

Black-Box Separations for Non-Interactive Commitments in a Quantum World

Kai-Min Chung^{1 *}, Yao-Ting Lin^{2 **}, and Mohammad Mahmoody^{3 ***}

¹ Academia Sinica, Taiwan

² UCSB, USA

³ University of Virginia, USA

Abstract. Commitments are fundamental in cryptography. In the classical world, commitments are equivalent to the existence of one-way functions. It is also known that the most desired form of commitments in terms of their round complexity, i.e., *non-interactive* commitments, *cannot* be built from one-way functions in a black-box way [Mahmoody-Pass, Crypto’12]. However, if one allows the parties to use quantum computation *and* communication, it is known that non-interactive commitments (to classical bits) are in fact possible [Koshiba-Odaira, Arxiv’11 and Bitansky-Brakerski, TCC’21].

We revisit the assumptions behind non-interactive commitments in a quantum world and study whether they can be achieved using quantum computation and *classical* communication based on a black-box use of one-way functions. We prove that doing so is impossible unless the Polynomial Compatibility Conjecture [Austrin et al. Crypto’22] is false. We further extend our impossibility to protocols with quantum decommitments. This complements the positive result of Bitansky and Brakerski [TCC’21], as they only required a classical decommitment message. Because non-interactive commitments can be based on injective one-way functions, assuming the Polynomial Compatibility Conjecture, we also obtain a black-box separation between one-way functions and injective one-way functions (e.g., one-way permutations) even when the construction and the security reductions are allowed to be quantum. This improves the separation of Cao and Xue [Theoretical Computer Science’21] in which they only allowed the *security reduction* to be quantum.

At a technical level, we prove that sampling oracles at random from “sufficiently large” sets (of oracles) will make them one-way against polynomial quantum-query adversaries who also get arbitrary polynomial-size quantum advice about the oracle. This gives a natural generalization of the recent results of Hhan et al. [Asiacrypt’19] and Chung et al. [FOCS’20].

* kmchung@iis.sinica.edu.tw. Supported in part by the NSTC QC project, under Grant no. NSTC 111-2119-M-001-004- and the Air Force Office of Scientific Research under award number FA2386-20-1-4066.

** yao-ting.lin@ucsb.edu. Part of the work was done when working at Academia Sinica.

*** mohammad@virginia.edu. Supported by NSF grants CCF-1910681 and CNS1936799.

Table of Contents

1	Introduction	2
1.1	Our Results	4
1.2	Technical Overview	6
1.3	Further Related Work	11
2	Preliminaries	12
2.1	Quantum Computation	12
2.2	Polynomial Compatibility Conjecture	14
2.3	The Donoho–Stark Uncertainty Principle	16
2.4	Non-Interactive Commitments	17
3	Non-uniform Hardness of Inverting Large Sets of Oracles	18
3.1	Oracle Puzzles with Advice	19
3.2	Multi-Instance Oracle Puzzles	20
3.3	Function-Inversion Oracle Puzzles	21
3.4	Proof of One-Wayness Under Quantum Advice	21
4	Quantum Black-Box Separation from One-Way Functions	22

1 Introduction

Commitment schemes are one of the most basic building blocks in the foundations of cryptography with a variety of applications. In a non-interactive commitment scheme, a sender Sen who holds a (say single bit) message b sends a *commitment* message com to a receiver Rec in such a way that the com acts as a secure vault hiding b ; this is formalized as follows. (1) The *hiding* property requires that com does not reveal anything about b to a computationally bounded receiver. (2) The *binding* property requires that after sending com , the sender is essentially bound to at most one $b \in \{0, 1\}$ and cannot change its mind afterwards. (3) The *completeness* of the scheme requires that the sender shall be able to convincingly reveal b using a *decommitment* message dec that functions like a password to the vault holding b .

In the classical setting, *interactive* commitments can be based on the *minimal* assumption that one-way functions exist [IL89, Nao90, HILL99]. However, when one wants to obtain the more desirable non-interactive variant, cryptographic primitives such as injective one-way functions [Yao82, GL89] and most public-key assumptions [LS19] have been shown to be sufficient. These constructions are black-box [IR89, RTV04, BBF13], in the sense that (1) they use the assumed primitive (e.g., in this case, one-way functions) as an oracle in their implementation, and (2) their security is proved by a reduction that treats the imagined adversary (breaking the construction) as an oracle as well. Moreover, it is known that OWFs cannot be used in a black-box way to obtain any of the primitives that are known to imply non-interactive commitments [MM11, IR89] or the non-interactive commitment itself [MP12].

Commitments in the quantum setting. With the rise of quantum computation in computer science and cryptography, questions that were previously considered to be well-understood in the classical setting are being revisited. In the quantum setting, we allow Sen , Rec to both run in *quantum* polynomial time, while the committed bit b is still classical. By default, the commitment and decommitment messages would also be quantum messages (but we would prefer them to be classical too, if possible). It has been shown [May97, LC97] that similarly to the classical setting, *some* form of computational intractability is necessary for commitments in the quantum setting, and (even interactive) commitments with statistical (hiding and binding) security cannot be achieved if adversaries are allowed to be quantum. When it comes to the *assumptions* behind commitments in the quantum setting, a sequence of works [DMS00, CLS01, KO09, LQWY14, Yan20] led to the perhaps surprising result that black-box constructions of non-interactive commitments (for classical messages) with various forms of binding properties that are meaningful in the quantum setting could in fact be constructed from (post-quantum) one-way functions [KO11, BB21] in the quantum setting.

The QCCC model: quantum computation and classical communication. The full advantage of quantum cryptography will rely on using both quantum computation as well as quantum communication. However, the internet is currently a classical communication medium, so it is much more desirable to design protocols that stick to classical communication as much as possible, even if they rely on local quantum computation (and hardness assumptions). In fact, the recent active line of work on classically verifying quantum power [Mah18, CCY20, ACGH20, BKVV20, Zha21, Bar21] also falls into this quantum-computation and classical-communication (QCCC) model. Thus, we revisit constructing non-interactive commitments from OWFs in the quantum setting as well and ask the following question.

Can we construct non-interactive commitments from (post-quantum) one-way functions using quantum computation and (only) classical communication?

In fact, one can study a relaxed version of the question above by limiting only *one* of the commitment or decommitment messages to be quantum. The OWF-based constructions of [KO11, BB21] and the candidate OWF-based construction of [YWLQ15] all used *quantum* messages. Among them, the work of [BB21] managed to make the decommitment message classical (while using a quantum commitment). Hence, their work suggests that perhaps using quantum commitment messages is the key to getting non-interactive commitments from OWFs. Therefore, a natural related question is whether one can obtain non-interactive commitments from OWFs in the quantum setting while limiting the *commitment* message to be classical and allowing the decommitment to be quantum. Note that commitment messages are *stored* for a longer time between the two phases of the commitment scheme, decommitment messages are revealed at the *very end*. Therefore, if only one of these messages is going to be quantum, it is perhaps preferred that the decommitment message is the quantum one.

Quantum binding vs. classical binding. We note that while hiding remains reasonably straightforward to define against quantum polynomial-time adversaries, binding is a

subtle property and could be defined in different ways. In fact, for statistically hiding commitments, it is *impossible* to achieve the same strong notions of binding, similar to the classical variant, in which the commitment message essentially binds the committed bit to be at most one value [May97, LC97, BB21]. As a result, in the statistically-hiding setting, we usually settle down for the weaker notion of “sum binding”, in which the probability p_b of opening successfully to b satisfies $p_0 + p_1 \leq 1 + \text{negl}(\kappa)$, where κ is the security parameter. For this setting, [Unr16a] proposed the alternative notion of “collapse-binding”. By only requiring computational hiding, [BB21] showed that a very close notion to the classical form of binding is in fact possible if one allows the receiver to only make a (partial) measurement right after the commitment message is sent. Our main question above is meaningful with respect to all these variants of binding. Therefore, as we will clarify, it is *not* crucial for the reader to follow these subtle differences, as our (negative) results can be stated with a *weaker* notion of binding in which the adversary has to successfully decommit into both values of 0, 1.

1.1 Our Results

At a high level, we answer our main question above negatively with respect to quantum black-box constructions [HY20], unless a recent conjecture about low-degree and low-influence polynomials is false [ACC⁺22]. The work of [ACC⁺22] showed that assuming this conjecture, one can break perfectly complete QCCC key-agreement protocols (with classical communication) in the Quantum Random Oracle Model [BDF⁺11] by asking only a polynomial number of queries to the oracle. As a result, they obtained black-box separations for perfectly complete key agreements in the QCCC model from OWFs. We first explain this conjecture and then will state our results based on it.

The Polynomial Compatibility Conjecture (PCC). Suppose $f = \sum_{S \subseteq [N]} \alpha_S \prod_{i \in S} x_i$ is a polynomial over Boolean variables $x_i \in \{\pm 1\}, i \in [N]$ and real coefficients $\alpha_S \in \mathbb{R}, S \subseteq [N]$. The degree of f is $\max_{\alpha_S \neq 0} |S|$ and $\|f\|_2 = \mathbb{E}_{\mathbf{x} \leftarrow \{\pm 1\}^N} [f(\mathbf{x})^2]$. The influence of x_i on f is $\text{Inf}_i(f) = \sum_{S \ni i} \alpha_S^2$, and for a distribution $\mathbf{F}$ over such polynomials, we let $\text{Inf}_i(\mathbf{F}) = \mathbb{E}_{f \leftarrow \mathbf{F}} [\text{Inf}_i(f)]$ be the *expected* influence. The PCC (for the group $\mathbb{Z}_2$) states that for sufficiently small $\delta(d) = 1/\text{poly}(d)$, if $\mathbf{F}, \mathbf{G}$ are distributions over polynomials of degree d over variables $x_1, \dots, x_N \in \{\pm 1\}$, $\|\cdot\|_2$ -norm equal to 1, and expected influences $\text{Inf}_i(\mathbf{F}), \text{Inf}_i(\mathbf{G}) \leq \delta(d)$ for all $i \in [N]$, there exist $f \in \text{supp}(\mathbf{F}), g \in \text{supp}(\mathbf{G})$ and $\mathbf{x} \in \{\pm 1\}^N$ such that $f(\mathbf{x}) \cdot g(\mathbf{x}) \neq 0$. The work of [ACC⁺22] gave some evidence for the validity of the PCC by proving a weaker statement than the PCC in which the influences are exponentially $\exp(-d)$ small.⁴

We prove the following conditional black-box separation for non-interactive commitments in the quantum setting. Our result holds even for a “weak” variant of binding that is necessary for all the proposed forms of binding in the quantum setting. In particular, we say that a malicious sender breaks the weak binding if it can come up with a commitment message com and a pair of decommitment messages $(\text{dec}_0, \text{dec}_1)$ such that using dec_b allows the sender to successfully decommit com to b .

⁴ The PCC bears some similarities to a conjecture by Aaronson and Ambanis [AA09] that also deals with polynomials with a low degree and low influence and which is also proved for exponentially small influences. See [ACC⁺22] for more discussions and comparisons.

Theorem 1.1 (Black-box separation of QCCC commitments from OWFs). *Assuming the Polynomial Compatibility Conjecture, there is no black-box construction of non-interactive commitments from (post-quantum) one-way functions in the QCCC model. Moreover, this holds even if the decommitment message is allowed to be quantum.*

Theorem 1.1 complements the *positive* result of [BB21], in which they show that there is a commitment scheme with a quantum commitment and a classical decommitment based on post-quantum OWFs. In other words, our work (conditionally) shows that one cannot trade a quantum commitment message with a quantum decommitment message and still use post-quantum OWFs when constructing non-interactive commitments from OWFs in a black-box way.

Corollary: separating injective OWFs from OWFs in the quantum setting. Injective one-way functions (e.g., one-way permutations) with classical input/outputs imply non-interactive commitments in a black-box way [GL89]. Therefore, Theorem 1.1 also implies the corollary that assuming the PCC, black-box construction of injective one-way functions from general one-way functions does not exist, even if the construction is allowed to use quantum computation. The work of [CX21] proved such a separation only when the *security reduction* is quantum, but our result extends to fully quantum constructions (assuming the PCC).

Corollary: separating NICs from pseudorandom states. The recent works of [AQY22, MY22] suggest that as opposed to the classical setting, OWFs might *not* be necessary for non-interactive commitments in the quantum setting: they could be constructed from “pseudorandom states” [BS20], which are weaker than OWFs [Kre21]. The work of [JLS18] showed that “pseudorandom states” (PRSs) can be based on OWFs in a black-box way. Therefore, as a corollary of our result, we obtain the separation between NICs and PRSs in the CCQD model (i.e., the model in which the commitment message is classical, but the decommitment is allowed to be quantum). We point out that the construction in [MY22] requires quantum commitment messages (but classical decommitment messages), and the construction in [AQY22] is interactive.

We will explain the key ideas behind the proof of Theorem 1.1 in Section 1.2. Before doing so we highlight one key technical tool that we develop along with the proof of Theorem 1.2 and believe to be of independent interest.

When are randomized oracles quantum one-way? It is known that if $\mathcal{H}_n$ denotes the set of all functions from $\{0, 1\}^n$ to $\{0, 1\}^n$, then a random oracle $f \leftarrow \mathcal{H}_n$ is one way against polynomial-time adversaries who even get arbitrary polynomial-size advice about the random oracle [IR89, GT00]. This classical result holds even if the adversary can ask quantum superposition queries to the (random *permutation*) oracle [NABT15], and even if the auxiliary information about the random oracle is quantum [HXY19, CGLQ20]. We revisit this phenomenon in a more general setting and ask the following question. What happens if the oracle $f: \{0, 1\}^n \mapsto \{0, 1\}^n$ is *not* completely random, yet it is sampled at random from a “large” set of oracles $\mathcal{F} \subseteq \mathcal{H}_n$. We give a concrete bound on how large $\mathcal{F}$ needs to be to make a random $f \leftarrow \mathcal{F}$ one-way against efficient non-uniform quantum adversaries that also receive quantum auxiliary advice about f .

Theorem 1.2 (One-wayness of oracles under quantum auxiliary input). *Suppose $\mathcal{H}_n$ denotes the set of all functions from $\{0, 1\}^n$ to $\{0, 1\}^n$, and that*

$$|\mathcal{F}| \geq 2^{-\frac{2^n}{n^{\omega(1)}}} \cdot |\mathcal{H}_n|$$

for a set of functions $\mathcal{F}$. Then a randomly selected $f \leftarrow \mathcal{F}$ will be one-way against quantum adversaries who ask $\text{poly}(n)$ quantum queries to f and receive at most $\text{poly}(n)$ bits of quantum advice about f .

See Section 3 for a more quantitative and general statement.

At a very high level, we use Theorem 1.2 to prove Theorem 1.1 by picking $\mathcal{F}$ to model a *large* set of oracles that can be used by a *cheating receiver*, while the advice about each oracle is a pair of decommitments to $b = 0, 1$. The security reduction of the supposedly black-box construction of non-interactive commitments from one-way functions would then lead to an adversary who inverts $f \leftarrow \mathcal{F}$ using $\text{poly}(n)$ number of queries and $\text{poly}(n)$ bits of advice, which is a contradiction due to Theorem 1.2.

1.2 Technical Overview

In this section, we explain some of the key ideas behind the proofs of Theorems 1.1 and 1.2 and the links between these two results. Our starting point for proving Theorem 1.1 is the black-box separation of non-interactive commitments from one-way functions in the classical setting [MP12]. We first sketch the argument for the classical case and then explain the challenges that arise in the quantum setting.

Recap of the proof for the classical setting. An approach for proving a black-box separation between primitives $\mathcal{Q}$ and OWFs is as follows. We show that the primitive $\mathcal{Q}$ can be broken relative to a *random oracle* h by asking only a polynomial number of queries.⁵ However, when we want to separate commitments from OWFs (even in the classical setting), we cannot simply use random oracles as mentioned above. The reason is that the random oracle can indeed be used to obtain *injective* one-way functions (with high probability), which in turn do imply non-interactive commitments. That is why, in [MP12], the oracle h used for the separation is chosen from a more subtle distribution: h is chosen *either* at random, *or* from a “partially fixed” random oracle. The key idea is to show that *at least* one of these two oracle distributions leads to breaking the commitment scheme while one-way functions exist relative to both oracle distributions. In particular, each randomized oracle corresponds to one of the parties of the commitment scheme to be the cheater.

- *Cheating receiver Rec^* relative to a random oracle.* Let h be a (fully) random oracle $h: \{0, 1\}^n \mapsto \{0, 1\}^n$, where $n = \kappa$ is the security parameter. Suppose Sen is committing to a *random* $b \in \{0, 1\}$ and sends message com to the receiver. Then, let $\text{Rec}^*(\text{com})$ be a cheating receiver who tries to learn the oracle answer to any query $x \in \{0, 1\}^n$ such that x has been asked by Sen with probability at least ε , for a parameter $\varepsilon = 1/\text{poly}(\kappa)$. Such queries were called “ ε -heavy” in [BM09].

⁵ E.g., one can re-interpret the proofs of [IR89, BM09] to fall into this framework.

and it was shown that regardless of com , there are (on average) at most d/ε of them if d is the number of oracle queries by Sen . Suppose the partial oracle $\mathcal{L}$ contains all the query-answer pairs learned by Rec^* . If Rec^* could now guess the random committed bit b (information-theoretically) with probability $1/2 + 1/\text{poly}(\kappa)$, it means Rec^* has succeeded in its attack *in the random oracle* model. In this case, we would be done with the separation; the reason is that the security reduction S of the black-box construction shall now be able to use Rec^* and invert the random oracle h with non-negligible probability, which is in fact impossible because the combined algorithm S^{Rec^*} is still asking polynomially many queries to h .

- *Cheating sender Sen^* relative to a non-random oracle.* Now, suppose the above attack by Rec^* does *not* succeed. In this case, we show that one can construct a cheating sender strategy Sen^* along with a fixed triple $(\text{com}, \text{dec}_0, \text{dec}_1)$ and a distribution $\mathbf{h}$ over the oracles with the following.

- com is a commitment message and dec_b is a decommitment message for b .
- $\mathbf{h}$ is a distribution over oracles that are random everywhere except on a $\text{poly}(n)$ -size subset of the input domain $\{0, 1\}^n$.
- The honest receiver accepts both $(\text{com}, \text{dec}_b)$, $b \in \{0, 1\}$ relative to all $h \leftarrow \mathbf{h}$.

If one can demonstrate the existence of the above triple $(\text{com}, \text{dec}_0, \text{dec}_1, \mathbf{h})$, it again implies that black-box construction of Sen, Rec from one-way functions are impossible: the security reduction S shall again be able to use Sen^* and invert the partially fixed random oracle $h \leftarrow \mathbf{h}$, but that is again impossible because the oracle $\mathbf{h}$ is only partially fixed, and partially fixed random oracles are also one-way.

The reason that such $(\text{com}, \text{dec}_0, \text{dec}_1, \mathbf{h})$ exists is as follows. Since we assumed that Rec^* had failed in its own attack above, conditioned on $(\text{com}, \mathcal{L})$, both bits $b = 0, 1$ are equally likely to be the truly committed bit. Therefore, if we further condition on $b = 0$ or $b = 1$, no 3ε -heavy queries would exist outside $\mathcal{L}$ (because both of the events $b = 0, 1$ have probability about $1/2$). Now, if we sample the view of the Sen twice, one conditioned on $b = 0$ and one conditioned on $b = 1$, two things happen: (1) we obtain decommitments two $\text{dec}_0, \text{dec}_1$ for $b = 0, 1$, and (2) we obtain *partial functions* h_0, h_1 that denote (only) the queries asked by Sen while committing to 0, 1 to generate com and $\text{dec}_0, \text{dec}_1$. Due to the lack of heavy queries in both of these sampling processes, the partial oracles h_0, h_1 will be *disjoint* with (high) probability $1 - O(d\varepsilon)$, and so they can be combined into a *single* partial oracle $h_{0,1} = h_0 \cup h_1$. Together with the partial oracle $\mathcal{L}$, $h_{\text{fixed}} = h_{0,1} \cup \mathcal{L}$ will shape the fixed part of the random oracles $\mathbf{h}$ that is useful for the cheating sender.

New challenges in the quantum setting. When we move to the setting in which the honest parties are quantum, several steps of the argument above will break down. We go over these issues one by one and explain the ideas for resolving them.

1. *Quantum analogue of learning heavy queries.* Since Sen^* can ask *quantum* queries to its oracle h (that is supposedly a random oracle), it no longer makes sense to use the classical ε -heavy query learners. However, the recent work of [ACC⁺22] showed how to extend this technique (by relying on ideas inspired by Zhandry's compressed oracle technique [Zha19]) to the quantum setting as follows. The re-

ceiver shall consider the sender's computation and the oracle all in a purified⁶ way, while the oracle's answers are represented in the Fourier basis. This way, any query x that has at least ε chance of having a *nonzero* answer in the Fourier conditioned on the commitment message com , will be considered *quantum ε -heavy*. The intuition is that being zero in the Fourier basis is (almost) the same as not being read by anyone, and hence remaining uniform. Note that the heavy queries are classical. It can also be shown, just like in the classical setting, that the total number of quantum heavy queries is $O(d/\varepsilon)$ where d is the number of quantum queries by the sender.

2. *Quantum analogue of partially fixed oracles.* In the classical setting, we could fix the partial oracle $h_{\text{fixed}} = h_{0,1} \cup \mathcal{L}$ that is consistent with two fixed openings $\text{dec}_0, \text{dec}_1$ as well as the learned and pick the rest of the oracle at random. However, in the quantum setting, it is no longer well-defined to refer to the “oracle queries asked by the sender” (i.e., $h_{0,1}$) as a partial oracle. That is because we cannot “record” the sender's queries, due to the quantum nature of its algorithm. Below we explain how to resolve this challenge. It turns out that resolving this challenge is even harder to resolve for protocols with quantum decommitments, so we will first go over the easier case of protocols in the QCCC model, before discussing the classical commitment quantum decommitment (CCQD) case.

Finding compatible oracles. Suppose after the sender runs out of learning quantum ε -heavy queries, $|\phi_0\rangle, |\phi_1\rangle$ are the marginal quantum states of the sender and the oracle for the two cases of $b = 0, 1$, the same commitment string com , and the set of fixed oracle answers in $\mathcal{L}$. First, we show that if the Polynomial Compatibility Conjecture (PCC) of [ACC⁺22] holds, then there is *at least one* oracle h (in the computational basis) that is consistent with both quantum states $|\phi_0\rangle, |\phi_1\rangle$. To do this, first, assume that the decommitment messages are classical. This means one can find two ensembles $\mathbf{H}_0, \mathbf{H}_1$ of quantum states for the oracle registers such that (1) $\mathbf{H}_b$ denotes quantum states for the oracle compatible with committing to b , and (2) $\mathbf{H}_0, \mathbf{H}_1$ can be uniquely modeled using distributions $\mathbf{F}_0, \mathbf{F}_1$ over degree- d polynomials (where d is the number of oracle queries of the sender) of influence at most 3ε , and (3) the oracles compatible with decommitting to b are the non-zero points of the polynomials sampled from $\mathbf{F}_b$. Therefore, by the PCC, there are indeed samples $f_0 \leftarrow \mathbf{F}_0, f_1 \leftarrow \mathbf{F}_1$, and an oracle h such that $f_b(h) \neq 0$ for both $b \in \{0, 1\}$. This means that the oracle h is compatible with two decommitments $\text{dec}_0, \text{dec}_1$ into both $b = 0, 1$ with respect to the same commitment com .

Boosting to many compatible oracles. Having only one compatible oracle h that allows opening com successfully into both $b \in \{0, 1\}$ using $\text{dec}_0, \text{dec}_1$ is not enough for proving the black-box separation, as h might be easy to invert. In particular, we need to show that such compatible oracle h can be found while it is also one-way. Our first idea for achieving this goal is that since the two polynomials f_0, f_1 have degree d , their (non-zero) product also has a degree at most $2d$. Therefore, we can

⁶ In the context of quantum information theory, purifying a quantum process means delaying all intermediate measurements to the end at the cost of introducing additional qubits. So that the whole computation remains a pure state until the final measurement.

use a variant of the Schwartz-Zippel lemma to conclude that at least 2^{-2d} fraction of all oracles h will satisfy $f_0(h)f_1(h) \neq 0$. When the group defining the oracle is not $\mathbf{Z}_2^n$, we can no longer apply the Schwartz-Zippel lemma, as the two functions $f_0(h)f_1(h)$ will not be low-degree polynomials, yet we derive a similar result using a generalization of the Schwartz-Zippel lemma known as the *Donoho–Stark support-size uncertainty principle* [DS89].

So far, we have shown that in the case of classical communication (including classical decommitments), the PCC implies that there is a *large* set of oracles $\mathcal{F}$ such that every $h \leftarrow \mathcal{F}$ is compatible with the commitment com followed by *both* decommitments $\text{dec}_0, \text{dec}_1$ into 0, 1. It remains to show that a random sample $h \leftarrow \mathcal{F}$ is hard to invert by $\text{poly}(n)$ -query *quantum* adversaries. This idea is implicit in [HXY19] (about the one-wayness of random oracles under quantum queries and classical auxiliary information) and generalizes to the case of sampling an oracle from a large set of oracles as well. However, this approach does not work when we want to attack protocols with *quantum* decommitment messages, as that requires working with *quantum* auxiliary information.

For classical decommitments, we rely on the fact that we can sample decommitments to 0, 1 and create an oracle (distribution) that is consistent with both. When decommitment messages are quantum, we can no longer measure the sender’s registers to create cheating strategies, because the decommitment messages are *quantum* and need to be kept as such. Therefore, we need to modify the approach above. Let $\mathcal{F}_b$ be the set of oracles (in the computational basis) that are compatible with an opening into b . Since measuring (or not measuring) the sender’s own registers (that will be used to produce the decommitment message) will *not* change the set $\mathcal{F}_b$, we first *pretend* that such measurement is happening to define two ensembles of oracles and use the PCC again to argue that the set $\mathcal{F} = \mathcal{F}_0 \cap \mathcal{F}_1$ contains at least an $\approx 1/d^{\Theta(n \cdot d)}$ fraction of the oracles that map $\{0, 1\}^n$ to $\{0, 1\}^n$.

To finish the proof, we need to prove two things: (1) for each $h \in \mathcal{F}$, there exist a pair of *quantum* decommitments $(\text{dec}_0, \text{dec}_1)$ that successfully decommit into 0, 1 with respect to commitment message com , and (2) picking a random oracle $h \leftarrow \mathcal{F}$ will lead to h that is hard to invert by polynomial-query adversaries. Item (1) is rather straightforward, due to the fact that $f \in \mathcal{F}_0$ (resp. $f \in \mathcal{F}_1$) are already defined to be the set of oracles that are compatible with at least one decommitments to 0 (resp. 1) with respect to com . However, Item (2) is now more challenging to prove when the decommitments are quantum messages. That is because, the security reduction S , now has access to $f \in \mathcal{F}$ as well as a pair of “advice” $(\text{dec}_0, \text{dec}_1)$ which can be seen as a piece of *quantum auxiliary information* about f , and so we would need to prove the one-wayness of the oracle $f \leftarrow \mathcal{F}$ against adversaries with quantum auxiliary information about f . Below, we focus on explaining the ideas for proving this specific one-wayness as an independent problem of its own.

Functions sampled from large sets are one-way for adversaries with quantum advice. As explained above, Theorem 1.1 reduces to Theorem 1.2, which states that any sufficiently large subset $\mathcal{F}$ of all oracles $\mathcal{H}_n = \{h \mid h: \{0, 1\}^n \mapsto \{0, 1\}^n\}$, a randomly selected function $h \leftarrow \mathcal{H}_n$ is “one-way” against any adversary who asks $\text{poly}(n)$ quan-

tum queries and gets $\text{poly}(n)$ bits of quantum advice about f . At a high level, we prove this result through a reduction to carefully chosen results from [CGLQ20].

Below we first recall the results in [CGLQ20], which can be used to prove the hardness of *completely random* oracles against quantum adversaries with quantum advice, and then show how their approach can be adapted to our case, in which a function is sampled from a large set (rather than all) of functions.

- *Non-uniform one-wayness of fully random functions* [CGLQ20]. Here we describe the approach of [CGLQ20] for proving non-uniform hardness of fully random functions. We then describe how the components of the proof of [CGLQ20] can be adapted for our setting. Let $h: \{0, 1\}^n \mapsto \{0, 1\}^n$ be a function. Consider a classical adversary A who receives S bit classical advice $\alpha = \alpha(h)$ about h and can ask T queries to h and manages to invert h with probability $\geq \varepsilon$, i.e.,

$$\Pr_{A, h, x} [A^h(\alpha, h(x)) \in h^{-1}(h(x))] \geq \varepsilon,$$

where the probability is over the randomness of A and the random choices of the *completely random* h and x . Now, consider a different attacking algorithm B that can ask $k \cdot T$ queries to h , but it does not receive any advice. However, the job of B is harder, as it needs to solve a *multi-instance version* of the inversion problem as follows: B is asked to invert k challenges $h(x_1), \dots, h(x_k)$. For each of these k challenges, the chance of inverting them is $O\left(\frac{kT}{2^n}\right)$ by a lazy-evaluation argument. Interestingly, as it is shown in [CGLQ20], and one can show that for the k -instance version, the success probability of any such algorithm B will decrease exponentially in k ,

$$\Pr_{B, h, x_1, \dots, x_k} [B^h(h(x_1), \dots, h(x_k)) \text{ inverts } h(x_1), \dots, h(x_k)] \leq O\left(\frac{kT}{2^n}\right)^k. \quad (1)$$

Now, we relate the success probability of A to that of B for bounding ε . First, using A we construct a new algorithm B' with S -bit classical advice and kT queries to h . B' simply uses the single copy of advice given to A and runs A to invert each $h(x_i)$ independently for all $i \in [k]$. To analyze the success probability of B' , we use the following argument. By an averaging argument, with probability at least $\varepsilon/2$ over the choice of h , A can invert them successfully with probability $\varepsilon/2$. Denote the set of such “good” functions by $\mathcal{G}$. Then, we have

$$\Pr_{B', x_1, \dots, x_k} [B'^h(\alpha, h(x_1), \dots, h(x_k)) \text{ inverts } h(x_1), \dots, h(x_k) \mid h \in \mathcal{G}] \geq \left(\frac{\varepsilon}{2}\right)^k.$$

Now, let B use B' and simply guess the advice α and use B' ; we have,

$$O\left(\frac{kT}{2^n}\right)^k \geq \Pr[B \text{ guesses } \alpha \text{ correctly}] \cdot \Pr_h[h \in \mathcal{G}] \quad (2)$$

$$\cdot \Pr_{B, x_1, \dots, x_k} [B^h(h(x_1), \dots, h(x_k)) \text{ inverts } h(x_1), \dots, h(x_k) \mid h \in \mathcal{G}] \quad (3)$$

$$\geq 2^{-S} \cdot \left(\frac{\varepsilon}{2}\right)^{k+1} \geq 2^{-S} \cdot \left(\frac{\varepsilon}{2}\right)^k. \quad (4)$$

By choosing k large enough $k = \tilde{O}(S)$ we obtain the desired bound of $\varepsilon = \tilde{O}(ST/2^n)$ between the adversary's number of queries, advice length, and its (small) chance of success. The main magic in the above argument is to leverage the exponential drop in success probability of the multi-instance game as shown in Equation (1) and to absorb the loss caused by guessing the advice.

Even though the above sketch was for the case of classical advice, in which guessing the advice is rather easy to analyze, as it was shown in [CGLQ20], a similar argument can be used for “guessing” quantum advice as well and use the above blueprint for proving one-wayness of a truly random function against adversaries with quantum queries and quantum advice.

- *Non-uniform one-wayness of functions sampled from large sets.* We now explain how the outline above can be adapted to the setting where we work with an oracle h that is sampled from a large enough sets of oracles $\mathcal{F}$, instead of picking h completely at random. To achieve our results, we have a simple but extremely useful observation as follows. Let's start by assuming an algorithm A can invert $f(x)$ with probability ε , when $f \leftarrow \mathcal{F}$ and $x \in \{0, 1\}^n$ are chosen uniformly at random. Using a similar averaging argument as the one above, we can still obtain a set of functions $\mathcal{G} \subseteq \mathcal{F}$ such that $|\mathcal{G}| \geq \varepsilon/2 \cdot |\mathcal{F}|$ such that A has success probability $\varepsilon/2$ conditioned on $h \in \mathcal{G}$. Going forward, the calculations in Equation (2) break down. In particular, we previously had $\Pr[h \in \mathcal{G}] \geq \varepsilon/2$, while we now have $\Pr[h \in \mathcal{G}] \geq \rho \cdot \varepsilon/2$, where ρ is the fraction of $|\mathcal{F}|$ in the set of all functions.

Our key observation is that, although ρ is very small, we prove that it is not *too* small. Therefore, the loss in the calculation of Equation (2) can be compensated by picking k *even larger* than before. In particular, recall that increasing k can bound the success probability of adversary to be *exponentially* small in k . Therefore, by picking k large enough in a careful way, we can recover an argument similar to the case of “all oracles” as outlined above.

1.3 Further Related Work

Here we discuss further related works that were not mentioned above already.

Quantum black-box separations. Hosoyamada and Yamakawa [HY20] initiated *quantum* black-box separations by formalizing quantum black-box constructions (for primitives with non-interactive security games) and ruling out the possibility of basing collision-resistant hash functions on one-way functions. Subsequently, [CX21] ruled out classical black-box constructions of post-quantum one-way permutations from post-quantum OWFs. The work of [ACC⁺22] ruled out quantum black-box constructions of perfectly complete key agreements from OWFs in the QCCC model. The work of [DLS22] ruled out quantum black-box reductions for proving the Fiat-Shamir heuristic, even in the presence of quantum shared entanglements.

Other assumptions than OWFs. The recent work of [BCQ22] showed that sampling statistically-far computationally-indistinguishable pairs of (mixed) quantum states, as a primitive, is a minimal assumption for many quantum primitives such as commitments, oblivious transfer, and secure multiparty computation.

In the classical setting, [BOV03] showed how to derandomize Naor’s 2-message commitment scheme that is based on OWFs and obtain a scheme that is non-interactive at the cost of introducing extra (derandomization-related) assumptions [NW94].

Post-quantum security. Our focus here is on commitments in which parties are quantum. Another line of work studies the post-quantum security of classical constructions [Unr12, Unr16a, Unr16b]. Another exciting recent line of work studies constructing stronger cryptographic primitives (such as oblivious transfer) from the minimal assumption that post-quantum OWFs exist [CDMS04, BCKM21, GLSV21].

2 Preliminaries

Notation. By κ we denote the security parameter. We use bold letters (e.g., $\mathbf{f}$) to denote random variables and distributions. We use calligraphic letters (e.g., $\mathcal{X}$) to denote sets. We use $\mathcal{Y}^{\mathcal{X}}$ to denote the set of all functions from $\mathcal{X}$ to $\mathcal{Y}$.

Throughout this work, we use the standard bra-ket notation (e.g., $|\psi\rangle$) for quantum objects. For the basics of quantum computation, we refer readers to [NC10].

2.1 Quantum Computation

An *oracle-aided quantum algorithm* $A^{(\cdot)}$ is a quantum algorithm with superposition query access to oracles. For any $d \geq 0$, an oracle quantum algorithm that makes d queries to oracles can be specified by a sequence of unitaries $U_0, \dots, U_d$, where the queries are executed between each unitary. Throughout this work, for any oracle $h : \mathcal{X} \rightarrow \mathcal{Y}$, we additionally define the range of the oracle $\mathcal{Y}$ to be an additive abelian group. In particular, by O_h we denote the query operator that maps the state $|x, y\rangle$ to $|x, y + h(x)\rangle$, where the addition is associated with the corresponding abelian group. The algorithm also has access to the inverse of the query operator $O_h^\dagger$.

In the *quantum random oracle model* (QROM for short) [BDF⁺11], a random function $h : \{0, 1\}^* \mapsto \{0, 1\}^\kappa$ is sampled in the beginning. Every party in the protocol (including honest parties and adversaries) has quantum access to h . If an algorithm in the QROM asks at most d queries to the oracle, we call it a d -query algorithm.

Zhandry [Zha19] showed that the purified random oracle is perfectly indistinguishable from the (standard) quantum random oracle. Since the sampling of the oracle commutes with the operators of the algorithm accessing the oracle, it can be deferred to the end. Here, we consider a more general setting. Consider an algorithm A with classical input that accesses quantum random oracle and outputs classical transcripts (classical leakage) during its computation (e.g., during an interactive protocol). Inspired by Zhandry’s work, we consider the *purified view* of such algorithms in the QROM. By the deferred measurement principle [NC10], all measurements of A can be replaced by unitaries if we introduce additional qubits for recording those measurement outcomes. After this modification, roughly speaking, by the purified view of A we mean the quantum

state obtained by executing A from scratch in a coherent way, in which the sampling of the oracle and intermediate measurements are deferred. The formal definition follows.

Definition 2.1 (Purified view of algorithms with classical leakage in the QROM).

Let A be an algorithm with quantum access to a random oracle $h : \mathcal{X} \rightarrow \mathcal{Y}$ that takes as classical input b chosen randomly from some set $\mathcal{B}$ with probability p_b , and (possibly) outputs classical transcripts c (perhaps produced in several steps) during its computation. Suppose A consists of a sequence of unitaries and query operators (but no measurements). For ease of notation, we represent A as a sequence $V_1, \dots, V_n$, where n is the size of A and each V_i is either a unitary operator or a query operator⁷. Let B be the input register, W be the workspace register, C be the transcript register, and H be the oracle register consisting of H_x for all $x \in \mathcal{X}$ while the content of each H_x stores $h(x)$. The purified view of A , denoted by $|\psi_n\rangle$, is defined as

$$|\psi_n\rangle := V_n V_{n-1} \dots V_1 |\psi_0\rangle,$$

where

$$|\psi_0\rangle := \frac{1}{\sqrt{|\mathcal{Y}|^{|\mathcal{X}|}}} \sum_{b \in \mathcal{B}, h \in \mathcal{Y}^{\mathcal{X}}} \sqrt{p_b} |b\rangle_B |0\rangle_W |h\rangle_H |0\rangle_C.$$

Purified view of the sender in commitments. We now apply Definition 2.1 to senders in commitments as follows. Let A be the sender's algorithm; the register B stores the input bit $b \in \{0, 1\}$, the register C stores the classical commitment message com , and part of the register W stores the classical (resp. quantum) decommitment message in the QCCC (resp. CCQD) model.

In a seminal work [Zha19], Zhandry observes that any d -query algorithm in the QROM has a *sparse* Fourier representation. In this work, we closely follow the rephrased version based on [ACC⁺22] for our use.

Definition 2.2 (Non-zero queries in Fourier basis). Let $\mathcal{Y}$ be a finite abelian group and $\hat{\mathcal{Y}}$ be the dual group. For any $\hat{h} \in \hat{\mathcal{Y}}^{\mathcal{X}}$, we define the size of $\hat{h}$ to be

$$|\hat{h}| := |\{x : x \in \mathcal{X}, \hat{h}(x) \neq \hat{0}\}|.$$

Definition 2.3 (The computational and the Fourier basis). Let $\mathcal{Y}$ be a finite abelian group with cardinality M . Let $\{|y\rangle\}_{y \in \mathcal{Y}}$ be an orthonormal basis of $\mathbb{C}^M$. We refer to this basis as the computational basis. Let $\hat{\mathcal{Y}}$ be the dual group which is known to be isomorphic to $\mathcal{Y}$. Recall that a member $\hat{y} \in \hat{\mathcal{Y}}$ is a character function (i.e., a function from $\mathcal{Y}$ to the multiplicative group of non-zero complex numbers). The Fourier basis $\{|\hat{y}\rangle\}_{\hat{y} \in \hat{\mathcal{Y}}}$ of $\mathbb{C}^M$ is defined as

$$|\hat{y}\rangle = \frac{1}{\sqrt{M}} \sum_y \hat{y}(y)^* |y\rangle \quad \text{and} \quad |y\rangle = \frac{1}{\sqrt{M}} \sum_{\hat{y}} \hat{y}(y) |\hat{y}\rangle.$$

⁷ Since A takes b as input, each V_i is defined to be a controlled-unitary with the control bit b .

Lemma 2.4 (Sparse representation [Zha19], rephrased). *For any d -query algorithm A with classical leakage in the $QROM$ with the oracle $h : \mathcal{X} \rightarrow \mathcal{Y}$, the purified view of A can be written as a (normalized) quantum state in the form of*

$$|\psi\rangle = \sum_{w,c,\hat{h}:|\hat{h}|\leq d} \alpha_{w,c,\hat{h}} |w\rangle_W |\hat{h}\rangle_H |c\rangle_C,$$

where W, H , and C , in order, denote the workspace of A , the oracle register, and the register recording the classical leakage.

When $\mathcal{Y}$ is a product of groups, i.e., $\mathcal{Y} = \mathcal{Y}_o^k$ for some integer $k \geq 1$ and abelian group $\mathcal{Y}_o$, then we immediately have the following corollary.

Corollary 2.5. *For any d -query algorithm A with classical leakage in the $QROM$ with the oracle $h : \mathcal{X} \rightarrow \mathcal{Y}_o^k$, the purified view of A can be written as a normalized quantum state in the form of*

$$|\psi\rangle = \sum_{w,c,\hat{h}_o:|\hat{h}_o|\leq dk} \alpha_{w,c,\hat{h}_o} |w\rangle_W |\hat{h}_o\rangle_H |c\rangle_C,$$

where $\hat{h}_o \in \hat{\mathcal{Y}}_o^{\mathcal{X}}$ and W, H and C , in order, denotes the workspace of A , the oracle register, and the register recording the classical leakage.

Definition 2.6 (Oracle support). *For any quantum state $|\phi\rangle = \sum_{w,\hat{h}} \alpha_{w,\hat{h}} |w\rangle_W |\hat{h}\rangle_H$ defined on an arbitrary register W and the oracle register H , define the oracle support in the Fourier basis of $|\phi\rangle$ as*

$$\widehat{\text{supp}}^H(|\phi\rangle) := \{\hat{h} \mid \exists w: \alpha_{w,\hat{h}} \neq 0\}.$$

Let $\hat{h}_{\max}^H(|\psi\rangle)$ denote the function $\hat{h} \in \widehat{\text{supp}}(|\phi\rangle)$ that has the largest size $|\hat{h}|$ (if such function is not unique, by default we pick the lexicographically first one). The definition extends naturally when the register W does not exist.

Definition 2.7 (Quantum ε -heavy queries [ACC⁺22]). *For any $x \in \mathcal{X}$, define the projector*

$$\Pi_x := \sum_{\hat{y} \in \hat{\mathcal{Y}} \setminus \{\hat{0}\}} |\hat{y}\rangle \langle \hat{y}|_{H_x}.$$

Given a quantum state $|\phi\rangle$ over registers W and H , the quantum heaviness of any $x \in \mathcal{X}$ is defined as

$$w(x) := \|\Pi_x |\phi\rangle\|^2,$$

i.e., the quantum heaviness of x is the probability of obtaining a non- $\hat{0}$ outcome while measuring H_x in the Fourier basis. We call x a quantum ε -heavy query if $w(x) \geq \varepsilon$.

2.2 Polynomial Compatibility Conjecture

In this section, we formally describe the Polynomial Compatibility Conjecture (PCC) of [ACC⁺22]. There are two equivalent formulations of this conjecture; one is based on low-degree polynomials, and the other is based on quantum states.

To keep the notation clean in this subsection, we identify $\mathcal{X}$ with $[N]$.

The Polynomial Formulation. Recall that for any $f : \mathcal{Y}^N \rightarrow \mathbb{C}$, it can be written in terms of its Fourier transform

$$f(\mathbf{x}) = \sum_{\chi \in \hat{\mathcal{Y}}^N} \hat{f}(\chi) \prod_{i=1}^N \chi_i(x_i),$$

where $\mathbf{x} = x_1 \parallel \dots \parallel x_N$. The *degree* of a character $\chi \in \hat{\mathcal{Y}}^N$ is $\deg(\chi) = |\{i \in [N] \mid \chi_i \neq \hat{0}\}|$, and the degree of f is $\deg(f) = \max\{\deg(\chi) \mid \hat{f}(\chi) \neq 0\}$. The ℓ_2 -norm of a function f is defined as $\|f\|_2 := \sqrt{\mathbb{E}_{\mathbf{x} \leftarrow \mathcal{Y}^N} |f(\mathbf{x})|^2}$. We say that f is *normalized* if $\|f\|_2 = 1$. The *influence* of variable i on f is $\text{Inf}_i(f) = \sum_{\chi \in \hat{\mathcal{Y}}^N, \chi_i \neq \hat{0}} |\hat{f}(\chi)|^2$.

Conjecture 2.8 (Polynomial Compatibility). There exists a finite abelian group $\mathcal{Y}$ and a function $\delta(d) = 1/\text{poly}(d)$ such that the following holds for all d, N . Let $\mathbf{F}$ and $\mathbf{G}$ be two distributions of functions from $\mathcal{Y}^N$ to $\mathbb{C}$ ⁸ such that the following holds for all $f \in \text{supp}(\mathbf{F})$ and $g \in \text{supp}(\mathbf{G})$.

- **Unit ℓ_2 norm:** f and g have ℓ_2 -norm 1.
- **d -degrees:** $\deg(f) \leq d$ and $\deg(g) \leq d$.
- **δ -influences on average:** For all $i \in [N]$, we have $\mathbb{E}_{f \leftarrow \mathbf{F}}[\text{Inf}_i(f)] \leq \delta$ and $\mathbb{E}_{g \leftarrow \mathbf{G}}[\text{Inf}_i(g)] \leq \delta$, where $\delta = \delta(d)$.

Then, there is an $f \in \text{supp}(\mathbf{F})$, $g \in \text{supp}(\mathbf{G})$, and $\mathbf{x} \in \mathcal{Y}^N$ such that $f(\mathbf{x}) \cdot g(\mathbf{x}) \neq 0$.

Here we describe an equivalence between quantum states and polynomials. In Section 4, we first use the formulation of quantum states. After proving that the states possess certain properties, we will convert the states into polynomials by Lemma 2.12, which enables us to apply Conjecture 2.8. For completeness, we provide relevant definitions below; we refer readers to Sections 4 and 5 in [ACC⁺22] for more details.

Definition 2.9 (($\mathcal{Y}, \delta, d, N$)-state). Let H be a register over the Hilbert space $\mathbb{C}^{\mathcal{Y}^{\mathcal{X}}}$, where $|\mathcal{X}| = N$. A quantum state $|\psi\rangle$ over registers W and H is a $(\mathcal{Y}, \delta, d, N)$ -state if it satisfies the following two conditions:

- **d -sparsity:** $|\hat{h}_{\max}^H(|\psi\rangle)| \leq d$. In other words, for any measurement of the registers H in the Fourier basis, the oracle support in the Fourier basis (as defined in Definition 2.6) is at most d (note that this is regardless of the basis in which we measure the register W).
- **δ -lightness:** For every $x \in \mathcal{X}$, it holds that $w(x) \leq \delta$.

Definition 2.10 (State polynomial). For a (normalized) quantum state $|\psi\rangle$ over the register H , the state polynomial of $|\psi\rangle$ is the function $f_\psi : \mathcal{Y}^N \rightarrow \mathbb{C}$ defined by

$$f_\psi(h) = |\mathcal{Y}|^{N/2} \cdot \langle \psi | h \rangle = \sum_{\chi \in \hat{\mathcal{Y}}^N} \langle \psi | \chi \rangle \prod_{i=1}^N \chi_i(h_i). \quad (5)$$

Note that $\|f_\psi\|_2 = 1$.

⁸ As shown in [ACC⁺22], regardless of the image being $\mathbb{R}$ or $\mathbb{C}$, the conjectures are equivalent up to a constant factor in δ . For convenience, we use the version with $\mathbb{C}$.

Definition 2.11 (State polynomial distribution). For a (normalized) quantum state $|\psi\rangle$ over registers W, H , the state polynomial distribution of $|\psi\rangle$ is the distribution $\mathbf{F}_\psi$ over (normalized) functions f which is sampled by measuring W in the computational basis and then taking the (normalized) state polynomial corresponding to the residual collapsed state over the register H . Explicitly, if $|\psi\rangle_{WH} = \sum_{w, \hat{h}} \alpha_{w\hat{h}} |w\rangle_W |\hat{h}\rangle_H$, then the support set of $\mathbf{F}_\psi$ consists of the state polynomial f_{ψ_w} of the normalized state $|\psi_w\rangle := \sum_{\hat{h}} \alpha_{w\hat{h}} |\hat{h}\rangle_H / \left\| \sum_{\hat{h}} \alpha_{w\hat{h}} |\hat{h}\rangle_H \right\|$ for each w . The probability of each f_{ψ_w} is defined to be $\left\| \sum_{\hat{h}} \alpha_{w\hat{h}} |\hat{h}\rangle_H \right\|^2$.

Lemma 2.12. Let $\mathbf{F}_\psi$ be the state polynomial distribution of an arbitrary $(\mathcal{Y}, \delta, d, N)$ -state $|\psi\rangle$. Then the following holds.

1. **Unit ℓ_2 norm:** f has ℓ_2 -norm 1 for every $f: \mathcal{Y}^N \rightarrow \mathbb{C}$ in the support set of $\mathbf{F}_\psi$.
2. **d -degrees:** $\deg(f) \leq d$ for every $f: \mathcal{Y}^N \rightarrow \mathbb{C}$ in the support set of $\mathbf{F}_\psi$.
3. **δ -influences on average:** For all $i \in [N]$, we have $\mathbb{E}_{f \leftarrow \mathbf{F}_\psi} [\text{Inf}_i(f)] \leq \delta$.

2.3 The Donoho–Stark Uncertainty Principle

We now explain the Donoho–Stark support-size uncertainty principle [DS89]. For our purpose, we use the following rephrased version from [WW21]. Informally, the uncertainty principle states that one cannot *simultaneously* obtain high-precision information of a state in the computational and Fourier basis. Consider the purified oracle as a motivating example. The oracle register in the Fourier basis starts with the all-zero state, while it is uniformly random in the computational basis. This phenomenon can be interpreted as the following: the algorithm knows the oracle with perfect precision in the Fourier basis while having absolutely no precision in the computational basis. Lemma 2.13 below provides a trade-off between the achievable precision in the computational and Fourier bases in terms of the size of supports.

Lemma 2.13 (Theorem 3.1 in [WW21]). Let $\mathcal{Y}$ be a finite abelian group. If $f: \mathcal{Y} \rightarrow \mathbb{C}$ is a non-zero function and $\hat{f}: \hat{\mathcal{Y}} \rightarrow \mathbb{C}$ denotes its Fourier transform, then

$$|\text{supp}(f)| \cdot |\text{supp}(\hat{f})| \geq |\mathcal{Y}|.$$

Corollary 2.14. Given $f_0, f_1: \mathcal{Y}^{\mathcal{X}} \rightarrow \mathbb{C}$ such that $\deg(f_0), \deg(f_1) \leq d$, we have

$$|\text{supp}(f_0) \cap \text{supp}(f_1)| \geq \frac{|\mathcal{Y}|^{|\mathcal{X}|}}{O(d|\mathcal{X}|^{2d}|\mathcal{Y}|^{2d})}.$$

Proof. Let $f := f_0 \cdot f_1$. It's easy to see that $\mathbf{x} \in \text{supp}(f)$ if and only if $\mathbf{x} \in \text{supp}(f_0) \cap \text{supp}(f_1)$. Since the degree of each f_0 and f_1 is at most d , their Fourier expansion can be written as

$$f_b(\mathbf{x}) = \sum_{\chi \in \hat{\mathcal{Y}}^N: \deg(\chi) \leq d} \hat{f}_b(\chi) \prod_{i=1}^N \chi_i(x_i)$$

where $b \in \{0, 1\}$.

Therefore, in the Fourier expansion of f , the characters with non-zero coefficients are of degree at most $2d$. Then the size of $\text{supp}(\hat{f})$ is at most the number of characters of degree at most $2d$. Namely,

$$|\text{supp}(\hat{f})| \leq \sum_{i=0}^{2d} \binom{|\mathcal{X}|}{i} (|\mathcal{Y}| - 1)^i \leq (2d + 1) \cdot (|\mathcal{X}||\mathcal{Y}|)^{2d}.$$

Together with Lemma 2.13, this finishes the proof. $\square$

2.4 Non-Interactive Commitments

Below we define non-interactive commitments.

Models. By QCCC we refer to the *quantum-computation classical-communication* model in which all the communications (including the commitment and decommitment messages) are classical. By CCQD we refer to the *classical-commitment quantum-decommitment* model, which is only defined for commitment schemes.

We now define non-interactive commitments with an extremely weak notion of binding. To break the weak binding, the adversary needs to prepare two decommitments for both $b = 0, 1$ such that both will be accepted if used during the decommitment. Using this notion makes our negative result stronger.

Definition 2.15 (Non-interactive weakly-binding commitments in CCQD model). A non-interactive commitment in the CCQD model consists of two quantum algorithms Sen, Rec . On input $b \in \{0, 1\}$, the sender $\text{Sen}(b, 1^\kappa)$ starts with $\text{poly}(\kappa)$ zero registers, $\text{poly}(\kappa)$ qubits of advice, and produces classical commitment message com and quantum decommitment message dec . The receiver (who also has $\text{poly}(\kappa)$ zero registers) receives $(\text{com}, b, \text{dec})$ and either accepts or rejects.

- **Completeness.** $\Pr[\text{Rec}(\text{com}, b, \text{dec}) = 1 \mid b \leftarrow \{0, 1\}, (\text{com}, \text{dec}) \leftarrow \text{Sen}(b)] = 1$.
- **Hiding.** We say Rec^* breaks hiding with advantage ε , if by picking $b \leftarrow \{0, 1\}$ at random, $\text{Rec}^*(\text{com})$ can correctly guess b with probability $(1 + \varepsilon)/2$. We call Sen hiding, if for every $\text{poly}(\kappa)$ -size quantum circuit Rec^* the advantage of Rec^* is at most $\text{negl}(\kappa)$.
- **Weak binding.** We say $(\text{com}, \text{dec}_0, \text{dec}_1)$ breaks the weak binding, if

$$\Pr[\text{Rec}(\text{com}, b, \text{dec}_b) = 1] = 1 \text{ for both } b \in \{0, 1\}.$$

We say that Rec has weak binding, if for all sequence $\{(\text{com}_\kappa, \text{dec}_{0,\kappa}, \text{dec}_{1,\kappa})_\kappa\}$ where $\text{com}_\kappa, \text{dec}_{0,\kappa}, \text{dec}_{1,\kappa}$ are of lengths at most $\text{poly}(\kappa)$, for all but finitely many κ , $(\text{com}_\kappa, \text{dec}_{0,\kappa}, \text{dec}_{1,\kappa})$ does not break the weak binding of Rec .

When the decommitment messages in a CCQD scheme are also classical, we say the resulting scheme is in the QCCC (quantum-computation classical-communication) model.

Note that in the definition above, we are implicitly working with poly -size (non-uniform) adversaries in our notion of weak binding. That is because a non-uniform

adversary might simply know the best way to open into both cases of 0, 1 without computational limitations. Having said that, even if we further weaken the security and ask for a *uniform* polynomial-time adversaries, it will not make a difference for a black-box separation (of an assumption behind non-interactive commitments). The reason is that the definition of black-box constructions (see below) requires the security reduction to work whenever it is given any *oracle* adversary regardless of its complexity.

Definition 2.16. A quantum black-box construction of weakly-binding non-interactive commitments from (length preserving) one-way functions is a pair of uniform QPT oracle-aided quantum algorithms (G, S) as follows.

- For every abelian group $\mathcal{Y}$ and every $f: \mathcal{Y}^\kappa \mapsto \mathcal{Y}^\kappa$, the oracle-aided quantum algorithm $G^f = (G_S^f, G_R^f)$ implements a quantum commitment scheme (both for the sender and receiver).
- For every abelian group $\mathcal{Y}$, for every $f: \mathcal{Y}^\kappa \mapsto \mathcal{Y}^\kappa$, and any oracle adversary $A = (A_h, A_b)$ who breaks the hiding or the weak binding of G^f , the algorithm $S^{f,A}$ inverts f with a non-negligible probability. In particular, S consists of two algorithms $S = (S_h, S_b)$, and there is a function $\delta = \text{poly}(\varepsilon/\kappa)$ such that: (1) if A_h ε -breaks the hiding of G_S^f , then S_h^{f,A_h} inverts f with probability $\delta = \text{poly}(\varepsilon/\kappa)$, and (2) if $A_b = (\text{com}_\kappa, \text{dec}_{0,\kappa}, \text{dec}_{1,\kappa})$ breaks the weak binding of G_R^f , then S_b^{f,A_b} inverts f with probability $\delta = \text{poly}(1/\kappa)$.

Remark 2.17. First, restricting the OWFs to have the same input and output spaces is without loss of generality. Because according to the definition of black-box reduction, the construction of the commitment scheme should work for *any* OWF. Hence, toward a contradiction, it's sufficient to show that the commitment scheme is impossible to be constructed from some *specific* OWF in a black-box way.

Next, we note that in the quantum setting, the quantum oracle access to a classical function depends on the underlying abelian group. By default, we assume $\mathcal{Y} = \mathbb{Z}_2$ and $f: \{0, 1\}^\kappa \mapsto \{0, 1\}^\kappa$ simply uses $\mathbb{Z}_2^\kappa$ as the group used for writing the answers in the registers (by adding them in $\mathbb{Z}_2^\kappa$). However, when we say a black-box construction from OWFs exists, it means that there is a version of the construction for *any* abelian group G (of constant size) instead of $\mathbb{Z}_2$, in which case the one-way function would look like $f: G^\kappa \mapsto G^\kappa$. Moreover, there are finite groups of any order, so assuming the input and output spaces of the OWFs have group structure is also without loss of generality.

3 Non-uniform Hardness of Inverting Large Sets of Oracles

In this section, we analyze a variant of the standard random functions inversion game in which the function is uniformly chosen from a specific set of functions instead of the set of all functions. In particular, we formalize and prove Theorem 1.2 in this section.

We consider the adversaries which are given classical or quantum advice and have quantum query access to the oracle. Arguments implicit in [HXY19] can be used for obtaining similar results but only for classical advice. Our proof, however, uses definitions and technical tools from [CGLQ20], and even in the case of classical advice we can obtain sharper bounds (than those obtained by arguments implicit in [HXY19]).

3.1 Oracle Puzzles with Advice

Definition 3.1 (Oracle algorithm with advice). An (S, T) -oracle-algorithm $A = (A_1, A_2)$ with (oracle-dependent) advice consists of two procedures:

- $|\alpha\rangle \leftarrow A_1(f)$, which is an arbitrary function of the oracle f , and outputs an S -qubit quantum state $|\alpha\rangle$;
- $|\text{ans}\rangle \leftarrow A_2^f(|\alpha\rangle, \text{ch})$, which is a computationally unbounded algorithm that takes advice $|\alpha\rangle$, a challenge ch , makes at most T quantum queries to f , and outputs an answer $|\text{ans}\rangle$, which we measure in the computational basis to obtain a classical answer ans if needed.

Furthermore, we distinguish the following cases:

- If the output of A_1 is classical, we call it a quantum algorithm with classical advice or an (S, T) -algorithm in the AI-QOM (auxiliary input quantum oracle model);
- If the output of A_1 is quantum, we call it a quantum algorithm with quantum advice or an (S, T) -algorithm in the QAI-QOM (quantum auxiliary input quantum oracle model);
- If $S = 0$, we call it a quantum algorithm without advice, or an algorithm in the QOM (quantum oracle model).

In the following interactive setting, the two terms “algorithm” and “adversary” will be used interchangeably.

Definition 3.2 (Oracle puzzle). An oracle puzzle $G = (\text{Chal}, \mathbf{f})$ is specified by a challenger $\text{Chal} = (\text{Samp}, \text{Ver})$ and a distribution $\mathbf{f}$ over oracles. In the beginning, an oracle is sampled $f \leftarrow \mathbf{f}$ and

- $\text{ch} \leftarrow \text{Samp}^f(r)$ is a deterministic classical algorithm that takes randomness r as input and outputs a classical challenge ch .
- $\text{Ver}^f(r, \text{ans})$ is a deterministic classical algorithm that takes as the input ans and outputs a decision b indicating whether the puzzle is won by the adversary.

For every algorithm with advice, i.e., $A = (A_1, A_2)$, we define

$$A_{\text{win}}^f := \text{Ver}^f(r, A_2^f(A_1(f), \text{Samp}^f(r)))$$

to be the binary variable indicating whether A wins the oracle puzzle.

We define the security loss in the AI-QOM, QAI-QOM of an oracle puzzle $G = (\text{Chal}, \mathbf{f})$ to be

$$\delta = \delta(S, T) := \sup_A \Pr_{f \leftarrow \mathbf{f}, r, A} [A_{\text{win}}^f = 1],$$

where A in the probability denotes the randomness of the (quantum) algorithm, and supremum is taken over all (S, T) -adversaries A in the AI-QOM/QAI-QOM respectively. We say an oracle puzzle G is $(1 - \delta)$ -secure if its security loss is at most δ .

In particular, we focus on a class of oracle puzzles in which the adversary can verify the answer by itself.

Definition 3.3 (Publicly-verifiable security game). We call an oracle puzzle to be publicly-verifiable with verification time T_{Ver} , if $\text{Ver}^f(r, \cdot) = \widetilde{\text{Ver}}^f(\text{ch}, \cdot)$ for some deterministic classical algorithm $\widetilde{\text{Ver}}^f$ where ch is determined by r and T_{Ver} is the upper bound on the number of f queries for computing $\widetilde{\text{Ver}}^f(\text{ch}, \cdot)$.

3.2 Multi-Instance Oracle Puzzles

Definition 3.4 (Multi-instance oracle puzzle). For any oracle puzzle $G = (\text{Chal}, f)$ and any positive integer $k \geq 1$, we define the multi-instance oracle puzzle $G^{\otimes k} = (\text{Chal}^{\otimes k}, f)$, where $\text{Chal}^{\otimes k}$ is given as follows

- For $i \in [k]$, do:
 1. Sample fresh randomness r_i ;
 2. Compute $\text{ch}_i \leftarrow \text{Chal.Samp}^f(r_i)$ and send it to the adversary;
 3. Give the adversary oracle access to f until the adversary submits a quantum state $|\text{ans}_i\rangle$;
 4. Let $\{P_0, P_1\}$ be a projective measurement where P_1 defines all ans 's such that $\text{Ver}(r, \text{ans}) = 1$ and $P_0 = I - P_1$. Measure $|\text{ans}_i\rangle$ in $\{P_0, P_1\}$ to get the quantum state $|\text{ans}'_i\rangle$ and store the result in $b_i \in \{0, 1\}$;
 5. Send $|\text{ans}'_i\rangle$ back to the adversary;
- Output $b_1 \wedge b_2 \wedge \dots \wedge b_k$;

Definition 3.5 (Multi-instance adversary). A (k, S, T) -adversary with advice for a multi-instance oracle puzzle $G^{\otimes k} = (\text{Chal}^{\otimes k}, f)$ consists of $A = (A_1, A_2)$, where the interaction between $A_2(|\alpha\rangle)$ and $\text{Chal}^{\otimes k}$ is defined as follows:

- $|\alpha\rangle \leftarrow A_1(f)$, which is an arbitrary (unbounded) function of f and outputs an S -qubit quantum state $|\alpha\rangle$ for A_2 ;
- For each $i \in [k]$,
 1. A_2 is given a challenge ch_i and the oracle access to f from $\text{Chal}^{\otimes k}$;
 2. A_2 makes at most T queries to f and prepares $|\text{ans}_i\rangle$;
 3. A_2 sends $|\text{ans}_i\rangle$ to $\text{Chal}^{\otimes k}$ and gets $|\text{ans}'_i\rangle$ back;
- Finally, $\text{Chal}^{\otimes k}$ outputs a bit b .

In particular, if $S = 0$, we also call it a (k, T) -adversary (without advice), or a (k, T) -algorithm in the QOM. In the rest of the section, we sometimes use such notation when it is clear from the context.

For any A which is a (k, S, T) -adversary with advice, we define $A_{\text{win}}^{\otimes k, f}$ to be the binary variable indicating whether A wins the multi-instance oracle puzzle.

We say a multi-instance oracle puzzle $G^{\otimes k}$ is $(1 - \delta)$ -secure⁹ in the QOM if for any (k, T) -adversary A (without advice),

$$\Pr_{f, A, \text{Chal}^{\otimes k}} [A_{\text{win}}^{\otimes k, f} = 1] \leq \delta^k = \delta(k, T)^k,$$

where A in the probability denotes the randomness of the algorithm, $\text{Chal}^{\otimes k}$ in the probability denotes the randomness of the challenger.

⁹ Actually, the security loss here is at most δ^k instead of δ . We follow this convention for ease of the presentation.

3.3 Function-Inversion Oracle Puzzles

Definition 3.6 (Function inversion oracle puzzle). The oracle puzzle $G_{\text{InvSet},N,M,R} = (\text{Chal}, \mathbf{f})$ parameterized by integers $R, N, M \geq 0$ is defined as follows:

- $\mathbf{f}$ is a uniform distribution over $\mathcal{F} \subseteq [M]^{[N]}$ such that $|\mathcal{F}|$ is at least M^{N-R} .
- $\text{Samp}^{\mathbf{f}}$ chooses x from $[N]$ uniformly at random, and outputs $\text{ch} = \mathbf{f}(x)$.
- $\text{Ver}^{\mathbf{f}}(x, x')$ outputs 1 if $\mathbf{f}(x) = \mathbf{f}(x')$.

Notice that $G_{\text{InvSet},N,M,R}$ is publicly-verifiable with $T_{\text{Ver}} = 1$. When $R = 0$, as a special case, the oracle puzzle corresponds to the standard random functions inversion game denoted by $G_{\text{InvAll},N,M}$.

In particular, [CGLQ20] prove the security of multi-instance oracle puzzle $G_{\text{InvAll},N,M}^{\otimes k}$ against (k, T) -adversaries in the quantum random oracle model (QROM). The formal statements are presented as follows.

Lemma 3.7 (Lemma 5.2 in [CGLQ20]). $G_{\text{InvAll},N,M}^{\otimes k}$ is $(1 - \delta(k, T))$ -secure¹⁰ in the QROM, where

$$\delta(k, T) = O\left(\frac{kT + T^2}{\min\{N, M\}}\right).$$

3.4 Proof of One-Wayness Under Quantum Advice

The following lemma reduces the multi-instance oracle puzzle $G^{\otimes k}$ against a (k, T) -adversary (without advice) to the (single-instance) oracle puzzle G against an (S, T) -adversary (with quantum advice).

Lemma 3.8 (Corollary 4.14 in [CGLQ20]). There exists a universal constant $c > 0$ such that the following holds. Given a publicly-verifiable oracle puzzle G with verification time T_{Ver} . Given an (S, T) -adversary A (with quantum advice) for G with winning probability δ , there exists a (k, T') -adversary A' (without advice) for the multi-instance oracle puzzle $G^{\otimes k}$ with winning probability at least $\delta' \geq 2^{-\ell S} \cdot (\delta/4)^{k+1}$ for any positive integer $k \geq 1$, where $T' = 2\ell(T + T_{\text{Ver}})$ and $\ell = c \cdot \log(k + 1)/\delta$.

Fact 3.9 (Fact 4.15 in [CGLQ20]) Given any real $C \geq 0, D \geq 2$. If $k_0 = C + D + 14$ and $k = 2k_0 \log k_0$, then we have $k \geq C \log(k + 1) + D$.

Now, we are ready to prove the function inversion oracle puzzle $G_{\text{InvSet},N,M,R}$ is secure against an (S, T) -adversary in the QAI-QOM.

Theorem 3.10. For any integer $R \geq 0$, the oracle puzzle $G_{\text{InvSet},N,M,R}$ is $(1 - \delta(S, T))$ -secure in the QAI-QOM, where

$$\delta(S, T) = \tilde{O}\left(\sqrt[3]{\frac{(S + R \log M) \cdot T + T^2}{\min\{N, M\}}}\right).$$

In particulate, if $S(\kappa) = \text{poly}(\kappa)$, $T(\kappa) = \text{poly}(\kappa)$, $R(\kappa) = \text{poly}(\kappa)$, $N = 2^{\Theta(\kappa)}$, and $M = 2^{\Theta(\kappa)}$, the security loss $\delta(\kappa)$ will be negligible in κ .

¹⁰ Recall that by our convention, the security loss is at most $\delta(k, T)^k$.

Proof. Suppose there exists an (S, T) -adversary A for $G_{\text{InvSet}, N, M, R} = (\text{Chal}, \mathbf{f})$ with winning probability $\delta = \delta(S, T)$. Then, by Lemma 3.8, there exists a (k, T') -adversary A' for $G_{\text{InvSet}, N, M, R}^{\otimes k}$ with winning probability at least $\delta' \geq 2^{-\ell S} \cdot (\delta/4)^{k+1}$ for any $k \geq 1$, where $T' = 2\ell(T + T_{\text{Ver}})$ and $\ell = c \cdot \log(k+1)/\delta$.

Here, we construct an adversary A'' for $G_{\text{InvAll}, N, M}^{\otimes k}$ by using A' as a black box. When A'' receives the challenge $f(x)$, it simply runs $A'^f(f(x))$ and outputs whatever $A'^f(f(x))$ outputs. The winning probability of A'' , denoted by δ'' , is at least

$$\begin{aligned} \delta'' &\geq \Pr[f \in \text{supp}(\mathbf{f})] \cdot \Pr[A'^f(f(x)) \text{ wins} \mid f \in \text{supp}(\mathbf{f})] \\ &\geq M^{-R} \cdot \delta' \geq 2^{-\ell S - R \log M} \cdot (\delta/4)^{k+1}, \end{aligned}$$

where $\text{supp}(\mathbf{f})$ denotes the support of $\mathbf{f}$.

By the definition of multi-instance security of $G_{\text{InvAll}, N, M}^{\otimes k}$, for all $k \geq 1$ we have

$$\delta(k, T')^k \geq \delta'' \geq 2^{-\ell S - R \log M} \cdot (\delta/4)^{k+1} \geq 2^{-\ell S - R \log M} \cdot (\delta_0/4) \cdot (\delta/4)^k,$$

where $1/N \leq \delta_0 \leq \delta$ is the winning probability of an adversary that outputs a random answer without advice or making any query.

Pick $k_0 = \frac{c}{\delta} S + R \log M + \log(1/\delta_0) + 16$ and $k = 2k_0 \log k_0$. By Fact 3.9, let $C = \frac{c}{\delta} S$ and $D = \log(1/\delta_0) + 2 + R \log M$, we have $k \geq C \log(k+1) + D = c \log(k+1) S/\delta + \log(1/\delta_0) + 2 + R \log M$.

Therefore, we have

$$\begin{aligned} \delta(k, T')^k &\geq 2^{-\ell S - R \log M} \cdot (\delta_0/4) \cdot (\delta/4)^k \\ &= 2^{-c \cdot \log(k+1) S/\delta} \cdot 2^{-\log(1/\delta_0) - 2 - R \log M} \cdot (\delta/4)^k \\ &\geq (\delta/8)^k \end{aligned}$$

or equivalently

$$\delta \leq 8\delta(k, T'),$$

where $k = \tilde{O}(S/\delta + R \log M)$ and $T' = \tilde{O}(T + T_{\text{Ver}})/\delta$.

By Lemma 3.7, it holds that

$$\delta \leq 8\delta(k, T') = \tilde{O} \left(\frac{(\frac{S}{\delta} + R \log M) \cdot \frac{T}{\delta} + \frac{T^2}{\delta^2}}{\min\{N, M\}} \right)$$

which leads to

$$\delta = \tilde{O} \left(\sqrt[3]{\frac{(S + R \log M) \cdot T + T^2}{\min\{N, M\}}} \right).$$

□

4 Quantum Black-Box Separation from One-Way Functions

In this section, assuming Conjecture 2.8 is true, we show that there is no black-box construction of non-interactive commitments (with perfect completeness) from OWFs

in the CCQD model. We emphasize that *all* known constructions of NICs that we are aware of have perfect completeness. The following theorem formalizes, and in fact generalizes, Theorem 1.1. In particular, Theorem 1.1 stated the result for the QCCC model (in which both messages are classical), while Theorem 4.1 allows the model to be CCQD, which lets the decommitment message to be quantum.

Theorem 4.1 (Black-box separation of CCQD commitments from OWFs). *Assuming Conjecture 2.8, there is no quantum black-box construction of non-interactive commitments in the CCQD model from one-way functions.*

We need the following notion characterizing the cardinality of sets of functions.

Definition 4.2 (α -flat distributions). *For $\alpha \in [0, 1]$, a distribution $\mathbf{f}$ over functions from $\mathcal{X}$ to $\mathcal{Y}$ is called an α -flat distribution if the size of the support is at least an α fraction of $\mathcal{Y}^{\mathcal{X}}$, i.e., $|\text{supp}(\mathbf{f})|/|\mathcal{Y}^{\mathcal{X}}| \geq \alpha$, and $\mathbf{f}$ is uniform over its support set.*

Next, we introduce a useful lemma from [ACC⁺22] that will help us argue about the efficiency of our attacks.

Lemma 4.3 (Efficiently learning quantum-heavy queries [ACC⁺22]). *Let A be an algorithm that asks at most d quantum queries to the random oracle $h : \mathcal{X} \rightarrow \mathcal{Y}$ and outputs a classical message com . For any $0 < \varepsilon < 1$, there exists a deterministic learning algorithm that learns a list $\mathcal{L}$ of (classical) query-answer pairs from the random oracle (i.e., a partial function), such that the following two conditions hold.*

1. *Efficiency of the learner: $\mathbb{E}[|\mathcal{L}|] \leq d/\varepsilon$, where the expectation is over the randomness of the oracle and the algorithm A .*
2. *Learning quantum heavy queries: When the learner stops and learns a list $\mathcal{L}$, there is no $x \notin \mathcal{Q}_{\mathcal{L}}$ that is quantum ε -heavy in the purified view of A conditioned on knowing $\mathcal{L}$ and com , where $\mathcal{Q}_{\mathcal{L}}$ denotes the domain of $\mathcal{L}$.*

The rest of the section is dedicated to proving Theorem 4.1. For readability and simplicity of the presentation, we first assume the abelian group associated with the random oracle to be $\mathbb{Z}_2^\kappa$ and Conjecture 2.8 holds for $\mathbb{Z}_2$. For the general case in which Conjecture 2.8 holds for some abelian group $\mathcal{Y}_\circ$, we instead pick the OWFs in Definition 2.16 as $f : \mathcal{Y}_\circ^\kappa \mapsto \mathcal{Y}_\circ^\kappa$. The following analysis still holds by replacing $\mathbb{Z}_2$ with $\mathcal{Y}_\circ$.

We will use the following lemma as the key to our proof of Theorem 4.1.

Lemma 4.4. *If Conjecture 2.8 is true, then for any quantum-black-box implementation of non-interactive commitments from oracle $f : \{0, 1\}^\kappa \mapsto \{0, 1\}^\kappa$ in which the sender asks d quantum oracle queries, there are cheating strategies $\text{Sen}^*, \text{Rec}^*$ such that at least one of the following holds.*

1. *Rec^* asks d oracle queries such that: if the f is a random oracle, then Rec^* has a non-negligible distinguishing advantage in breaking the hiding property of the commitment scheme.*

2. There is a $2^{-\text{poly}(\kappa)}$ -flat distribution $\mathbf{f}$ over the oracles such for all $f \leftarrow \mathbf{f}$, there exists (an auxiliary information) $(\text{com}, \text{dec}_0, \text{dec}_1)$ such that com is classical and $\text{dec}_0, \text{dec}_1$ are quantum messages and $(\text{com}, \text{dec}_0, \text{dec}_1)$ breaks the weak binding of the scheme relative to f .¹¹

Proof. Suppose (Sen, Rec) is a quantum-black-box implementation of non-interactive commitment from one-way functions in the CCQD model.

Construction 4.5 (The cheating receiver Rec^* with parameter ε) Let d be the number of oracle queries asked by the sender. Given the commitment com which commits to a random bit $b \in \{0, 1\}$, the description of the cheating receiver Rec^* is as follows:

1. Let A in Lemma 4.3 be Sen in which the sender commits to a random bit b . The output of A will be the commitment com . The cheating receiver Rec^* runs the learning algorithm in Lemma 4.3 over A with the parameter $\varepsilon = \frac{1}{10\delta(d\kappa)}$, where $\delta(\cdot)$ is the function defined in Conjecture 2.8 for $\mathcal{Y} = \mathbb{Z}_2$.
2. The cheating receiver Rec^* outputs the more likely input bit $b \in \{0, 1\}$ according to the purified view (i.e., conditioned on com and the learned classical queries $\mathcal{L}$ of the oracle) as its own output bit.

If the conditional distribution of input bit b has already been noticeably biased after the learning algorithm, then Rec^* would have a decent chance of breaking the hiding. Let $\mathcal{E}$ be the event that the distinguishing advantage

$$\frac{1}{2} |\Pr[b = 0 \mid \text{com}, \mathcal{L}] - \Pr[b = 1 \mid \text{com}, \mathcal{L}]|$$

is non-negligible holds. Then we either have $\Pr[\mathcal{E}] > 1/\kappa$ or $\Pr[\bar{\mathcal{E}}] \geq 1 - 1/\kappa$. If the former holds, it implies that Rec^* has a non-negligible distinguishing advantage and thus the proof is done. Therefore, we assume that we are in the latter case. By Lemma 4.3 and an averaging argument, the number of queries asked by Rec^* satisfies

$$\mathbb{E}[|\mathcal{L}| \mid \bar{\mathcal{E}}] \leq \frac{\mathbb{E}[|\mathcal{L}|]}{\Pr[\bar{\mathcal{E}}]} \leq \frac{1.01d}{\varepsilon}$$

for sufficiently large κ . Then by Markov's inequality, we have

$$\Pr\left[|\mathcal{L}| \geq \kappa^2 \cdot \frac{1.01d}{\varepsilon} \mid \bar{\mathcal{E}}\right] \leq \frac{1}{\kappa^2}.$$

Putting things together, we conclude that with probability at least $1 - O(1/\kappa^2)$, all of the following events hold:

- Rec^* is efficient: Rec^* asks at most $1.01\kappa^2 d/\varepsilon = \text{poly}(d, \kappa)$ queries.
- No quantum ε -heavy query left: for all $x \notin \mathcal{Q}_{\mathcal{L}}$, $w(x) < \varepsilon$ where $w(\cdot)$ is defined in Definition 2.7.
- $b = 0, 1$ are almost as likely: $|\Pr[b = 0 \mid \text{com}, \mathcal{L}] - \Pr[b = 1 \mid \text{com}, \mathcal{L}]| = \text{negl}(\kappa)$.

¹¹ One can think of $(\text{com}, \text{dec}_0, \text{dec}_1)$ as a cheating sender Sen^* .

Let $\mathcal{G}$ denote the event that all the above three events hold.

Next, assuming that Rec^* fails, we describe the cheating sender Sen^* as follows.

Construction 4.6 (The cheating sender and the flat distribution) *Now, we describe the cheating sender's strategy Sen^* and a corresponding α -flat distribution $\mathbf{f}$.*

1. *The cheating sender Sen^* samples $(\text{com}, \mathcal{L})$ according to the first step of the cheating receiver Rec^* in Construction 4.5.*

Before proceeding to the next step, we introduce some notations. Consider the purified view $|\Phi_{\text{com}, \mathcal{L}}\rangle$ of the honest sender of the commitment conditioned on the (classical) commitment message com and the list $\mathcal{L}$. Let $|\Phi_{0, \text{com}, \mathcal{L}}\rangle$ and $|\Phi_{1, \text{com}, \mathcal{L}}\rangle$ be the purified views further conditioned on b being 0 and 1. That is,

$$|\Phi_{\text{com}, \mathcal{L}}\rangle = \sqrt{\Pr[b=0 \mid \text{com}, \mathcal{L}]} |\Phi_{0, \text{com}, \mathcal{L}}\rangle + \sqrt{\Pr[b=1 \mid \text{com}, \mathcal{L}]} |\Phi_{1, \text{com}, \mathcal{L}}\rangle.$$

Let $\mathcal{X}' := \mathcal{X} \setminus \mathcal{Q}_{\mathcal{L}}$ and $N' := |\mathcal{X}'| = |\mathcal{X}| - |\mathcal{L}|$. Let H' be the oracle register corresponding to $\mathcal{X}'$. Note that conditioning on the list $\mathcal{L}$, the content of the oracle registers corresponding to $\mathcal{Q}_{\mathcal{L}}$ is fixed. So they are not entangled with H' . By abusing notation, for $b \in \{0, 1\}$, we also denote by $|\Phi_{b, \text{com}, \mathcal{L}}\rangle$ the state obtained by discarding the registers corresponding to $\mathcal{Q}_{\mathcal{L}}$. Let $\mathbf{F}_0$ be the state polynomial distribution of $|\Phi_{0, \text{com}, \mathcal{L}}\rangle$. Define $\mathbf{F}_1$ similarly.

2. *Find $f_0 \in \text{supp}(\mathbf{F}_0)$, $f_1 \in \text{supp}(\mathbf{F}_1)$ such that $f_0 \cdot f_1$ is not constant zero. If no such functions exist, then abort.*

Let $\mathcal{F}'$ be the set of all $h' \in \mathcal{Y}^{\mathcal{X}'}$ such that $(f_0 \cdot f_1)(h) \neq 0$, i.e.,

$$\mathcal{F}' := \{h' \in \mathcal{Y}^{\mathcal{X}'} \mid (f_0 \cdot f_1)(h) \neq 0\}.$$

The α -flat distribution $\mathbf{f}$ will be uniform over the set $\mathcal{F} \subseteq \mathcal{Y}^{\mathcal{X}}$ which contains all functions in $\mathcal{F}'$ combined with $\mathcal{L}$, i.e.,

$$\mathcal{F} := \{h \in \mathcal{Y}^{\mathcal{X}} \mid \exists h' \in \mathcal{F}' : h = h' \cup \mathcal{L}\}.$$

3. *The cheating sender Sen^* sends com as the commitment and uses the oracle-dependent quantum advice dec_b to decommit com into $b \in \{0, 1\}$.*

Suppose $\mathcal{G}$ occurs in the rest of the proof. Before using Lemma 2.12 to relate quantum states with polynomials, we first show that the purified views satisfy certain properties. First, by Corollary 2.5, the purified views satisfy

$$|\hat{h}_{\max}^{H'}(|\Phi_{b, \text{com}, \mathcal{L}}\rangle)| \leq d \cdot \kappa$$

for $b \in \{0, 1\}$, where the degree is defined over $\mathbb{Z}_2$. Next, notice that after the first step, none of the conditional probability of each input is greater than $2/3$ for sufficiently large κ . That is, both probabilities $\Pr[b=0 \mid \text{com}, \mathcal{L}]$ and $\Pr[b=1 \mid \text{com}, \mathcal{L}]$ are between $1/3$ and $2/3$. Therefore, given that the purified view $|\Phi_{\text{com}, \mathcal{L}}\rangle$ has no quantum ε -heavy query in $\mathcal{X}'$, we can conclude that both $|\Phi_{0, \text{com}, \mathcal{L}}\rangle$ and $|\Phi_{1, \text{com}, \mathcal{L}}\rangle$ have no quantum 3ε -heavy query in $\mathcal{X}'$. By our choice of ε , we have $3\varepsilon \leq \delta(d\kappa)$. Consequently, we have both $|\Phi_{0, \text{com}, \mathcal{L}}\rangle$ and $|\Phi_{1, \text{com}, \mathcal{L}}\rangle$ are $(\mathbb{Z}_2, \delta(d\kappa), d\kappa, N')$ -states. By Lemma 2.12, every $f: \mathbb{Z}_2^{N'} \rightarrow \mathbb{C}$ in the support set of $\mathbf{F}_0$ satisfies the following properties.

1. **Unit ℓ_2 norm:** f has ℓ_2 -norm 1.
2. **$d\kappa$ -degrees:** $\deg(f) \leq d\kappa$.
3. **δ -influences on average:** For all $i \in [N']$, we have $\mathbb{E}_{f \leftarrow \mathbf{F}_0}[\text{Inf}_i(f)] \leq \delta(d\kappa)$.

The same conditions hold for $\mathbf{F}_1$ as well. Assuming Conjecture 2.8 holds for $\mathcal{Y} = \mathbb{Z}_2$, there must exist $f_0 \in \text{supp}(\mathbf{F}_0)$ and $f_1 \in \text{supp}(\mathbf{F}_1)$ such that $f_0 \cdot f_1 \neq 0$.

Finally, we show that the cardinality of $\mathcal{F}$ is large. By Corollary 2.14, it holds that the size of $\mathcal{F}'$ satisfies

$$\frac{|\mathcal{F}'|}{|\mathcal{Y}^{\mathcal{X}'}|} \geq \frac{1}{O(d\kappa|\mathcal{X}'|^{2d\kappa}|\mathcal{Y}|^{2d\kappa})}.$$

Furthermore, note that the size of the sub-domain $\mathcal{Q}_{\mathcal{L}}$ fixed by $\mathcal{L}$ satisfies $|\mathcal{L}| \leq 100d/\varepsilon = \text{poly}(\kappa, d)$. Therefore, it holds that

$$\frac{|\mathcal{F}|}{|\mathcal{Y}^{\mathcal{X}}|} = \frac{|\mathcal{F}'|}{|\mathcal{Y}^{\mathcal{X}'}|} \geq |\mathcal{Y}|^{-|\mathcal{L}|} \cdot \frac{1}{O(d\kappa|\mathcal{X}'|^{2d\kappa}|\mathcal{Y}|^{2d\kappa})} = 2^{-\text{poly}(\kappa, d)} = 2^{-\text{poly}(\kappa)},$$

which means the uniform distribution over $\mathcal{F}$ is a $2^{-\text{poly}(\kappa)}$ -flat distribution. $\square$

Finally, we use Lemma 4.4 to prove Theorem 4.1.

Proof of Theorem 4.1. Suppose there exists a black-box construction (G, S) of non-interactive commitments from OWF $f: \{0, 1\}^\kappa \mapsto \{0, 1\}^\kappa$ (as in Definition 2.16). By Lemma 4.4, at least one of the following holds.

1. Let f be a random oracle. There exist Rec^* and S_h^{f, Rec^*} such that S_h^{f, Rec^*} breaks the one-wayness of f by asking $\text{poly}(\kappa)$ queries to f . However, then one can combine the algorithms S and Rec^* as a single algorithm that inverts a random oracle f with non-negligible probability by asking $\text{poly}(\kappa)$ queries to it. This contradicts the known optimality of Grover search [BBBV97].
2. There exist Sen^* and S_b^{f, Sen^*} such that S_b^{f, Sen^*} breaks the one-wayness of f , where f has a $2^{-R(\kappa)}$ -flat distribution with respect to Sen^* , where $R(\kappa) = \text{poly}(\kappa)$. In detail, for each query asked to Sen^* , it outputs polynomially many classical bits as the commitment and polynomially many qubits as for the two decommitment. By assumption, S_b^{f, Sen^*} asks only a polynomial number of queries to both f and Sen^* , but the answer by Sen^* is already fixed and so not worth asking them more than once. The answers that Sen^* provides could be interpreted as polynomial-size quantum advice about the oracle f that is passed down to the security reduction S . Putting things together, we conclude that S_b^{f, Sen^*} is an algorithm that inverts f with non-negligible probability by asking $S(\kappa) = \text{poly}(\kappa)$ number of queries and having $T(\kappa) = \text{poly}(\kappa)$ many bits of quantum advice about f . However, this contradicts the one-wayness of f as proven in Theorem 3.10.

$\square$

References

- AA09. Scott Aaronson and Andris Ambainis. The need for structure in quantum speedups. *arXiv preprint arXiv:0911.0996*, 2009. 4
- ACC⁺22. Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *Annual International Cryptology Conference (CRYPTO)*, 2022. <https://ia.cr/2022/218>. 4, 7, 8, 11, 13, 14, 15, 23
- ACGH20. Gorjan Alagic, Andrew M Childs, Alex B Grilo, and Shih-Han Hung. Non-interactive classical verification of quantum computation. In *Theory of Cryptography Conference*, pages 153–180. Springer, 2020. 3
- AQY22. Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. *International Cryptology Conference CRYPTO*, 2022. 5
- Bar21. James Bartusek. Secure quantum computation with classical communication. Cryptology ePrint Archive, Report 2021/964, 2021. <https://ia.cr/2021/964>. 3
- BB21. Nir Bitansky and Zvika Brakerski. Classical binding for quantum commitments. In *Theory of Cryptography Conference*, pages 273–298. Springer, 2021. 3, 4, 5
- BBBV97. Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM journal on Computing*, 26(5):1510–1523, 1997. 26
- BBF13. Paul Baecker, Christina Brzuska, and Marc Fischlin. Notions of black-box reductions, revisited. In Kazuo Sako and Palash Sarkar, editors, *Advances in Cryptology – ASIACRYPT 2013, Part I*, volume 8269 of *Lecture Notes in Computer Science*, pages 296–315, Bangalore, India, December 1–5, 2013. Springer, Heidelberg, Germany. 2
- BCKM21. James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 467–496, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany. 12
- BCQ22. Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. Cryptology ePrint Archive, Paper 2022/1181, 2022. <https://eprint.iacr.org/2022/1181>. 12
- 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, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 41–69, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany. 4, 12
- BKVV20. Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. Simpler proofs of quantumness. *arXiv preprint arXiv:2005.04826*, 2020. 3
- BM09. Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle puzzles are optimal - an $O(n^2)$ -query attack on any key exchange from a random oracle. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 374–390, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Heidelberg, Germany. 6
- BOV03. Boaz Barak, Shien Jin Ong, and Salil P. Vadhan. Derandomization in cryptography. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 299–315, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany. 12
- BS20. Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology –*

- CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 417–440, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany. [5](#)
- CCY20. Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. Classical verification of quantum computations with efficient verifier. In *Theory of Cryptography Conference*, pages 181–206. Springer, 2020. [3](#)
- CDMS04. Claude Crépeau, Paul Dumais, Dominic Mayers, and Louis Salvail. Computational collapse of quantum state with application to oblivious transfer. In Moni Naor, editor, *TCC 2004: 1st Theory of Cryptography Conference*, volume 2951 of *Lecture Notes in Computer Science*, pages 374–393, Cambridge, MA, USA, February 19–21, 2004. Springer, Heidelberg, Germany. [12](#)
- CGLQ20. Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. Tight quantum time-space tradeoffs for function inversion. In *61st Annual Symposium on Foundations of Computer Science*, pages 673–684, Durham, NC, USA, November 16–19, 2020. IEEE Computer Society Press. [5](#), [10](#), [11](#), [18](#), [21](#)
- CLS01. Claude Crépeau, Frédéric Légaré, and Louis Salvail. How to convert the flavor of a quantum bit commitment. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 60–77, Innsbruck, Austria, May 6–10, 2001. Springer, Heidelberg, Germany. [3](#)
- CX21. Shujiao Cao and Rui Xue. Being a permutation is also orthogonal to one-wayness in quantum world: Impossibilities of quantum one-way permutations from one-wayness primitives. *Theoretical Computer Science*, 855:16–42, 2021. [5](#), [11](#)
- DLS22. Frédéric Dupuis, Philippe Lamontagne, and Louis Salvail. Fiat-shamir for proofs lacks a proof even in the presence of shared entanglement. *Cryptology ePrint Archive*, Report 2022/435, 2022. <https://eprint.iacr.org/2022/435>. [11](#)
- DMS00. Paul Dumais, Dominic Mayers, and Louis Salvail. Perfectly concealing quantum bit commitment from any quantum one-way permutation. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 300–315, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany. [3](#)
- DS89. David L Donoho and Philip B Stark. Uncertainty principles and signal recovery. *SIAM Journal on Applied Mathematics*, 49(3):906–931, 1989. [9](#), [16](#)
- GL89. Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, WA, USA, May 15–17, 1989. ACM Press. [2](#), [5](#)
- GLSV21. Alex B Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in minicrypt. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 531–561. Springer, 2021. [12](#)
- GT00. Rosario Gennaro and Luca Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *41st Annual Symposium on Foundations of Computer Science*, pages 305–313, Redondo Beach, CA, USA, November 12–14, 2000. IEEE Computer Society Press. [5](#)
- HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. [2](#)
- HXY19. Minki Hhan, Keita Xagawa, and Takashi Yamakawa. Quantum random oracle model with auxiliary input. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019, Part I*, volume 11921 of *Lecture Notes in Computer Science*, pages 584–614, Kobe, Japan, December 8–12, 2019. Springer, Heidelberg, Germany. [5](#), [9](#), [18](#)

- HY20. Akinori Hosoyamada and Takashi Yamakawa. Finding collisions in a quantum world: Quantum black-box separation of collision-resistance and one-wayness. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part I*, volume 12491 of *Lecture Notes in Computer Science*, pages 3–32, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany. 4, 11
- IL89. Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th Annual Symposium on Foundations of Computer Science*, pages 230–235, Research Triangle Park, NC, USA, October 30 – November 1, 1989. IEEE Computer Society Press. 2
- IR89. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, WA, USA, May 15–17, 1989. ACM Press. 2, 5, 6
- JLS18. Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Annual International Cryptology Conference*, pages 126–152. Springer, 2018. 5
- KO09. Takeshi Koshihara and Takanori Odaira. Statistically-hiding quantum bit commitment from approximable-preimage-size quantum one-way function. In *Workshop on Quantum Computation, Communication, and Cryptography*, pages 33–46. Springer, 2009. 3
- KO11. Takeshi Koshihara and Takanori Odaira. Non-interactive statistically-hiding quantum bit commitment from any quantum one-way function. *arXiv preprint arXiv:1102.3441*, 2011. 3
- Kre21. William Kretschmer. Quantum pseudorandomness and classical complexity. *arXiv preprint arXiv:2103.09320*, 2021. 5
- LC97. Hoi-Kwong Lo and Hoi Fung Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78(17):3410, 1997. 3, 4
- LQWY14. Dongdai Lin, Yujuan Quan, Jian Weng, and Jun Yan. Quantum bit commitment with application in quantum zero-knowledge proof. Cryptology ePrint Archive, Report 2014/791, 2014. <https://eprint.iacr.org/2014/791>. 3
- LS19. Alex Lombardi and Luke Schaeffer. A note on key agreement and non-interactive commitments. Cryptology ePrint Archive, Report 2019/279, 2019. <https://eprint.iacr.org/2019/279>. 2
- Mah18. Urmila Mahadev. Classical verification of quantum computations. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 259–267. IEEE, 2018. 3
- May97. Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical review letters*, 78(17):3414, 1997. 3, 4
- MM11. Takahiro Matsuda and Kanta Matsuura. On black-box separations among injective one-way functions. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 597–614, Providence, RI, USA, March 28–30, 2011. Springer, Heidelberg, Germany. 2
- MP12. Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments - on the power of black-box vs. non-black-box use of primitives. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 701–718, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany. 2, 6
- MY22. Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments without one-way functions. *International Cryptology Conference CRYPTO*, 2022. 5
- NABT15. Aran Nayebi, Scott Aaronson, Aleksanders Belovs, and Luca Trevisan. Quantum lower bound for inverting a permutation with advice. *Quantum Information & Computation*, 15(11-12):901–913, 2015. 5

- Nao90. Moni Naor. Bit commitment using pseudo-randomness. In Gilles Brassard, editor, *Advances in Cryptology – CRYPTO’89*, volume 435 of *Lecture Notes in Computer Science*, pages 128–136, Santa Barbara, CA, USA, August 20–24, 1990. Springer, Heidelberg, Germany. 2
- NC10. Michael A Nielsen and Isaac L Chuang. Quantum computation and quantum information. *Cambridge University Press*, 2010. 12
- NW94. Noam Nisan and Avi Wigderson. Hardness vs randomness. *Journal of computer and System Sciences*, 49(2):149–167, 1994. 12
- RTV04. Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *TCC 2004: 1st Theory of Cryptography Conference*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20, Cambridge, MA, USA, February 19–21, 2004. Springer, Heidelberg, Germany. 2
- Unr12. Dominique Unruh. Quantum proofs of knowledge. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 135–152. Springer, 2012. 12
- Unr16a. Dominique Unruh. Collapse-binding quantum commitments without random oracles. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016, Part II*, volume 10032 of *Lecture Notes in Computer Science*, pages 166–195, Hanoi, Vietnam, December 4–8, 2016. Springer, Heidelberg, Germany. 4, 12
- Unr16b. Dominique Unruh. Computationally binding quantum commitments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 497–527, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany. 12
- WW21. Avi Wigderson and Yuval Wigderson. The uncertainty principle: variations on a theme. *Bulletin of the American Mathematical Society*, 58(2):225–261, 2021. 16
- Yan20. Jun Yan. General properties of quantum bit commitments. *Cryptology ePrint Archive*, Paper 2020/1488, 2020. <https://eprint.iacr.org/2020/1488>. 3
- Yao82. Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press. 2
- YWLQ15. Jun Yan, Jian Weng, Dongdai Lin, and Yujuan Quan. Quantum bit commitment with application in quantum zero-knowledge proof. In *International Symposium on Algorithms and Computation*, pages 555–565. Springer, 2015. 3
- Zha19. Mark Zhandry. How to record quantum queries, and applications to quantum indistinguishability. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 239–268, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany. 7, 12, 13, 14
- Zha21. Jiayu Zhang. Succinct blind quantum computation using a random oracle. In *STOC ’21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, pages 1370–1383, 2021. 3