

Hard Quantum Extrapolations in Quantum Cryptography

Luowen Qian¹

Justin Raizes²

Mark Zhandry¹

¹NTT Research, Inc.

²Carnegie Mellon University

Abstract

Although one-way functions are well-established as the minimal primitive for classical cryptography, a minimal primitive for quantum cryptography is still unclear. Universal extrapolation, first considered by Impagliazzo and Levin (1990), is hard if and only if one-way functions exist. Towards better understanding minimal assumptions for quantum cryptography, we study the quantum analogues of the universal extrapolation task. Specifically, we put forth the classical→quantum extrapolation task, where we ask to extrapolate the rest of a bipartite pure state given the first register measured in the computational basis. We then use it as a key component to establish new connections in quantum cryptography: (a) quantum commitments exist if classical→quantum extrapolation is hard; and (b) classical→quantum extrapolation is hard if any of the following cryptographic primitives exists: quantum public-key cryptography (such as quantum money and signatures) with a classical public key or 2-message quantum key distribution protocols.

For future work, we further generalize the extrapolation task and propose a fully quantum analogue. We show that it is hard if quantum commitments exist, and it is easy for quantum polynomial space.

1 Introduction

Modern cryptography works by reducing the security of complicated cryptosystems down to the hardness of solving simpler problems. These reductions also have the side benefit that it enables us to modularize cryptographic constructions and theorems. A celebrated example developed through a sequence of such transformations is the fact that one-way functions — functions that are easy to compute but hard to invert — are sufficient to realize a large swath of symmetric key cryptography [GGM86, LR88, Rom90, HILL99]. On the other hand, one-way functions are inherent to the vast majority of complexity-based cryptography [IL89, Gol90]. This makes the one-way function abstraction a bedrock of modern cryptography, and cryptographers frequently refer to the assumed existence of one-way functions as the “minimal” assumption in cryptography.

The similar question for quantum cryptography is just as important but is a lot less studied for now. Interestingly, one-way functions are no longer minimal: recent works have demonstrated that using quantum information, cryptography can be based on assumptions that appear strictly milder than even a one-way function [Kre21, AQY22, MY22, KQST23]. This leads to the following fundamental question:

Is there a quantum analog of a one-way function, that is both inherent to essentially all of complexity-based cryptography, while also being sufficient to build useful cryptosystems?

In recent years, there have been some progress towards answering this question but it does not have a satisfactory answer yet. To explain this, let us first recall how classical cryptographic security is usually formalized. Broadly, we can divide them into two categories:

- In a **decision-style** security game, the adversary’s goal is to distinguish between two experiments, such as distinguishing whether the encrypted message is 0 or 1, or whether the game is in the real world or the ideal world.

Decision-style quantum security assumptions are fairly well understood. Specifically, security games of this style can be captured by the notion of an EFI pair, and any such EFI pairs implies a number of cryptographic tasks, such as commitments [Yan22, BCQ23].

- In a **search-style** security game, the adversary’s goal is simply to produce certain message to satisfy some predicate, such as forging a signature or a quantum money state.

Search-style quantum security assumptions are a lot less studied. One glaring issue is that public-key quantum money, a very natural quantum cryptographic primitive, is not known to be related to the other quantum cryptographic primitives. In fact, public-key quantum money and its strengthening, quantum lightning [Zha21], are only known from extremely strong assumptions like post-quantum obfuscation or in ideal oracle models, yet we do not even know if any cryptographic assumption such as EFI pairs are necessary for them.

A few recent works have studied what we call a **quantum**→**classical** search-style assumption, where the adversary is given a potentially quantum input, and must produce some classical output. An example is a quantum one-way state generator [MY24], where informally a keyed mixed state can be efficiently prepared from its classical key but it is hard to find its key given the state. It was shown that quantum one-way state generators with inefficient verification algorithms are existentially equivalent to EFI pairs [KT24a, BJ24]. As a consequence of this, EFI pairs are implied by primitives like secret-key quantum money, given the scheme has a classical secret. However, public-key quantum money may not have a classical secret, and quantum lightning must not have a classical secret. Thus, this does not give us a way to construct EFI pairs from public-key quantum money.

Going to more general search-style assumptions, very little is known about the search-style assumption where the adversary is supposed to output some quantum state, or even a **classical**→**quantum** search-style assumption such as public-key quantum money. To make progress towards addressing the fundamental question above, we can ask a more concrete question:

Is it possible to build EFI pairs from public-key quantum money?

1.1 Classical→Quantum Extrapolation

To understand how EFI pairs, or equivalently, commitments may be built from public-key cryptography, let us take a step back and think about how a classical analogue of this is proved, for example, how we can construct a classical commitment scheme or a one-way function from a classical signature scheme. The key middle step of the classical proof [IL90] is the hardness of the *universal extrapolation* task. In particular, it is shown that secure signature schemes imply hardness of universal extrapolation, which in turn implies the existence of (distributional) one-way functions.

Inspired by this template, we study quantum analogues of this extrapolation task, which we emphasize is a search-style assumption. We first define classical→quantum extrapolation which, given

an efficient quantum pure state, asks to extrapolate the rest of the quantum state conditioned/post-selected on measuring the first half of the state (in the computational basis).

Definition 1.1 (Classical→Quantum Extrapolation, informal). *A classical→quantum extrapolation problem is specified by a circuit Gen that produces a pure state, which can be written as*

$$|\text{Gen}\rangle = \sum_s \alpha_s |s\rangle_{\text{A}} \otimes |\psi_s\rangle_{\text{B}}$$

for some $\alpha_s \geq 0$ and unit vectors $|\psi_s\rangle$. We say (a uniform family of) Gen is hard if for every QPT adversary Adv (potentially with auxiliary input), its output has a negligible overlap with the correct state $|\psi_s\rangle$ given the classical part s , i.e.

$$\mathbb{E}[\text{Tr}(|\psi_s\rangle\langle\psi_s| \text{Adv}(s))] = \text{Tr}\left(\sum_s \alpha_s^2 |\psi_s\rangle\langle\psi_s| \text{Adv}(s)\right)$$

is negligible.

This is a natural generalization of the classical extrapolation task for distributions to pure states, but instead of asking for the conditional distribution on B conditioned on s , we use a stronger requirement of preparing the pure (unmeasured) state $|\psi_s\rangle$.

Using similar proof ideas as classically, we can straightforwardly establish the following.

Proposition 1.2. *There exists a hard classical→quantum extrapolation task if any of the following exists:*

- *Public-key quantum money scheme (with a classical public key).*
- *Public-key quantum signature scheme (with a classical public key).*
- *2-message quantum key distribution that is only unpredictably-secure.*

Our main theorem then finishes the proof by establishing the following.

Theorem 1.3. *If there exists a hard classical→quantum extrapolation task, then quantum commitment schemes exist.*

In fact, our commitment construction trivially adapts to *cloneable*→quantum extrapolation, where we only require the challenge A register to be cloneable instead of strictly classical. As a consequence, we can construct commitments from public-key quantum money scheme with a cloneable quantum public key as well. Interested readers should refer to the formal treatment later.

We remark that these results are optimal in a few perspectives. First, asking for the public key to be cloneable is barely a restriction: if one would like to reap the benefits of having a public key infrastructure, it is important for the authority to be able to efficiently clone and distribute public keys without all the users being online to produce new copies of the public key. Second, it is known that uncloneable public-key signatures (with bounded number of copies of the public key) are in fact statistically possible [GC01]. Third, 3-message quantum key distribution with unpredictability security is also statistically possible.

Remark 1.4. *It might be tempting to think that some form of distributional one-wayness should follow from hard classical \rightarrow quantum extrapolation tasks. Classically, consider a sampler of a hard-to-extrapolate distribution $\text{Samp}(r) \rightarrow (a, b)$. Then the truncated sampler $f(r) \rightarrow a$ is immediately a distributionally one-way function. However, since a quantum algorithm can build a pure state from scratch, it is unclear if any meaningful form of distributional one-wayness can be constructed here. From this perspective, it appears that extrapolation may be a more useful abstraction than distributional one-wayness.*

1.2 Quantum Extrapolation

Given the result above, it is natural to wonder if some version of the following generalization is true: any search-style assumption implies EFI pairs, even if both the inputs and outputs are potentially quantum. With that said, to the best of our knowledge, all natural examples (binding security of commitments, or soundness of zero knowledge arguments) already imply EFI pairs.

Hence, we put forth a candidate *quantum input, quantum output* search task. Intuitively, the task is to convert a bipartite pure state into its canonical purification. We show that this task is implied by commitment schemes and on the other hand, can be solved in quantum polynomial space. We do not know how to use it to construct other primitives such as commitments, and we provide a discussion of the difficulties in extending our construction in Section 2.3.

Theorem 1.5 (Informal). *If quantum bit commitments exist, then there exists an efficiently preparable state with Schmidt decomposition*

$$\sum_i \alpha_i |A_i\rangle \otimes |B_i\rangle,$$

such that it is hard to coherently “strongly” map $|B_i\rangle$ to $|A_i^\rangle$ on average over α_i . On the other hand, this task is possible in pureUnitaryPSPACE for any efficiently preparable state.*

Discussion: quantum minimalism. Mark Zhandry’s talk on “Quantum Minimalism” [Zha23] highlighted various features that make one-way functions “minimal” for classical cryptography. We recall these features here and highlight the fact that quantum extrapolation satisfy the first 7 proposed features to roughly the same extent as classically. Thus, hard quantum extrapolations might be a good candidate for the minimal quantum cryptographic assumption from this perspective if we can build useful cryptography from these.

Feature 1: (Somewhat) trivially implied by most general primitives As discussed above, this matches one-way functions since we can somewhat trivially prove hardness of quantum extrapolation from commitments, which is implied by most general primitives.

Feature 2: Trivially and robustly implied by most concrete assumptions This holds for all concrete assumptions that we can think of.

Feature 3: Simple to define Modulo the complex conjugate condition which is necessary in order to satisfy other features, this is essentially true.

Feature 4: Falsifiable This is probably not true in general. However, in its defense, its classical analogues, universal extrapolation hardness or distributional one-way functions, are not known to be falsifiable either. (If falsifiable assumptions like commitments are implied by this, then this feature is not as essential.)

Feature 5: Search problem Quantum extrapolation is defined as a search problem.

Feature 6: Trivial combiners and universal constructions A trivial combiner can be constructed by simply concatenating the respective parts of the two states: this thus implies a universal construction via standard techniques.

Feature 7: Minimal correctness requirements This holds since the only requirement on the state is that it should be efficiently preparable.

Feature 8: Can build crypto? We leave as interesting open question whether the hardness of quantum extrapolation can be used to construct EFI pairs or other cryptographically useful hardness.

2 Technical Overview

2.1 Why a New Approach?

Before we discuss our commitment construction, one might wonder if it might be simpler to generalize the construction of commitments (or EFI pairs) from quantum one-way state generators [KT24a, BJ24] to this setting. After all, their techniques can handle quantum \rightarrow classical search-style assumption and we only need to do the other way around.

Upon closer inspection, it appears that unfortunately their techniques crucially rely on having access to a classical secret. More specifically, their construction of EFI pair on a high level still goes through the classical construction of hardcore predicates from one-way functions, and this in turn builds on the fact that the secret is classical.

To see this, recall that the construction essentially leaks some random inner product about the secret and argues that such a leakage appears computationally random to the adversary. However, the presence of a quantum secret shuts down this approach due to the lack of a quantum analogue of the Goldreich–Levin hardcore bit. In fact, some natural formulation of such an analogue is flat out impossible. For example, assume the quantum secret is a (computationally) Haar random state. Then a non-trivial leakage would instead be statistically independent of the secret, whereas the GL hardcore bit is statistically determined by the secret.

Another naïve approach is to try to apply the construction of one-way functions from hardness of universal extrapolation to the classical \rightarrow quantum extrapolation setting. Unfortunately, this construction also applies randomness extractors on the secret and we run into a similar difficulty as above.

2.2 Commitments from Classical \rightarrow Quantum Extrapolation

Given the known approaches of constructing EFI pairs do not seem to work, it appears that a drastically different approach is needed. The initial idea is that since the binding security game of the commitment scheme is already a search-style assumption and is known to be existentially equivalent to EFI pairs, it might be easier to construct a computationally *binding* commitment instead. Perhaps we can carefully craft a commitment scheme so that breaking its binding would directly correspond to solving the extrapolation problem. This idea was also used to construct commitments from secretly-verifiable statistically-invertible one-way state generators [MY24].

To construct a commitment, we will use the canonical commitment scheme template [Yan22]. As a starting point, imagine a sender who manufactures $|\text{Gen}\rangle$ in the opening register \mathbf{D} and copies s to the commitment register \mathbf{C} . This results in the state

$$\sum_s \alpha_s |s\rangle_{\mathbf{C}} \otimes |s, \psi_s\rangle_{\mathbf{D}},$$

which we will use as the commitment and the decommitment register if we want to commit to bit 1.

Now imagine if we could also “obliviously” quantum sample the challenge distribution, that is, we can prepare the following pure state

$$\sum_s \alpha_s |s\rangle_{\mathbf{C}} \otimes |s\rangle_{\mathbf{D}}$$

for the case of committing to 0. For example, if the reduced density matrix on s is maximally mixed, then this is essentially asking to prepare the maximally entangled state, which can be efficiently done.

However, note that if this were indeed possible, we would be done as the commitment would be computationally binding. If a malicious committer commits to 0 and wishes to change his mind to decommit to 1, he would essentially be forced to craft the state $|\psi_s\rangle$ given only s .

Removing quantum sampling assumption. Unfortunately, not every efficient distribution can be quantum sampled unless, for example, $\text{SZK} \subseteq \text{BQP}$ [AT07]. Thus we want to also handle the case where s can be arbitrarily distributed, which could be the case for a general quantum money scheme.

A first idea is that maybe starting from the 0-commitment state, we can “effectively” remove $|\psi_s\rangle$ from the committer’s view (or equivalently, its register \mathbf{D}). Naïvely we could simply assign the state $|\psi_s\rangle$ to the \mathbf{C} register. However, the resulting commitment scheme would be not hiding, so some care has to be taken.

The next idea, then, is to maybe employ some encryption scheme to do this. More specifically, our second attempt is as follows:

$$\begin{aligned} |\text{com}_0\rangle &= \mathbb{E}_k \left[\sum_s \alpha_s |s, k\rangle_{\mathbf{C}} \otimes |s, \text{Enc}(k, \psi_s)\rangle_{\mathbf{D}} \right], \\ |\text{com}_1\rangle &= \mathbb{E}_k \left[\sum_s \alpha_s |s, k\rangle_{\mathbf{C}} \otimes |s, \text{Enc}(k, 0), \psi_s\rangle_{\mathbf{D}} \right], \end{aligned}$$

for some (unitary) encryption scheme Enc .

Now, what encryption scheme would work? The most natural candidate, quantum one-time pad, turns out not to work: if $|\psi_s\rangle$ are all encoded in Hadamard basis then the resulting scheme would again not be hiding. More specifically, consider the state $1 \otimes |+\rangle^{\otimes n}$. Then (ignoring the $|+\rangle^{\otimes n}$ part in com_1 which is irrelevant for hiding)

$$\begin{aligned} |\text{com}_0\rangle &= \mathbb{E}_{k \in \{0,1\}^{2n}} [|k\rangle_{\mathbf{C}} \otimes |\text{Enc}(k, |+\rangle^{\otimes n})\rangle_{\mathbf{D}}] = \left(|+\rangle_{\mathbf{C}_X} \otimes \frac{|0+\rangle_{\mathbf{C}_Z \mathbf{D}} + |1-\rangle_{\mathbf{C}_Z \mathbf{D}}}{\sqrt{2}} \right)^{\otimes n}, \\ |\text{com}_1\rangle &= \mathbb{E}_{k \in \{0,1\}^{2n}} [|k\rangle_{\mathbf{C}} \otimes |\text{Enc}(k, 0)\rangle_{\mathbf{D}}] = \left(|+\rangle_{\mathbf{C}_Z} \otimes \frac{|00\rangle_{\mathbf{C}_X \mathbf{D}} + |11\rangle_{\mathbf{C}_X \mathbf{D}}}{\sqrt{2}} \right)^{\otimes n}. \end{aligned}$$

Thus we see that these two density matrices have statistical distance close to 1 as $n \rightarrow \infty$.

Nevertheless, we make the crucial observation: this attack appears to rely on the secret quantum state being encoded in a specific basis. If we attempt to encrypt, instead of only in two bases, in exponentially many bases, then the probability that we hit such a bad basis would be negligible. To make this efficient, let us say that we just pick one of bases to encrypt uniformly at random. Since adding random phases to a basis is equivalent to measuring it, this gives us the following random-measurement commitment scheme.

Construction 2.1 (Informal). *Consider an (exponentially large) family of complete measurement bases $M = \{|\phi_{k,x}\rangle\}$ indexed by k and measurement outcome x . Given a quantum state $|\psi\rangle$, we use the notation $|M_k(\psi)\rangle \otimes |M_k(\psi)\rangle$ to informally denote (coherently) measuring $|\psi\rangle$ in basis k and then cloning the outcome into the second register. More formally, we use this notation to denote the following pure state*

$$|M_k(\psi)\rangle \otimes |M_k(\psi)\rangle := \sum_x \langle \phi_{k,x} | \psi \rangle |x\rangle \otimes |x\rangle.$$

Then the random-measurement commitment with bases M is the following construction:

$$\begin{aligned} |\text{com}_0\rangle &= \mathbb{E}_k \left[\sum_s \alpha_s |s, k, M_k(\psi_s)\rangle_{\text{C}} \otimes |s, k, M_k(\psi_s)\rangle_{\text{D}} \right], \\ |\text{com}_1\rangle &= \mathbb{E}_k \left[\sum_s \alpha_s |s, k, M_k(0)\rangle_{\text{C}} \otimes |s, k, M_k(0), \psi_s\rangle_{\text{D}} \right]. \end{aligned}$$

In the encryption perspective, this is equivalently using a random encryption scheme indexed by some public randomness k , and secret random phases are added to the basis specified by k . Moving forward, we will use the measurement perspective since it is more convenient for the analysis.

We now need to argue that the commitment is hiding and binding and neither is obvious at this point. Intuitively, hiding should hold as our intuition above indicates that a sufficiently random set of bases should evade the bad distinguishing attack with overwhelming probability. On the other hand, giving the adversary a complete measurement of the state in a probably useless basis should not help the adversary to swap the state out with the zero state — or in the encryption perspective, the encryption should be secure enough that the adversary could only build the state from scratch instead of extracting useful information from the encryption.

Weak statistical hiding. As observed above, the family of measurement bases must be sufficiently large, and different bases need to be “independent” enough to avoid the attack above. Given these criteria, we consider a few candidates: (1) mutually unbiased bases (MUBs); (2) unitary t -designs; (3) stripped down t -designs such as binary phase bases [MPST24]. It turns out that both MUBs and t -designs for $t \geq 2$ work, while 1-designs and binary phase bases do not necessarily work. For comparison, MUBs have very good randomness complexity (n bits of randomness for hiding n qubits) while 2-designs or random Cliffords are more efficient (can be applied in quasi-linear time). For interested readers, we explain the counterexamples in Section A.

On a high level, the hiding proof can be reduced the following question: given any two mixed states ρ_0, ρ_1 that could be statistically far; is it true that the states after measuring them in a random basis k , $(\rho_0^{(k)}, k), (\rho_1^{(k)}, k)$ are somewhat statistically close? This also might be an independently interesting question on its own beyond this application: traditionally, tomography asks for a set of

measurements that preserves the quantum information; here, we ask for a set of measurements that destroys quantum information with respect to statistical distance.

Crucially, we need to bound the statistical distance *even if the basis choice k is revealed*. This is problematic for the proof of t -designs, since t -designs generally do not guarantee anything if the secret key k is revealed: everyone can check that the specific unitary corresponding to k is applied instead of a Haar random unitary. However, there is a simple trick we can apply: the statistical distance we need to bound above can be equivalently expressed as $\mathbb{E}_k[TD(\rho_0^{(k)}, \rho_1^{(k)})]$, as the distribution over k is identical. Furthermore, here the key is no longer given to the distinguisher. Even though trace distance is not a polynomial (due to the absolute value), we can nevertheless bound it by an appropriate degree-2 polynomial and then invoke the security of 2-designs. In the end, we show that the average statistical distance under 2-designs is upper bounded by $\|\rho_0 - \rho_1\|_2 / 2$. A naïve upper bound would be $\|\rho_0 - \rho_1\|_2 / 2 \leq \|\rho_0 - \rho_1\|_1 / 2 = TD(\rho_0, \rho_1)$ which yields a trivial bound. However, by more carefully leveraging Jordan–Hahn decomposition of $\rho_0 - \rho_1$, we are able to show that $\|\rho_0 - \rho_1\|_2 / 2 \leq TD(\rho_0, \rho_1) / \sqrt{2}$. Curiously, it appears that even measurements under Haar random unitaries would still leave behind $1/2$ statistical distance.

To prove MUBs work, we leverage a useful fact by Ivanovic [Iva81, Iva92] that any density matrix can be decomposed into the sum of its post-measurement states over all MUBs; furthermore, each pair of states are orthogonal with respect to the Hilbert–Schmidt inner product after appropriate renormalization. Utilizing the orthogonality, we can translate the average (L^1) statistical distance to (L^2) Hilbert–Schmidt distance to invoke the decomposition, and then back to statistical distance using the same trick above without losing too much. In the end, we can compute that the average statistical distance under MUBs is at most $1/\sqrt{2}$ as well.

We refer interested readers to the later sections for a more detailed proof.

Computational binding. To show (honest) binding, we show that any adversary who is given register D of $|\text{com}_0\rangle$ and successfully opens to 1 can be used to solve the classical→quantum extrapolation problem. In the classical→quantum extrapolation game, the reduction is essentially given a mixed state $\sum_s \alpha_s^2 |s\rangle\langle s|$. In order to meaningfully invoke the binding adversary, the reduction must produce the reduced density matrix on D of com_0 , which looks like

$$\sum_{k,s} \alpha_s^2 |s\rangle\langle s| \otimes |k\rangle\langle k| \otimes \rho_{k,s} \otimes |\vec{0}\rangle\langle \vec{0}|$$

when C is traced out, where $\rho_{k,s}$ denotes the mixed state resulting from measuring $|\psi_s\rangle$ in the basis k . Observe that since the canonical-form commitment receiver is going to project on to the $|\text{com}_1\rangle$ state, the adversary’s output must (at the very least) be the following form:

$$\sum_{k,s} \alpha_s^2 |s\rangle\langle s| \otimes |k\rangle\langle k| \otimes |\psi_s\rangle\langle \psi_s| \otimes \rho'_{k,s}$$

for some $\rho'_{k,s}$. This is true since the committer already holds a copy of k and s . Therefore, as long as we can mimic the distribution on k and $\rho_{k,s}$, the reduction would work.

Recall our intuition above that $\rho_{k,s}$ should disclose very little information so that it would not help the adversary to prepare $|\psi_s\rangle$. This suggests that maybe the reduction can just put there maybe a maximally mixed state and hope that the adversary would not notice. However, recall from above that $\rho_{k,s}$ is only *weakly* hiding. In other words, it may contain some information about

$|\psi_s\rangle$. It is not clear that the reduction can generate such a state without already violating the hardness of the extrapolation problem.

Our solution is to simply generate $\rho_{k,0}$, a measurement of $|\vec{0}\rangle$ in the basis k , and feed it to the binding adversary instead. In other words, we are again measuring all zeroes and expect the adversary to not see a difference. This turns out to work. Although $\rho_{k,0}$ only matches $\rho_{k,s}$ to a limited degree, it *does* match register **C** from $|\text{com}_0\rangle$. Since register **C** in $|\text{com}_1\rangle$ would contain $\rho_{k,s}$ instead of $\rho_{k,0}$, the support states where $\rho_{k,s}$ significantly differs from $\rho_{k,0}$ precisely cover the case where the commitment is *statistically binding*. Thus, the assumed binding adversary's advantage must be entirely on the portion of $\rho_{k,s}$ which matches $\rho_{k,0}$. In order for the substitution to work, the receiver's view of $\rho_{k,0}$ in $|\text{com}_0\rangle$ must be symmetric with the sender's view of $\rho_{k,s}$ in $|\text{com}_1\rangle$. Therefore if the choice of basis k appears in one view, it must appear in both.

2.3 Quantum Extrapolation

Here, we discuss the more general task of quantum extrapolation, where an adversary operates on register **B** of the Schmidt decomposed state

$$\sum_i \alpha_i |A_i\rangle_{\text{A}} \otimes |B_i\rangle_{\text{B}}.$$

For the sake of exposition, we omit conjugations of the states and assume it is hard to coherently map $|B_i\rangle$ to $|A_i\rangle$ on average over α_i . Technically, the conjugation is necessary for ensuring that the solution is well-defined, since the Schmidt decomposition of a state is not necessarily unique. See the technical sections for a formal treatment.

Implications from Public-Key Quantum Money and Commitments. It is not hard to see that both public-key quantum money and commitments imply hard quantum extrapolation tasks. In the case of public-key quantum money, we may regard the serial number as $|B_i\rangle$ and the banknote as $|A_i\rangle$; thus, any adversary mapping $|B_i\rangle \mapsto |A_i\rangle$ counterfeits banknotes using their serial numbers.

The case of commitments is only slightly more complicated. A perfectly hiding commitment in canonical quantum form has a Schmidt decomposition

$$\begin{aligned} |\text{com}_0\rangle &= \sum_i \alpha_i |c_i\rangle_{\text{C}} \otimes |d_{i,0}\rangle_{\text{D}}, \\ |\text{com}_1\rangle &= \sum_i \alpha_i |c_i\rangle_{\text{C}} \otimes |d_{i,1}\rangle_{\text{D}}. \end{aligned}$$

Note that because the commitment is perfectly hiding, the left hand side of the two decompositions are the same. Thus, if an adversary could coherently map $|d_{i,0}\rangle \mapsto |c_i\rangle \mapsto |d_{i,1}\rangle$, they could open $|\text{com}_0\rangle$ to $|\text{com}_1\rangle$ by operating only on the opening register, breaking computational binding.

Handling statistical hiding is essentially the same reduction, but the analysis is more difficult. Essentially, we need to argue that statistical hiding implies that the two target states are close. We prove this by establishing that the closeness of the target states is captured by the Holevo fidelity and then utilizing a tight connection between Holevo fidelity and trace distance [Hol72].

Difficulties in Building Commitments. It is natural to wonder whether our construction from cloneable \rightarrow quantum extrapolation can be extended to the more general quantum extrapolation task.

A natural attempt is to prepare the extrapolation state in the opening register D, then “destroy” either $|A_i\rangle$ or $|B_i\rangle$, depending on the message bit, by measuring it and copying it in a random basis to the commitment register C. This results in a commitment which looks like the following:

$$\begin{aligned} |\text{com}_0\rangle &= \mathbb{E}_k \left[\sum_i \alpha_i |k, M_k(A_i)\rangle_{\text{C}} \otimes |k, B_i, M_k(A_i)\rangle_{\text{D}} \right], \\ |\text{com}_1\rangle &= \mathbb{E}_k \left[\sum_i \alpha_i |k, M_k(B_i)\rangle_{\text{C}} \otimes |k, A_i, M_k(B_i)\rangle_{\text{D}} \right]. \end{aligned}$$

In order to open $|\text{com}_1\rangle$ to $|\text{com}_0\rangle$, it would suffice to map $|B_i\rangle \mapsto |A_i\rangle$ and let the statistical hiding take care of the differences between $M_k(A_i)$ and $M_k(B_i)$. Unfortunately, it is not clear whether this mapping is *necessary* to break binding. In particular, the commitment register only contains partial information about $|A_i\rangle$, so the adversary is not necessarily bound to any particular i when breaking binding. Instead, it could potentially map $|B_i\rangle$ onto a superposition of $|A_j\rangle$ ’s.

For a more concrete example, consider the state $\sum_x |x\rangle \otimes H|x\rangle = \sum_{xy} (-1)^{x \cdot y} |x\rangle \otimes |y\rangle$ and the fixed measurement of measuring in the computational basis (note that our binding proof before works for any family of measurement bases). Then the resulting commitment has identical commitment states: $\sum_{xy} (-1)^{x \cdot y} |x\rangle \otimes |y, x\rangle = \sum_{xy} (-1)^{x \cdot y} |y\rangle \otimes |x, y\rangle$. This means that the action to break binding is the identity unitary, whereas we would expect the adversary to apply the Hadamard gates.

2.4 Related Works

Both our construction and the commitment construction from one-way state generator [KT24a, BJ24] (can) make use of unitary t -designs. Interestingly, the similarity is superficial since the purposes are opposite. (As a historical note, we actually arrived at this construction before we became aware of [KT24a].) In our work, we use 2-designs to reduce statistical distance, whereas in their work, 3-designs are indirectly used by classical shadow to preserve information about the quantum state since they are tomographically complete.

After we have announced our results, a more recent work by Dakshita Khurana and Kabir Tomer [KT24b] independently constructs commitment schemes from a state puzzle, which is equivalent to our definition of a hard classical \rightarrow quantum extrapolation task. They also prove an amplification theorem for state puzzles, thus commitments can also be built from weak state puzzles, or equivalently, weakly-hard classical \rightarrow quantum extrapolation tasks. However, their techniques would not immediately extend if we instead have a hard cloneable \rightarrow quantum extrapolation task. To see this, their main idea is to view this hardness as a hard state synthesis problem, and construct hard one-way puzzles (or classical \rightarrow classical extrapolation tasks) from this hardness using similar techniques as for solving state synthesis with a classical oracle; afterwards, commitments are built using prior work [KT24a]. In comparison, we construct a commitment directly from hard classical \rightarrow quantum extrapolation tasks and the construction trivially extends to cloneable bases as well.

3 Commitments

We recall the definition of a canonical non-interactive quantum bit commitment from [Yan22].

Definition 3.1 (Quantum Bit Commitment Syntax). *A canonical (non-interactive) quantum bit commitment is specified by a family of unitaries $\{\text{Com}_\lambda\}_{\lambda \in \mathbb{N}}$ which acts on two registers \mathbf{C} and \mathbf{D} . It consists of two stages:*

- **Commit.** *To commit to a bit b , the sender prepares the state $|b\rangle \otimes |\vec{0}\rangle$ in register (\mathbf{C}, \mathbf{D}) , then applies Com_λ to it to obtain the state $|\text{com}_b\rangle_{\mathbf{C}, \mathbf{D}}$.¹ It sends register \mathbf{C} to the receiver as the commitment register and keeps \mathbf{D} as the opening register.*
- **Open.** *To open the commitment, the sender sends register \mathbf{D} to the receiver. The receiver applies $\text{Com}_\lambda^\dagger$ to register (\mathbf{C}, \mathbf{D}) , then measures the register in the computational basis. If the measurement result is of the form $b||\vec{0}$ for a bit b , the receiver outputs b . Otherwise they output \perp .*

Note that a canonical commitment inherently enjoys completeness; the receiver will always output b as the result of opening a commitment to b , because its measurement for output b is exactly a projection onto the state the sender prepares.

A quantum bit commitment must also satisfy hiding and a notion of binding. We consider honest-binding, which informally guarantees that no adversary given an honestly-prepared commitment to 0 can open it to 1 instead, and vice-versa. It is known that honest-binding for canonical form commitments suffices for stronger binding security [Yan22].

Definition 3.2 (Honest Binding). *A commitment scheme $\{\text{Com}_\lambda\}_{\lambda \in \mathbb{N}}$ is computationally (resp. statistically) ϵ -honest-binding if for every sufficiently large security parameter λ , every auxiliary input $|\psi\rangle$ in register \mathbf{A} , and every polynomial-time (resp. physically) realizable unitary U operating on register (\mathbf{A}, \mathbf{D}) ,*

$$\text{Tr}[(I \otimes |\text{com}_1\rangle\langle\text{com}_1|) U (|\psi\rangle\langle\psi| \otimes |\text{com}_0\rangle\langle\text{com}_0|) U^\dagger] \leq \epsilon$$

When $\epsilon = \text{negl}$, we simply refer to the commitment as honest-binding.

It is known that binding on $0 \rightarrow 1$ as defined above also implies binding on $1 \rightarrow 0$, where the adversary's task is instead transforming an honest commitment to the 1 bit into one to the 0 bit.

Definition 3.3 (Commitment Hiding). *A commitment is computationally (resp. statistically) ϵ -hiding if the states*

$$\frac{\text{Tr}_{\mathbf{D}}(|\text{com}_0\rangle\langle\text{com}_0|)}{\text{Tr}_{\mathbf{D}}(|\text{com}_1\rangle\langle\text{com}_1|)}$$

are ϵ -computationally (resp. ϵ -statistically) indistinguishable for every sufficiently large λ . When $\epsilon = \text{negl}$, we simply refer to the commitment as hiding.

While the security presented here may appear weak, it is known that canonical form commitments satisfying honest binding and hiding is sufficient for constructing commitments with stronger security.

¹To simplify the notation, we write $|\vec{0}\rangle$ to denote enough $|0\rangle$ states to finish filling the register.

4 Statistical Hiding via Random Measurements

We give two methods to coherently “destroy” the information in a quantum state by dividing it into two registers so that each register on its own contains only a limited amount of information about the original state. Concretely, consider measuring a n -qubit register \mathcal{R} with respect to a basis B_i which is drawn from a set of bases \mathcal{B} , then outputting both the measurement result and the choice of basis i . We prove that for certain sets of bases, the expected distance between the distributions induced by this procedure on any two states (averaged over the choice of basis, which equals the total trace distance) is bounded away from 1.

We first prove hiding when the set of bases \mathcal{B} is mutually unbiased. To do this, we recall the definition of mutually unbiased bases and their properties.

Consider an N -dimensional quantum state. A complete measurement is a set of rank-1 projectors $\{P_i\}_{i \in [N]}$ such that $\text{Tr}(P_i P_j) = \delta_{ij}$. Two complete measurements $\{P_i\}_{i \in [N]}$ and $\{Q_i\}_{i \in [N]}$ are mutually unbiased if $\text{Tr}(P_i Q_j) = 1/N$ for all i, j . A maximum set of mutually unbiased bases (MUBs) is a set of $(N+1)$ complete measurements $\{\{P_i^{(r)}\}_{i \in [N]}\}_{r \in [N+1]}$ that are mutually unbiased. For the qubit-string case ($N = 2^n$), it is known how to measure any complete measurement in a maximum set of MUBs in time $O(n^3)$ [DPS04, Section 5.1].

All Hermitian N -dimensional matrices form an N^2 -dimensional real vector space under the Hilbert–Schmidt inner product $\langle X, Y \rangle := \text{Tr}(XY)$. This also induces the Hilbert–Schmidt norm $\|X\| = \sqrt{\text{Tr}(X^2)}$.

Lemma 4.1 ([Iva81, Iva92]). *When a maximum set of MUBs is known, a density matrix W has the following orthogonal decomposition*

$$W = \frac{I}{N} + \sum_r \left(W^{(r)} - \frac{I}{N} \right) \quad (1)$$

with respect to the Hilbert–Schmidt inner product, where $W^{(r)} := \sum_i P_i^{(r)} W P_i^{(r)}$ is the projection onto the r -th MUB. Furthermore, each term is mutually orthogonal even when considering two different density matrices.

Leveraging these, we prove the following lemma, which intuitively states that the statistical distance would be somewhat hidden under most MUB measurements. Notably, this is true even if the distinguisher knows the measurement basis.

Lemma 4.2. *Let N be a prime power and let \mathcal{B} be any set of $p \leq N+1$ MUBs. For any two density matrices W_0 and W_1 , the expected trace distance between the post-measurement state of W_0 and W_1 under a random $r \leftarrow \mathcal{B}$ is*

$$\mathbb{E}_{r \leftarrow \mathcal{B}} \left[\text{TD}(W_0^{(r)}, W_1^{(r)}) \right] \leq \sqrt{\frac{N}{2p}} \cdot \text{TD}(W_0, W_1).$$

In particular, if we instantiate $p = N$, then we get statistical distance to be at most $1/\sqrt{2}$.

Proof. Leveraging the orthogonal decomposition from (1), we can see that

$$\begin{aligned}
\|W_0 - W_1\|^2 &= \left\| \left(\frac{I}{N} - \frac{I}{N} \right) + \sum_{r \in [N+1]} \left(\left(W_0^{(r)} - \frac{I}{N} \right) - \left(W_1^{(r)} - \frac{I}{N} \right) \right) \right\|^2 \\
&= \sum_{r \in [N+1]} \left\| \left(W_0^{(r)} - \frac{I}{N} \right) - \left(W_1^{(r)} - \frac{I}{N} \right) \right\|^2 \\
&= \sum_{r \in [N+1]} \left\| W_0^{(r)} - W_1^{(r)} \right\|^2 \\
&\geq \sum_{r \in \mathcal{B}} \left\| W_0^{(r)} - W_1^{(r)} \right\|^2,
\end{aligned}$$

where the second equality follows from the fact that all the cross terms are 0. Dividing both sides by p , we get

$$\begin{aligned}
\mathbb{E}_r \left[\left\| W_0^{(r)} - W_1^{(r)} \right\|^2 \right] &\leq \frac{1}{p} \|W_0 - W_1\|^2 \\
&= \frac{1}{p} \|W_+ - W_-\|^2 \\
&= \frac{\|W_+\|^2 + \|W_-\|^2}{p} \\
&\leq \frac{\|W_+\|_1^2 + \|W_-\|_1^2}{p} \\
&= \frac{\|W_0 - W_1\|_1^2}{2p},
\end{aligned}$$

where the second equality is considering the Jordan–Hahn decomposition of $W_0 - W_1$, which gives two orthogonal PSD matrices W_+, W_- such that $\text{Tr}(W_+) = \text{Tr}(W_-) = \frac{1}{2} \|W_0 - W_1\|_1$. This is because $\text{Tr}(W_+) + \text{Tr}(W_-) = \text{Tr}(W_+ + W_-) = \|W_0 - W_1\|_1$ and $\text{Tr}(W_+) - \text{Tr}(W_-) = \text{Tr}(W_+ - W_-) = \text{Tr}(W_0 - W_1) = 0$.

Then the overall trace distance of doing a random MUB measurement is, by Cauchy–Schwarz,

$$\begin{aligned}
\frac{1}{2} \mathbb{E}_r \left[\left\| W_0^{(r)} - W_1^{(r)} \right\|_1 \right] &\leq \frac{1}{2} \sqrt{N} \cdot \mathbb{E}_r \left[\left\| W_0^{(r)} - W_1^{(r)} \right\| \right] \\
&\leq \frac{\sqrt{N}}{2} \cdot \sqrt{\mathbb{E}_r \left[\left\| W_0^{(r)} - W_1^{(r)} \right\|^2 \right]} \\
&\leq \sqrt{\frac{N}{2p}} \cdot \frac{\|W_0 - W_1\|_1}{2},
\end{aligned}$$

where the second inequality is due to Jensen’s inequality. \square

In Section A, we further discuss a few alternative choices of bases of measurement that work or do not work.

5 Commitments from Hard Classical→Quantum Extrapolation

5.1 Classical→Quantum Extrapolation

Here, we define a computational task which we call *classical→quantum extrapolation*. At a high level, a classical→quantum extrapolation problem is specified by an efficiently sampleable distribution over pairs $(s, |\psi_s\rangle)$ consisting of a classical string s and a quantum state $|\psi_s\rangle$. Given a random s , the adversary’s task is to extrapolate the quantum half of the pair $|\psi_s\rangle$. We say that this task is hard if no QPT adversary succeeds at this with noticeable probability over the choice of s .

Definition 5.1. *A cloneable→quantum extrapolation problem is specified by two efficient families of isometries $(\text{Gen}, \text{Clone})$ satisfying the following property.*

- (*Cloning correctness*) For every λ , let $\{|\text{chal}_i\rangle\}_i$ be the set of unit vectors such that $\text{Clone}_\lambda |\text{chal}_i\rangle = |\text{chal}_i\rangle |\text{chal}_i\rangle$.² Then the state generated by Gen_λ should admit the following decomposition $\text{Gen}_\lambda = \sum_i \alpha_i |\text{chal}_i\rangle_{\text{A}} \otimes |\psi_i\rangle_{\text{B}}$ for some $\alpha_i \geq 0$ and unit vector $|\psi_i\rangle$. In other words, $((I - \sum_i |\text{chal}_i\rangle\langle\text{chal}_i|) \otimes I) \text{Gen}_\lambda = 0$.

The goal of the task is on input $|\text{chal}_i\rangle$, produce $|\psi_i\rangle$ with non-negligible probability, or more formally, the probability that Adv with auxiliary input aux succeeds

$$\text{Tr} \left(\sum_i \alpha_i^2 |\psi_i\rangle\langle\psi_i| \text{Adv}(|\text{chal}_i\rangle\langle\text{chal}_i| \otimes \text{aux}) \right)$$

is non-negligible.³

In particular, if $\text{Clone} |x\rangle = |x\rangle |x\rangle$ is simply cloning in the computational basis, then we consider this as a classical→quantum extrapolation task.

Here, we make the distinction between efficiently cloneable basis and classical (or telegraphable) basis since efficiently cloneable basis might be a larger class [NZ24].

5.2 Hardness from Cryptographic Primitives

Hard classical→quantum extrapolation is implied by many natural quantum primitives, such as public-key quantum money and signatures with classical verification keys. Morally, such hard extrapolation task is simply the impossibility to sample the secret quantum state conditioned on the public classical verification key; the reduction simply formalizes this idea using standard techniques. In slightly more details,

- In the case of public-key quantum money (or even a mini-scheme), no QPT adversary who is given the public verification key vk should be able to construct a state $|\$ \rangle$ which passes $\text{Verify}(\text{vk}, |\$ \rangle)$ with noticeable probability; otherwise, they could “clone” or forge banknotes just by looking at the verification key. Since it is efficient to sample vk together with a passing banknote $|\$ \rangle$, any adversary for the classical→quantum extrapolation task defined by that

²By the fact that isometries preserve inner products, it is not hard to see that all vectors in this set are mutually orthogonal. However, this might not form a basis since the set might be empty.

³Here, we do not require that the target states $|\psi_i\rangle$ are orthogonal. For example, in public key quantum money, different serial numbers may recognize overlapping banknotes. The challenger can prepare the challenge mixed state $\sum_i \alpha_i^2 |\text{chal}_i\rangle\langle\text{chal}_i|$ in the adversary’s view by preparing Gen_λ , then cloning $|\text{chal}_i\rangle$ into an internal register.

sampling procedure forges banknotes with noticeable probability. This also extends to public-key quantum money where the public key is only cloneable instead of classical (and we get a hard cloneable→quantum extrapolation task instead), or if the money has a serial number.

- Similarly, we can construct a hard cloneable→quantum extrapolation task from signatures with cloneable verification keys. The task is again “inverting” the quantum secret state corresponding to the cloneable verification key that is used to sign additional messages. If we can create the quantum secret state with some noticeable overlap, then the forged signature would also pass verification with some noticeable probability as well.

Finally, we consider a slightly different task of quantum key distribution (QKD) with a bounded number of rounds. In QKD, two parties, Alice and Bob, have access to an insecure quantum channel and a classically authenticated channel⁴. In the protocol, they take turns to send messages over the two channels with Alice sending the first message. The malicious party, called the interceptor, can passively monitor the messages in the classical authenticated channel and can arbitrarily modify messages in the quantum channel.

There are many variants of QKD security considered in the literature. Here we consider a very weak security notion, which we note is a search-type assumption.

Definition 5.2. *Consider a QKD protocol. Let k_A, k_B, k_* be the keys Alice, Bob, and the interceptor produce correspondingly at the end of the protocol. We require the QKD protocol to be correct, meaning that without the presence of the interceptor, $\Pr[k_A = k_B \neq \perp]$ is negligibly close to 1.*

We say the QKD protocol is (statistically) unpredictably-secure if for all efficient (or unbounded, resp.) interceptors, $\Pr[k_ = k_A \neq \perp \vee k_* = k_B \neq \perp]$ is negligible.*

Fact 5.3 ([SP00, Protocol 2]). *There exists a 3-message QKD that is statistically unpredictably-secure. Furthermore, the only quantum message is sent in the first message.*

Fact 5.4 ([MW24]). *Assuming the existence of post-quantum one-way functions, there exists a 2-message QKD that is computationally unpredictably-secure.*

Fact 5.5. *1-message QKD cannot be unpredictably secure.*

Proof sketch. This follows from a straightforward observation that the interceptor could simply perform Bob’s honest action coherently and do a gentle measurement to extract Bob’s key. The measurement is gentle since Alice’s key is already determined after the first message, and by correctness, Bob should output the correct key with overwhelming probability. Finally, note that this attack is efficient. \square

Given these known results, it is natural to wonder if 2-message QKD can be statistically secure, and whether it requires computational assumptions. We observe that hard classical→quantum extrapolation task also follows from 2-message QKD and thus it requires computational assumptions somewhere between a post-quantum one-way function and a hard classical→quantum extrapolation task. In fact, we can prove something slightly stronger.

⁴Having access to a classically authenticated channel is also required classically: it is easy to see that without any authenticated channel the task is trivially impossible since the interceptor could always pretend to be Bob to Alice and pretend to be Alice to Bob. Public key infrastructure also does not get around this problem since it still assumes authenticated channels between the parties and the authority.

Proposition 5.6. *If there is an unpredictably-secure QKD where all but the last two messages are classical, then there exists a hard classical→quantum extrapolation task.*

Proof. Without loss of generality, let us say Alice receives the last message. Then consider the state immediately after Alice sends her quantum message, and the task is to produce the joint quantum state on Alice’s and Bob’s internal register and the message register conditioned on the classical transcript so far: this is a classical→quantum extrapolation task. Note that this is the state before any quantum communication reaches the other party, thus Bob’s internal registers must be unentangled from the rest conditioned on the classical transcript.

Suppose for contradiction that this is easy, then we construct the interceptor as follows:

- Upon receipt of Alice’s last classical message, discard her quantum message and produce a fresh copy of the joint quantum state and use its message register as the quantum message to send to Bob.
- Upon receipt of Bob’s messages, run Alice’s honest action to extract Bob’s key.

By correctness of the QKD protocol, whenever we succeed in producing the quantum state with some noticeable overlap (which happens with non-negligible probability by assumption), Bob will not abort the protocol and we can predict Bob’s key with roughly the same probability, breaking its security. \square

Note that as special cases, this proposition covers any 2-message QKD, or any 3-message QKD whose first message is classical.

5.3 A Quantum Bit Commitment Scheme

Let $\{\text{Rand}_k\}_{k \in \{0,1\}^{\text{poly}(n)}}$ be a family of efficiently implementable n -qubit “randomizing” unitaries, which satisfy

$$\mathbb{E}_{k \leftarrow \{0,1\}^{\text{poly}(n)}} \left[\text{TD} \left(|\psi_0\rangle\langle\psi_0|^{(k)}, |\psi_1\rangle\langle\psi_1|^{(k)} \right) \right] \leq \frac{1}{\sqrt{2}} \quad (2)$$

for any orthogonal states $|\psi_0\rangle$ and $|\psi_1\rangle$, where $|\psi\rangle\langle\psi|^{(k)}$ is mixed state resulting from applying Rand_k to $|\psi\rangle$ and measuring it in the standard basis.⁵ These exist by Theorem 4.2 or Theorem A.1. More explicitly, for $P_x^k := \text{Rand}_k |x\rangle\langle x| \text{Rand}_k^\dagger$, we define

$$|\psi\rangle\langle\psi|^{(k)} := \sum_{x \in \{0,1\}^n} P_x^k |\psi\rangle\langle\psi| P_x^k.$$

Roughly speaking, for any fixed state $|\psi\rangle$, with high probability over $k \leftarrow \{0,1\}^n$, the state $\text{Rand}_k |\psi\rangle$ is well-spread in the standard basis (i.e., measuring it gives an outcome that is *somewhat* uniform).

Let CRand denote the controlled Rand_k unitary

$$\text{CRand} := \sum_{k \in \{0,1\}^{\text{poly}(n)}} |k\rangle\langle k| \otimes \text{Rand}_k.$$

We use the notation $\text{CRand}_{A \rightarrow B}$ as shorthand for $\sum_k |k\rangle\langle k|_A \otimes (\text{Rand}_k)_B$, and we follow a similar convention for other controlled gates such as CNOT .

⁵The key may be much longer than the number of qubits operated on, for example if using random Cliffords.

Construction 5.7 (Weakly Hiding Commitment). *Let $(\text{Gen}, \text{Clone})$ be a cloneable \rightarrow quantum extrapolation task. Our weakly-binding quantum bit commitment works as follows.*

- To commit to a bit b :

1. The sender first prepares the following initial state (which is independent of b):

$$|\text{init}\rangle := |0\rangle_{C_0} \otimes |+\rangle_K \otimes \left(\sum_s \beta_s |\text{chal}_s\rangle_S |\psi_s\rangle_{M_0} \right) \otimes |0\rangle_{M_1},$$

where M_0, M_1 are both m qubit registers, and C_0 is the same as size (K, M_0) .

2. Next apply $(\text{Clone}_{S \rightarrow (S, C_1)} \cdot \text{CNOT}_{(K, M_b) \rightarrow C_0}) \cdot (\text{CRand}_{K \rightarrow M_b})$.
3. Finally, send the $C = (C_0, C_1)$ register.

- To decommit, send b along with all the remaining registers K, S, M_0, M_1 . The receiver verifies by checking that the state is the expected state

$$|\text{com}_b\rangle := (\text{Clone}_{S \rightarrow (S, C_1)} \cdot \text{CNOT}_{(K, M_b) \rightarrow C_0}) \cdot (\text{CRand}_{K \rightarrow M_b}) \cdot |\text{init}\rangle.$$

Theorem 5.8. *If $(\text{Gen}, \text{Clone})$ is a hard cloneable \rightarrow quantum extrapolation problem, then Theorem 5.7 is a $\frac{1}{\sqrt{2}}$ -statistically hiding and computationally binding commitment.*

Proof. We prove $\frac{1}{\sqrt{2}}$ -statistically hiding in Theorem 5.11 and prove computational binding in Theorem 5.12. \square

We compile the weakly-statistically hiding construction to be statistically hiding using standard XOR amplification. In detail:

Construction 5.9 (Full Commitment). *The commitment scheme works as follows.*

- To commit to a bit b :

1. The sender samples a string $x \leftarrow \{0, 1\}^\lambda$ such that the parity of x is b .
2. For each bit x_i of x , the sender commits to x_i using Theorem 5.7.

- To decommit, send b, x and the openings of the commitments to each bit of x . The verifier verifies by checking that for each $i \in [n]$, commitment i validly opens to x_i according to Theorem 5.7, then checks that the parity of x is b .

Corollary 5.10. *If $(\text{Gen}, \text{Clone})$ is a hard cloneable \rightarrow quantum extrapolation problem, then Theorem 5.9 is a statistically hiding and computationally binding commitment.*

Proof. Statistical hiding amplification follows from the fact that trace distance is amplified exponentially under XOR compositions [Wat02, Lemma 2]. Computational binding holds since in order to change the final bit revealed, the adversary must flip at least one of the λ bits committed using the original commitment scheme; by randomly guessing the index, we complete the reduction to breaking the binding security of the original commitment scheme. \square

5.4 Statistical Hiding

Claim 5.11. *Theorem 5.7 is $\frac{1}{\sqrt{2}}$ -statistically hiding.*

Proof. The state of (\mathbf{C}, \mathbf{R}) after a commitment to b is:

$$\begin{aligned} |\text{com}_0\rangle &:= \sum_{s,k,x} \beta_s 2^{-n/2} |k, \text{chal}_s, x\rangle_{\mathbf{C}} \otimes (|k, \text{chal}_s\rangle \otimes |x\rangle\langle x| \text{Rand}_k |\psi_s\rangle \otimes |0\rangle)_{\mathbf{D}}, \\ |\text{com}_1\rangle &:= \sum_{s,k,x} \beta_s 2^{-n/2} |k, \text{chal}_s, x\rangle_{\mathbf{C}} \otimes (|k, \text{chal}_s\rangle \otimes |\psi_s\rangle \otimes |x\rangle\langle x| \text{Rand}_k |0\rangle)_{\mathbf{D}}. \end{aligned}$$

After tracing out register \mathbf{D} , the mixed states on register \mathbf{C} are

$$\begin{aligned} \rho_0 &:= \sum_s \beta_s^2 |\text{chal}_s\rangle\langle \text{chal}_s| \otimes \sum_{k \in \{0,1\}^n} 2^{-n} |k\rangle\langle k| \otimes \sum_{x \in \{0,1\}^n} |\langle x | \text{Rand}_k |\psi_s\rangle|^2 |x\rangle\langle x|, \\ \rho_1 &:= \sum_s \beta_s^2 |\text{chal}_s\rangle\langle \text{chal}_s| \otimes \sum_{k \in \{0,1\}^n} 2^{-n} |k\rangle\langle k| \otimes \sum_{x \in \{0,1\}^n} |\langle x | \text{Rand}_k |0\rangle|^2 |x\rangle\langle x|. \end{aligned}$$

Denote the projections of $|\psi_s\rangle$ and $|0\rangle$, respectively, onto measurement in the basis Rand_k as

$$\begin{aligned} \rho_{0,s,k} &:= \sum_{x \in \{0,1\}^n} |\langle x | \text{Rand}_k |\psi_s\rangle|^2 |x\rangle\langle x|, \\ \rho_{1,s,k} &:= \sum_{x \in \{0,1\}^n} |\langle x | \text{Rand}_k |0\rangle|^2 |x\rangle\langle x|. \end{aligned}$$

Then the trace distance between ρ_0 and ρ_1 is bounded as

$$\frac{1}{2} \|\rho_0 - \rho_1\|_1 \leq \frac{1}{2} \sum_s \beta_s^2 \sum_{k \in \{0,1\}^n} 2^{-n} \|\rho_{0,s,k} - \rho_{1,s,k}\|_1 \quad (3)$$

$$\begin{aligned} &= \sum_s \beta_s^2 \mathbb{E}_{k \in \{0,1\}^n} \left[\frac{1}{2} \|\rho_{0,s,k} - \rho_{1,s,k}\|_1 \right] \\ &\leq \sum_s \beta_s^2 \frac{1}{\sqrt{2}} \\ &= \frac{1}{\sqrt{2}}. \end{aligned} \quad (4)$$

(3) follows from the convexity of the trace norm and its invariance under tensor product with the same state; (4) follows from (2). \square

5.5 Computational Binding

Claim 5.12. *If $(\text{Gen}, \text{Clone})$ is a hard cloneable \rightarrow quantum extrapolation problem, then Theorem 5.7 is computationally binding.*

Proof. Suppose we have an adversary Adv that can map $|\text{com}_0\rangle$ to $|\text{com}_1\rangle$ by acting only on the decommitment register $(\mathbf{K}, \mathbf{S}, \mathbf{M}_0, \mathbf{M}_1)$.

We'll use this to give a solver for the cloneable→quantum extrapolation problem. Let's think of the cloneable→quantum extrapolation problem as follows. The challenger prepares

$$|\chi_{\text{start}}\rangle := \sum_s \beta_s |\text{chal}_s\rangle_S |\psi_s\rangle_M \otimes |\text{chal}_s\rangle_{S'}$$

and gives us the S' register. To win the game, it suffices to turn this into the pure state

$$|\chi_{\text{end}}\rangle := \sum_s \beta_s |\text{chal}_s\rangle_S |\psi_s\rangle_M \otimes |\text{aux}_s\rangle_A |\psi_s\rangle_{M'}$$

for some aux_s since the M' register will also pass the verification.

The reduction, given the (S', M') register of $|\chi_{\text{start}}\rangle$, initializes a K register to the uniform superposition $|+\rangle^n$, M_1 registers to $|0\rangle$, and then applies $(\text{Clone}_{(K, S, M') \rightarrow C}) \cdot (\text{CRand}_{K \rightarrow M'})$. It applies the binding adversary on registers K, S', M', M_1 . It then coherently checks that the registers K, S', M_1 agrees⁶ with C and abort if not (in other words only the decision is measured). Finally, it returns register M' to the challenger.

Random measurements. In general, the result of applying Rand_k can be expressed as

$$\begin{aligned} \text{Rand}_k |\psi_s\rangle &= \sum_x \alpha_{k,s,x} |x\rangle, \\ \text{Rand}_k |0\rangle &= \sum_x \alpha_{k,x} |x\rangle \end{aligned}$$

for some complex coefficient α 's. Using these, we can write the commitment states as

$$\begin{aligned} |\text{com}_0\rangle &= \sum_{k,s,x} 2^{-n/2} \beta_s \alpha_{k,s,x} |k, \text{chal}_s, x\rangle_C \otimes |k, \text{chal}_s, x, 0\rangle_D, \\ |\text{com}_1\rangle &= \sum_{k,s,x} 2^{-n/2} \beta_s \alpha_{k,x} |k, \text{chal}_s, x\rangle_C \otimes |k, \text{chal}_s, \psi_s, x\rangle_D. \end{aligned}$$

Here, C contains the commitment (sent to the receiver) and D contains the opening (kept by the sender).

Binding adversary's attack. Say the adversary consists of a unitary U_A and an auxiliary quantum input $|\text{aux}\rangle_A$. It applies U_A to registers C and A . We assume the adversary always prepares a state $|k, s, \phi_{k,s,x}, x\rangle$ in register D along with potentially another private register. This is without loss of generality since not preparing a state like this form can only decrease the binding attack advantage since when the challenger projects onto $|\text{com}_1\rangle$ it necessarily projects k, s, x to the correct value; furthermore, in our reduction, we also perform this check so the amplitude outside would not affect the reduction either. More formally, we can account for this by considering the residual state to be potentially sub-normalized when the adversary does not comply. The output of the adversary is

$$|\varphi_A\rangle = \sum_{k,s,x,y} 2^{-n/2} \beta_s \alpha_{k,s,x} |k, \text{chal}_s, x\rangle_C \otimes |k, \text{chal}_s, y, x\rangle_D \otimes |\text{aux}_{k,s,y,x}\rangle_A.$$

⁶This can be done efficiently for any cloneable basis. In particular, take any efficient unitary completion of Clone and we simply compute its inverse and check if the auxiliary registers return to $|0\rangle$.

The state $|\phi_{k,s,x}\rangle := \sum_y |y\rangle |\mathbf{aux}_{k,s,y,x}\rangle$ might be subnormalized but the remaining component will never contribute (in either direction) to the binding advantage. From this we can see that the adversary's action can be seen as first swapping $M_0 M_1$ and then prepare a (possibly subnormalized) quantum state on the registers $M_0 A$ controlled on k, s, x .

We now consider the state after projecting M_0 to $|\psi_s\rangle$ to be

$$|\varphi'_A\rangle = \sum_{k,s,x} 2^{-n/2} \beta_s \alpha_{k,s,x} |k, s, x\rangle_C \otimes |k, \mathbf{chal}_s, \psi_s, x\rangle_D \otimes |\mathbf{aux}_{k,s,x}\rangle_A,$$

where $|\mathbf{aux}_{k,s,x}\rangle := \sum_y \langle\psi_s|y\rangle |\mathbf{aux}_{k,s,y,x}\rangle$. Note that this projection again contains both the commitment receiver's projection as well as the extrapolation verification projection. As a consequence, we can for simplicity consider the overlap with this state, which will be a lower bound on the success probability.

Additional notations. Let $\Pr[s] = \beta_s^2$ be the probability of sampling a challenge s . Define $p_{k,s}(x) = |\alpha_{k,s,x}|^2$ to be the probability that a measurement of $|\psi_s\rangle$ in the basis k outputs x . Similarly define $q_k(x) = |\alpha_{k,x}|^2$. $w(k, s, x)$ denotes the probability that the adversary, given $|k, \mathbf{chal}_s, x\rangle$, successfully produces $|\psi_s\rangle$. In other words,

$$w(k, s, x) = \|\mathbf{aux}_{k,s,x}\|^2.$$

Reduction success probability. Given the notation above, the reduction's success probability in the cloneable \rightarrow quantum extrapolation game is

$$\sum_{k,s} 2^{-n} \Pr[s] \sum_x q_k(x) w(k, s, x) = \mathbb{E}_{k,s} \left[\sum_x q_k(x) w(k, s, x) \right],$$

since again, in the reduction we also trimmed out the useless amplitudes.

We now show that the reduction's success probability in the cloneable \rightarrow quantum extrapolation game is at least the adversary's probability of success in the binding attack. To that end, we

compute the binding success probability to be

$$\begin{aligned} \left\| \langle \text{com}_1 | \otimes I \rangle | \varphi'_A \rangle \right\|^2 &= \left\| 2^{-n} \sum_{k,s,x} \beta_s^2 \alpha_{k,s,x} \overline{\alpha_{k,x}} | \text{aux}_{k,s,x} \rangle \right\|^2 \\ &\leq \left(2^{-n} \sum_{k,s,x} \beta_s^2 |\alpha_{k,s,x} \overline{\alpha_{k,x}}| \cdot \| | \text{aux}_{k,s,x} \rangle \| \right)^2 \end{aligned} \quad (5)$$

$$\begin{aligned} &= \mathbb{E}_{k,s} \left[\sum_x |\alpha_{k,s,x} \alpha_{k,x}| \sqrt{w(k,s,x)} \right]^2 \\ &\leq \mathbb{E}_{k,s} \left[\left(\sum_x |\alpha_{k,s,x} \alpha_{k,x}| \sqrt{w(k,s,x)} \right)^2 \right] \end{aligned} \quad (6)$$

$$\begin{aligned} &\leq \mathbb{E}_{k,s} \left[\left(\sum_x |\alpha_{k,s,x}|^2 \right) \left(\sum_x |\alpha_{k,x}|^2 w(k,s,x) \right) \right] \\ &= \mathbb{E}_{k,s} \left[\sum_x q_k(x) w(k,s,x) \right]. \end{aligned} \quad (7)$$

(5) follows from triangle inequality; (6) follows from Jensen's inequality; (7) follows from Cauchy-Schwarz inequality. Therefore, our reduction succeeds with the same probability as the binding adversary's advantage, completing the proof. \square

6 Quantum Extrapolation

In this section, we define (fully) quantum extrapolation tasks and give some initial observations.

Definition 6.1 (Quantum Extrapolation). *Let Gen be a family of efficiently preparable bipartite pure states on registers \mathbf{A}, \mathbf{B} . Let ρ be the reduced density matrix on \mathbf{A} . For every λ , define the target state to be its canonical purification*

$$|T_\lambda\rangle := (\sqrt{\rho} \otimes I) |\Phi\rangle$$

where $|\Phi\rangle_{\mathbf{A}\mathbf{B}'} = \sum_i |i\rangle |i\rangle$ is the unnormalized maximally entangled state where \mathbf{B}' has same dimension as \mathbf{A} . Gen is a hard quantum extrapolation problem if there exists a polynomial p such that for every auxiliary input aux and polynomial-time channel $V_{\mathbf{B}\mathbf{C} \rightarrow \mathbf{B}'}$,

$$F\left(|T_\lambda\rangle \langle T_\lambda|_{\mathbf{A}, \mathbf{B}'}, V_{\mathbf{B}\mathbf{C} \rightarrow \mathbf{B}'}(\text{Gen}_\lambda() \otimes \text{aux}_\lambda)\right) \leq 1 - \frac{1}{p(\lambda)}.$$

Readers might notice that this is defined differently from Theorem 1.5. We now explain why the two definitions are equivalent, or why Theorem 6.1 can be viewed as an extrapolation problem.

Lemma 6.2. *Let*

$$|\text{Gen}\rangle = \sum_i \alpha_i |A_i\rangle_{\mathbf{A}} \otimes |B_i\rangle_{\mathbf{B}}$$

be any of $|\text{Gen}\rangle$'s Schmidt decompositions. Then the target state

$$|T\rangle = \sum_i \alpha_i |A_i\rangle_A \otimes |A_i^*\rangle_{B'}$$

where $|A_i^*\rangle$ is the complex conjugate of $|A_i\rangle$.

Proof. By the Schmidt decomposition, we can see that $\rho = \sum_i \alpha_i^2 |A_i\rangle\langle A_i|$. Since $\{A_i\}_i$'s form an orthonormal basis, $\sqrt{\rho} = \sum_i \alpha_i |A_i\rangle\langle A_i|$ and $|\Phi\rangle = \sum_i |A_i\rangle \otimes |A_i^*\rangle$. The lemma follows by plugging these into the definition of $|T\rangle$. \square

Corollary 6.3. *For any Gen , there exists an isometry U such that $U_{B \rightarrow B'} |\text{Gen}\rangle = |T\rangle$.*

Proof. Consider any Schmidt decomposition as above. Take any U such that $U |B_i\rangle = |A_i^*\rangle$ for all i . This is well defined since $|B_i\rangle$'s and $|A_i^*\rangle$'s form a basis. \square

The reader may notice that the Schmidt decomposition of a state is not necessarily unique. Indeed, if we asked the adversary to map $|B_i\rangle$ to $|A_i\rangle$ directly, the optimal action would depend on the specific Schmidt decomposition that we consider. However, by requiring the adversary to instead map it to $|A_i^*\rangle$ instead, we ensure uniqueness of the target state $|T\rangle$, and thus that the quantum extrapolation problem is well-defined.

Let us briefly compare this task with classical \rightarrow quantum extrapolation. On a high level, there are three important differences:

1. Obviously, the input or the challenge is now a quantum state instead of classical bitstrings.
2. We also require the solver to solve the problem coherently, in other words, the solver cannot measure (or remember) which i is given if he wishes to succeed. This is different from classical \rightarrow quantum extrapolation where the solver is free to measure the input. This condition is important to ensure the uniqueness of the action as otherwise the action would depend on the specific Schmidt decomposition that we consider.
3. Finally, we also require the problem to be “strongly” solved, i.e. solving the task to arbitrary inverse polynomial error; whereas we can accept any non-negligible overlap for classical \rightarrow quantum extrapolation.⁷

Before proceeding, let us first establish the robustness of the quantum extrapolation task. Informally, we show that if the instance specified by Gen is slightly disturbed, then the target state will not change drastically either. As a corollary, the same optimal action will still do pretty good even if a slightly disturbed input is given instead. This is not obvious since the target state depends on Gen non-linearly. In particular, the square-root superoperator is not linear.

To prove this, we first need the following lemma on Holevo fidelity $F_H(\rho, \sigma) := \text{Tr}(\sqrt{\rho}\sqrt{\sigma})$.

Lemma 6.4 ([Hol72]). *Let ρ, σ be two density matrices. Then $1 - \sqrt{F_H(\rho, \sigma)} \leq TD(\rho, \sigma) \leq \sqrt{1 - F_H(\rho, \sigma)}$.*

Lemma 6.5. *If two quantum extrapolation tasks $|\langle \text{Gen} | \text{Gen}' \rangle| = \sqrt{1 - \varepsilon}$ have a large overlap, then their target states has noticeable overlap as well: $\langle T | T' \rangle \geq (1 - \sqrt{\varepsilon})^2$.*

⁷This difference would not be as important if one could show an amplification theorem for this task. Known techniques [BQSY24] do not apply since it is inefficient to project onto the target state.

Proof. Let ρ, σ be the reduced density matrices of Gen and Gen' respectively. By Uhlmann's theorem,

$$\sqrt{F}(\rho, \sigma) \geq |\langle \text{Gen} | \text{Gen}' \rangle| = \sqrt{1 - \varepsilon}. \quad (8)$$

Since the target states are the canonical purifications of the bipartite states, we compute the overlap of the two target states

$$\begin{aligned} \langle T' | T \rangle &= \langle \Phi | (\sqrt{\sigma} \otimes I) (\sqrt{\rho} \otimes I) | \Phi \rangle \\ &= \text{Tr}(\sqrt{\sigma} \sqrt{\rho}) \\ &= F_H(\rho, \sigma) \\ &\geq (1 - TD(\rho, \sigma))^2 \end{aligned} \quad (9)$$

$$\geq (1 - \sqrt{1 - F(\rho, \sigma)})^2 \quad (10)$$

$$\geq (1 - \sqrt{\varepsilon})^2. \quad (11)$$

(9) is by Theorem 6.4; (10) is by Fuchs–van de Graaf inequality; and finally, (11) is by (8). \square

6.1 Relations to Other Hardness Assumptions

In this section, we begin by showing that quantum extrapolation easily generalizes classical→quantum extrapolation tasks, and then move onto construction from commitments.

Proposition 6.6. *If hard classical→quantum extrapolation exists, then hard quantum extrapolation exists.*

Proof. In classical→quantum extrapolation, the challenger efficiently prepares the state

$$U |\vec{0}\rangle = \sum_{x \in \{0,1\}^\lambda} \alpha_x |A_x, x\rangle_{\mathbf{A}} \otimes |x\rangle_{\mathbf{B}}$$

and asks the adversary to find $|A_x\rangle$ given x .⁸ Observe that the complex conjugate of this state can also be prepared efficiently by conjugating each gate in U , producing

$$U^* |\vec{0}\rangle = \left(U |\vec{0}\rangle \right)^* = \sum_{x \in \{0,1\}^\lambda} \alpha_x |A_x^*, x\rangle_{\mathbf{A}} \otimes |x\rangle_{\mathbf{B}}. \quad (12)$$

Since both $|A_x^*, x\rangle$'s as well as $|x\rangle$'s are orthonormal, (12) must be a Schmidt decomposition.

Assume for contradiction that quantum extrapolation were easy, then there exists a QPT adversary operating on register \mathbf{B} of $U^* |\vec{0}\rangle$ which produces a state with $\frac{1}{2}$ fidelity to

$$\sum_{x \in \{0,1\}^\lambda} \alpha_x |A_x^*, x\rangle_{\mathbf{A}} \otimes |A_x, x\rangle_{\mathbf{B}}.$$

We now show that this adversary, when applied to register \mathbf{B} of $U |\vec{0}\rangle$, produces a state with noticeable fidelity to

$$\sum_{x \in \{0,1\}^\lambda} \alpha_x |A_x, x\rangle_{\mathbf{A}} \otimes |A_x, x\rangle_{\mathbf{B}}.$$

⁸Here, the extra copy of x in \mathbf{A} acts as a delayed measurement of x , which purifies the classical→quantum extrapolation experiment.

Thus, this contradicts the hardness of classical→quantum extrapolation.

Since $|A_x, x\rangle$ is orthogonal to $|A_{x'}, x'\rangle$ for any $x \neq x'$, there exists a unitary M mapping $|A_x, x\rangle \mapsto |A_x^*, x\rangle$. Applying M to register A maps $U|\vec{0}\rangle$ to $U^*|\vec{0}\rangle$. Then, since

$$\begin{aligned} (I_A \otimes V_{B,C}) \left(U|\vec{0}\rangle \otimes |\text{aux}\rangle_C \right) &= (M_A^\dagger M_A \otimes V_{B,C}) \left(U|\vec{0}\rangle \otimes |\text{aux}\rangle_C \right) \\ &= (M_A^\dagger \otimes I_{B,C}) (I_A \otimes V_{B,C}) (M_A \otimes I_{B,C}) \left(U|\vec{0}\rangle \otimes |\text{aux}\rangle_C \right) \\ &= (M_A^\dagger \otimes I_{B,C}) (I_A \otimes V_{B,C}) \left(U^*|\vec{0}\rangle \otimes |\text{aux}\rangle_C \right), \end{aligned}$$

the adversary produces a state with fidelity $\frac{1}{2}$ to

$$(M_A^\dagger \otimes I_{B,C}) \sum_{x \in \{0,1\}^\lambda} \alpha_x |A_x^*, x\rangle_A \otimes |A_x, x\rangle_B \otimes I_C = \sum_{x \in \{0,1\}^\lambda} \alpha_x |A_x, x\rangle_A \otimes |A_x, x\rangle_B \otimes I_C.$$

Since this state is always accepted by the classical→quantum extrapolation challenger, the reduction must succeed with probability at least $\frac{1}{2}$ as well, which is non-negligible. \square

We next prove that the hardness is also implied by commitments. This technically subsumes the proposition above, but the reduction is slightly more involved.

Lemma 6.7 (Triangle-Like Inequality for Fidelity [NC10]). *For any three states ρ , σ , and τ , the following inequality holds:*

$$\arccos(F(\rho, \tau)) \leq \arccos(F(\rho, \sigma)) + \arccos(F(\sigma, \tau)).$$

Theorem 6.8. *If quantum bit commitments exist, then hard quantum extrapolation exists.*

Proof. Without loss of generality, we can simply consider this for statistically hiding commitment schemes, which can be constructed from general quantum bit commitments [Yan22, BCQ23]. Consider the commitment states $|\text{Com}_0\rangle, |\text{Com}_1\rangle$ when committing to 0, 1 respectively. Let $\rho_b := \text{Tr}_D(|\text{Com}_b\rangle\langle\text{Com}_b|)$ be the commitment message for bit b . By statistical hiding, we have that $TD(\rho_0, \rho_1) \leq .01$ for all sufficiently large security parameters.

Let the target states $|T_b\rangle = (\sqrt{\rho_b} \otimes I) |\Phi\rangle$. On a high level, we will construct an adversary that turns $|\text{Com}_0\rangle$ to $|T_0\rangle$ using quantum extrapolation, and then pretend that $|T_0\rangle \approx |T_1\rangle$ by statistical hiding, and then undo the quantum extrapolation to turn $|T_1\rangle$ to $|\text{Com}_1\rangle$, breaking the binding property.

Assume for contradiction that quantum extrapolation were easy, then there exist efficient quantum channels \mathcal{A}_b that only act on the D register, such that $\text{Tr}(|T_b\rangle\langle T_b| \mathcal{A}_b(|\text{Com}_b\rangle\langle\text{Com}_b|)) \geq .99$ for both $b = 0, 1$. Note that this also implies the ability to efficiently invert the extrapolation: there exists efficient \mathcal{A}'_b that only acts on the D register, such that $\text{Tr}(|\text{Com}_b\rangle\langle\text{Com}_b| \mathcal{A}'_b(|T_b\rangle\langle T_b|)) \geq .99$ for both $b = 0, 1$.⁹

The adversary for breaking computational binding is simply $\mathcal{A}'_1 \circ \mathcal{A}_0$. We now analyze the success probability. After running \mathcal{A}_0 , we must have some state that has fidelity $\geq .99$ with $|T_0\rangle$. Since $|\langle T_1 | T_0 \rangle|^2 = |\langle \Phi | (\sqrt{\rho_1} \otimes I)(\sqrt{\rho_0} \otimes I) | \Phi \rangle|^2 = |\text{Tr}(\sqrt{\rho_1} \sqrt{\rho_0})|^2 = |F_H(\rho_0, \rho_1)|^2 \geq (1 - TD(\rho_0, \rho_1))^4 \geq .96$ by Theorem 6.4, this state must have fidelity $\geq .9$ with $|T_1\rangle$ by Theorem 6.7. Using Theorem 6.7 again, we get that after running \mathcal{A}'_1 , the output state has fidelity $\geq .8$ with $|\text{Com}_1\rangle$. Therefore, we break computational binding with noticeable probability for all sufficiently large security parameters, a contradiction. \square

⁹One essentially runs the unitary purification of $\mathcal{A}_b(|\text{Com}_b\rangle\langle\text{Com}_b|)$ for some sufficiently high-fidelity \mathcal{A}_b , swap its output register with the input and then uncompute the purified \mathcal{A}_b .

6.2 Polynomial-Space Upper Bound

Finally, we give an upper bound on how one could solve any quantum extrapolation task in (quantum) polynomial space, more specifically `pureUnitaryPSPACE` [MY23, Definition 2.7]. This suggests that it might be hard to unconditionally establish the hardness of a quantum extrapolation task with an explicit classical description.

Theorem 6.9. *For any quantum extrapolation task specified by `Gen`, it can be solved in quantum polynomial space and exponential time.*

Proof sketch. The algorithm essentially reduces to a special case of a succinct Uhlmann instance, which can be solved in quantum polynomial space [BEM⁺23]. To solve a quantum extrapolation task specified by `Gen`, it is equivalent to consider the fidelity-1 Uhlmann transformation problem from $|\text{Gen}\rangle$ to $|T\rangle$, or equivalently put, consider breaking the binding of a (not necessarily efficient) commitment scheme where the two commitment states are $|\text{Gen}\rangle$ and $|T\rangle$. Recall that the target state is simply the canonical purification of $|\text{Gen}\rangle$, i.e. consider the reduced density matrix $\rho := \text{Tr}_B(|\text{Gen}\rangle\langle\text{Gen}|)$, then $|T\rangle = (\sqrt{\rho} \otimes I)|\Phi\rangle$ where $|\Phi\rangle_{AB}$ is the unnormalized maximally entangled state.

We prove this by invoking Algorithmic Uhlmann’s Theorem [MY23, Theorem 7.4], which states that we can (approximately) perform the Uhlmann transformation problem in polynomial space given that we know how to (approximately) prepare each of the two states in polynomial space. Therefore, it suffices to show how to prepare the state $|T\rangle$ in polynomial space, or equivalently, to compute the amplitudes of $|T\rangle$ in `PSPACE`; and this follows from the fact that we can space-efficiently approximate the block encoding of ρ [MY23, Lemma 3.3], the square root superoperator [MY23, Lemma 3.15], as well as the product of two operators [MY23, Lemma 3.5], and from there it is easy to compute the amplitudes from the block encoding. \square

Acknowledgements

We thank Fermi Ma for his extended insightful discussions. We also thank John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, and Henry Yuen for their initial discussions on the quantum extrapolation task. Finally, we thank Tomoyuki Morimae for the reference on the 3-message QKD protocol [SP00].

This work was done in part when LQ was at Boston University and supported by DARPA under Agreement No. HR00112020023.

References

- [AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. “Pseudorandom (Function-Like) Quantum State Generators: New Definitions and Applications”. In: *Theory of Cryptography*. 2022, pp. 237–265. DOI: 10.1007/978-3-031-22318-1_9 (cit. on p. 28).
- [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 p. 1).

- [AT07] Dorit Aharonov and Amnon Ta-Shma. “Adiabatic Quantum State Generation”. In: *SIAM Journal on Computing* 37.1 (2007), pp. 47–82. DOI: 10.1137/060648829 (cit. on p. 6).
- [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. DOI: 10.4230/LIPIcs.ITCS.2023.24 (cit. on pp. 2, 24).
- [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. 25).
- [BJ24] Rishabh Batra and Rahul Jain. “Commitments are equivalent to one-way state generators”. In: *2024 IEEE 65th Symposium on Foundations of Computer Science (FOCS)*. 2024 (cit. on pp. 2, 5, 10).
- [BQSY24] John Bostanci, Luowen Qian, Nicholas Spooner, and Henry Yuen. “An Efficient Quantum Parallel Repetition Theorem and Applications”. In: *Proceedings of the 56th Annual ACM SIGACT Symposium on Theory of Computing*. 2024. DOI: 10.1145/3618260.3649603 (cit. on p. 22).
- [BS19] Zvika Brakerski and Omri Shmueli. “(Pseudo) Random Quantum States with Binary Phase”. In: *Theory of Cryptography*. 2019, pp. 229–250. DOI: 10.1007/978-3-030-36030-6_10 (cit. on p. 28).
- [CLLW16] Richard Cleve, Debbie Leung, Li Liu, and Chunhao Wang. “Near-Linear Constructions of Exact Unitary 2-Designs”. In: *Quantum Information and Computation* 16.9–10 (2016), pp. 721–756. DOI: 10.5555/3179473.3179474 (cit. on p. 29).
- [DPS04] Ivan Damgård, Thomas Pedersen, and Louis Salvail. “On the Key-Uncertainty of Quantum Ciphers and the Computational Security of One-Way Quantum Transmission”. In: *Advances in Cryptology - EUROCRYPT 2004*. 2004, pp. 91–108. DOI: 10.1007/978-3-540-24676-3_6 (cit. on p. 12).
- [GC01] Daniel Gottesman and Isaac Chuang. *Quantum Digital Signatures*. 2001. arXiv: quant-ph/0105032 (cit. on p. 3).
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. “How to construct random functions”. In: *J. ACM* 33.4 (1986), pp. 792–807. DOI: 10.1145/6490.6503 (cit. on p. 1).
- [Gol90] Oded Goldreich. “A note on computational indistinguishability”. In: *Information Processing Letters* 34.6 (1990), pp. 277–281. DOI: 10.1016/0020-0190(90)90010-U (cit. on p. 1).
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. “A Pseudo-random Generator from any One-way Function”. In: *SIAM Journal on Computing* 28.4 (1999), pp. 1364–1396. DOI: 10.1137/S0097539793244708 (cit. on p. 1).
- [Hol72] Alexander S. Holevo. “On quasiequivalence of locally normal states”. In: *Theoretical and Mathematical Physics* 13.2 (1972), pp. 1071–1082. DOI: 10.1007/BF01035528 (cit. on pp. 9, 22).

- [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. 1).
- [IL90] Russell Impagliazzo and Leonid A. Levin. “No better ways to generate hard NP instances than picking uniformly at random”. In: *Proceedings [1990] 31st Annual Symposium on Foundations of Computer Science*. 1990, 812–821 vol.2. DOI: 10.1109/FSCS.1990.89604 (cit. on p. 2).
- [Iva81] I D Ivanovic. “Geometrical description of quantal state determination”. In: *Journal of Physics A: Mathematical and General* 14.12 (1981), p. 3241. DOI: 10.1088/0305-4470/14/12/019 (cit. on pp. 8, 12).
- [Iva92] I D Ivanovic. “An inequality for the sum of entropies of unbiased quantum measurements”. In: *Journal of Physics A: Mathematical and General* 25.7 (1992), p. L363. DOI: 10.1088/0305-4470/25/7/014 (cit. on pp. 8, 12).
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. “Pseudorandom Quantum States”. In: *Advances in Cryptology – CRYPTO 2018*. 2018, pp. 126–152. DOI: 10.1007/978-3-319-96878-0_5 (cit. on p. 28).
- [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. DOI: 10.4230/LIPIcs.TQC.2021.2 (cit. on p. 1).
- [KT24a] Dakshita Khurana and Kabir Tomer. “Commitments from Quantum One-Wayness”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024, pp. 968–978. DOI: 10.1145/3618260.3649654 (cit. on pp. 2, 5, 10).
- [KT24b] Dakshita Khurana and Kabir Tomer. *Founding Quantum Cryptography on Quantum Advantage, or, Towards Cryptography from #P-Hardness*. 2024. arXiv: 2409.15248 (cit. on p. 10).
- [LR88] Michael Luby and Charles Rackoff. “How to Construct Pseudorandom Permutations from Pseudorandom Functions”. In: *SIAM Journal on Computing* 17.2 (1988), pp. 373–386. DOI: 10.1137/0217022 (cit. on p. 1).
- [MPSY24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. *Simple constructions of linear-depth t -designs and pseudorandom unitaries*. 2024. arXiv: 2404.12647 (cit. on p. 7).
- [MW24] Giulio Malavolta and Michael Walter. “Robust Quantum Public-Key Encryption with Applications to Quantum Key Distribution”. In: *Advances in Cryptology – CRYPTO 2024*. 2024, pp. 126–151. DOI: 10.1007/978-3-031-68394-7_5 (cit. on p. 15).
- [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 p. 1).

- [MY23] Tony Metger and Henry Yuen. “stateQIP = statePSPACE”. In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. 2023, pp. 1349–1356. DOI: 10.1109/FOCS57990.2023.00082. arXiv: 2301.07730v2 (cit. on p. 25).
- [MY24] Tomoyuki Morimae and Takashi Yamakawa. “One-Wayness in Quantum Cryptography”. In: *19th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2024)*. Vol. 310. 2024. DOI: 10.4230/LIPIcs.TQC.2024.4 (cit. on pp. 2, 5).
- [NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. DOI: 10.1017/CB09780511976667 (cit. on p. 24).
- [NZ24] Barak Nehoran and Mark Zhandry. “A Computational Separation Between Quantum No-Cloning and No-Telegraphing”. In: *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Vol. 287. 2024, 82:1–82:23. DOI: 10.4230/LIPIcs.ITCS.2024.82 (cit. on p. 14).
- [Rom90] John Rompel. “One-way functions are necessary and sufficient for secure signatures”. In: *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*. 1990, pp. 387–394. DOI: 10.1145/100216.100269 (cit. on p. 1).
- [SP00] Peter W. Shor and John Preskill. “Simple Proof of Security of the BB84 Quantum Key Distribution Protocol”. In: *Physical Review Letters* 85 (2 2000), pp. 441–444. DOI: 10.1103/PhysRevLett.85.441 (cit. on pp. 15, 25).
- [Wat02] John Watrous. “Limits on the power of quantum statistical zero-knowledge”. In: *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*. 2002, pp. 459–468. DOI: 10.1109/SFCS.2002.1181970 (cit. on p. 17).
- [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. 2, 6, 10, 11, 24).
- [Zha21] Mark Zhandry. “Quantum Lightning Never Strikes the Same State Twice. Or: Quantum Money from Cryptographic Assumptions”. In: *Journal of Cryptology* 34.1 (2021), p. 6. DOI: 10.1007/s00145-020-09372-x (cit. on p. 2).
- [Zha23] Mark Zhandry. “Quantum Minimalism”. Minimal Complexity Assumptions for Cryptography Workshop, Simons Institute. 2023. URL: <https://mzhandry.github.io/talks.html> (cit. on p. 4).

A Alternative choices of bases

A.1 A Counterexample for Binary Phase Bases

Without loss of generality, we are going to consider a basis to be described by a unitary followed by a complete measurement in the standard basis. In this language, a set of bases can be equivalently expressed as a set of unitaries.

One natural choice of basis, inspired by the binary phase construction of pseudorandom states [JLS18, BS19, AGQY22], is to first apply a (pseudo-)random ± 1 diagonal matrix and then measure in the Hadamard basis. Consider the counterexample where we have two states

$|\pm 0 \cdots 0\rangle = \frac{1}{\sqrt{2}}(|0^n\rangle \pm |10^{n-1}\rangle)$. (In other words, state b is $\frac{1}{\sqrt{2}}(|0^n\rangle + (-1)^b |10^{n-1}\rangle)$.) After applying a random diagonal matrix, these states are still either $|\pm 0 \cdots 0\rangle$ or $|\mp 0 \cdots 0\rangle$, and thus perfectly distinguishable when measured in the Hadamard basis.

A.2 A Counterexample for Pauli Bases

Consider two states $|0^n\rangle$ and $|1^n\rangle$. Applying a random Pauli and then measuring in the standard basis will not work in this case since with overwhelming probability at least $n/2 - o(n)$ qubits will be measured in the standard basis, thus the measurement outcome can be distinguished with overwhelming probability. This shows that 1-design cannot work since Pauli group forms a 1-design.

A.3 Statistical Hiding with 2-Designs

In this section, we are going to show that using a unitary 2-design suffices, and we can efficiently sample a unitary 2-design over n qubits uniformly from $2^{5n} - 2^{3n}$ elements and it can be implemented in quasi-linear time using Clifford gates [CLLW16].

Lemma A.1. *Let D_2 be a unitary 2-design over dimension N and D_∞ be the Haar random unitary distribution. For any two mixed states ρ_0 and ρ_1 , the expected trace distance after a random measurement drawn from D_2 is*

$$\mathbb{E}_{U \sim D_2} [\text{TD}(\rho_0^{(U)}, \rho_1^{(U)})] \leq \sqrt{\frac{N}{2(N+1)}} \cdot \text{TD}(\rho_0, \rho_1)$$

where $\rho^{(U)} := \sum_{x=0}^{N-1} |x\rangle\langle x| U \rho U^\dagger |x\rangle\langle x|$ is the projection onto the basis $\{U|x\rangle\}_x$.

Proof. Consider the normalized Jordan–Hahn decomposition of $\rho_0 - \rho_1 = \text{TD}(\rho_0, \rho_1) \cdot (P_+ - P_-)$ for two orthogonal density matrices P_+, P_- . Then we compute

$$\begin{aligned} & \mathbb{E}_{U \sim D_2} \left[\text{TD} \left(\sum_x |x\rangle\langle x| U \rho_0 U^\dagger |x\rangle\langle x|, \sum_x |x\rangle\langle x| U \rho_1 U^\dagger |x\rangle\langle x| \right) \right] \\ &= \mathbb{E}_{U \sim D_2} \left[\frac{1}{2} \sum_x \left| \langle x| U \rho_0 U^\dagger |x\rangle - \langle x| U \rho_1 U^\dagger |x\rangle \right| \right] \\ &= \mathbb{E}_{U \sim D_2} \left[\frac{1}{2} \sum_x \left| \langle x| U (\rho_0 - \rho_1) U^\dagger |x\rangle \right| \right] \\ &= \frac{N}{2} \text{TD}(\rho_0, \rho_1) \cdot \mathbb{E}_{U \sim D_2} \left[\left| \langle 0| U (P_+ - P_-) U^\dagger |0\rangle \right| \right] \\ &\leq \frac{N}{2} \text{TD}(\rho_0, \rho_1) \cdot \sqrt{\mathbb{E}_{U \sim D_2} \left[(\langle 0| U (P_+ - P_-) U^\dagger |0\rangle)^2 \right]} \\ &= \frac{N}{2} \text{TD}(\rho_0, \rho_1) \cdot \sqrt{\mathbb{E}_{U \sim D_\infty} \left[(\langle 0| U (P_+ - P_-) U^\dagger |0\rangle)^2 \right]}. \end{aligned}$$

where the inequality is due to Jensen's inequality.

It remains to calculate the last term. To that end, let $P_+ - P_- = \sum_i p_i |\psi_i\rangle\langle\psi_i|$ be its eigendecomposition since it is Hermitian. Plugging this in, we get that

$$\begin{aligned}
& \mathbb{E}_{U \sim D_\infty} \left[\left(\langle 0| U (P_+ - P_-) U^\dagger |0\rangle \right)^2 \right] \\
&= \mathbb{E}_{U \sim D_\infty} \left[\left(\langle 0| U \left(\sum_i p_i |\psi_i\rangle\langle\psi_i| \right) U^\dagger |0\rangle \right)^2 \right] \\
&= \mathbb{E}_{U \sim D_\infty} \left[\sum_i p_i^2 \left(\langle 0| U |\psi_i\rangle\langle\psi_i| U^\dagger |0\rangle \right)^2 + \sum_{i \neq j} p_i p_j \langle 0| U |\psi_i\rangle\langle\psi_i| U^\dagger |0\rangle \langle 0| U |\psi_j\rangle\langle\psi_j| U^\dagger |0\rangle \right] \\
&= \sum_i p_i^2 \cdot \frac{2}{N(N+1)} + \sum_{i \neq j} p_i p_j \cdot \frac{1}{N(N+1)} \tag{13}
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{N(N+1)} \cdot \sum_i \left(p_i^2 + p_i \sum_j p_j \right) \\
&= \frac{1}{N(N+1)} \cdot \sum_i p_i^2 \tag{14}
\end{aligned}$$

$$\leq \frac{2}{N(N+1)}. \tag{15}$$

(13) holds since by unitary invariance, it is equivalent to consider $\langle 0| U |\psi_i\rangle$ as the overlap between a Haar random state $|\psi\rangle$ and $|0\rangle$, thus $\mathbb{E}_{U \sim D_\infty} \left[\left(\langle 0| U |\psi_i\rangle\langle\psi_i| U^\dagger |0\rangle \right)^2 \right] = \mathbb{E}[|\langle 0|\psi\rangle|^4]$, where $|\psi\rangle$ is a Haar random state; on the other hand, $\mathbb{E}[|\langle 0|\psi\rangle|^4] = \frac{2}{N(N+1)}$ can be computed via the probability of measuring $|0\rangle \otimes |0\rangle$ on the maximally mixed state over the symmetric subspace Sym_2^N , which has dimension $\binom{N+1}{2}$. Similarly, $\mathbb{E}_{U \sim D_\infty} [\langle 0| U |\psi_i\rangle\langle\psi_i| U^\dagger |0\rangle \langle 0| U |\psi_j\rangle\langle\psi_j| U^\dagger |0\rangle] = \mathbb{E}[|\langle 0|\psi\rangle \langle 1|\psi\rangle|^2] = \frac{1}{N(N+1)}$. (14) holds since $\sum_j p_j = \text{Tr}(P_+ - P_-) = 0$. (15) holds since $\sum_i p_i^2 = \|P_+ - P_-\|_2^2 = \|P_+\|_2^2 + \|P_-\|_2^2 \leq \|P_+\|_1^2 + \|P_-\|_1^2 = 2$.

Putting everything together, we obtain that it is upper bounded by $TD(\rho_0, \rho_1) \cdot \sqrt{\frac{N}{2(N+1)}}$. \square

Higher t -Designs. Interestingly, the bound gets only a bit better than $1/\sqrt{2}$ if we use Haar random unitaries. Consider $\rho_0 = |0\rangle\langle 0|$ and $\rho_1 = |1\rangle\langle 1|$. The calculation starts the same as the last section except that we start with using D_∞ instead of D_2 and we avoid the step involving Jensen's inequality, so we obtain that the trace distance is

$$\frac{N}{2} \cdot \mathbb{E} \left[\left| |\langle 0|\psi\rangle|^2 - |\langle 1|\psi\rangle|^2 \right| \right].$$

We note that when $N \rightarrow \infty$, each $|\langle x|\psi\rangle|^2$ is approximately an independent draw from the exponential distribution with rate $1/N$ (also known as the Porter–Thomas distribution). Let X, Y be

two such independent random variables with PDF p , then we get that it is approximately

$$\begin{aligned}\frac{N}{2} \cdot \mathbb{E}[|X - Y|] &= N \cdot \mathbb{E}_{X \geq Y}[X - Y] \\ &= N \cdot \int_0^\infty p(y) \int_y^\infty p(x)(x - y) dx dy \\ &= \frac{1}{2}.\end{aligned}$$

This implies that going to higher t -designs does not seem to lead to a more asymptotically efficient construction.