

Unconditionally secure quantum commitments with preprocessing

Luowen Qian*

Abstract

We demonstrate how to build computationally secure commitment schemes with the aid of quantum auxiliary inputs without unproven complexity assumptions. Furthermore, the quantum auxiliary input can be prepared either (1) efficiently through a trusted setup similar to the classical common random string model, or (2) strictly between the two involved parties in uniform exponential time. Classically this remains impossible without first proving $P \neq NP$.

1 Introduction

In this work, we reexamine constructing quantum commitments, a central task in quantum cryptography [Yan22, BCQ23, BEM⁺23]. While commitments statistically (or information theoretically) secure against both parties are impossible [May97, LC97], recent works have demonstrated that computationally secure ones are possible under complexity assumptions [BCQ23, BEM⁺23, Bra23] that are potentially milder than $P \neq NP$ [Kre21, AQY22, MY22, KQST23, LMW23]. This suggests that achieving unconditionally computationally secure quantum cryptography might not be susceptible to the barriers that apply to classical cryptography.

Cryptography by default considers security against non-uniform adversaries where they, in addition to running in polynomial time, get some advice at the beginning from either an inefficient preprocessing phase or some residual information from another protocol execution. Auxiliary-input cryptography, studied since the 1990s [OW93], is a non-uniform version of cryptography where every party in the protocol also gets access to a copy of some public information that might not be efficiently preparable.

In this work we consider quantum auxiliary input, meaning that every party receives copies of the same quantum pure state as input. Therefore, in some sense, this quantum auxiliary information can be thought of as a “cryptographic magic state”, extremely similar to magic states that occur in quantum fault tolerance, where a piece of quantum state can augment the computational power of a less powerful circuit family.

Our main theorem constructs quantum auxiliary-input commitments without unproven complexity assumptions, improving from the prior work of Chailloux, Kerenidis, and Rosgen [CKR16] who established the same result assuming the unproven complexity separation of $QIP \not\subseteq QMA$.¹ While

*Boston University. Email: luowenq@bu.edu

¹There are two different variants of auxiliary-input security considered in the literature: in this work, we focus on the strong variant where the adversary’s success probability is small for all but finitely many auxiliary inputs; however, there is also a weaker variant where we only require the adversary’s success probability to be small for infinitely many auxiliary inputs [OW93]. Note that classical auxiliary input *weakly* secure quantum commitments can be built assuming $QCZK \not\subseteq BQP$ [BCQ23], classical auxiliary-input (strong) quantum commitment can be built assuming $QCZK \not\subseteq QMA$, and finally standard quantum commitments without auxiliary input can be built assuming various other unproven assumptions.

we share some techniques, our construction uses the sparse pseudorandom ensemble constructed with a probabilistic method instead of the hardness of a QIP-complete problem.

Theorem 1. *There exists a computationally-hiding statistically-binding non-interactive quantum commitment scheme with quantum auxiliary input. Furthermore, the quantum auxiliary input has an exponential-size classical description that can be sampled uniformly in exponential time.*

We note that a classical analogue of Theorem 1 would still imply $P \neq NP$ [IL89]. In fact, as we show in Theorem 4, this is true even if we consider the model where all parties get access to a single (possibly inefficient) sampling oracle, and each sample is private to the requested party. This model is even stronger than having access to classical auxiliary input, since the oracle could just output that fixed string with probability 1. (Considering every party having access to the same randomized advice can be simply replaced with a fixed classical advice by an averaging argument.) Thus in some sense, our result is even stronger than Raz’s result of $QIP/qpoly = IP/rpoly = ALL$ [Raz05] as the power of the quantum advice does not come from the advice being inherently randomized and this randomness being private to each party.

Removing trust. One might be skeptical of the security since naïvely it appears that the parties need to assume they can trust the quantum auxiliary input given to them. One way to fix this problem is to ask the skeptical party to simply inspect the classical description of the magic state and verify that the commitment built with it is secure. This takes at most a doubly exponential time since we can cast it as a QMA problem (Proposition 9).

In Section 4, we show how to achieve computationally secure commitments without any trusted auxiliary inputs through a less inefficient preprocessing. Specifically, in the scheme we only need to ask both the committer and the receiver to perform a uniform exponential-time preprocessing phase.

Removing inefficiency. Complementary to that, we also show how to do commitments in a *completely efficient* setting with a weakly trusted setup in Section 3. Specifically, we adapt our construction into a trusted setup model where a trusted third party efficiently generates a few copies of the quantum auxiliary information for every party in the protocol to use before the protocol begins. However, unlike Theorem 1, here the pure state distributed to every party is not deterministic. Furthermore, the scheme could in fact be statistically secure against all parties if the number of copies distributed is restricted.

A stronger (in terms of trust) setup model is the secret parameter model [Ps05], where we need to trust the setup to sample two correlated strings for each party. While statistically secure classical non-interactive commitments are possible in the secret parameter model, our model appears to be meaningfully weaker than the secret parameter model. In particular, we can invoke our classical impossibility (Theorem 4) to argue that statistically secure commitments remain impossible in a classical sampling analogue of our model, since an unbounded adversary could always solve any NP problem.

Polynomially bounded adversaries are physical, probably. A possible concern regarding the claimed “unconditional” nature of this result is that the security relies on the “assumption” that the adversary is polynomially bounded during the execution of the protocol. We address that concern by noting that it is possible to reduce this assumption to a physical assumption. On one hand, physical assumptions and more generally modeling assumptions (the mapping between real world

and mathematics) are unavoidable in any form of provably secure cryptography. On the other hand, this suggests a new win-win philosophy, since either we can have secure cryptography or we are able to discover exciting new physics.

More specifically, there is most likely a fundamental physical limit to the density of quantum information before collapsing into a black hole given by the Bekenstein bound (see [Llo00] for an exposition of this). Since any malicious computation must be done in polynomial time, and thus space by the no superluminal signaling principle, it follows from extended Church–Turing thesis (for quantum information tasks) that any polynomial-time physical computation can be described by a polynomial-size quantum circuit. Therefore, we can force adversary to be polynomially bounded by simply limiting the protocol execution time. Even if the adversary could somehow leverage black holes to perform useful computations, we could still monitor the energy density nearby to make sure that this does not happen. We leave further materializing this idea to future work.

Additional applications. We note that once we have commitments then stronger applications are also within reach, as we further detail in Appendix B (see also [BCQ23]). One notable application is to use it to perform secure multiparty computations.

For a more concrete example, consider the classical Yao’s Millionaires’ Problem [Yao82], except that now it’s a Trillionaires’ Problem! This means that two trillionaires want to figure out who is richer without revealing anything else. Also since they are trillionaires, their entire families and businesses are somehow also on the line, so if cheating is detected then chaos would ensue. Furthermore, they also have the world’s bestest cryptanalyst to break any computational hardness assumption should it be necessary and possible, and they are willing to use exponential-time preprocessing as a small sacrifice.

Secure multiparty computations with preprocessing is *the perfect* solution for this problem! They first spend exponential time to set up and during the protocol execution, they ensure that the other party finish in time and no cheating occurs. So if nothing bad occurs, then both of them can be satisfied knowing that their secrets are safe. Additionally, using certified everlasting transfer [BK23], they can achieve everlasting security by having a third party referee, who is trusted to be uninterested in spending exponential resources recovering the input, to certifiably delete the remaining information.

Open problems. One undesirable feature of our scheme is that there is no good way to get additional copies of the magic state. In the trusted setup model, it is possible to redo the setup every once in a while to preserve statistical security and efficiency, but otherwise generating new copies still takes exponential time, although it is possible to build dedicated hardware to pipeline this process. A natural question is whether we can construct quantum cryptography with quantum auxiliary information that is efficiently clonable or reusable. While there are families of quantum states that is efficiently clonable but cannot be prepared in uniform polynomial time relative to a quantum oracle [NZ23], even having a standard model candidate is open since the quantum oracle there is a cloning oracle. The biggest issue with reusing our schemes is that for each invocation of the commitment scheme, on average $\Theta(\lambda)$ copies of the magic state is transferred to the other party, and there does not appear to be a way to certifiably retrieve these states. Of course, the ultimate goal would be to construct these without inefficient preprocessing or having trusted setup at all.

Concurrent work. Near the completion of this work, I became aware of a parallel work done by Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa [MNY23]. In particular, both works have the same constructions of quantum auxiliary-input commitments (Theorem 1) and statistically secure commitments with a weakly trusted (stateful) setup (Corollary 6 and Proposition 7). They have an additional statistical commitment construction with stateless setup, however, the downside is that this can only be secure given that at most a bounded number of copies is generated. Complementary to the classical impossibility of Theorem 4, they also give an impossibility of classical auxiliary-input (quantum) commitments in a weak setup model. They in addition have the application of quantum auxiliary-input zero-knowledge proofs for NP (with negligible simulation security). They also point out an observation by Fermi Ma that the cube root security loss of Proposition 2 can be improved to square root if we instead augment the [GK92] argument with a matrix Chernoff bound like was done in [LMW23].

Notably, their work point out that the quantum auxiliary-input commitment cannot be immediately used to construct simulation-secure commitments through the [BCKM21] compiler due to the use of Watrous rewinding in the simulator there. After the discussions with [MNY23], I present the resolution of this in Appendix B by showing how to adapt [BCKM21] to recover ε -simulation security, which still suffices to recover almost all applications of a standard quantum commitment.

2 Quantum auxiliary-input commitment

We formally define non-interactive commitments and handle its subtleties for later applications in Appendix A. To prove our main theorem, we first need the following result on the quantum non-uniform security of pseudorandomness.

Proposition 2 ([CGLQ20a, Liu23]). *For a random function $H : [N] \rightarrow [M]$, the best quantum non-uniform circuit of size S (potentially depending on H) can distinguish its output from a random element from $[M]$ with advantage at most $12 \cdot \sqrt[3]{\frac{S}{N}}$ (averaged over the choice of H).*

While these works consider the more general case where quantum algorithms that could additionally make queries to H , the same proofs also give a polynomial upper bound for standard query-less algorithms. For completeness, we show this precise bound in Appendix C.

In particular, this implies the existence of a good function for which this is true against size $S - 1$, since we can always use an extra bit to make the bias always have the same sign for all H . In other words, for any $(S - 1)$ -size algorithm A' , there is an S -size algorithm A such that

$$\left| \mathbb{E}_H \left[\mathbb{E}_x [A'(H(x))] - \mathbb{E}_y [A'(y)] \right] \right| = \mathbb{E}_H \left[\left| \mathbb{E}_x [A(H(x))] - \mathbb{E}_y [A(y)] \right| \right],$$

where A simply runs A' and XOR its output with the extra advice bit. With some small inverse exponential security loss, we can also use Markov's inequality to argue that even sampling a function uniformly at random would satisfy this with overwhelming probability.

Instantiating the above with $S = 2^\lambda + 1$, $N = 2^{5\lambda}$ and $M = 2^{6\lambda}$, we get the following corollary, which generalizes the classical (inefficient) sparse pseudorandom ensemble construction of Goldreich and Krawczyk [GK92] to the post-quantum setting².

²The naïve attempt of quantizing [GK92] proof cannot seem to handle quantum advice. For readers familiar with

Corollary 3 (Exponentially secure sparse pseudorandom ensemble). *There exists a pseudorandom ensemble $H : \{0, 1\}^{5\lambda} \rightarrow \{0, 1\}^{6\lambda}$ against all 2^λ -size quantum circuits A with quantum auxiliary information, whose security $|\mathbb{E}_x[A(H(x))] - \mathbb{E}_y[A(y)]| \leq 2^{-\lambda}$ for all $\lambda \geq 11$.*

Proof of Theorem 1. Let $\{H(x)\}_{x \in \{0,1\}^{5\lambda}}$ be the exponentially secure sparse pseudorandom ensemble as above. We first specify the quantum auxiliary input

$$|M_\lambda\rangle := 2^{-5\lambda/2} \sum_{x \in \{0,1\}^{5\lambda}} |x\rangle \otimes |H(x)\rangle,$$

a pure state of 11λ qubits that can be prepared using a circuit of size $\tilde{O}(2^{5\lambda})$.

We now specify the protocol. The sender, to commit to 0, simply sends the second half of $|M_\lambda\rangle$; and to commit to 1, sends the second half of a maximally entangled state

$$|\Psi_\lambda\rangle := 2^{-3\lambda} \sum_{y \in \{0,1\}^{6\lambda}} |y\rangle \otimes |y\rangle;$$

to later decommit, the rest of the state is sent. The receiver, upon receiving the entire state and the bit b can efficiently test by performing a SWAP test between the received state and the correct state of either $|M_\lambda\rangle$ or $|\Psi_\lambda\rangle$. Finally, a λ -fold parallel repetition is applied to this construction, meaning that the committer commits to the same bit λ times in parallel, and the receiver checks that all commitment decommits to the same bit.

For binding, we note that our construction is identical to that of [CKR16, Section 4] except for the choice of two pure states used for the commitments. In our case, since the two reduced density matrices on the commitment register are classical, so the fidelity between them is

$$\left(\sum_y \sqrt{\frac{\Pr[H(x) = y]}{2^{6\lambda}}} \right)^2 \leq 2^{-\lambda} \tag{1}$$

by applying Cauchy–Schwarz on y 's in the image of H . The rest follows the same proof as [CKR16, Proposition 4.4], thus we get that our commitment after taking λ -fold parallel repetition is statistical sum binding.

For hiding, the two reduced density matrices given to the hiding adversary exactly corresponds to the security game of the pseudorandom ensemble. Finally taking a parallel repetition also preserves hiding by a standard hybrid argument. \square

We remark that this construction can be straightforwardly “derandomized” using a (post-quantum) pseudorandom generator, and thus eliminating the inefficiency at the cost of assuming the security of the pseudorandom generator.

In fact, if we are only aiming for a commitment scheme with quantum auxiliary input and do not require the classical description to be computable in exponential time, any non-trivial computationally indistinguishable pair of quantum states suffices. The construction is essentially the same except that the receiver does a SWAP test with the corresponding pure state for decommitments to either bits.

that work, a natural strategy is to use union bound over all quantum algorithms/advice using an ε -net, however, this fails since the size of the ε -net is doubly exponential, and thus eliminating the single exponential concentration we get from Hoeffding's bound.

Classical impossibility. For convenience, we focus on non-interactive statistically-binding commitments for the classical case. This is a fair comparison since such a commitment with classical auxiliary information does exist assuming existence of one-way functions, by applying averaging argument to the first message of Naor commitment [Nao91]. We leave improving this impossibility to future work. The main insight is that having sample access to the oracle suffices to break the security.

Theorem 4. *If there exists a computationally-hiding statistically-binding non-interactive classical commitment scheme where all parties share access to a sampling oracle, then $\text{NP} \not\subseteq \text{P/poly}$.*

Proof. For simplicity, we consider that committer and receiver each gets a private sample s, r respectively from the sampling oracle; to commit to bit b , the committer computes a deterministic function $\text{Commit}(s, b) = \tau$ and sends τ to the receiver; to decommit, the committer sends s, b and the receiver computes a deterministic function $\text{Reveal}(s, b, \tau, r) \in \{\top, \perp\}$ indicating accept or reject the revealed bit b . This is without loss of generality since we can always have the oracle give multiple samples per query and pad the oracle with uniform random bits so that the only source of private randomness is from the oracle.

Assume $\text{NP} \subseteq \text{P/poly}$, we show how to efficiently break the hiding. Let each sample be ℓ bits long. The malicious receiver, on input τ and $(\ell + 1)$ independent samples $\vec{r} := r_0, \dots, r_\ell$ from the sampling oracle, predicts the commitment is to 0 if there exists s_0 such that $\text{Reveal}(s_0, 0, \tau, r_i) = \top$ for all i , and predicts to 1 otherwise. This is also an NP language since Reveal is efficient. Then in the case when the committer commits to 0, by completeness with overwhelming probability over τ , $\Pr_r[\text{Reveal}(s_0, 0, \tau, r)]$ is negligibly close to 1, so this receiver always predicts 0 except with negligible probability. On the other hand, when the committer commits to 1, by statistical binding, with overwhelming probability over τ , for every s_0 we have that $\Pr_r[\text{Reveal}(s_0, 0, \tau, r)]$ is negligible, and by independence of samples, $\Pr_{\vec{r}}[\forall i : \text{Reveal}(s_0, 0, \tau, r_i)] \leq 2^{-\ell-1}$ for all sufficiently large ℓ , thus by union bound, $\Pr_{\vec{r}}[\exists s_0 \forall i : \text{Reveal}(s_0, 0, \tau, r_i)] \leq \frac{1}{2}$. This gives a hiding adversary with advantage at least $\frac{1}{2} - \text{negl}$. \square

3 Commitment in the unclonable common random state model

The common random string (CRS) model is a commonly considered relaxation of the standard trustless model where the only trust in the setup is that a classical string is uniformly sampled and then published. This model was first proposed in the context of non-interactive zero knowledge (NIZK) [BFM88] since interesting NIZK is impossible in the trustless model.

In this work, we introduce a quantum analogue of the CRS model that we call the unclonable common random state ($|\text{UCRS}\rangle$) model. In the $|\text{UCRS}\rangle$ model, the only trust in the setup is that a random pure state is drawn from a state distribution, and then many copies of that pure state are made available to all parties. Furthermore, there should be a way to efficiently generate any polynomial (but a priori unbounded) number of the common state. We emphasize that unclonability only indicates a lack of the cloning functionality of the common state but the scheme's security does not rely on the common state being unclonable — in fact in this model, any malicious party is allowed to inefficiently process the classical description of the random state before the protocol begins.

Going back to the commitment construction of Theorem 1, we note that the magic state is simply a uniform superposition query on a function H , and even using a random function H is

secure except with inverse exponential probability. Therefore, that construction is indeed also secure in the $|\text{UCRS}\rangle$ model, so the only missing piece of the puzzle is to show efficient sampling of the magic state. This is tricky since the classical description requires an exponential size and there are a doubly exponential number of possible states to sample from.

Nevertheless, we show that there is an efficient way to statefully sample these states. To do this, we invoke Zhandry’s compressed oracle technique [Zha19].

Lemma 5 ([Zha19]). *There exists a stateful simulation oracle CStO that perfectly simulate any number of (quantum) queries to a random function $H : \{0, 1\}^n \rightarrow \{0, 1\}^m$ for any n, m . Furthermore, the t -th query can be processed in time polynomial in mt , and the state (also called the database) after t queries consists of $(m + n + 1) \cdot t$ qubits.*

Since the simulation is perfect, the commitment constructed in Theorem 1, after replacing H with a truly random function, still works. The simulation is also efficient since both m and t are polynomial in λ . Thus we arrive at the following corollary.

Corollary 6. *There exists a computationally-hiding statistically-binding non-interactive quantum commitment scheme in the $|\text{UCRS}\rangle$ model.*

Interestingly, the hiding of the commitment in fact becomes statistical if at most a polynomial number of states are given out to the adversary, and the proof of this is deferred to Appendix C. Therefore, we get a completely statistically secure commitment if we in addition trust the setup to not generate too many copies of the magic state.

Proposition 7. *Assuming the receiver has at most P copies of the magic state on H , then the commitment scheme of Corollary 6 is $\left(8\sqrt{2} \cdot \sqrt{\frac{P}{N}}\right)$ -statistically hiding.*

Another interesting consequence of this is that for our scheme, we cannot trust either party to distribute the magic state. If the receiver chooses the magic state for the committer, then computational hiding can be trivially broken by picking a bad magic state like the all zero state. If the committer chooses the magic state for the receiver, then Proposition 7 shows that this scheme is in fact statistically hiding and thus not statistical binding. In fact, it is not even computationally binding as we show below.

Proposition 8. *Take the commitment scheme from Theorem 1 but instead have the committer choose H as a random function. Then this commitment scheme has computational sum binding error of at least $1 - O(t/\sqrt{N})$ even after taking t -fold parallel repetition using the same random function.*

Proof sketch. We sketch how to efficiently break sum binding, even if the scheme is repeated t times in parallel. Consider a binding adversary that commits to 0 honestly using the compressed oracle. Certainly by sending the decommitment registers honestly, the receiver would accept 0 with probability 1. We now show how to decommit to 1 with probability $1 - O(t/\sqrt{N})$. First, we measure the x for every magic state (including the ones held by the decommitter are also measured, which is okay since the receiver does not touch those registers when checking decommitment to 1 so they are essentially traced out) and abort if not, which happens with probability at most $O(t^2/N)$ by collision probability. Otherwise, every magic state holds a distinct x . For each fold, the decommitment register contains some x_i : we apply StdDecomp_{x_i} and search in the database where the entry x_i occurs and send the corresponding image register as decommitment to 1, which

is maximally entangled with the corresponding commitment register. To make this into a malicious committer that does not measure the x 's held by the receiver, we note that this measurement is only used to conditionally abort the committer, so the success probability of the same adversary except that it never aborts is still $1 - O(t/\sqrt{N})$ by gentle measurement. \square

4 Eliminating trust with preprocessing

We first give a PromiseQMA upper bound on the complexity of verifying the computational insecurity of H up to a constant multiplicative loss. Therefore, we can check if a function H is secure in doubly exponential time, or even exponential time if $\text{BQP} = \text{QMA}$.

Proposition 9. *The language L , consisting of all functions H such that the commitment constructed in Theorem 1 using H is insecure, is in PromiseQMA.*

Proof. We prove this by constructing a PromiseQMA verifier. The verifier, on input H (of length $N \log M$) and a witness $|C\rangle$ (a distinguisher quantum circuit of length S which is polynomial in $|H|$), samples a random bit b and runs the universal quantum circuit on either $(|C\rangle, H(x))$ for a random $x \in [N]$ or $(|C\rangle, y)$ for a random $y \in [M]$, and accepts if the universal quantum circuit predicts b correctly. We set completeness to be $\frac{1}{2} + 2^{-\lambda}$ and soundness to be $\frac{1}{2} + 2^{-\lambda-1}$, so the gap is $2^{-\lambda-1}$ which is inverse polynomial in $|H|$.

If the output of H is pseudorandom, then we have that $H \notin L$ as desired. On the other hand, if there exists an S -sized witness for H that distinguishes with advantage higher than $2 \cdot 2^{-\lambda}$, then $H \in L$. \square

We now give an alternative *unconditional* approach to eliminate trust on the magic state in uniform exponential time.

Theorem 10. *There exists a computationally-hiding statistically-binding non-interactive quantum commitment scheme with an exponential time preprocessing phase.*

Proof. We again adapt the commitment from Theorem 1 by adding a preprocessing phase to generate the magic states for both parties. More specifically, the sender samples a random function H on their own, then send the classical description of H to the receiver. Afterwards, they generate multiple copies of the magic state for that function on their own before the protocol begins. (It is important that they generate the quantum magic state on their own to not run into the impossibility of Proposition 8 above.)

Computational hiding is preserved since in this case H is honestly generated, and the proof of Theorem 1 actually suffices to show that the commitment is statistically binding for any H : specifically, the fidelity computed in (1) is bounded by $2^{-\lambda}$ for any H . \square

Acknowledgements

The author thanks Ran Canetti, William Kretschmer, Qipeng Liu, as well as Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa [MNY23] for their invaluable feedback on an earlier draft of this work. Special gratitude is owed to Qipeng Liu for proposing to augment Theorem 1 with compressed oracles (Corollary 6), and Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa [MNY23] for their helpful discussions contributing to the development of Appendix B. The author is supported by DARPA under Agreement No. HR00112020023.

References

- [AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. “Cryptography from Pseudorandom Quantum States”. In: *Advances in Cryptology – CRYPTO 2022*. 2022, pp. 208–236. DOI: 10.1007/978-3-031-15802-5_8 (cit. on pp. 1, 11, 15).
- [BCKM21] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. “One-Way Functions Imply Secure Computation in a Quantum World”. In: *Advances in Cryptology – CRYPTO 2021*. 2021, pp. 467–496. DOI: 10.1007/978-3-030-84242-0_17 (cit. on pp. 4, 13, 15).
- [BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. “On the Computational Hardness Needed for Quantum Cryptography”. In: *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Vol. 251. 2023, 24:1–24:21. DOI: 10.4230/LIPIcs.ITCS.2023.24 (cit. on pp. 1, 3, 11).
- [BEM⁺23] John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. *Unitary Complexity and the Uhlmann Transformation Problem*. 2023. arXiv: 2306.13073 (cit. on p. 1).
- [BFM88] Manuel Blum, Paul Feldman, and Silvio Micali. “Non-Interactive Zero-Knowledge and Its Applications”. In: *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*. 1988, pp. 103–112. DOI: 10.1145/62212.62222 (cit. on p. 6).
- [BK23] James Bartusek and Dakshita Khurana. “Cryptography with Certified Deletion”. In: *Advances in Cryptology – CRYPTO 2023*. 2023, pp. 192–223. DOI: 10.1007/978-3-031-38554-4_7 (cit. on p. 3).
- [Bra23] Zvika Brakerski. “Black-Hole Radiation Decoding Is Quantum Cryptography”. In: *Advances in Cryptology – CRYPTO 2023*. 2023, pp. 37–65. DOI: 10.1007/978-3-031-38554-4_2 (cit. on p. 1).
- [CGLQ20a] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. “Tight Quantum Time-Space Tradeoffs for Function Inversion”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 673–684. DOI: 10.1109/FOCS46700.2020.00068 (cit. on pp. 4, 16).
- [CGLQ20b] Kai-Min Chung, Siyao Guo, Qipeng Liu, and Luowen Qian. *Tight Quantum Time-Space Tradeoffs for Function Inversion*. 2020. arXiv: 2006.05650v2 (cit. on p. 16).
- [CKR16] André Chailloux, Iordanis Kerenidis, and Bill Rosgen. “Quantum commitments from complexity assumptions”. In: *computational complexity* 25.1 (2016), pp. 103–151. DOI: 10.1007/s00037-015-0116-5 (cit. on pp. 1, 5).
- [FUYZ22] Junbin Fang, Dominique Unruh, Jun Yan, and Dehua Zhou. “How to Base Security on the Perfect/Statistical Binding Property of Quantum Bit Commitment?” In: *33rd International Symposium on Algorithms and Computation (ISAAC 2022)*. Vol. 248. 2022, 26:1–26:12. DOI: 10.4230/LIPIcs.ISAAC.2022.26 (cit. on pp. 12, 13).
- [GK92] Oded Goldreich and Hugo Krawczyk. “Sparse pseudorandom distributions”. In: *Random Structures & Algorithms* 3.2 (1992), pp. 163–174. DOI: 10.1002/rsa.3240030206 (cit. on p. 4).

- [IL89] Russell Impagliazzo and Michael Luby. “One-way functions are essential for complexity based cryptography”. In: *30th Annual Symposium on Foundations of Computer Science*. 1989, pp. 230–235. DOI: 10.1109/SFCS.1989.63483 (cit. on p. 2).
- [KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. “Quantum Cryptography in Algorithmica”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 1589–1602. DOI: 10.1145/3564246.3585225 (cit. on p. 1).
- [Kre21] William Kretschmer. “Quantum Pseudorandomness and Classical Complexity”. In: *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*. Vol. 197. 2021, 2:1–2:20. DOI: 10.4230/LIPIcs.TQC.2021.2 (cit. on p. 1).
- [LC97] Hoi-Kwong Lo and H. F. Chau. “Is Quantum Bit Commitment Really Possible?” In: *Phys. Rev. Lett.* 78 (17 1997), pp. 3410–3413. DOI: 10.1103/PhysRevLett.78.3410 (cit. on p. 1).
- [Liu23] Qipeng Liu. “Non-uniformity and Quantum Advice in the Quantum Random Oracle Model”. In: *Advances in Cryptology – EUROCRYPT 2023*. 2023, pp. 117–143. DOI: 10.1007/978-3-031-30545-0_5 (cit. on pp. 4, 16).
- [Llo00] Seth Lloyd. “Ultimate physical limits to computation”. In: *Nature* 406.6799 (2000), pp. 1047–1054. DOI: 10.1038/35023282 (cit. on p. 3).
- [LMW23] Alex Lombardi, Fermi Ma, and John Wright. *A one-query lower bound for unitary synthesis and breaking quantum cryptography*. 2023. arXiv: 2310.08870 (cit. on pp. 1, 4).
- [May97] Dominic Mayers. “Unconditionally Secure Quantum Bit Commitment is Impossible”. In: *Phys. Rev. Lett.* 78 (17 1997), pp. 3414–3417. DOI: 10.1103/PhysRevLett.78.3414 (cit. on p. 1).
- [MNY23] Tomoyuki Morimae, Barak Nehoran, and Takashi Yamakawa. *Unconditionally Secure Commitments with Quantum Auxiliary Inputs (manuscript)*. 2023 (cit. on pp. 4, 8).
- [MY22] Tomoyuki Morimae and Takashi Yamakawa. “Quantum Commitments and Signatures Without One-Way Functions”. In: *Advances in Cryptology – CRYPTO 2022*. 2022, pp. 269–295. DOI: 10.1007/978-3-031-15802-5_10 (cit. on pp. 1, 13).
- [Nao91] Moni Naor. “Bit commitment using pseudorandomness”. In: *Journal of Cryptology* 4.2 (1991), pp. 151–158. DOI: 10.1007/BF00196774 (cit. on p. 6).
- [NZ23] Barak Nehoran and Mark Zhandry. *A Computational Separation Between Quantum No-cloning and No-teleportation*. 2023. arXiv: 2302.01858 (cit. on p. 3).
- [OW93] Rafail Ostrovsky and Avi Wigderson. “One-way functions are essential for non-trivial zero-knowledge”. In: *The 2nd Israel Symposium on Theory and Computing Systems*. 1993, pp. 3–17. DOI: 10.1109/ISTCS.1993.253489 (cit. on p. 1).
- [Ps05] Rafael Pass and abhi shelat. “Unconditional Characterizations of Non-interactive Zero-Knowledge”. In: *Advances in Cryptology – CRYPTO 2005*. 2005, pp. 118–134. DOI: 10.1007/11535218_8 (cit. on p. 2).

- [Raz05] Ran Raz. “Quantum information and the PCP theorem”. In: *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*. 2005, pp. 459–468. DOI: 10.1109/SFCS.2005.62 (cit. on p. 2).
- [Wat06] John Watrous. “Zero-Knowledge against Quantum Attacks”. In: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*. 2006, pp. 296–305. DOI: 10.1145/1132516.1132560 (cit. on p. 15).
- [Yan22] Jun Yan. “General Properties of Quantum Bit Commitments (Extended Abstract)”. In: *Advances in Cryptology – ASIACRYPT 2022*. 2022, pp. 628–657. DOI: 10.1007/978-3-031-22972-5_22 (cit. on pp. 1, 11).
- [Yao82] Andrew Chi-Chih Yao. “Protocols for secure computations”. In: *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. 1982, pp. 160–164. DOI: 10.1109/SFCS.1982.88 (cit. on p. 3).
- [Zha19] Mark Zhandry. “How to Record Quantum Queries, and Applications to Quantum Indifferentiability”. In: *Advances in Cryptology – CRYPTO 2019*. 2019, pp. 239–268. DOI: 10.1007/978-3-030-26951-7_9 (cit. on p. 7).

A Non-interactive quantum commitments

In this section, we use $x \in S$ as the auxiliary information for a certain set S and implicitly its length $|x|$ as the security parameter. Thus if x is a classical unary string, then it is a standard uniform commitment scheme. In order to prevent degeneracy, we require that S must contain arbitrarily long bitstrings or quantum states: for any integer n , there exists $x \in S$ such that $|x| \geq n$. We refer the readers to [BCQ23] for quantum information and cryptography background.

Definition 11. *A non-interactive commitment scheme is a pair of efficient quantum algorithms Commit and Reveal where Commit(x, b) produces a bipartite state ρ over the commitment register C and the decommitment register D , and Reveal(x, ρ) outputs either the committed bit b' or a rejection symbol \perp . Furthermore, $\Pr[\text{Reveal}(x, \text{Commit}(x, b)) \neq b]$ is negligible for any $b = 0, 1$ and $x \in S$.*

We say it is computationally (or statistically) hiding if the C registers of Commit($x, 0$) and Commit($x, 1$) are computationally (or statistically, respectively) indistinguishable.

We say it is statistically (or computationally) sum binding if for any state ρ (over C, D and a private register M) and any possibly inefficient (or efficient, respectively) unitary U not acting on register C , $p_0 + p_1 - 1$ is negligible, where p_b is the probability that the receiver accepts bit b . In particular, $p_0 := \Pr[\text{Reveal}(x, \rho_{CD}) = 0]$ and $p_1 := \Pr[\text{Reveal}(x, (U\rho)_{CD}) = 1]$.

Finally, we say it is in canonical form [Yan22], if Reveal takes the following form:

1. Perform a rank-1 projection on CD and output 0 if it succeeds.
2. Perform another orthogonal rank-1 projection on CD and output 1 if it succeeds.
3. Output \perp .

Since statistical sum binding is traditionally unwieldy to use in applications, the work of Ananth, Qian, and Yuen defines an extractor-based binding definition [AQY22, Definition 6] (hereby called extractable binding in order to distinguish), which on a high level states that for any malicious committer, there is an extractor that can extract the committed bit from the receiver’s view after the commitment phase in an imperceptible way, as long as the receiver does not touch those registers

until the reveal phase. More specifically, it is defined as an indistinguishability between a real experiment and an ideal experiment. The real experiment simply performs the commitment normally and outputs the revealed bit along with the committer's residual state. The ideal experiment is defined as follows: after a normal execution of the commit phase, the extractor is ran on the receiver's view to obtain a trit $b \in \{0, 1, \perp\}$ and a post-measurement view; after the reveal phase, the output is set to the revealed bit b' along with the committer's residual view if $b' \in \{b, \perp\}$, otherwise the output is set to a special symbol \perp indicating extraction failure.

For canonical form commitments, it is known that many variants of binding are equivalent, including statistical sum binding and extractable binding [FUYZ22]. It is unclear how to make commitments constructed in this work into canonical form due to the presence of quantum auxiliary information, so for completeness we show how to extend the equivalence to general non-interactive schemes.

We call a commitment scheme having a projective **Reveal** if **Reveal** is a projective measurement on the receiver's view. For example, any post-quantum commitment with a deterministic **Reveal** algorithm or any canonical-form commitment has a projective **Reveal** but our schemes do not due to the swap test. Without loss of generality we can always generically make a commitment have a projective **Reveal** via Stinespring dilation. In particular, we need to purify the **Reveal** algorithm so that it is a unitary on registers **CD** and potentially some auxiliary register **A** (which holds x and potentially some zeroes), followed by a complete measurement on a qutrit in the auxiliary register to obtain the output. Furthermore, we ask the receiver to prepare the auxiliary register **A** immediately after the commit phase. This change is also imperceptible from the committer's perspective.

We remark that due to a technicality, the equivalence cannot hold without giving the extractor access to **A**. Consider the following unnatural counterexample, a non-interactive (classical) commitment scheme without a projective **Reveal** that satisfies sum binding but not extractable binding. To commit to bit b , the committer simply send a mode bit 1 followed by b ; and the decommitment message is empty. The **Reveal** algorithm checks that if the mode bit is 1 then output b , and if the mode bit is 0 then output a random bit. Intuitively this scheme is binding since no matter which mode the malicious committer uses, he cannot later change the bit in any way since the decommitment message is empty; and indeed it is straightforward to see that it satisfies sum binding. However, a malicious committer can cause the extractor to fail since without access to the random bit register later used to sample the output of **Reveal**, it is impossible for the extractor to predict the bit that would be revealed later.

Theorem 12. *Statistical sum binding is equivalent to extractable binding for any commitment scheme with a projective **Reveal**, thus we can build an extractable binding commitment scheme from any statistical sum binding commitment scheme (and vice versa).*

Proof. Extractable binding implies sum binding is straightforward: $p_0 + p_1 \leq 1$ in the ideal experiment where the bit is extracted and guaranteed to be correct, thus the indistinguishability between the real and the ideal experiments implies sum binding.

For the other direction, fix any sender where the overall state before executing **Reveal** is ρ . Without loss of generality we assume ρ is pure by taking the purification into committer's private register. Let $|\rho_0\rangle$ be the state post-selected on the event that $\text{Reveal}(|\rho\rangle) = 0$ and $|\rho_1\rangle$ be the state post-selected on the event that $\text{Reveal}(|\rho\rangle) = 1$, or 0 if such post-selection is not possible. $|\rho_0\rangle$ and $|\rho_1\rangle$ are orthogonal by the fact that **Reveal** is projective. Consider a canonical-form variant of **Reveal** where the rank-1 projections are given by $|\rho_0\rangle\langle\rho_0|$ and $|\rho_1\rangle\langle\rho_1|$. We claim that this variant is

still sum binding since any sum binding adversary succeeding for this variant would also succeed in breaking sum binding against the original `Reveal`. Then the result of Fang, Unruh, Yan, and Zhou [FUYZ22] (also in [MY22, Appendix B]) implies that there is an extractor for this modified scheme, in particular, the extractor performs the optimal distinguishing measurement between these two states on registers `AC`. This extractor also works for the original scheme since the extractor almost perfectly project the state onto either $|\rho_0\rangle$ or $|\rho_1\rangle$ by statistical sum binding and gentle measurement, and the two `Reveal` algorithms behave identically in the subspace spanned by these states. \square

B Simulation security

Simulation security captures the security of a primitive using the real-ideal world paradigm more precisely than the game-based security definitions (which are usually used for hardness assumptions) and is the default security notion in the context of zero knowledge and secure multiparty computations. In this appendix, we discuss how to further augment our commitment scheme following the template of Bartusek, Coladangelo, Khurana, and Ma [BCKM21] so that it satisfies simulation security.

The simulation security for a bit commitment is morally trying to capture the following ideal world: (1) in the commit phase, the committer sends a bit b to the ideal functionality; (2) in the reveal phase, the committer asks the ideal functionality to open and the ideal functionality sends the bit b to the receiver. More specifically, the security against receiver is called equivocality, which states that the commitment can be simulated in a way that b is only determined at the beginning of the reveal phase. The security against committer is called extractability, which states that the commitment can be simulated in a way that b can be extracted from the committer after the commit phase completes.

ε -simulation security for inefficient preprocessing. We first consider augmenting the base protocol (Theorem 1) with simulation security. We note some caveats before proceeding.

The first caveat is that we allow the simulator to take a few copies of the auxiliary input state as additional inputs. Intuitively, this means that a malicious party can come out of the protocol obtaining a few extra copies of the magic state. Indeed, if we look at the construction of Theorem 1, the receiver after an honest interaction gains one copy of the magic state from the committer for each fold of repetition, and there does not seem to be a way for the committer to certifiably retrieve the state back. We believe that this weakening is still meaningful and non-trivial since (1) the magic state is supposed to be public knowledge anyways (everyone should have many copies), and (2) the simulation security still guarantees that the “real” input is hidden from the other party.

The second caveat is that we only achieve ε -simulation security (with quantum auxiliary information $|aux\rangle$), which states that there is an efficient simulator S and some polynomial t such that for every adversary A (that is possibly entangled with the distinguisher) and every ε , the view outputted by $(I \otimes S)(A, |aux\rangle^{\otimes t(1/\varepsilon)})$ is distinguishable from the real view except with advantage no more than ε , where S does not touch the distinguisher’s private register. We stress that the protocol itself is independent of ε . Furthermore, this still suffices for almost any game-based application that needs simulation security although with a larger polynomial security loss. To see this, for example, suppose an adversary can break a downstream game-based security with some non-negligible probability p , then we can set our overall simulation error to be $p/2$ to reach a contradiction. On the other hand,

to prove a downstream ε -simulation security, we can similarly pick a smaller ε for each fold and invoke hybrid argument.

The main technical ingredient we need is to implement reflection unitary for an arbitrary initial state, which in this case could contain some inefficient auxiliary information. This is proven in the following lemma, which constructs such an algorithm by using a generalized SWAP test.

Lemma 13 (Approximate state reflection). *For any pure state $|\psi\rangle$, let \mathcal{R}_ψ be the unitary channel for unitary $R_\psi := I - 2|\psi\rangle\langle\psi|$. Then there is a uniformly efficient channel $\tilde{\mathcal{R}}$ such that $\tilde{\mathcal{R}}(\cdot, |\psi\rangle\langle\psi|^{\otimes n})$ is $\sqrt[4]{\frac{64}{n+1}}$ -close to \mathcal{R}_ψ in diamond norm for all $\psi, n \geq 0$.*

Proof. We describe the algorithm as follows. We denote the registers as X_0, \dots, X_n with X_0 being the input register and the rest being initialized to $|\psi\rangle$.

1. Initialize a uniform superposition $|+\rangle_N := \frac{1}{\sqrt{n+1}} \sum_{i=0}^n |i\rangle_N$.
2. Controlled on N being $|i\rangle$, swap X_0 and X_i .
3. Controlled on N being $|+\rangle$, apply phase -1 .
4. Uncompute step 2.
5. Trace out everything except X_0 .

We begin analyzing the algorithm by considering pure state inputs. If the input is $|\psi\rangle$ then phase -1 is correctly applied since steps 2 and 4 do not affect the state. If the input is some orthogonal state $|\phi\rangle$, let $|\phi_i\rangle$ denote the state where X_i is $|\phi\rangle$ and everywhere else is $|\psi\rangle$. Then after step 3, we get the state

$$\frac{1}{\sqrt{n+1}} \sum_{i=0}^n |\phi_i\rangle (|i\rangle - 2|+\rangle) = \frac{1}{\sqrt{n+1}} \sum_i |\phi_i\rangle |i\rangle - \frac{2}{n+1} \sum_{i,j} |\phi_i\rangle |j\rangle.$$

Therefore after step 4, we get

$$|\tilde{\phi}\rangle := \left(1 - \frac{2}{\sqrt{n+1}}\right) |\phi_0\rangle |+\rangle - \frac{2}{n+1} \sum_{i \neq j} |\phi_i\rangle |j\rangle.$$

We compare this state with the expected output and get that the overlap

$$\langle\langle\phi_0| \langle +| | \tilde{\phi}\rangle = 1 - \frac{2}{\sqrt{n+1}} \left(1 + \frac{n}{n+1}\right) \geq 1 - \frac{4}{\sqrt{n+1}}.$$

By decomposing a general pure state $|x\rangle = \sqrt{p}|\psi\rangle + \sqrt{1-p}|\phi\rangle$ and let $|\psi_n\rangle := |\psi\rangle^{\otimes n} |+\rangle$, we have that

$$\langle\langle x | R_\psi^\dagger \otimes \langle\psi_n| | \tilde{x}\rangle \geq 1 - \frac{4}{\sqrt{n+1}} \quad (2)$$

as well by combining the two cases above. Furthermore, (2) generalizes to a larger entangled pure state by simply applying (2) linearly to the Schmidt decomposition. For diamond norm, it suffices to consider any input state where the overall entangled state is pure (since we can without loss of generality purify the state for the distinguisher), thus we have that the trace distance between $(I \otimes R_\psi) |x\rangle \otimes |\psi_n\rangle$ and $|\tilde{x}\rangle$ is at most

$$\sqrt{1 - \left(1 - \frac{4}{\sqrt{n+1}}\right)^2} \leq \sqrt[4]{\frac{64}{n+1}}.$$

This completes the proof since trace distance cannot increase after the operation of tracing out the auxiliary registers holding $|\psi_n\rangle$, which is CPTP. \square

Theorem 14. *ε -simulation secure commitment schemes with quantum auxiliary input exist. Furthermore, it can be built from any non-interactive extractable-binding computationally-hiding commitment scheme with quantum auxiliary input.*

Proof. This follows the same construction and proof strategy as [BCKM21, AQY22] except for one change. In particular, Watrous rewinding [Wat06] was used for a total of λ times in the equivocal simulator [BCKM21, Section 4.1], which involved applying a unitary $I - 2|0\rangle\langle 0|$ on a certain private register of the simulator’s. The purpose of this unitary was to check whether this register returned to all zero state. In our context, this register would be initialized to $|0\rangle \otimes |aux\rangle^{\otimes t}$ instead for a suitably large zero register and some polynomial $t(\lambda)$, and similarly we need to reflect around this state in order for the analysis to go through³.

We establish ε -equivocality as follows. We first consider a simulator that runs the [BCKM21] equivocal simulator with access to an inefficient reflection oracle: this gives a negligible simulation error, and thus it is at most $\varepsilon/2$ for all sufficiently large λ . We now instantiate this inefficient oracle with Lemma 13 where $n = \lceil 1024(\lambda/\varepsilon)^4 \rceil - 1$ and $|\psi\rangle = |0\rangle \otimes |aux\rangle^{\otimes t}$, then by a standard hybrid argument we arrive that the overall simulation error is at most ε . Furthermore, the number of copies of $|aux\rangle$ used is $O(t(\lambda/\varepsilon)^4)$ which is polynomial.

The rest of the proof follows as [BCKM21] using a similar trick of running the inner simulator with a polynomially smaller error. \square

Simulation security in the unclonable common random state model. We conclude by remarking on the simulation security of the commitment scheme with trusted setup from Corollary 6. First of all, we can still apply the [BCKM21] transformation, but in this case we would get negligible simulation security since with access to the compressed database register, the simulator can efficiently test/uncompute the magic state.

However, we note that the argument of Proposition 8 gives an adversary that breaks honest binding when the adversary could control the sampling of the random function using compressed oracles, and thus an analogous argument can show that the commitment of Corollary 6 is in fact already negligibly equivocal even without any further modification to the scheme.

Using similar ideas this commitment is probably also negligibly extractable as well, however, the argument is more complicated since in this case we need to be able to extract any malicious committer (this is unlike the equivocal case where the only freedom a passive but malicious receiver has is to distinguish the views, so it suffices to simply show simulation correctness). We leave formalizing this to future work.

C Post-quantum pseudorandomness

We view a quantum query-less circuit with auxiliary input of total size S as an S -qubit input fed to a universal quantum circuit, which itself is independent of the random function H . (Indeed we

³One naïve idea to fix this is to ask the simulator to only reflect the zero part and ignore the auxiliary information. However, this does not work since we can calculate and see that such a rewinding algorithm (without any further changes) would not work for an adversary that picks its challenge by measuring the auxiliary information it receives from the simulator.

can without loss of generality even take S to be the number of qubits that actually depend on the function H .)

The first polynomial upper bound for this problem was established by Chung, Guo, Liu, and Qian [CGLQ20a] and was subsequently improved by Liu [Liu23]. We follow the second work in this proof. We first recall a game G in a P -BF-QROM [Liu23] to be the following:

0. A random function $H : [N] \rightarrow [M]$ is sampled uniformly at random.
1. The adversary starts by making P (quantum) queries to H , and then we postselect on measuring its first qubit being 1 (abort if it is not possible). This postselection may affect its residual state as well as the conditional distribution of the random function.
2. The challenger then samples a random classical challenge to the adversary, using T_{samp} queries.
3. The adversary produces a response, using T queries.
4. The challenger outputs a bit indicating accept or reject, using T_{verify} queries.

We say G is $\nu(P, T)$ -secure in the P -BF-QROM if any adversary with T queries cannot make the challenger accept with probability higher than ν . In our case, the security game of a pseudorandom ensemble (or PRG) against a query-less adversary corresponds to $T_{\text{samp}} = 1$ and $T = T_{\text{verify}} = 0$: the challenger flips a random bit and either sends a pseudorandom $H(x)$ (using a single query) or a random y , and asks the adversary to predict the bit.

Similarly, an (S, T) non-uniform quantum adversary plays the same security game, except that in step 1 it can do an arbitrary amount of queries but is not allowed to do post-selection, and its output (to be used later in step 3) is restricted to at most S qubits.

Lemma 15 ([Liu23, Lemma 4], with the coefficient from [CGLQ20b, Proof of Lemma 5.13]). *The PRG game has $\nu(P, T) = \frac{1}{2} + 4\sqrt{2} \cdot \sqrt{\frac{P+T^2}{N}}$ in the P -BF-QROM.*

Theorem 16 ([Liu23, Theorem 5]). *Any game G that has security ν in the P -BF-QROM has security*

$$\delta(S, T) \leq \min_{\gamma > 0} \{ \nu(P/\gamma, T) + \gamma \}$$

against (S, T) non-uniform adversaries in QROM, where $P = S(T + T_{\text{verify}} + T_{\text{samp}})$.

Proof of Proposition 2. Combining Lemma 15 and Theorem 16, we find that for any non-uniform algorithm A of size S (that potentially depends on H),

$$\begin{aligned} \left| \mathbb{E}_H \left[\mathbb{E}_x [A(H(x))] - \mathbb{E}_y [A(y)] \right] \right| &= 2 \left| \delta(S \cdot 1, 0) - \frac{1}{2} \right| \\ &\leq 2 \min_{\gamma > 0} \left\{ \sqrt{2^5 \cdot \frac{S}{\gamma N}} + \gamma \right\} \\ &= 12 \cdot \sqrt[3]{\frac{S}{N}}, \end{aligned}$$

showing the bound above. □

Proof of Proposition 7. Since each copy of the magic state can be efficiently prepared through a single quantum query to H , any distinguishing adversary is a valid query-less adversary in the P -BF-QROM, and thus we arrive at the proposition by invoking Lemma 15. □