text{else.} \end{cases}$$

We define $X_{<i} = (X_1, \ldots X_{i-1})$, and

$$Z_i = \mathbb{E}\left[X_i | X_{<i}\right], \tag{15}$$
$$\hat{X}_i = Z_i - X_i, \text{ and} \tag{16}$$
$$\hat{Z}_i = \sum_{j=1}^{i} \hat{X}_j. \tag{17}$$

[4] Following established conventions, we denote both the query input register and certain random variables using the letter $X$. Abuse of notation is avoided as the random variables have subscripts and context should help the reader avoid confusion.

Note that conditioning on $X_{<i}$, $Z_{<i}$, $\hat{X}_i$ and $\hat{Z}_{<i}$ are all equivalent. Let the value of the standard cumulative mean be denoted as

$$\mu' = \sum_{i=1}^{n} \mathbb{E}[X_i \mid X_{<i}] = \sum_{i=1}^{n} Z_i. \tag{18}$$

**Lemma 8.** *We have the following probability bound*

$$\left\| (P_\xi)_{\tilde{D}} |\psi\rangle \right\|^2 \geq 1 - \exp\left( -\frac{\delta'^2 \underline{\mu}'^2}{2m} \right) =: 1 - \varepsilon'$$

*Here,*

$$\underline{\mu}' := 2^{-l} k \left[ N - 1 - \gamma \left( 1 + 4 \cdot 2^\ell + \log\left( \frac{4\gamma}{\log k} \right) \right) - 4\sqrt{2^\ell \gamma N} \right]$$

*and*

$$\xi = (1 - \delta')\underline{\mu}'.$$

*Proof.* As in the proof of Lemma 5, we use the fact that

$$\left\| (P_\xi)_{\tilde{D}} |\psi\rangle \right\|^2 = \Pr\left[ \sum_{i=1}^{m} X_i \geq \xi \right] = 1 - \Pr\left[ \sum_{i=1}^{m} X_i < \xi \right]$$

Note that conditioning on $X_{<i}$, $Z_{<i}$, $\hat{X}_i$ and $\hat{Z}_{<i}$ are all equivalent. We get

$$\begin{aligned}
\mathbb{E}\left[ \hat{X}_i \mid \hat{Z}_{<i} \right] &= \mathbb{E}[Z_i - X_i \mid X_{<i}] \\
&= \mathbb{E}[\mathbb{E}\left[X_i | X_{<i}\right] - X_i | \, X_{<i}] \\
&= \mathbb{E}\left[X_i | X_{<i}\right] - \mathbb{E}\left[X_i | X_{<i}\right] \\
&= 0, \\
\implies \mathbb{E}\left[ \hat{Z}_i \mid \hat{Z}_{<i} \right] &= \hat{Z}_{i-1} + \mathbb{E}\left[ \hat{X}_i \mid \hat{Z}_{<i} \right] = \hat{Z}_{i-1}
\end{aligned}$$

Hence $\hat{Z}_i$ is a martingale. Additionally, $\hat{Z}_0 = 0$ and $|\hat{Z}_i - \hat{Z}_{i-1}| = |\hat{X}_i| \leq 1$. Applying Lemma 4 we thus get

$$\begin{aligned}
\Pr\left[ \sum_{j=1}^{m} X_j \leq \mu' - \varepsilon \right] &= \Pr\left[ \sum_{j=1}^{m} X_j \leq \sum_{j=1}^{m} Z_j - \varepsilon \right] \\
&= \Pr\left[ 0 - \hat{Z}_m \leq -\varepsilon \right] \\
&\leq \exp\left( -\frac{\varepsilon^2}{2m} \right).
\end{aligned}$$

Using the lower bound on the cummulative mean from Lemma 9 we obtain the desired bound,

$$\Pr\left[ \sum_{j=1}^{m} X_j \leq \underline{\mu}' - \varepsilon \right] \leq \exp\left( -\frac{\varepsilon^2}{2m} \right).$$

□

We continue by providing a lower bound on the cumulative mean.

**Lemma 9.** *The cummulative mean defined in Eq. (18) is bounded as*

$$\mu' \geq \underline{\mu}' = 2^{-l}k \left[ N - 1 - \gamma \left( 1 + 4 \cdot 2^{\ell} + \log \left( \frac{4\gamma}{\log k} \right) \right) - 4\sqrt{2^{\ell}\gamma N} \right].$$

The proof proceeds by exploiting Lemma 7 and elementary bounds, and can be found in Section A in the Appendix.

**QROM Extractability Bound** We are now finally ready to prove our main QROM extractability result for deterministic commitments. For brevity, define $\tilde{\mathcal{E}}_f = \text{Comp}\mathcal{E}_f\text{Comp}$ and $\tilde{\mathcal{E}}_s = \text{Comp}\mathcal{E}_s\text{Comp}$. We make the convention that any projectors act on register $D$ unless specified otherwise.

Recall that the argument is a sandwich between two tail bounds pulling in opposite directions. For a fixed commitment vector $a_0$, we want to bound $\|\tilde{E}_f|\phi_{succ}\rangle\|^2$, the probability that the prover succeeds but the extractor fails. (Recall the (subnormalized) state $|\phi_{succ}\rangle = P_s|\phi\rangle$ projected onto prover success. The key quantity is a threshold $\eta$ on the number of database registers in state $|0^{\ell}\rangle$. On one hand, Lemma 5 gives an *upper* bound: unconditionally, the probability of seeing $\geq \eta$ such registers is exponentially small (call it $\varepsilon''$), because the database registers start independent and $\eta$ is chosen well above the mean $\mu = 2^{-\ell}kN$. On the other hand, Lemma 11 - which collects the decomposition, symmetrization, and martingale argument of Section 3 - gives a matching *lower* bound: conditioned on both prover success and extractor failure, seeing $\geq \eta$ registers in state $|0^{\ell}\rangle$ is likely, with probability at least $(1-\varepsilon')\|\tilde{E}_f|\phi_{succ}\rangle\|^2$, up to small error terms from the compressed-basis tail $\varepsilon_{\gamma,k}$. Combining the two via the triangle inequality, together with a bound on the cross term $\langle\phi_{succ}|(\mathbb{1} - \tilde{E}_f)P_\eta\tilde{E}_f|\phi_{succ}\rangle$ from Lemma 11, yields $\varepsilon'' \geq (1-\varepsilon')\|\tilde{E}_f|\phi_{succ}\rangle\|^2$, from which

$$\|\tilde{E}_f|\phi_{succ}\rangle\|^2 \leq \frac{\varepsilon}{1-\varepsilon}$$

follows immediately, where $\varepsilon = \max(\varepsilon', \varepsilon'')$. Note that both $\varepsilon''$ and $\varepsilon'$ are negligible in $k$, making the extraction error negligible.

**Lemma 10 (Extractability – Deterministic Commitments).** *Let $\mathcal{P}$ be a dishonest prover for the Fischlin transform of a special-sound $\Sigma$-protocol for a relation $R$ with unique responses. For some constant $c$, let the challenge space be of size $N = c \cdot 2^{\ell} \cdot \log(k)$. Let further $2^{\frac{1}{c}} \leq k \leq 2^{\frac{2^{\ell}}{256 \cdot c}}$ and $l \geq 14$. The probability of the event where the prover succeeds and the extractor fails is bounded by*

$$\mathbb{P}[\textit{Extractor fails and Prover Succeeds}] = ||\tilde{\mathcal{E}}_f P_s|\phi\rangle||^2 \leq \varepsilon/(1-\varepsilon) = \text{negl}(k),$$

*where* $\varepsilon \leq 3\exp\left[-\frac{k}{128 \cdot c \cdot 2^{\ell} \cdot \log k}\right] + 7\exp\left[-\frac{k}{8 \cdot 2^{\ell}}\right]$.

Note that the parameter restrictions allow picking $l = \Theta(\log\log k)$ which indeed makes the error negligible in $k$.

*Proof.* Using Lemma 5, we bound

$$\begin{aligned}
\varepsilon'' &\geq ||P_\eta\,|\phi\rangle\,||^2 \\
&\geq ||P_s P_\eta|\phi\rangle||^2 = ||P_\eta P_s|\phi\rangle||^2 \\
&= ||P_\eta(\tilde{\mathcal{E}}_s + \tilde{\mathcal{E}}_f)|\phi_{\text{succ}}\rangle||^2 \\
&\geq ||P_\eta\tilde{\mathcal{E}}_s|\phi_{\text{succ}}\rangle||^2 + ||P_\eta\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 - 2\,\text{Re}\,\langle\phi_{\text{succ}}|\,\tilde{\mathcal{E}}_s P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle \\
&\geq ||P_\eta\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 - 2\,\text{Re}\,\langle\phi_{\text{succ}}|\,\tilde{\mathcal{E}}_s P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle \\
&\geq (1-\varepsilon')\|\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 - 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}} - 2\,\text{Re}\,\langle\phi_{\text{succ}}|\,\tilde{\mathcal{E}}_s P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle \\
&\quad (\text{here } \exp\left[-\frac{\delta'^2\mu'}{2m}\right] = \varepsilon' \text{as in lemma 11, applying lemma 11}) \\
&\geq (1-\varepsilon')\|\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 - 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}} - \sqrt{2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}}\ (\text{due to (19) below})
\end{aligned}$$

In the second line, we have used that $P_\eta$ and $P_s$ are both diagonal in the standard basis on $D$. The bound on the inner product expression is obtained as

$$\begin{aligned}
&\text{Re}\langle\phi_{\text{succ}}|\tilde{\mathcal{E}}_s P_\eta\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle \leq \langle\phi_{\text{succ}}|\,(\mathbb{I}-\tilde{\mathcal{E}}_f)\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle + \langle\phi_{\text{succ}}|\,(\mathbb{I}-\tilde{\mathcal{E}}_f)(P_\eta\tilde{\mathcal{E}}_f - \tilde{\mathcal{E}}_f)\,|\phi_{\text{succ}}\rangle \\
&\leq ||(\mathbb{I}-\tilde{\mathcal{E}}_f)|||(P_\eta\tilde{\mathcal{E}}_f - \tilde{\mathcal{E}}_f)|\phi_{\text{succ}}\rangle|| \leq \sqrt{2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}}, \qquad (19)
\end{aligned}$$

where the last inequality is a gentle-measurement-type bound which we formulate as a specialized lemma, Lemma 12, in Section B in the Appendix for convenience.

$$\begin{aligned}
&\implies \mathbb{P}[\text{Extractor fails and Prover Succeeds}] = ||\tilde{\mathcal{E}}_f P_s|\phi\rangle||^2 \\
&\leq \frac{\varepsilon'' + \sqrt{2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} + 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}}{1-\varepsilon'} \leq \frac{\varepsilon}{1-\varepsilon}
\end{aligned}$$

$$\begin{aligned}
\text{where } \varepsilon &= \max\{\varepsilon', \varepsilon'' + \sqrt{2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} + 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}\} \\
&= \varepsilon'' + \sqrt{2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} + 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}} \\
&\leq \varepsilon'' + \sqrt{2\varepsilon'} + \sqrt{8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} + 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}\ (\text{ as } \sqrt{a^2+b^2} \leq |a| + |b|\ ) \\
&\leq \varepsilon'' + 2\sqrt{\varepsilon'} + 7\sqrt{\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} \\
&\leq 3\exp\left[-\frac{k}{128\cdot c\cdot 2^\ell\cdot\log k}\right] + 7\exp\left[-\frac{k}{8\cdot 2^\ell}\right].
\end{aligned}$$

The last inequality is obtained by crudely balancing errors, which is done in Lemma 13 in the Appendix. □

**Lemma 11 (Using the martingale bounds).**

$$||P_\eta\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 \geq (1-\varepsilon')||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 - 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}$$

*where $\varepsilon_{\gamma,(1-\gamma)k}$ is tail bound as in lemma 6, $\underline{\mu}'$ is as in Lemma 8,*

$$\eta - (1-2\gamma)k = (1-\delta')\underline{\mu}' \quad \textit{and} \quad \varepsilon' = \exp\left(-\frac{\delta'^2\underline{\mu}'^2}{2m}\right)$$

*Proof.* Consider $\tilde{\mathcal{E}}_f P_s |\phi\rangle = \tilde{\mathcal{E}}_f |\phi_{\text{succ}}\rangle$ and

$$\begin{aligned}|\tilde{\phi}'\rangle &= \frac{\tilde{\mathcal{E}}_f |\phi_{\text{succ}}\rangle}{\|\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle\|} \\ &= \sqrt{\Gamma}\sum_{\mathbf{x}} \sqrt{p_{\mathbf{x}}}\,|\mathbf{x}\rangle_O\,|\psi^{\mathbf{x}}_{\text{norm}}\rangle_{ED} + \frac{|\delta'_{\gamma,k}\rangle}{\|\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle\|}.\end{aligned}$$

Here we have used Eqs. (8), (9) and (11) and defined

$$\begin{aligned}|\psi^{\mathbf{x}}_{\text{norm}}\rangle &= \frac{|\psi^{\mathbf{x}}\rangle}{\|\,|\psi^{\mathbf{x}}\rangle\,\|}, \text{ the probability distribution} \\ p_{\mathbf{x}} &= \frac{\|\,|\psi^{\mathbf{x}}\rangle\,\|^2}{\sum \mathbf{x}\|\,|\psi^{\mathbf{x}}\rangle\,\|^2} \text{ and } \Gamma = \frac{\sum \mathbf{x}\|\,|\psi^{\mathbf{x}}\rangle\,\|^2}{\|\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle\|^2}.\end{aligned}$$

We bound

$$\begin{aligned}\|P_\eta\,|\tilde{\phi}'\rangle\,\|^2 &\geq \Gamma\mathbb{E}_{\mathbf{x}\leftarrow p}\left[\|P_\eta\,|\psi^{\mathbf{x}}_{\text{norm}}\rangle\,\|^2\right] - 2\sqrt{\varepsilon_{\gamma,k}} \\ &= \Gamma\mathbb{E}_{\mathbf{x}\leftarrow p}\left[\left\|P_\eta \frac{1}{\|\,|\psi^{\mathbf{x}}_{\text{sym}}\rangle\,\|}\,|\psi^{\mathbf{x}}_{\text{sym}}\rangle\right\|^2\right] - 2\sqrt{\varepsilon_{\gamma,k}}\end{aligned}$$

where we have used Eq. (10). Here, $\mathbf{x} \leftarrow p$ denotes sampling $\mathbf{x}$ from the probability distribution $p$. Again using the simplifying notation $|\psi\rangle = \frac{1}{\||\psi^{\mathbf{x}}_{\text{sym}}\rangle\|}|\psi^{\mathbf{x}}_{\text{sym}}\rangle$ we bound

$$\|P_\eta\,|\psi\rangle_{ED}\,\|^2 \geq \|\left(P_{(1-2\gamma)k}\right)_{D_{\mathsf{hi}(\mathbf{x})}}\left(P_{\eta-(1-2\gamma)k}\right)_{D_{\mathsf{hi}(\mathbf{x})^c}}|\psi\rangle_{ED}\,\|^2$$

as choosing particular registers to be in $|0^\ell\rangle$ makes the norm smaller. Defining

$$|\psi'\rangle = \left(P_{\eta-(1-2\gamma)k}\right)_{D_{\mathsf{hi}(\mathbf{x})^c}}|\psi\rangle_{ED}$$

we have that

$$\|\,|\psi'\rangle\,\|^2 \geq 1 - \Pr\left[\sum_{i=1}^m X_i \leq \eta - (1-2\gamma)k\right] \geq 1 - \varepsilon'$$

by applying Lemma 8 with $\xi = \eta - (1-2\gamma)k$, and

$$|\psi'\rangle = \sum_{\substack{S\subset\mathsf{hi}(\mathbf{x}) \\ s=|S|\geq(1-\gamma)k}} |\tilde{\psi}^{\mathbf{x}}_S\rangle_{ZD_{S^C}}\left(\text{Comp}|0^\ell\rangle\right)^{\otimes s}_{D_S}$$

for some states $\left|\tilde{\psi}_S^{\mathbf{x}}\right\rangle_{ZD_{S^C}}$ due to Eq. (8). Applying Lemma 6 with $(1-\gamma)k$ in place of $k$ we get

$$|\psi'\rangle = \sum_{\substack{S\subset \mathsf{hi}(\mathbf{x}) \\ s=|S|\geq(1-2\gamma)k}} \left|\tilde{\tilde{\psi}}_S^{\mathbf{x}}\right\rangle_{ZD_{S^C}} |0^\ell\rangle_{D_S}^{\otimes s} + |\delta'_{\gamma,(1-\gamma)k}\rangle \tag{20}$$

for some states $\tilde{\tilde{\psi}}^{\mathbf{x}}$, where

$$\| \, |\delta_{\gamma,(1-\gamma)k}\rangle \, \|^2 \leq \varepsilon_{\gamma,(1-\gamma)k}.$$

It follows that

$$\begin{aligned} \| \left(P_{(1-2\gamma)k}\right)_{D_{\mathsf{hi}(\mathbf{x})}} \left(P_{\eta-(1-2\gamma)k}\right)_{D_{\mathsf{hi}(\mathbf{x})^c}} |\psi\rangle_{ED} \, \|^2 &= \| \left(P_{(1-2\gamma)k}\right)_{D_{\mathsf{hi}(\mathbf{x})}} |\psi'\rangle \, \|^2 \\ &\geq 1-\varepsilon'-2\varepsilon_{\gamma,(1-\gamma)k}. \end{aligned}$$

In summary, we have

$$\|P_\eta \, |\psi\rangle_{ED} \, \|^2 \geq 1-\varepsilon'-2\varepsilon_{\gamma,(1-\gamma)k}$$

And thus

$$\begin{aligned} ||P_\eta \tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 &\geq \|\tilde{\mathcal{E}}_f \, |\phi_{\text{succ}}\rangle \, \|^2(1-\varepsilon'-2\varepsilon_{\gamma,(1-\gamma)k}) - 2\sqrt{\varepsilon_{\gamma,k}} \\ &\geq \|\tilde{\mathcal{E}}_f \, |\phi_{\text{succ}}\rangle \, \|^2(1-\varepsilon') - 4\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}. \end{aligned}$$

□

### 3.2 Extractability – lifting to general prover.

We now go on to apply our bound for a prover outputting a proof with a fixed value for $\mathbf{a}$ to obtain a bound for a general prover. To this end, we prove a bound on the probability that a general prover succeeds but the extractor fails, in terms of the probability that the same happens for a fixed vector of commitments $\mathbf{a}$.

**Theorem 2.** *Let $\mathcal{P}$ be a dishonest prover for the Fischlin transform of a special-sound* $\Sigma$-protocol *for a relation $R$ with unique responses, in the QROM, making no more than $q$ quantum queries to the random oracle. We have*

$$\Pr_{(x,\pi,w)\leftarrow\langle\mathcal{P},\mathcal{E}\rangle}[\mathcal{V}(x,\pi)=1] = \Pr_{(x,\pi)\leftarrow\mathcal{P}^H}[\mathcal{V}(x,\pi)=1] \tag{21}$$

*and*

$$\Pr_{(x,\pi,w)\leftarrow\langle\mathcal{P},\mathcal{E}\rangle}[\mathcal{V}(x,\pi)=1 \wedge (x,w)\notin R] \leq 4(q+k)^2\cdot\delta, \tag{22}$$

*where $\delta$ is such that for any prover $\mathcal{P}'$ and any $\mathbf{a}_0\in\mathcal{M}^k$,*

$$\Pr_{(x,\pi=(\mathbf{a},\mathbf{c},\mathbf{z}),w)\leftarrow\langle\mathcal{P}',\mathcal{E}\rangle}[\mathbf{a}=\mathbf{a}_0 \wedge \mathcal{V}(x,\pi)=1 \wedge (x,w)\notin R] \leq \delta. \tag{23}$$

Combining with Lemma 10, we get our main result.

**Corollary 1 (straight-line extractability of the Fischlin transform in the QROM).** *Let $\Sigma$ be a $\Sigma$-protocol for a relation $R$ with special soundness and unique responses. If the challenge space is of size $N = c \cdot 2^\ell \cdot \log(k)$ for some constant $c$ and parameters $k, l$ of the Fischlin transform, $2^{\frac{1}{c}} \leq k \leq 2^{\frac{2^\ell}{256 \cdot c}}$, $l \geq 14$, and $l = O(\log\log k)$ then $\mathsf{Fis}[\Sigma]$ is a proof of knowledge with straight-line extractability. The extractor's simulation of the quantum-accesible random oracle can be made perfect ($\varepsilon_{\text{sim}} = 0$) and for a $q$-query adversarial prover,*

$$\varepsilon_{\text{ex}} \leq q^2 \cdot \text{negl}(k)$$

A quick note on instantiation. Given a $\Sigma$-protocol $\hat{\Sigma}$ with constant challenge space of size $\hat{N}$ we can pick $k$ (not too small) first. Now we can choose $c$ and $l = \lceil \log\log k + \log c + 8 \rceil$, then set $r = \lceil \log_{\hat{N}}(c \cdot 2^\ell \cdot \log(k)) \rceil$ to allow for the desired challenge space size, and define the $\Sigma$-protocol $\Sigma$ as the $r$-fold parallel repetition of $\hat{\Sigma}$ with the challenge space artificially restricted to size $c \cdot 2^\ell \cdot \log(k)$.[5] The combination of $\Sigma$, $k, l, c$ now fulfils the prerequisites of Corollary 1. As $k$ is a lower bound for the security parameter, this calculation shows that our result supports a parameter region that allows asymptotic security (and thus arbitarily high concrete levels of security).

*Proof (of Theorem 2).* The statement in Equation (21) is immediate by the perfect correctness of the compressed oracle.

We begin by constructing a family of provers that either output a valid proof with a fixed commitment vector $\hat{\mathbf{a}}$, or a dummy symbol. Construct prover $\tilde{\mathcal{P}}_{\hat{\mathbf{a}}}$ as follows.

1. Run $(x, \pi = (\mathbf{a}, \mathbf{c}, \mathbf{z})) \leftarrow \mathcal{P}^H$.
2. Run $b \leftarrow \mathcal{V}^H(x, \pi)$.
3. If $b = 1$ and $\mathbf{a} = \hat{\mathbf{a}}$, output $(x, \pi)$, else, output $\perp$.

Let $\mathcal{P}_{\hat{\mathbf{a}}}$ be a version of $\tilde{\mathcal{P}}_{\hat{\mathbf{a}}}$ with all measurements delayed to the end, i.e., where verification is run coherently. Observe that $\mathcal{P}_{\hat{\mathbf{a}}}$ makes at most $q' = q + k$ queries. Without loss of generality, $\mathcal{P}_{\hat{\mathbf{a}}}$ starts with some initial state $|\phi_0\rangle$, applies $\left(UO^H\right)^{q'}$ and measures the output register $O$. We now consider $\mathcal{P}_{\hat{\mathbf{a}}}$ as run by the extractor, i.e., it gets query access to the compressed oracle instead of a standard quantum-accessible random oracle. Denote by $\Xi = |\mathcal{M}|^k \cdot k \cdot N \cdot |\mathcal{Z}|$ the cardinality of the domain of the random oracle. We decompose the final state after the unitary part of $\mathcal{P}_{\hat{\mathbf{a}}}$ as

$$\begin{aligned}
&\left(UO^H\right)^{q'} |\phi_0\rangle_{XYE} \, |+^\ell\rangle_D^{\otimes \Xi} \\
&= \left(UO^H(\mathbb{1} - |\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|_{X_\mathcal{M}})\right)^{q'} |\phi_0\rangle_{XYE} \, |+^\ell\rangle_D^{\otimes \Xi} \\
&\quad + \sum_{i=1}^{q'} \left(UO^H\right)^{q'-i+1} |\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|_{X_\mathcal{M}} \left(UO^H(\mathbb{1} - |\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_\mathcal{M}}\right)^{i-1} |\phi_0\rangle_{XYE} \, |+^\ell\rangle_D^{\otimes \Xi} .
\end{aligned}$$

[5] We ignore rounding here, which changes $c$ slightly.

Here, $E$ is the adversary's working register. Note that by construction, $\mathcal{P}_{\hat{\mathbf{a}}}$ always queries an input starting with the $\mathbf{a}$ in its output proof (if it outputs one). Therefore we have

$$P_s \left(UO^H(\mathbb{1} - |\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_{\mathcal{M}}}\right)^{q'} |\phi_0\rangle_{XYE} |+^{\ell}\rangle_D^{\otimes \Xi} = 0.$$

Now note that by construction,

$$|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|_{X_{\mathcal{M}}} \left(UO^H(\mathbb{1} - |\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_{\mathcal{M}}}\right)^{i-1} |\phi_0\rangle_{XYE} |+^{\ell}\rangle_D^{\otimes \Xi} = |\tilde{\psi}_i\rangle_{XYED^c_{\hat{\mathbf{a}}}} |+^{\ell}\rangle_{D_{\hat{\mathbf{a}}}}^{\otimes([k]\cdot N\cdot|\mathcal{Z}|)},$$

for some sub-normalized state $|\tilde{\psi}_i\rangle$, as before each query, the input register is projected away from $|\hat{\mathbf{a}}\rangle$.[6] Let $|\psi_i\rangle = \| |\tilde{\psi}_i\rangle \|^{-1} |\tilde{\psi}_i\rangle$. We now define, for each pair $(\hat{\mathbf{a}}, i) \in \mathcal{M}^k \times [q]$, a new prover $\mathcal{P}^H_{\hat{\mathbf{a}},i}$ as follows.

1. Prepare $|\psi_i\rangle$ as follows:
   (a) Initialize a compressed oracle for a random function $H'$ identically distributed as $H$.
   (b) Run $\mathcal{P}^{H'}_{\hat{\mathbf{a}}}$ until right before the $i$-th query, but for each query, apply the binary measurement whether the query input starts with $\hat{\mathbf{a}}$. If it does, abort and restart from (a).
   (c) Finally, apply the same measurement to the query input register. If it returns "not $\hat{\mathbf{a}}$", abort and restart from (a).
2. Run $(x, \pi) \leftarrow \mathcal{P}^{H''}_{\hat{\mathbf{a}}}$ from its $i$-th query, starting with the prepared joint state $|\psi_i\rangle |+^{\ell}\rangle^{\otimes([k]\cdot N\cdot|\mathcal{Z}|)}$ of prover and $H'$ oracle database, where

$$H''(\mathbf{a}, i, c, z) = \begin{cases} H(\mathbf{a}, i, c, z) & \text{if } \mathbf{a} = \hat{\mathbf{a}} \\ H'(\mathbf{a}, i, c, z) & \text{else.} \end{cases}$$

   Here, it is understood that $H'$ is still instantiated with the compressed oracle simulated by the prover $\mathcal{P}^H_{\hat{\mathbf{a}},i}$, while $H$ is the oracle the prover $\mathcal{P}^H_{\hat{\mathbf{a}},i}$ has access to.
3. Output $(x, \pi)$.

This prover might not be time-efficient. This is not a problem, however, as all extractability statements hold for provers only (at most) fulfilling a query bound.

[6] Sub-normalization follows from $\sum_i \| |\tilde{\psi}_i\rangle \|^2 \leq 1$ which holds by orthogonality of the decomposition.

We can now relate the probability that $\mathcal{P}$ succeeds and the extractor fails to the provers $\mathcal{P}^H_{\hat{\mathbf{a}},i}$ as follows,

$$
\begin{aligned}
&\Pr_{(x,\pi,w)\leftarrow\langle\mathcal{P},\mathcal{E}\rangle}[\mathcal{V}(x,\pi)=1\wedge(x,w)\notin R]\\
=&\sum_{\hat{\mathbf{a}}}\left\|\mathcal{E}_f P_s\sum_{i=1}^{q'}\left(UO^H\right)^{q'-i+1}|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|_{X_{\mathcal{M}}}\left(UO^H(\mathbb{1}-|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_{\mathcal{M}}}\right)^{i-1}|\phi_0\rangle_{XYE}\,|+^{\ell}\rangle_D^{\otimes\Xi}\right\|^2\\
\leq&\sum_{\hat{\mathbf{a}}}\left(\sum_{i=1}^{q'}\Big\|\mathcal{E}_f P_s\left(UO^H\right)^{q'-i+1}\right.\\
&\qquad\left.|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|_{X_{\mathcal{M}}}\left(UO^H(\mathbb{1}-|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_{\mathcal{M}}}\right)^{i-1}|\phi_0\rangle_{XYE}\,|+^{\ell}\rangle_D^{\otimes\Xi}\Big\|\right)^2\\
\leq&\delta\sum_{\hat{\mathbf{a}}}\left(\sum_{i=1}^{q'}\||\tilde{\psi}_i\rangle\|\right)^2,
\end{aligned}
$$

where the first inequality is the triangle inequality and the second inequality results from assumption (23) applied to the provers $\mathcal{P}^H_{\hat{\mathbf{a}},i}$. It remains to bound the sum of norms on the right-hand side. Define

$$
\varepsilon_{\hat{\mathbf{a}},i}=\left\||\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|\,(UO)^{i-1}\,|\phi_0\rangle\,|+^{\ell}\rangle_D^{\otimes\Xi}\right\|^2.
$$

We can now bound, using Cauchy-Schwarz in the first step,

$$
\begin{aligned}
&\frac{1}{q'}\left(\sum_{i=1}^{q'}\||\tilde{\psi}_i\rangle\|\right)^2\leq\sum_{i=1}^{q'}\||\tilde{\psi}_i\rangle\|^2\\
=&\sum_{i=1}^{q'}\langle\phi_0|_{XYE}\,\langle+^{\ell}|_D^{\otimes\Xi}\left((\mathbb{1}-|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_{\mathcal{M}}}\left(O^H\right)^{\dagger}U^{\dagger}\right)^{i-1}\\
&\qquad|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|_{X_{\mathcal{M}}}\left(UO^H\left(\mathbb{1}-|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|\right)_{X_{\mathcal{M}}}\right)^{i-1}|\phi_0\rangle_{XYE}\,|+^{\ell}\rangle_D^{\otimes\Xi}\\
=&1-\langle\phi_0|_{XYE}\,\langle+^{\ell}|_D^{\otimes\Xi}\left((\mathbb{1}-|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|)_{X_{\mathcal{M}}}\left(O^H\right)^{\dagger}U^{\dagger}\right)^{q'}\\
&\qquad\left(UO^H\left(\mathbb{1}-|\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|\right)_{X_{\mathcal{M}}}\right)^{q'}|\phi_0\rangle_{XYE}\,|+^{\ell}\rangle_D^{\otimes\Xi}\\
\leq&4\sum_{i=1}^{q'}\varepsilon_{\hat{\mathbf{a}},i}.
\end{aligned}
$$

Here, we have used the quantum union bound, Lemma 2, with projectors $P_i = \left(\left(O^H\right)^\dagger U^\dagger\right)^{i-1} (\mathbb{1} - |\hat{\mathbf{a}}\rangle\langle\hat{\mathbf{a}}|) (UO)^{i-1}$ in the last line. We thus finally get

$$\begin{aligned}\Pr_{(x,\pi,w)\leftarrow\langle\mathcal{P},\mathcal{E}\rangle}[\mathcal{V}(x,\pi)=1 \wedge (x,w)\notin R] &\leq \delta \sum_{\hat{\mathbf{a}}} \left(\sum_{i=1}^{q'} \left\||\tilde{\psi}_i\rangle\right\|\right)^2 \\ &\leq 4q' \cdot \delta \sum_{\hat{\mathbf{a}}} \sum_{i=1}^{q'} \varepsilon_{\hat{\mathbf{a}},i} \\ &= 4q'^2 \cdot \delta,\end{aligned}$$

where we have used the normalization of $|\phi_0\rangle$ in the last line.

### 3.3 Main Result

We now state the main result on QROM security of Fischlin's Transform.

**Theorem 3 (Post-Quantum Security of the Fischlin Transform).** *Let $\Sigma$ be a $\Sigma$-protocol for a relation $R$ with special soundness, unique responses, and commitment entropy. Let the parameters $k$, $\ell$, $c$ of the Fischlin transform* Fis$[\Sigma]$ *satisfy $N = c \cdot 2^\ell \cdot \log k$, $2^{1/c} \leq k \leq 2^{2^\ell/(256c)}$, $\ell \geq 14$, and $\ell = O(\log\log k)$. Then* Fis$[\Sigma]$ *is a non-interactive zero-knowledge proof of knowledge with straight-line extractability in the* QROM*, concretely satisfying:*

1. (Perfect simulation.) *The extractor simulates the quantum-accessible random oracle perfectly, i.e. $\varepsilon_{\text{sim}} = 0$.*
2. (Straight-line extractability.) *For any $q$-query malicious prover $P^*$, the extraction error satisfies*
$$\varepsilon_{\text{ex}} \;\leq\; q^2 \cdot \text{negl}(k).$$
3. (Zero-knowledge.) Fis$[\Sigma]$ *is computationally indistinguishable from simulated proofs for any quantum polynomial-time distinguisher.*

*Proof.* Perfect simulation and straight-line extractability follow from Lemma 10 and Corollary 1; zero-knowledge from Theorem 4 in Section D in the Appendix.

## 4 Acknowledgements

We have used AI tools for spell checking, grammar and other English-language-related improvements. CM thanks Gorjan Alagic for useful discussions. CM and JS acknowledge support by the Independent Research Fund Denmark via a DFF Sapere Aude grant (IM-3PQC, grant ID 10.46540/2064-00034B). The authors thank the anonymous reviewers of the CRYPTO 2026 program committee for useful feedback.

## References


AHJ+23. Carlos Aguilar Melchor, Andreas Hülsing, David Joseph, Christian Majenz, Eyal Ronen, and Dongze Yue. SDitH in the QROM. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VII*, volume 14444 of *LNCS*, pages 317–350, Guangzhou, China, December 4–8, 2023. Springer, Singapore, Singapore.

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

BBB+25. Carsten Baum, Ward Beullens, Lennart Braun, Cyprien Delpech de Saint Guilhem, Michael Klooß, Christian Majenz, Shibam Mukherjee, Emmanuela Orsini, Sebastian Ramacher, Christian Rechberger, Lawrence Roy, and Peter Scholl. Shorter, tighter, FAESTer: Optimizations and improved (QROM) analysis for VOLE-in-the-head signatures. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VI*, volume 16005 of *LNCS*, pages 124–156, Santa Barbara, CA, USA, August 17–21, 2025. Springer, Cham, Switzerland.

BDF+11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69, Seoul, South Korea, December 4–8, 2011. Springer Berlin Heidelberg, Germany.

CFHL21. Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 598–629, Zagreb, Croatia, October 17–21, 2021. Springer, Cham, Switzerland.

CMS19a. Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 1–29, Nuremberg, Germany, December 1–5, 2019. Springer, Cham, Switzerland.

CMS19b. Alessandro Chiesa, Pratyush Mishra, and Nicholas Spooner. On succinct arguments and the quantum random oracle model. In *Theory of Cryptography Conference (TCC 2019)*, volume 11891 of *Lecture Notes in Computer Science*, pages 37–67. Springer, 2019.

DFM20. Jelle Don, Serge Fehr, and Christian Majenz. Security of the fiat–shamir transformation in the quantum random-oracle model. In *Advances in Cryptology – CRYPTO 2020*, volume 12171 of *Lecture Notes in Computer Science*, pages 356–383. Springer, 2020.

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, Paper 2019/190, 2019.

DFMS22a. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Efficient NIZKs and signatures from commit-and-open protocols in the QROM. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 729–757, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland.

DFMS22b. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part III*, volume 13277 of *LNCS*, pages 677–706, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland.

Fis05. Marc Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 152–168, Santa Barbara, CA, USA, August 14–18, 2005. Springer Berlin Heidelberg, Germany.

GHHM21. Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the QROM. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part I*, volume 13090 of *LNCS*, pages 637–667, Singapore, December 6–10, 2021. Springer, Cham, Switzerland.

HJMN24. Andreas Hülsing, David Joseph, Christian Majenz, and Anand Kumar Narayanan. On round elimination for special-sound multi-round identification and the generality of the hypercube for MPCitH. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 373–408, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland.

Ks22. Yashvanth Kondi and abhi shelat. Improved straight-line extraction in the random oracle model with applications to signature aggregation. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part II*, volume 13792 of *LNCS*, pages 279–309, Taipei, Taiwan, December 5–9, 2022. Springer, Cham, Switzerland.

MU05. Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 1st edition, 2005.

OV22. Ryan O'Donnell and Ramgopal Venkateswaran. The quantum union bound made easy. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 314–320. SIAM, 2022.

Pas03. Rafael Pass. On deniability in the common reference string and random oracle model. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 316–337, Santa Barbara, CA, USA, August 17–21, 2003. Springer Berlin Heidelberg, Germany.

RT25. Lior Rotem and Stefano Tessaro. Straight-line knowledge extraction for multi-round protocols. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 95–127, Santa Barbara, CA, USA, August 17–21, 2025. Springer, Cham, Switzerland.

Unr15. Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784, Sofia, Bulgaria, April 26–30, 2015. Springer Berlin Heidelberg, Germany.

Zha19. Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 239–268, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.

# Appendix

## A Deferred proof of Lemma 9

*Proof (of Lemma 9).* Let $|\beta\rangle_{\tilde{D}_{[i+1:m']}}$ be the post-measurement state after the first $i-1$ subregisters of $\tilde{D}$ have been measured and fixed outcomes $y_1, \ldots y_{i-1}$ have been obtained. Then $|\beta\rangle \in W_{n-i+1}^{m-i+1}$. When measuring the first remaining register of $|\beta\rangle$, the probability of obtaining outcome $y_i = 0^\ell$ is thus lower-bounded as

$$\Pr[X_i = 1 | y_1, \ldots, y_{i-1}] \geq \frac{2^{-l}(n-i+1)}{(m-i+1)} - 2^{1-l/2}\sqrt{\frac{(n-i+1)(m-n)}{(m-i+1)^2}}$$

by Lemma 7 and therefore in particular

$$\begin{aligned}\mathbb{E}[X_i | X_{<i}] = \Pr[X_i = 1 | X_{<i}] &\geq \frac{2^{-l}(n-i+1)}{(m-i+1)} - 2^{1-l/2}\sqrt{\frac{(n-i+1)(m-n)}{(m-i+1)^2}} \\ &\geq \frac{2^{-l}(n-i)}{(m-i)} - 2^{1-l/2}\sqrt{\frac{(n-i)(m-n)}{(m-i)^2}}.\end{aligned}$$

We can now lower-bound now for any $n' \leq n$,

$$\begin{aligned}\mu' &= \sum_{i=1}^{n} \mathbb{E}[X_i | X_{<i}] \\ &\geq \sum_{i=1}^{n'} \mathbb{E}[X_i | X_{<i}] \\ &\geq \sum_{i=1}^{n'} \left(\frac{n-i}{m-i}\right) 2^{-l} - 2^{1-l/2}\sqrt{\frac{(n-i)(m-n)}{(m-i)^2}} \\ &= \sum_{i=1}^{n'} \left(1 - \frac{m-n}{m-i}\right) 2^{-l} - 2^{1-l/2}\sqrt{\frac{(n-i)(m-n)}{(m-i)^2}}. \\ &= n' - 2^{-l}(m-n)\sum_{i=1}^{n'} \frac{1}{m-i} - \sum_{i=1}^{n'} 2^{1-l/2}\sqrt{\frac{(n-i)(m-n)}{(m-i)^2}}.\end{aligned}$$

We have the following inequalities for $n' < m$ on the above terms in the sum:

1. On the second term: $\displaystyle\sum_{i=1}^{n'} \frac{1}{m-i} \leq \int_1^{n'} \frac{dx}{m-x} = \log\left(\frac{m-1}{m-n'}\right),$
2. On the third term: $\displaystyle\sum_{i=1}^{n'} \sqrt{\frac{(n-i)}{(m-i)^2}} \leq \sum_{i=1}^{n'} \sqrt{\frac{(m-i)}{(m-i)^2}}$
$\displaystyle\leq \int_1^{n'} \frac{dx}{\sqrt{m-x}} = 2\left[\sqrt{(m-1)} - \sqrt{(m-n')}\right]$

Hence we get:

$$\begin{aligned}
\mu' &\geq 2^{-l}\left[n' - (m-n)\log\left(\frac{m-1}{m-n'}\right)\right] - 2^{2-l/2}\sqrt{m-n}\left[\sqrt{(m-1)} - \sqrt{(m-n')}\right] \\
&\geq 2^{-l}\left[n' + (m-n)\log\left(\frac{m-n'}{m-1}\right)\right] - 2^{2-l/2}\sqrt{m(m-n)} \\
&= 2^{-l}\left[k(N-1-\gamma) - 2^{2+l}k\gamma + (k\gamma)\log\left(\frac{\gamma(1+2^{2+l})}{2^{\ell}\log k - 1}\right)\right] - 2^{2-l/2}\sqrt{m(m-n)} \\
&\geq 2^{-l}k\left[(N - (1+\gamma+2^{2+l}\gamma) + (\gamma)\log\left(\frac{4\gamma}{\log k}\right)\right] - 2^{2-l/2}\sqrt{k(N-1)\gamma k} \\
&\geq 2^{-l}k\left[(N - 1 - \gamma\left(1 + 2^{2+l} + \log\left(\frac{4\gamma}{\log k}\right)\right) - 2^{2+l/2}\sqrt{N\gamma}\right]
\end{aligned}$$

Here, we have chose $n'$ so that we do not unnecessarily lower-bound non-negative terms by negative ones,

$$\begin{aligned}
&\left[\frac{n-i}{m-i}\right]2^{-l} - \frac{2^{1-l}\sqrt{(n-i)(m-n)}}{m-i} \geq 0 \\
\implies &(n-i)2^{-l} \geq 2^{1-l/2}\sqrt{(n-i)(m-n)} \\
\implies &(n-i)^2 2^{-2l} \geq 2^{2-l}(n-i)(m-n) \\
\implies &(n-i) \geq 2^{2+l}(m-n) \\
\implies &i \leq n - 2^{2+l}(m-n) =: n'. \qquad (24)
\end{aligned}$$

□

We remark that the parameter restrictions of Lemma 10 imply $n' \geq 0$, so the truncation in the proof is non-empty.

## B A gentle measurement lemma

**Lemma 12 (Gentle measurement with additive error).** $||(P_\eta\tilde{\mathcal{E}}_f - \tilde{\mathcal{E}}_f)|\phi_{\text{succ}}\rangle||^2 \leq 2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}$.

*Proof. Directly expanding:*

$$\begin{aligned}
||(P_\eta\tilde{\mathcal{E}}_f - \tilde{\mathcal{E}}_f)|\phi_{\text{succ}}\rangle||^2 &= ||P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 + ||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 - 2\langle\phi_{\text{succ}}|\tilde{\mathcal{E}}_f P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle \\
&= ||P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 + ||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 - 2||P_\eta\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 \\
&\leq ||P_\eta||^2||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 + ||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 - 2(1-\varepsilon')||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 \\
&\quad + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}\textit{(due to lemma 11 )} \\
&\leq (1 + 1 - 2(1-\varepsilon'))||\tilde{\mathcal{E}}_f\,|\phi_{\text{succ}}\rangle\,||^2 + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}} \\
&= 2\varepsilon'||\tilde{\mathcal{E}}_f|\phi_{\text{succ}}\rangle||^2 + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}} \\
&\leq 2\varepsilon' + 8\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}
\end{aligned}$$

□

## C Bounding Lemmas

The following Lemmas gives numerical analysis for the bounds.

**Lemma 13.** *For some constant c, let the challenge space be of size* $N = c \cdot 2^{\ell} \cdot \log(k)$. *Let further* $\gamma = 4 \cdot 2^{-l}$, $2^{\frac{1}{c}} \leq k \leq 2^{\frac{2^{\ell}}{256 \cdot c}}$ *and* $l \geq 14$. *Then*

$$\varepsilon'' + 2\sqrt{\varepsilon'} + 7\sqrt{\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} \leq 3\exp\left[-\frac{k}{128 \cdot c \cdot 2^{\ell} \cdot \log k}\right] + 7\exp\left[-\frac{k}{8 \cdot 2^{\ell}}\right] = \mathrm{negl}(k).$$

*Proof.* Substituting the formulas for the parameters we get

$$\begin{aligned}\bar{\varepsilon} &\coloneqq \varepsilon'' + 2\sqrt{\varepsilon'} + 7\sqrt{\sqrt{\varepsilon_{\gamma,(1-\gamma)k}}} \\ &\leq \exp\left[-\frac{\delta^2\mu}{3}\right] + 2\exp\left[-\frac{\delta'^2\underline{\mu'}}{2^{\ell+1}}\right] + 7\exp\left[-\left(\gamma - 2 \cdot 2^{-l}\right)(1-\gamma)k/8\right] \qquad (25)\end{aligned}$$

using Lemma 14 below.

We simplify

$$\bar{\varepsilon} \leq \exp\left[-\frac{\delta^2\mu}{3}\right] + 2\exp\left[-\frac{\delta'^2\underline{\mu'}}{2^{\ell+1}}\right] + 7\exp\left[-2^{-l}k/8\right]$$

where we have used $1 - \gamma \geq 1/2$.

We bound

$$\begin{aligned}\underline{\mu'} &= 2^{-l}k\left[N - 1 - \gamma\left(1 + 2^{2+l} + \log\left(\frac{4\gamma}{\log k}\right)\right) - 2^{2+l/2}\sqrt{N\gamma}\right] \\ &= \mu - 2^{-l}k\left[1 + \gamma\left(1 + 2^{2+l} + \log\left(\frac{4\gamma}{\log k}\right)\right) + 2^{2+l/2}\sqrt{N\gamma}\right] \leq \mu \qquad (26)\end{aligned}$$

as $\log\left(\frac{4\gamma}{\log k}\right) \geq 4 - 2l \geq -4 \cdot 2^{\ell}$ by the assumptions $k \leq 2^{2^{\ell}}$ and $l \geq 14$.

For simplicity, we set $\delta\mu = \delta'\underline{\mu'}$. We thus get

$$\begin{aligned}\delta'^2\underline{\mu'} &= \frac{\delta^2\mu^2}{\underline{\mu'}} \\ &\geq \delta^2\mu. \qquad (27)\end{aligned}$$

using Eq. (26).

We have $\eta = \mu(1+\delta)$ (from applying Lemma 5) and $\eta - (1-2\gamma)k = \underline{\mu'}(1-\delta')$ (see Lemma 11). We thus get

$$\begin{aligned}\mu(1+\delta) - (1-2\gamma)k &= \underline{\mu'}(1-\delta') \\ \Leftrightarrow \delta &= \frac{(1 - 8 \cdot 2^{-l})k - (\mu - \underline{\mu'})}{2\mu}. \qquad (28)\end{aligned}$$

We bound

$$\begin{aligned}\mu - \underline{\mu}' &= 2^{-l}k\left[1+\gamma\left(1+2^{2+l}+\log\left(\frac{4\gamma}{\log k}\right)\right)+2^{2+l/2}\sqrt{N\gamma}\right]\\ &\leq 2^{-l}k\left[1+4\cdot 2^{-l}\left(1+4\cdot 2^{\ell}\right)+8\sqrt{N}\right]\\ &\leq 2^{-l/2}k\left[21\cdot 2^{-l/2}+8\sqrt{c\cdot\log k}\right]\end{aligned}$$

We can use the assumption on $k$,

$$8\sqrt{c\cdot\log k}\leq 8\sqrt{\frac{2^{\ell}}{256}}\leq\frac{1}{2}2^{\ell/2},$$

so we get

$$\begin{aligned}\mu-\underline{\mu}' &\leq 2^{-l/2}k\left[21\cdot 2^{-l/2}+8\sqrt{c\cdot\log k}\right]\\ &\leq\frac{3k}{4},\end{aligned}$$

using $l\geq 14$. Plugging into Eq. (28) we get

$$\begin{aligned}\delta &\geq\frac{\left(\frac{1}{4}-8\cdot 2^{-l}\right)k}{2\mu}\\ &\geq\frac{1}{8c\log k}.\end{aligned}$$

Here we have used $8\cdot 2^{-l}\leq 2^{-11}\leq 1/8$ due to the assumption $l\geq 14$. We can now use this to get

$$\begin{aligned}\delta^2\mu &\geq\frac{2^{-l}kN}{64c^2\log^2 k}\\ &=\frac{2^{-l}k\cdot c\cdot 2^{\ell}\cdot\log k}{64c^2\log^2 k}\\ &=\frac{k}{64c\log k}.\end{aligned}$$

Finally, we obtain

$$\begin{aligned}\bar{\varepsilon} &\leq\exp\left[-\frac{\delta^2\mu}{3}\right]+2\exp\left[-\frac{\delta'^2\underline{\mu}'}{2^{\ell+1}}\right]+7\exp\left[-2^{-l}k/8\right]\\ &\leq\exp\left[-\frac{\delta^2\mu}{3}\right]+2\exp\left[-\frac{\delta^2\mu}{2^{\ell+1}}\right]+7\exp\left[-2^{-l}k/8\right]\\ &\leq 3\exp\left[-\frac{\delta^2\mu}{2^{\ell+1}}\right]+7\exp\left[-2^{-l}k/8\right]\\ &\leq 3\exp\left[-\frac{k}{128\cdot c\cdot 2^{\ell}\cdot\log k}\right]+7\exp\left[-2^{-l}k/8\right].\end{aligned}$$

□

We prove the deferred lemma.

**Lemma 14.** *For some constant c, let the challenge space be of size* $N = c \cdot 2^\ell \cdot \log(k)$. *Let further* $\gamma = 4 \cdot 2^{-l}$, $2^{\frac{1}{c}} \leq k \leq 2^{\frac{2^\ell}{256 \cdot c}}$ *and* $l \geq 14$. *Then*

$$\exp\left(-\frac{\delta'^2 \underline{\mu'}^2}{2m}\right) \leq \exp\left(-\frac{\delta'^2 \underline{\mu'}}{2^\ell},\right).$$

*Proof.* We bound

$$\begin{aligned}\exp\left(-\frac{\delta'^2 \underline{\mu'}^2}{2m}\right) &\leq \exp\left(-\frac{\delta'^2 \underline{\mu'}^2}{2kN}\right) \\ &\leq \exp\left(-\frac{\delta'^2 \underline{\mu'}}{2^\ell},\right),\end{aligned}$$

where the first inequality is straightforward from $m = k(N-1) \leq kN$, and the second inequality is obtained using the definition of $\underline{\mu'}$ and the assumptions on the parameters,

$$\begin{aligned}\frac{\underline{\mu'}}{N} &= \frac{2^{-l}k\left[N - 1 - \gamma\left(1 + 4 \cdot 2^\ell - \log\left(\frac{\log k}{4\gamma}\right)\right) - 4\sqrt{2^\ell \gamma N}\right]}{kN} \\ &\geq 2^{-l}\left[1 - \frac{(1 + \gamma\left(1 + 4 \cdot 2^\ell\right) + 4\sqrt{2^\ell \gamma N})}{N}\right] \\ &\geq 2^{-l}\left[1 - \frac{(1 + 4 \cdot 2^{-l}\left(1 + 4 \cdot 2^\ell\right) + 4\sqrt{4N})}{N}\right] \\ &\geq 2^{-l}\left[1 - \frac{(21 + 8\sqrt{N})}{N}\right] \\ &\geq 2^{-l}\left[1 - \frac{29}{\sqrt{c \log k \cdot 2^\ell}}\right] \\ &\geq 2^{-l}\left[1 - \frac{29}{\sqrt{2^\ell}}\right] \\ &\geq 2^{-l}\left[1 - 29 \cdot 2^{-7}\right] \\ &\geq \frac{1}{2}2^{-l}\end{aligned}$$

□

## D Proof of Zero Knowledge

We prove that the Fischlin transform is zero-knowledge in the quantum random oracle model (QROM). Our proof follows standard techniques for adaptive oracle reprogramming in the QROM [GHHM21]. We begin by defining the

zero-knowledge simulator. Let $\mathsf{Sim}_\Sigma$ denote the honest-verifier zero-knowledge simulator for the underlying $\Sigma$-protocol as in Definition 6 .

**Definition 11 (Simulator $\mathsf{Sim}_{\mathsf{Fis}}$).** *The simulator $\mathsf{Sim}^H_{\mathsf{Fis}}$ proceeds as follows.*

***Setup.*** *Sample a random function*

$$\tilde{H} : [k] \times \mathcal{C} \to \{0,1\}^\ell.$$

*For each $i \in [k]$, sample $c_i \leftarrow^{\$} \mathcal{C}$ such that $\tilde{H}(i, c_i) = 0^\ell$. If no such value exists, abort.*

***Proof Generation.*** *For each $i \in [k]$, run*

$$(a_i, z_i) \leftarrow \mathsf{Sim}_\Sigma(x, c_i).$$

*Set $\mathbf{a} = (a_i)_{i\in[k]}$ and similarly $\mathbf{c}$, $\mathbf{z}$.*

***Oracle Programming.*** *For any query $(\mathbf{a}, i, c, z)$, it answers as follows:*

$$H'(\mathbf{a}, i, c, z) = \begin{cases} \tilde{H}(i, c) & \text{if } V(a_i, c, z) = 1 \\ H(\mathbf{a}, i, c, z) & \text{otherwise.} \end{cases}$$

***Output.*** *Define the proof*

$$\pi = (\mathbf{a}, \mathbf{c}, \mathbf{z}).$$

*and output $(\pi, H')$.*

With this simulator in hand, we can prove zero-knowledge.

**Theorem 4.** *For a $\Sigma$-protocol with commitment entropy, the Fischlin transform is zero-knowledge in the QROM.*

*Proof.* The simulator clearly runs in expected polynomial time.

*Hybrid Argument.* We prove indistinguishability by a sequence of hybrid experiments.

**Hybrid $H_0$.** The honest prover executes the Fischlin protocol using an unmodified quantum random oracle.

**Hybrid $H_1$.** The prover generates commitments $(a_i)_{i\in[k]}$ honestly. The oracle is reprogrammed on all accepting transcripts with prefix $(\mathbf{a}, i)$ as in Definition 11. Challenges $c_i$ are sampled uniformly, and responses $z_i$ are computed honestly.

By the adaptive reprogramming lemma of [GHHM21] (Lemma 1), $H_1$ is computationally indistinguishable from $H_0$, except with probability negligible in the min-entropy of the commitment vector.

**Hybrid $H_1'$.** As in $H_1$, but challenges $c_i$ are sampled conditioned on $\tilde{H}(i, c_i) = 0^\ell$.

Since $\tilde{H}$ is uniformly random, this conditioning affects the distribution only negligibly. Hence $H_1'$ is indistinguishable from $H_1$.

**Hybrid $H_2$.** As in $H_1'$, but responses are generated using the simulator $\mathsf{Sim}_\Sigma$ instead of the honest prover.

By the honest-verifier zero-knowledge property of the $\Sigma$-protocol, $H_2$ is indistinguishable from $H_1'$ up to advantage at most $k\epsilon$, where $\epsilon$ is the distinguishing advantage for a single repetition.

Hybrid $H_2$ corresponds exactly to the output of $\mathsf{Sim}_{\mathsf{Fis}}$. Since each transition introduces only negligible distinguishing advantage, the simulator output is computationally indistinguishable from a real proof. This establishes zero knowledge in the QROM. □