On the Feasibility of Unclonable Encryption, and More

Prabhanjan Ananth[*] Fatih Kaleoglu[†] Xingjian Li[‡] Qipeng Liu[§]
Mark Zhandry[¶]

Abstract

Unclonable encryption, first introduced by Broadbent and Lord (TQC'20), is a one-time encryption scheme with the following security guarantee: any non-local adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ cannot simultaneously distinguish encryptions of two equal length messages. This notion is termed as unclonable indistinguishability. Prior works focused on achieving a weaker notion of unclonable encryption, where we required that any non-local adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ cannot simultaneously recover the entire message m. Seemingly innocuous, understanding the feasibility of encryption schemes satisfying unclonable indistinguishability (even for 1-bit messages) has remained elusive.

We make progress towards establishing the feasibility of unclonable encryption.

- We show that encryption schemes satisfying unclonable indistinguishability exist unconditionally in the quantum random oracle model.
- Towards understanding the necessity of oracles, we present a negative result stipulating that a large class of encryption schemes cannot satisfy unclonable indistinguishability.
- Finally, we also establish the feasibility of another closely related primitive: copy-protection for single-bit output point functions. Prior works only established the feasibility of copy-protection for multi-bit output point functions or they achieved constant security error for single-bit output point functions.

1 Introduction

Quantum information ushers in a new era for cryptography. Cryptographic constructs that are impossible to achieve classically can be realized using quantum information. In particular, the no-cloning principle of quantum mechanics has given rise to many wonderful primitives such as quantum money [Wie83] and its variants [AC12, Zha21, RS22], tamper detection [Got02], quantum copy-protection [Aar09], one-shot signatures [AGKZ20], single-decryptor encryption [GZ20, CLLZ21], secure software leasing [AL21], copy-detection [ALL+21] and many more.

[*]University of California, Santa Barbara. Email: prabhanjan@cs.ucsb.edu
[†]University of California, Santa Barbara. Email: kaleoglu@ucsb.edu
[‡]Tsinghua University. Email: lixj18@mails.tsinghua.edu.cn
[§]Simons Institute for the Theory of Computing. Email: qipengliu0@gmail.com
[¶]NTT Research & Princeton University. Email: mzhandry@gmail.com

Unclonable Encryption. Of particular interest is a primitive called unclonable encryption, introduced by Broadbent and Lord [BL20]. Roughly speaking, unclonable encryption is a one-time secure encryption scheme with *quantum* ciphertexts having the following security guarantee: any adversary given a ciphertext, modeled as a quantum state, cannot produce two (possibly entangled) states that both encode some information about the original message. This is formalized in terms of a splitting game.

A splitting adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ first has $\mathcal{A}$ receive as input an encryption of m_b, for two messages m_0 and m_1. $\mathcal{A}$ then outputs a bipartite state to $\mathcal{B}$ and $\mathcal{C}$. $\mathcal{B}$ and $\mathcal{C}$ additionally receive as input the classical decryption key and respectively output b_B and b_C. They win if $b = b_B = b_C$. Clearly, $\mathcal{A}$ could give $\mathcal{B}$ the entire ciphertext and $\mathcal{C}$ nothing, in which case $b_B = b$ but b_C would be independent of b, giving an overall winning probability of $1/2$. Security therefore requires that the splitting adversary wins with probability only negligibly larger than $1/2$. This security property, introduced by [BL20], is called *unclonable indistinguishability*. Unclonable indistinguishability clearly implies plain semantic security, as $\mathcal{A}$ could use any semantic security adversary to make a guess b_A for b, and then simply send b_A to $\mathcal{B}$ and $\mathcal{C}$, who set $b_B = b_C := b_A$.

Unclonable encryption is motivated by a few interesting applications. Firstly, unclonable encryption implies private-key quantum money. It is also useful for preventing storage attacks where malicious entities steal ciphertexts in the hope that they can decrypt them when the decryption key is compromised later. Recently, the works of [CMP20, AK21] showed that unclonable encryption implies copy-protection for a restricted class of functions with computational correctness guarantees.

Despite being a natural primitive, actually constructing unclonable encryption (even for 1-bit messages!) and justifying its security has remained elusive. Prior works [BL20, AK21] established the feasibility of unclonable encryption satisfying a weaker property simply called *unclonability*: this is modeled similar to unclonable indistinguishability, except that the message m encrypted is sampled uniformly at random and both $\mathcal{B}$ and $\mathcal{C}$ are expected to guess the entire message m. This weaker property is far less useful, and both applications listed above – preventing storage attacks and copy-protection – crucially rely on indistinguishability security. Moreover, unclonability does not on its own even imply plain semantic security, meaning the prior works must separately posit semantic security.

The following question has been left open from prior works:

Q1. Do encryption schemes satisfying unclonable indistinguishability, exist?

Copy-Protection for Point Functions. Copy-protection, first introduced by Aaronson [Aar09], is another important primitive closely related to unclonable encryption. Copy-protection is a compiler that converts a program into a quantum state that not only retains the original functionality but also satisfies the following property: a splitting adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ first has $\mathcal{A}$ receive as input a copy-protected state that can be used to compute a function f. $\mathcal{A}$ then outputs a bipartite state to $\mathcal{B}$ and $\mathcal{C}$. As part of the security guarantee, we require that both $\mathcal{B}$ and $\mathcal{C}$ should not be able to simultaneously compute f.

While copy-protection is known to be impossible for general unlearnable functions [AL21], we could still hope to achieve it for simple classes of functions. Of particular interest to us is the class of point functions. A single-bit output point function is of the form $f_y(\cdot)$: it takes as input x and outputs 1 if and only if $x = y$. One could also consider the notion of multi-bit output point functions, where the function outputs a large string, rather than 0 or 1.

Prior works [CMP20, AK21] either focus on constructing copy-protection for *multi*-bit output point functions or they construct copy-protection for single-bit output point functions with constant security, rather than optimal security, where the adversary can only do negligibly better than a trivial guess.

Yet another important question that has been left open from prior works is the following:

Q2. Does copy-protection for single-bit output point functions, with optimal security, exist?

As we will see later, the techniques used in resolving Q1 will shed light on resolving Q2. Hence, we focus on highlighting challenges in resolving Q1. The reader familiar with the challenges involved in constructing unclonable encryption could skip Section 1.1 and directly go to Section 1.2.

1.1 Achieving Unclonable Indistinguishability: Challenges

We need to achieve a *one-time* secure encryption scheme for *1-bit* messages satisfying unclonable indistinguishability: *how hard can this problem be?* Indeed one might be tempted to conclude that going from the weaker unclonability property to the stronger unclonable indistinguishability notion is a small step. The former is a search problem while the latter is a decision problem, and could hope to apply known search-to-decision reductions. As we will now explain, unfortunately this intuition is false, due both to the effects of quantum information and also to the fact that unclonable encryption involves multiple interacting adversaries.

- Recall that in an unclonable encryption scheme, the secret key is revealed to both $\mathcal{B}$ and $\mathcal{C}$. As a consequence, the secret information of any underlying cryptographic tool we use to build unclonable encryption could be revealed. For example, consider the following construction: to encrypt $m \in \{0,1\}$, compute $(r, \mathsf{PRF}(k,r) \oplus m)$, where $k \xleftarrow{\$} \{0,1\}^\lambda$ is the pseudorandom function key and $r \xleftarrow{\$} \{0,1\}^\lambda$ is a random tag. In the security experiment, the secret key, namely k, will be revealed to both $\mathcal{B}$ and $\mathcal{C}$. This restricts the type of cryptographic tools we can use to build unclonable encryption.

- Another challenge is to perform security reductions. Typically, we use the adversary to come up with a reduction that breaks a cryptographic game that is either conjectured to be or provably hard. However, this is tricky when there are two adversaries, $\mathcal{B}$ and $\mathcal{C}$. Which of the two adversaries do we use to break the underlying game? Suppose we decide to use $\mathcal{B}$ to break the game. For all we know, $\mathcal{A}$ could have simply handed over the ciphertext it received to $\mathcal{B}$ and clearly, $\mathcal{B}$ cannot be used to break the underlying game. Even worse, Alice can send a superposition of $\mathcal{B}$ getting the ciphertext and $\mathcal{C}$ receiving nothing v.s. $\mathcal{C}$ receiving the ciphertext and $\mathcal{B}$ getting nothing.

- Even if we somehow manage to achieve unclonable indistinguishability for 1-bit messages, it is a priori unclear how to achieve unclonable indistinguishability for multi-bit messages. In classical cryptography, the standard transformation goes from encryption of 1-bit messages to encryption of multi-bit messages via a hybrid argument. This type of argument fails in the setting of unclonable encryption. Let us illustrate why: suppose we encrypt a 2-bit message $m = m_1||m_2$ by encrypting 1-bit messages m_1 and m_2, denoted respectively by ρ_1 and ρ_2. This scheme is unfortunately insecure. An encryption of 11 can be (simultaneously) distinguished from an encryption of 00 by a non-local adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$: $\mathcal{A}$ can send ρ_1

to $\mathcal{B}$ and ρ_2 to $\mathcal{C}$. Since, both $\mathcal{B}$ and $\mathcal{C}$ receive the secret key, they can check whether the underlying message was 1 or 0.

- A recent result by Majenz, Schaffner and Tahmasbi [MST21] explores the difficulties in constructing unclonable encryption schemes. They show that any unclonable encryption scheme satisfying indistinguishability property needs to have ciphertexts, when represented as density matrices, with sufficiently large eigenvalues. As a consequence, it was shown that [BL20] did not satisfy unclonable-indistinguishability property. Any unclonable encryption scheme we come up with needs to overcome the hurdles set by [MST21].

We take an example below that concretely highlights some of the challenges explained above.

Example: Issues with using Extractors. For instance, we could hope to use randomness extractors. To encrypt a message m, we output $(\rho_x, c_r, \mathsf{Ext}(r,x) \oplus m)$, where ρ_x is an unclonable encryption of x satisfying the weaker unclonability property, c_r is a classical encryption of a random seed r, and Ext is an extractor using seed r. The intuition for this construction is that unclonable security implies that at least one of the two parties, say $\mathcal{B}$ cannot predict x, and therefore x has min-entropy conditioned on $\mathcal{B}$'s view. Therefore, $\mathsf{Ext}(r,x)$ extracts bits that are statistically random against $\mathcal{B}$, and thus completely hides m.

There are a few problems with this proposal. First, since $\mathcal{A}$ generates $\mathcal{B}$'s state and has access to the entire ciphertext, the conditional distribution of x given Bob's view will depend on c_r. This breaks the extractor application, since it requires r to be independent. One could hope to perform a hybrid argument to replace c_r with a random ciphertext, but this is not possible: $\mathcal{B}$ eventually learns the decryption key for c_r and would be able to distinguish such a hybrid. This example already begins to show how the usual intuition fails.

A deeper problem is that extractor definitions deal with a single party, whereas unclonable encryption has two recipient parties. To illustrate the issue, note that it is actually *not* the case that x has min-entropy against one of the parties: if $\mathcal{A}$ randomly sends the ciphertext to $\mathcal{B}$ or $\mathcal{C}$, each one of them can predict x with probability $1/2$, so the min-entropy is only 1. In such a case the extractor guarantee is meaningless. Now, in this example one can condition on the message $\mathcal{A}$ sends to $\mathcal{B}, \mathcal{C}$, and once conditioned it will in fact be the case that one of the two parties has high min-entropy. But other strategies are possible which break such a conditioning argument. For example, $\mathcal{A}$ could send messages that are in *superposition* v.s. $\mathcal{B}$ getting the ciphertext (and $\mathcal{C}$ nothing) v.s. $\mathcal{C}$ getting the ciphertext (and $\mathcal{B}$ nothing). By being in superposition, we can no longer condition on which party receives the ciphertext.

1.2 Our Results

We overcome the aforementioned challenges and make progress on addressing both questions Q1 and Q2. We start with our results on unclonable encryption before moving onto copy-protection.

Unclonable Encryption. For the first time, we establish the feasibility of unclonable encryption. Our result is in the quantum random oracle model. Specifically, we prove the following.

Theorem 1.1 (Informal)**.** *There exists an unconditionally secure one-time encryption scheme satisfying unclonable indistinguishability in the quantum random oracle model.*

Our construction is simple: we make novel use of coset states considered in recent works [CLLZ21]. However, our analysis is quite involved: among many other things, we make use of threshold projective implementation introduced by Zhandry [Zha21].

A recent work [AK21] showed a generic transformation from one-time unclonable encryption to public-key unclonable encryption[1]. By combining the above theorem with the generic transformation of [AK21], we obtain a public-key unclonable encryption satisfying the unclonable indistinguishability property.

Theorem 1.2 (Informal)**.** *Assuming the existence of post-quantum public-key encryption, there exists a post-quantum public-key encryption scheme satisfying the unclonable indistinguishability property in the quantum random oracle model.*

It is natural to understand whether we can achieve unclonable encryption in the plain model. Towards understanding this question, we show that a class of unclonable encryption schemes, that we call *deterministic* schemes, are impossible to achieve. By 'deterministic', we mean that the encryptor is a unitary U and the decryptor is $U^\dagger$. Moreover, the impossibility holds even if the encryptor and the decryptor are allowed to run in exponential time!

In more detail, we show the following.

Theorem 1.3 (Informal)**.** *There do not exist unconditionally secure deterministic one-time encryption schemes satisfying the unclonable indistinguishability property.*

In light of the fact that any classical one-time encryption scheme can be made deterministic without loss of generality[2], we find the above result to be surprising. An interesting consequence of the above result is an alternate proof that the conjugate encryption scheme of [BL20] does not satisfy unclonable indistiguishability[3]. This was originally proven by [MST21].

We can overcome the impossibility result by either devising an encryption algorithm that traces out part of the output register (in other words, performs non-unitary operations) or the encryption scheme is based on computational assumptions.

Copy-Protection for Point Functions. We also make progress on Q2. We show that there exists copy-protection for single-bit output functions with optimal security. Prior work by Coladangelo, Majenz and Poremba [CMP20] achieved a copy-protection scheme for single-bit output point functions that only achieved constant security.

We show the following.

Theorem 1.4 (Informal)**.** *There exists a copy-protection scheme for single-bit output point functions in the quantum random oracle model.*

While there are generic transformations from unclonable encryption to copy-protection for point functions explored in the prior works [CMP20, AK21], the transformations only work for multi-bit point functions. Our construction extensively makes use of the techniques for achieving unclonable encryption (Theorem 1.1). Our result takes a step closer in understanding the classes of functions for which the feasibility of copy-protection can be established.

[1] While their result demonstrates that the generic transformation preserves the unclonability property, we note that the same transformation preserves unclonable indistinguishability.

[2] We can always include the randomness used in the encryption as part of the secret key.

[3] It is easy to see why conjugate encryption of multi-bit messages is insecure. The insecurity of conjugate encryption of 1-bit messages was first established by [MST21] .

1.3 Organization

The rest of the paper is organized as follows. In Section 2, we cover all the necessary preliminaries, including Jordan's lemma, measuring success probability of a quantum adversary and the definitions of unclonable encryption schemes. Followed by Section 4, we recall coset states and their properties. We introduce a new game called "strengthened MOE games in the QROM" and prove security in this game. This part contains the main technical contribution of our paper. In Section 5, we build our unclonable encryption on the new property . In the final section (Section 6), we present our construction for copy-protection of single-output point functions. Finally, we talk about our impossibility result in Section 3.

1.4 Technical Overview

Attempts based on Wiesner States. We start by recalling the unclonable encryption scheme proposed by Broadbent and Lord [BL20]. The core idea is to encrypt a message m under a randomly chosen secret key x and encode x into an unclonable quantum state ρ_x. Intuitively, for any splitting adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$, there is no way for $\mathcal{A}$ to split ρ_x into two quantum states, such that no-communicating $\mathcal{B}$ and $\mathcal{C}$ can both recover enough information about x to decrypt $\mathsf{Enc}(x, m)$.

A well-known choice of no-cloning states is the Wiesner conjugate coding (or Wiesner states for short) [Wie83]. For a string $x = x_1 x_2 \cdots x_\lambda \in \{0,1\}^\lambda$, λ bases are chosen uniformly at random, one for each x_i. Let θ_i denote the basis for x_i. If θ_i is 0, x_i is encoded under the computational basis $\{|0\rangle, |1\rangle\}$; otherwise, x_i is encoded under the Hadamard basis $\{|+\rangle, |-\rangle\}$. The conjugate coding of x under basis θ is then denoted by $|x^\theta\rangle$. By knowing θ, one can easily recover x from the Wiesner state.

The unclonability of Wiesner states is well understood and characterized by *monogamy-of-entanglement games* (MOE games) in [TFKW13, BL20]. In the same paper, Broadbent and Lord show that no strategy wins the following MOE game[4] with probability more than 0.85^λ.

- A challenger samples uniformly at random $x, \theta \in \{0,1\}^\lambda$ and sends $|x^\theta\rangle$ to $\mathcal{A}$.
- $\mathcal{A}$ taking the input from the challenger, produces a bipartite state to $\mathcal{B}$ and $\mathcal{C}$.
- The non-communicating $\mathcal{B}$ and $\mathcal{C}$ then additionally receive the secret basis information θ and make a guess $x_\mathcal{B}, x_\mathcal{C}$ for x respectively.
- The splitting adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins the game if and only if $x_\mathcal{B} = x_\mathcal{C} = x$.

Figure 1: MOE Games for Wiesner States.

A natural attempt to construct unclonable encryption schemes is by composing a one-time pad with Wiesner states. A secret key is the basis information $\theta \in \{0,1\}^n$. An encryption algorithm takes the secret key θ and a plaintext m, it samples an $x \in \{0,1\}^n$ and outputs $m \oplus x$ together with the Wiesner conjugate coding of x, i.e. $|x^\theta\rangle$. On a high level, no split adversaries can both completely recover x, thus it is impossible for them to both recover the message m. However,

[4] This is a variant of MOE games discussed in [TFKW13]. We will use this notation throughout the paper.

such a scheme can never satisfy the stronger security: unclonable indistinguishability. Recall that unclonable indistinguishability requires either $\mathcal{B}$ or $\mathcal{C}$ can not distinguish whether the ciphertext is an encryption of message m_0 or m_1. Broadbent and Lord observe that although it is hard for $\mathcal{B}$ and $\mathcal{C}$ to recover the message completely, they can still recover half of the message and hence simultaneously distinguish with probability 1.

Towards unclonable indistinguishability, they introduce a random oracle $H : \{0,1\}^\lambda \times \{0,1\}^\lambda \to \{0,1\}^n$ in their construction (Figure 2). If an adversary can distinguish between $m_0 \oplus H(\alpha, x)$ and $m_1 \oplus H(\alpha, x)$, it must query $H(\alpha, x)$ at some point; hence, one can extract x from this adversary by measuring a random query. Following the same reasoning, one may hope to base the security (of Figure 2) on the MOE games (of Figure 1), by extracting x from both parties.

$\mathsf{Gen}(1^\lambda)$: on input λ, outputs uniformly random $(\alpha, \theta) \in \{0,1\}^{2\lambda}$.

$\mathsf{Enc}^H((\alpha,\theta), m)$: samples $x \in \{0,1\}^\lambda$, outputs $(|x^\theta\rangle, m \oplus H(\alpha, x))$.

$\mathsf{Dec}^H((\alpha,\theta), (|x^\theta\rangle, c))$: recovers x from $|x^\theta\rangle$, outputs $c \oplus H(\alpha, x)$.

Figure 2: Unclonable Encryption by Broadbent and Lord.

The above idea, thought intuitive, is hard to instantiate. It will require simultaneous extraction of the secret x from both $\mathcal{B}$ and $\mathcal{C}$. Since $\mathcal{B}$ and $\mathcal{C}$ can be highly entangled, a successful extraction of x on $\mathcal{B}$'s register may always result in an extraction failure on the other register. Broadbent and Lord use a "simultaneous" variant of the so-called "O2H" (one-way-to-hiding) lemma [Unr15] to prove their scheme satisfies unclonable indistinguishability for un-entangled adversaries $\mathcal{B}, \mathcal{C}$, or messages with constant length. The unclonable indistinguishability for general adversaries and message spaces remains quite unknown.

Even worse, Majenz, Schaffner, and Tahmasbi [MST21] show that there is an inherent limitation to this simultaneous variant of O2H lemma. They give an explicit example that shatters the hope of proving unclonable indistinguishability of the construction in [BL20] using this lemma.

Instantiating [BL20] using Coset States. Facing the above barrier, we may resort to other states possessing some forms of unclonability. One candidate is the so-called "coset states", first proposed by Vidick and Zhang [VZ21] in the context of proofs of quantum knowledge and later studied by Coladangelo, Liu, Liu, and Zhandry [CLLZ21] for copy-protection schemes.

A coset state is described by three parameters: a subspace $A \subseteq \mathbb{F}_2^\lambda$ of dimension $\lambda/2$ and two vectors $s, s' \in \mathbb{F}_2^\lambda$ denoting two cosets $A + s$ and $A^\perp + s'$[5]($A^\perp$ denotes the dual subspace of A); we write the state as $|A_{s,s'}\rangle$. Coset states have many nice properties, among those we only need the following:

1. Given $|A_{s,s'}\rangle$ and a classical description of subspace A, an efficient quantum algorithm can compute both s and s'.

[5]There are many vectors in $A + s$. In the rest of the discussion, we assume s is the lexicographically smallest vector in $A + s$. Similarly for s'.

2. No adversary can win the MOE game (Figure 3) for coset states with probability more than $\sqrt{e} \cdot (\cos(\pi/8))^\lambda$ (first proved in [CLLZ21]).

- A challenger samples uniformly at random a subspace $A \subseteq \mathbb{F}_2^\lambda$ of dimension $\lambda/2$ $s, s' \in \mathbb{F}_2^\lambda$ and sends $|A_{s,s'}\rangle$ to $\mathcal{A}$.
- $\mathcal{A}$ taking the input from the challenger, produces a bipartite state to $\mathcal{B}$ and $\mathcal{C}$.
- The non-communicating $\mathcal{B}$ and $\mathcal{C}$ then additionally receive a classical description of the subspace A and make a guess $s_\mathcal{B}, s'_\mathcal{B}, s_\mathcal{C}, s'_\mathcal{C}$ for s, s' respectively.
- The splitting adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins the game if and only if $s_\mathcal{B} = s_\mathcal{C} = s, s'_\mathcal{B} = s'_\mathcal{C} = s'$.

Figure 3: MOE Games for Coset States.

Readers may already notice the similarity between Wiesner states and coset states. If we substitute the basis information θ with A and the secret x with $s||s'$, we get coset states and their corresponding MOE games. Hence, we can translate the construction in [BL20] using the languages of coset states. A question naturally arises: if these two kinds of states are very similar, why does replacing Wiesner states with coset states even matter?

Indeed, they differ in one crucial place. Let us come back to Wiesner states. As shown by [Lut10] in the setting of private key quantum money, given $|x^\theta\rangle$ together with an oracle P_x[6] that outputs 1 only if input $y = x$, there exists an efficient quantum adversary that learns x without knowing θ. This further applies to the MOE games for Wiesner states: if $\mathcal{A}$ additionally gets oracle access to P_x, the MOE game is no longer secure.

MOE games for coset states remain secure if oracles for checking s and s' are given. More formally, let P_{A+s} be an oracle that outputs 1 only if the input $y \in A + s$, similarly for $P_{A^\perp+s'}$. No adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ can win the MOE games for coset states with more than some exponentially small probability in λ, even if $\mathcal{A}, \mathcal{B}, \mathcal{C}$ all query P_{A+s} and $P_{A^\perp+s'}$ polynomially many times. We call this game *MOE game for coset states with membership checking oracles*.

We now give our construction of unclonable encryption that satisfies unclonable indistinguishability in Figure 4. In our construction, we also get rid of the extra input α in [BL20] construction. We believe α can be similarly removed in their construction as well. Also, note that in our construction, we only require coset states and random oracles. The membership checking oracles will only be given to the adversary when we prove its security. Thus, we prove a stronger security guarantee (with membership checking oracle are given). Due to this, we can not prove the security of their construction using Wiesner states following the same idea; nonetheless, we do not know how to disprove it. We leave it as an interesting open question.

[6] [Lut10] showed that an algorithm breaks the money scheme, given oracle access to P_x^θ; P_x^θ outputs 1 if and only if input $y = x$ under basis specified by θ. One can change the algorithm so that it only needs P_x to break the money scheme.

$\mathsf{Gen}(1^\lambda)$: on input λ, outputs uniformly random subspace $A \subseteq \mathbb{F}_2^\lambda$ of dimension $\lambda/2$.

$\mathsf{Enc}^H(A, m)$: samples $s, s' \in \mathbb{F}_2^\lambda$[a], outputs $(|A_{s,s'}\rangle, m \oplus H(s, s'))$.

$\mathsf{Dec}^H(A, (|A_{s,s'}\rangle, c))$: recovers s, s' from the coset state, outputs $c \oplus H(s, s')$.

[a] We again require s, s' to be the lexicographically smallest vector in $A + s$ and $A^\perp + s'$.

Figure 4: Our Unclonable Encryption Scheme.

Basing Security on Reprogramming Games. Now we look at what property we require for coset states to establish unclonable indistinguishability. We will focus on the case $n = 1$ (length-1 messages) in this section. By a sequence of standard variable substitutions, unclonable indistinguishability of our scheme can be based on the following security game in the identical challenge mode (please refer to Figure 5), where each of $\mathcal{B}, \mathcal{C}$ tries to identify whether the oracle has been reprogrammed or not. We want to show any adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ only achieves successful probability $1/2 + \mathsf{negl}$. This ideal security matches the trivial attack: $\mathcal{B}$ gets the coset state and $\mathcal{C}$ makes a random guess, they win with probability $1/2$.

Note that in the above reprogramming game (Figure 5), $\mathcal{A}$ has no access to $H(s, s')$. This is different from unclonable indistinguishability games or MOE games. Nevertheless, we show that $\mathcal{A}$ never queries (s, s') and thus $H(s, s')$ does not help $\mathcal{A}$ and thus can be safely removed by introducing a small loss.

The security of the reprogramming games in the identical challenge mode can be reduced to the security in the independent challenge mode. A careful analysis of Jordan's lemma (Section 2.3) is required to show such a reduction. We believe that this reduction is non-trivial and we leave it to the last section in the overview.

The remaining is to show the security of the game in the independent challenge mode. Inspired by the work of [Zha20] which initiates the study of measuring success probability of a quantum program, we show there is an efficient procedure that operates locally on both the entangled adversaries $(\mathcal{B}, \mathcal{C})$ and outputs $(\mathcal{B}', p_{\mathcal{B}})$, $(\mathcal{C}', p_{\mathcal{C}})$ such that: (informally)

- $\mathcal{B}'$ and $\mathcal{C}'$ are un-entangled[7].
- The success probability of $\mathcal{B}'$ on guessing whether it has access to H_0 or H_1 is $p_{\mathcal{B}}$.
- The success probability of $\mathcal{C}'$ on guessing whether it has access to H_0 or H_1 is $p_{\mathcal{C}}$.
- The expectation of $p_{\mathcal{B}} \cdot p_{\mathcal{C}}$ is equal to $(\mathcal{B}, \mathcal{C})$'s success probability in the reprogramming game in the independent challenge mode.

The above estimation procedure requires to run $\mathcal{B}'$ and $\mathcal{C}'$ on H_0 and H_1. In other words, the procedure should be able to reprogram $H_{(s,s')\to\perp}$ on the input (s, s'). Since the procedure will be used in the reduction for breaking MOE games for coset states, it should not know s or s', but only

[7] $\mathcal{B}'$ and $\mathcal{C}'$ satisfy a weaker guarantee than being un-entangled. Informally, conditioned on any event of non-negligible chance on one's side, the other party still has success probability $p_{\mathcal{C}}$ (or $p_{\mathcal{B}}$, respectively). The same analysis applies to this weaker guarantee. For ease of presentation, we assume that they are un-entangled.

- H be a random oracle with binary range, $H : \mathbb{F}_2^\lambda \times \mathbb{F}_2^\lambda \to \{0,1\}$.
  Additionally, $\mathcal{A}, \mathcal{B}, \mathcal{C}$ get oracle access to P_{A+s} and $P_{A^\perp+s'}$.
- A challenger samples a coset state $|A_{s,s'}\rangle$ and sends $(|A_{s,s'}\rangle, H(s,s'))$ to $\mathcal{A}$.
- $\mathcal{A}$ taking the input from the challenger, has oracle access to $H_{(s,s')\to\perp}$ and produces a bipartite state to $\mathcal{B}$ and $\mathcal{C}$. Here $H_{(s,s')\to\perp}$ is the same as H except $H(s,s')$ is replaced with $\perp$[a].
- The non-communicating $\mathcal{B}$ and $\mathcal{C}$ then receive a classical description of the subspace A:
  - Let $H_0 := H$ be the original random oracle.
  - Let H_1 be identical to H, except the outcome on (s,s') is flipped.
  - (**Identical Challenge Mode**): Flip a coin b, both $\mathcal{B}$ and $\mathcal{C}$ get oracle access to H_b.
  - (**Independent Challenge Mode**): Flip two coins $b_\mathcal{B}, b_\mathcal{C}$, $\mathcal{B}$ has oracle access to $H_{b_\mathcal{B}}$ and $\mathcal{C}$ gets oracle access to $H_{b_\mathcal{C}}$.
- $\mathcal{B}, \mathcal{C}$ makes a guess b', b'' respectively.
- The adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins the game if and only if $b' = b'' = b$ (in the identical challenge mode), or $b' = b_\mathcal{B}$ and $b'' = b_\mathcal{C}$ (in the independent challenge mode).

[a] In the actual proof, $H(s,s')$ is replaced with a uniformly random u. Both approaches work.

Figure 5: Reprogramming Games for Coset States in the QROM

knows A and $P_{A+s}, P_{A^\perp+s'}$. Nonetheless, we show with the membership checking oracle, such reprogramming is possible. For example, H_1 can be reprogrammed as follows:

$$H_1 = \begin{cases} \neg H(s,s') & Q_s(z) = 1 \text{ and } Q_{s'}(z') = 1 \\ H_{(s,s')\to\perp}(z,z') & \text{Otherwise} \end{cases},$$

where Q_s is the point function that only outputs 1 on s, similarly for $Q_{s'}$. The remaining is to show Q_s (or $Q_{s'}$) can be instantiated by the classical description of A and P_{A+s} (or $P_{A^\perp+s'}$ respectively). Q_s can be implemented by (1) check if the input z is in $A+s$, (2) check if the input z is the lexicographically smallest in $A+s$. Step (1) can be done via P_{A+s}. Step (2) can be done by knowing A and some $z \in A+s$ (which is known from step (1)): one can check if there exists some lexicographically smaller z^* such that $(z - z^*) \in \mathsf{span}(A)$; this can be done efficiently, by enumerating each coordinate and doing Gaussian elimination. Thus, both Q_s and $Q_{s'}$ can be implemented.

Without membership checking oracle, we do not know how to reprogram the oracle, or run the above procedure. Thus the proof fails for Wiesner states.

Finally, we prove the security of reprogramming game in the independent challenge mode. If $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ has non-trivial success probability $1/2 + \gamma$ for some large γ, the above procedure must

output large $p_{\mathcal{B}}, p_{\mathcal{C}} > 1/2 + \gamma/2$ with non-negligible probability. If $\mathcal{B}'$ never queries H_0 or H_1 on (s, s'), the best probability it can achieve is $1/2$. Thus, by measuring a random query of $\mathcal{B}'$, we can extract s, s' with non-negligible probability. Similarly for $\mathcal{C}'$. This violates the MOE games for coset states with membership checking oracles, a contradiction. Therefore, the security of the reprogramming game in the independent mode is established.

Relating Identical Challenge Mode to Independent Challenge Mode. In the end, in this section, we discuss how to relate the reprogramming game in the identical challenge mode to that in the independent challenge mode. We refer the readers to the proof of Theorem 4.8 for further details.

We first elaborate on the above discussion for *independent challenge mode*. It helps us establish the language for the presentation of *identical challenge mode* and give a nice characterization of the state produced by Alice.

For a random choice of A, s, s' and oracles $H_{(s,s')\to\perp}$, let $|\sigma\rangle_{\mathbf{BC}}$ be the joint quantum state shared by Bob and Charlie after Alice's stage. We additionally define projections Π_b^B and Π_b^C for $b \in \{0, 1\}$:

- Π_0^B: Run Bob on its own register $\sigma[\mathbf{B}]$ with oracle access to H_0, project onto Bob outputting 0 and rewind;
- Π_1^B: Run Bob on $\sigma[\mathbf{B}]$ with oracle access to H_1, project onto Bob outputting 1 and rewind.

We can similarly define Π_0^C and Π_1^C. Namely, Π_b^B is the projection for Bob's success on H_b and Π_b^C is the projection for Charlie's success on H_b.

By definition, the success probability in the independent challenge mode is:

$$\mathsf{Tr}\left[\left(\frac{\Pi_0^B + \Pi_1^B}{2}\right) \otimes \left(\frac{\Pi_0^C + \Pi_1^C}{2}\right) |\sigma\rangle \langle\sigma|\right]. \tag{1}$$

Since $(\Pi_0^B + \Pi_1^B)/2$ is a POVM, let $\{|\phi_p\rangle\}_{p\in\mathbb{R}}$ be the set of eigenvectors with eigenvalues $p \in [0, 1]$[8]. Similarly, let $\{|\psi_q\rangle\}_{q\in\mathbb{R}}$ be the set of eigenvectors with eigenvalues $q \in [0, 1]$ for $(\Pi_0^C + \Pi_1^C)/2$. Therefore, we can write $|\sigma\rangle$ under the bases $\{|\phi_p\rangle\}$ and $\{|\psi_q\rangle\}$:

$$|\sigma\rangle = \sum_{p,q} \alpha_{p,q} |\phi_p\rangle |\psi_q\rangle .$$

The analysis in the last paragraph (for independent challenge mode) can show in this setting that, p and q cannot be simultaneously far away from the trivial guessing probability $1/2$, i.e., for any inverse polynomial ε,

$$\sum_{\substack{p:|p-1/2|>\varepsilon \\ q:|q-1/2|>\varepsilon}} |\alpha_{p,q}|^2 \approx 0.$$

In other words, $|\sigma\rangle$ is very close to the summation of the following subnormalized states:

$$|\sigma\rangle = \sum_{p:|p-1/2|\le\varepsilon} \alpha_{p,q} |\phi_p\rangle |\psi_q\rangle + \sum_{\substack{p:|p-1/2|>\varepsilon \\ q:|q-1/2|\le\varepsilon}} \alpha_{p,q} |\phi_p\rangle |\psi_q\rangle .$$

[8]There can be multiple eigenvectors with the same eigenvalues. In the overview, we assume that eigenvalues are unique.

Here we simply call the first subnormalized state as $|\sigma_{\mathcal{B}}^{\mathsf{bad}}\rangle$, denoting Bob can not behave in a significantly different way from random guessing; and call second subnormalized state as $|\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle$ for Charlie. We have $|\sigma\rangle = |\sigma_{\mathcal{B}}^{\mathsf{bad}}\rangle + |\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle$. Thus, 1 is bounded by at most $1/2 + \varepsilon$ for any inverse polynomial ε, concluding the security in the independent challenge mode.

The above analysis gives a characterization of $|\sigma\rangle$. Note that although the analysis is done assuming Alice, Bob and Charlie play the game in the independent challenge mode, it holds for the game in identical challenge mode as well.

Finally, we focus on the identical challenge mode. The success probability in the identical challenge mode is:

$$\mathsf{Tr}\left[\left(\frac{\Pi_0^B \otimes \Pi_0^C + \Pi_1^B \otimes \Pi_1^C}{2}\right)|\sigma\rangle\langle\sigma|\right]. \tag{2}$$

By plugging $|\sigma\rangle = |\sigma_{\mathcal{B}}^{\mathsf{bad}}\rangle + |\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle$ in the above formula, 2 is at most:

$$\frac{1}{2} + \varepsilon + \frac{1}{2}\left(\left|\langle\sigma_{\mathcal{B}}^{\mathsf{bad}}|\Pi_0^B \otimes \Pi_0^C|\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle\right| + \left|\langle\sigma_{\mathcal{B}}^{\mathsf{bad}}|\Pi_1^B \otimes \Pi_1^C|\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle\right|\right). \tag{3}$$

The only difference between 2 and 1 is the cross terms $\left|\langle\sigma_{\mathcal{B}}^{\mathsf{bad}}|\Pi_b^B \otimes \Pi_b^C|\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle\right|$, for $b \in \{0,1\}$. Perhaps surprisingly, we prove that the cross terms are **zero**. To show it, we prove a corollary of Jordan's lemma (see Corollary 2.4) that for any *two* projections Π_0, Π_1, let $|\phi_p\rangle$ be the set of eigenvectors for $(\Pi_0 + \Pi_1)/2$; if $p + q \neq 1$ and $p \neq q$, then their cross terms $\langle\phi_p|\Pi_0|\phi_q\rangle = \langle\phi_p|\Pi_1|\phi_q\rangle = 0$. Applying this corollary to 3, we can show that $\left|\langle\sigma_{\mathcal{B}}^{\mathsf{bad}}|\Pi_b^B \otimes \Pi_b^C|\sigma_{\mathcal{C}}^{\mathsf{bad}}\rangle\right| = 0$ for both $b \in \{0,1\}$. Therefore, we conclude the security in the identical challenge mode.

1.5 Related Work

Unclonable Encryption. Broadbent and Lord [BL20] demonstrated the feasibility of unclonable encryption satisfying the weaker unclonability property. They present two constructions. The first construction based on Wiesner states achieve 0.85^n-security (i.e., the probability that both $\mathcal{B}$ and $\mathcal{C}$ simultaneously guess the message is at most 0.85^n), where n is the length of the message being encrypted. Their second construction, in the quantum random oracle model, achieves $\frac{9}{2^n} + \mathsf{negl}(\lambda)$-security. In the same work, they show that any construction satisfying 2^{-n}-unclonability implies unclonable indistinguishability property. Following Broadbent and Lord, Ananth and Kaleoglu [AK21] construct public-key and private-key unclonable encryption schemes from computational assumptions. Even [AK21] only achieve unclonable encryption with the weaker unclonability guarantees.

Majenz, Schaffner and Tahmasbi [MST21] explore the difficulties in constructing unclonable encryption schemes. In particular, they show that any scheme achieving unclonable indistinguishability should have ciphertexts with large eigenvalues. Towards demonstrating a better bound for unclonability, they also showed inherent limitations in the proof technique of Broadbent and Lord.

Copy-Protection. Copy-protection was first introduced by Aaronson [Aar09]. Recently, Aaronson, Liu, Liu, Zhandry and Zhang [ALL+21] demonstrated the existence of copy-protection in the presence of classical oracles. Coladangelo, Majenz and Poremba [CMP20] showed that copy-protection for multi-bit output point functions exists in the quantum random oracle model. They

also showed that copy-protection for single-bit output point functions exists in the quantum random oracle model with constant security.

Ananth and La Placa [AL21] showed a conditional result that copy-protection for arbitrary unlearnable functions, without the use of any oracles, does not exist. Recently, Coladangelo, Liu, Liu and Zhandry [CLLZ21], assuming post-quantum indistinguishability obfuscation and one-way functions, demonstrated the first feasibility of copy-protection for a non-trivial class of functions (namely, pseudorandom functions) in the plain model. Another recent work by Broadbent, Jeffrey, Lord, Podder and Sundaram [BJL+21] studies copy-protection for a novel (but weaker) variant of copy-protection.

Acknowledgements

Qipeng Liu is supported in part by the Simons Institute for the Theory of Computing, through a Quantum Postdoctoral Fellowship and by the DARPA SIEVE-VESPA grant Np.HR00112020023. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

Mark Zhandry is supported in part by an NSF CAREER award.

2 Preliminaries

2.1 Basics

We will briefly introduce some basic notations in our work and some preliminaries on quantum computing in this section.

We denote by λ the security parameter. We write $\mathsf{poly}(\cdot)$ to denote an arbitrary polynomial and $\mathsf{negl}(\cdot)$ to denote an arbitrary negligible function. We say that an event happens with *overwhelming probability* if the probability is at least $1 - \mathsf{negl}(\lambda)$.

Readers unfamiliar with quantum computation and quantum information could refer to [NC10] for a comprehensive introduction.

Given Hilbert space $\mathcal{H}$, we write $\mathcal{S}(\mathcal{H})$ for the unit sphere set $\{x : ||x||_2 = 1\}$ in $\mathcal{H}$, $\mathcal{U}(\mathcal{H})$ for the set of unitaries acting on Hilbert space $\mathcal{H}$, $\mathcal{D}(\mathcal{H})$ for the set of density operators on $\mathcal{H}$. We write $\mathcal{H}_X$ to denote the Hilbert space associated with a quantum register X. Given two quantum states ρ, σ, we denote the (normalized) trace distance between them by

$$\mathsf{TD}(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_{\mathsf{tr}} .$$

We say that two states ρ, σ are δ-*close* if $\mathsf{TD}(\rho, \sigma) \leq \delta$.

A positive operator-valued measurement (POVM) on the Hilbert space $\mathcal{H}$ is defined as a set of positive semidefinite operators $\{E_i\}$ on $\mathcal{H}$ that satisfies $\sum_i E_i = I$. A projective measurement means the case that E_is are projectors.

A common technique in quantum computation is uncomputing [BBBV97]. A quantum algorithm could be modeled as a unitary U acting on some hilbert space $\mathcal{H}$, then perform measurement on output registers on without loss of generality. By uncomputation we mean that acting $U^\dagger$ on

the same hilbert space after the measurement. It is easy to examine that if the measurement outputs same result with overwhelming probability, the trace distance between the final state and the original state is negligible.

Quantum Oracle Algorithms A quantum oracle for a function f is defined as the controlled unitary O_f: $O_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$. We define a query to the quantum oracle as applying O_f on the given quantum state once.

We say that a quantum adversary $\mathcal{A}$ with access to oracle(s) is *query-bounded* if it makes at most $p(\lambda)$ queries to each oracle for some polynomial $p(\cdot)$.

2.2 Quantum Random Oracle Model (QROM)

This is the quantum analogue of Random Oracle Model, where we model a hash function H as a random classical function, and it can be accessed by an adversary in superposition, modeled by the unitary O_H.

The following theorem, paraphrased from [BBBV97], will be used for reprogramming oracles without adversarial detection on inputs which are not queried with large weight:

Theorem 2.1 ([BBBV97]). *Let $\mathcal{A}$ be an adversary with oracle access to $H : \{0,1\}^m \to \{0,1\}^n$ that makes at most T queries. Define $|\phi_i\rangle$ as the global state after $\mathcal{A}$ makes i queries, and $W_y(|\phi_i\rangle)$ as the sum of squared amplitudes in $|\phi_i\rangle$ of terms in which $\mathcal{A}$ queries H on input y. Let $\epsilon > 0$ and let $F \subseteq [0, T-1] \times \{0,1\}^m$ be a set of time-string pairs such that $\sum_{(i,y)\in F} W_y(|\phi_i\rangle) \leq \epsilon^2/T$.*

Let H' be an oracle obtained by reprogramming H on inputs $(i, y) \in F$ to arbitrary outputs. Define $|\phi'_i\rangle$ as above for H'. Then, $\mathsf{TD}(|\phi_T\rangle, |\phi'_T\rangle) \leq \epsilon/2$.

Note that the theorem can be straightforwardly generalized to mixed states by convexity.

2.3 More on Jordan's lemma

We first recall the following version of Jordan's lemma, adapted from [Reg05] and [Vid21]:

Lemma 2.2. *Let $w \in [0,1]$, $\mathcal{H}$ be a finite-dimensional Hilbert space and let Π_0, Π_1 be any two projectors in $\mathcal{H}$, then there exists an orthogonal decomposition of $\mathcal{H}$ into one-dimensional and two dimensional subspaces $\mathcal{H} = \oplus_i \mathcal{S}_i$ that are invariant under both Π_0 and Π_1; each $\mathcal{S}_i$ is spanned by one or two eigenvectors of $w\Pi_0 + (1-w)\Pi_1$.*

Whenever $\mathcal{S}_i$ is 2-dimensional, there is a basis for it in which Π_0 and Π_1 (restricting on $\mathcal{S}_i$) take the form:

$$\Pi_{0,\mathcal{S}_i} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \qquad \text{and} \qquad \Pi_{1,\mathcal{S}_i} = \begin{pmatrix} c_i^2 & c_i s_i \\ c_i s_i & s_i^2 \end{pmatrix},$$

where $c_i = \cos\theta_i$ and $s_i = \sin\theta_i$ for some principal angle $\theta_i \in [0, \pi/2]$.

Proof. The proof for the case $w = 1/2$ can be found in the references above, and the generalization is straightforward. □

We additionally show a relation between two eigenvalues in the same Jordan block.

Lemma 2.3. *For any two projectors Π_0, Π_1, let $\mathcal{S}_i$ be a 2-dimensional subspace in the above decomposition. Let $|\phi_0\rangle, |\phi_1\rangle$ be two eigenvectors of $w\Pi_0 + (1-w)\Pi_1$ that span $\mathcal{S}_i$ and λ_0, λ_1 be their eigenvalues. We have $\lambda_0 + \lambda_1 = 1$.*

Proof. Restricting on $\mathcal{S}_i$, we have:

$$\lambda_0 + \lambda_1 = \mathrm{Tr}\left[(\Pi_{0,\mathcal{S}_i} + \Pi_{1,\mathcal{S}_i})/2\right] = (1 + c_i^2 + s_i^2)/2 = 1.$$

□

Corollary 2.4. *For any two projectors Π_0, Π_1, let $|\phi_0\rangle$ and $|\phi_1\rangle$ be two eigenvectors of $w\Pi_0 + (1-w)\Pi_1$ with eigenvalues λ_0, λ_1. If $\lambda_0 + \lambda_1 \neq 1$ and $\lambda_0 \neq \lambda_1$, then*

$$\langle\phi_0|\Pi_0|\phi_1\rangle = \langle\phi_0|\Pi_1|\phi_1\rangle = 0.$$

Proof. If $\lambda_0 + \lambda_1 \neq 1$, by Lemma 2.3, $|\phi_0\rangle$ and $|\phi_1\rangle$ cannot be in the same Jordan block. Because $|\phi_0\rangle$ still belongs to the corresponding subspace $\mathcal{S}_0$ of its Jordan block after the action of Π_0, $\Pi_0|\phi_0\rangle$ is orthogonal to $|\phi_1\rangle$. Similarly, $\Pi_1|\phi_0\rangle$ is orthogonal to $|\phi_1\rangle$. □

2.4 Measuring Success Probability

In this section, we give preliminaries on how to measure success probability of quantum programs (with respect to a test distribution). Part of this section is taken verbatim from [ALL+21, CLLZ21]. Since this section will only be used for proving the strengthened monogamy-of-entanglement game of coset states in the quantum random oracle model (see Section 4), the reader can safely skip it to view our construction first, and return to this section when understanding the proof of the strengthened MOE game.

In classical cryptography, we are often interested in the success probability of a given program with respect to a test distribution. Assume that the test distribution is known to everyone and can be efficiently sampled, one can efficiently estimate the success probability of a given program within any inverse polynomial error. The estimating algorithm is fairly simple: just run the programs multiple times and output how many times the program succeeds. However, this method does not quite work when quantum programs are taken into account. One crucial reason is that the estimation algorithm only gets a single copy of the program. It is in general impossible to run the program multiple times without rewinding. However, rewinding a quantum program appears to be one of the difficulties in quantum cryptography. We refer the reader to [Zha20] for a more in-depth discussion.

Measure Probability. In [Zha20], Zhandry formalizes a measurement operator for estimating the success probability of a quantum program. This operator is inefficient to implement, but Zhandry also shows how to efficiently estimate the probability with large statistical confidence in the same work (following the idea in QMA amplification [MW05]). We will discuss the efficient measurement procedure later in this section.

The starting point is that a binary POVM specifies the probability distribution over outcomes $\{0,1\}$ ("success" or "failure") on any quantum program, but it does not uniquely determine the post-measurement state. Zhandry shows that, for any binary POVM $\mathcal{P} = (P, I-P)$, there exists a

nice projective measurement such that the post-measurement state is an eigenvector of P. In particular, Zhandry observes that there exists a projective measurement $\mathcal{E}$ which *measures* the success probability of a state with respect to $\mathcal{P}$. More precisely,

- $\mathcal{E}$ outputs a probability $p \in [0,1]$ from the set of eigenvalues of P. (We stress that $\mathcal{E}$ actually outputs a real number p).
- The post-measurement state upon obtaining outcome p is an *eigenvector* of P with eigenvalue p; it is also an eigenvector of $Q = I - P$ with eigenvalue $1 - p$.

Note that since $\mathcal{E}$ is projective, we are guaranteed that applying the same measurement again on the leftover state will yield the same outcome. Thus, what we obtain from applying $\mathcal{E}$ is a state with a "well-defined" success probability with respect to $\mathcal{P}$.

Furthermore, $\mathcal{E}$ is compatible with $\mathcal{P}$. In other words, one can safely measure the success probability of a program without disturbing the overall success probability. We now give the formal theorem statement.

Theorem 2.5 (Inefficient Measurement)**.** *Let $\mathcal{P} = (P, Q)$ be a binary outcome POVM. Let $\mathcal{D}$ be the set of eigenvalues of P. There exists a projective measurement $\mathcal{E} = \{E_p\}_{p \in \mathcal{D}}$ with index set $\mathcal{D}$ that satisfies the following: for every quantum state ρ, let ρ_p be the sub-normalized post-measurement state obtained after measuring ρ with respect to E_p. That is, $\rho_p = E_p \rho E_p$. We have,*

(1) For every $p \in \mathcal{D}$, ρ_p is an eigenvector of P with eigenvalue p;

(2) The probability of ρ when measured with respect to P is $\mathrm{Tr}[P\rho] = \sum_{p \in \mathcal{D}} \mathrm{Tr}[P\rho_p]$.

A measurement $\mathcal{E}$ which satisfies these properties is the measurement in the common eigenbasis of P and $Q = I - P$ (due to simultaneous diagonalization theorem, such common eigenbasis exists since P and Q commute). Let P have eigenbasis $\{|\psi_i\rangle\}$ with eigenvalues $\{\lambda_i\}$. Without loss of generality, let us assume ρ is a pure state $|\psi\rangle\langle\psi|$ and $\{\lambda_i\}$ has no duplicated eigenvalues. We write $|\psi\rangle$ in the eigenbasis of P: $|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$. Applying $\mathcal{E}$ will result in an outcome λ_i and a leftover state $|\psi_i\rangle$ with probability $|\alpha_i|^2$.

Looking ahead, we will write a quantum program under the eigenbasis of P in the proof of the strengthened MOE game.

Theorem 2.6 (Inefficient Threshold Measurement)**.** *Let $\mathcal{P} = (P, Q)$ be a binary outcome POVM. Let P have eigenbasis $\{|\psi_i\rangle\}$ with eigenvalues $\{\lambda_i\}$. Then, for every $\gamma \in (0,1)$ there exists a projective measurement $\mathcal{E}_\gamma = (E_{\leq\gamma}, E_{>\gamma})$ such that:*

(1) $E_{\leq\gamma}$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy $\lambda_i \leq \gamma$;

(2) $E_{>\gamma}$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy $\lambda_i > \gamma$.

Similarly, for every $\gamma \in (0, 1/2)$, there exists a projective measurement $\mathcal{E}'_\gamma = (\widetilde{E}_{\leq\gamma}, \widetilde{E}_{>\gamma})$ such that:

(1) $\widetilde{E}_{\leq\gamma}$ projects a quantum state into the subspace spanned by $\{|\psi_i\rangle\}$ whose eigenvalues λ_i satisfy $|\lambda_i - \frac{1}{2}| \leq \gamma$;

(2) $\widetilde{E}_{>\gamma}$ *projects a quantum state into the subspace spanned by* $\{|\psi_i\rangle\}$ *whose eigenvalues* λ_i *satisfy* $|\lambda_i - \frac{1}{2}| > \gamma$.

It is easy to see how to construct $\mathcal{E}_\gamma, \mathcal{E}'_\gamma$ from $\mathcal{E}$, e.g. by setting $\widetilde{E}_{\leq\gamma} = \sum_{i:|\lambda_i - 1/2| \leq \gamma} E_{\lambda_i}$. Note that for any quantum state ρ, $\mathrm{Tr}[\widetilde{E}_{>\gamma}\rho]$ is the weight over eigenvectors with eigenvalues λ that are γ away from $1/2$.

Efficient Measurement. The projective measurement $\mathcal{E}$ above is not efficiently computable in general. However, they can be approximated if the POVM is a mixture of projective measurements, as shown by Zhandry [Zha20], using a technique first introduced by Marriott and Watrous [MW05].

Consider the following procedure as a binary POVM $\mathcal{P} = (P, Q)$ acting on a quantum program ρ: samples a random challenge r, evaluates the program on r, and checks if the output is correct. This procedure can be viewed as (1). picking a uniformly random challenge r; (2). applying a *projective measurement* U_r. In this case, $P = \frac{1}{R}\sum_r U_r$ where R is the size of the challenge space. This POVM captures the situation where a challenger randomly samples a classical challenge and tests if a quantum program's classical outcome is correct on that challenge.

Below, we give the formal theorem statement about efficient approximated threshold measurement, which is adapted from Theorem 6.2 in [Zha20] and Lemma 3 in [ALL+21].

Theorem 2.7 (Efficient Threshold Measurement)**.** *Let* $\mathcal{P}_b = (P_b, Q_b)$ *be a binary outcome POVM over Hilbert space* $\mathcal{H}_b$ *that is a mixture of projective measurements for* $b \in \{1, 2\}$. *Let* P_b *have eigenbasis* $\{|\psi_i^b\rangle\}$ *with eigenvalues* $\{\lambda_i^b\}$. *For every* $\gamma_1, \gamma_2 \in (0, 1), 0 < \epsilon < \min(\gamma_1/2, \gamma_2/2, 1 - \gamma_1, 1 - \gamma_2)$ *and* $\delta > 0$, *there exist efficient binary-outcome quantum algorithms, interpreted as the POVM element corresponding to outcome 1,* $\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P}_b,\gamma}$ *such that for every quantum program* $\rho \in \mathcal{D}(\mathcal{H}_1) \otimes \mathcal{D}(\mathcal{H}_2)$ *the following are true about the product algorithm* $\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P}_1,\gamma_1} \otimes \mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P}_2,\gamma_2}$*:*

(0) Let $(E^b_{\leq\gamma}, E^b_{>\gamma})$ *be the inefficient threshold measurement in Theorem 2.6 for* $\mathcal{H}_b$.

(1) The probability of measuring 1 on both registers satisfies

$$\mathrm{Tr}\left[\left(\mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P}_1,\gamma_1} \otimes \mathsf{ATI}^{\epsilon,\delta}_{\mathcal{P}_2,\gamma_2}\right)\rho\right] \geq \mathrm{Tr}\left[\left(E^1_{>\gamma_1+\epsilon} \otimes E^2_{>\gamma_2+\epsilon}\right)\cdot\rho\right] - 2\delta.$$

(2) The post-measurement state ρ' *after getting outcome (1,1) is* 4δ*-close to a state in the support of* $\left\{|\psi_i^1\rangle\,|\psi_j^2\rangle\right\}$ *such that* $\lambda_i^1 > \gamma_1 - 2\epsilon$ *and* $\lambda_j^2 > \gamma_2 - 2\epsilon$.

(3) The running time of the algorithm is polynomial in the running time of P_1, P_2, $1/\epsilon$ *and* $\log(1/\delta)$.

Intuitively the theorem says that if a quantum state ρ has weight p on eigenvectors of (P_1, P_2) with eigenvalues greater than $(\gamma_1 + \epsilon, \gamma_2 + \epsilon)$, then the quantum algorithm will produce (with probability at least $p - 2\delta$) a post-measurement state which has weight $1 - 4\delta$ on eigenvectors with eigenvalues greater than $(\gamma_1 - 2\epsilon, \gamma_2 - 2\epsilon)$.

In this paper, we will work with indistinguishability games. Therefore, we will particularly be interested in the projective measurement that projects onto eigenvectors with eigenvalues away from $1/2$ (meaning its behavior is more than random guessing). For this reason, we will need the following symmetric version of Theorem 2.7:

Theorem 2.8 (Efficient Symmetric Threshold Measurement). *Let $\mathcal{P}_b = (P_b, Q_b)$ be a binary outcome POVM over Hilbert space $\mathcal{H}_b$ that is a mixture of projective measurements for $b \in \{1,2\}$. Let P_b have eigenbasis $\{|\psi_i^b\rangle\}$ with eigenvalues $\{\lambda_i^b\}$. For every $\gamma_1, \gamma_2 \in (0, 1/2), 0 < \epsilon < \min(\gamma_1/2, \gamma_2/2)$, and $\delta > 0$, there exist efficient binary-outcome quantum algorithms, interpreted as the POVM element corresponding to outcome 1, $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P}_b,\gamma}$ such that for every quantum program $\rho \in \mathcal{D}(\mathcal{H}_1) \otimes \mathcal{D}(\mathcal{H}_2)$ the following are true about the product algorithm $\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P}_1,\gamma_1} \otimes \mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P}_2,\gamma_2}$:*

(0) Let $(\widetilde{E}^b_{\leq\gamma_b}, \widetilde{E}^b_{>\gamma_b})$ be the inefficient threshold measurement in Theorem 2.6 for $\mathcal{H}_b$.

(1) The probability of measuring 1 on both registers satisfies

$$\mathrm{Tr}\left[\left(\mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P}_1,\gamma_1} \otimes \mathsf{SATI}^{\epsilon,\delta}_{\mathcal{P}_2,\gamma_2}\right)\rho\right] \geq \mathrm{Tr}\left[\left(\widetilde{E}^1_{>\gamma_1+\epsilon} \otimes \widetilde{E}^2_{>\gamma_2+\epsilon}\right)\cdot\rho\right] - 2\delta.$$

(2) The post-measurement state ρ' after getting outcome (1,1) is 4δ-close to a state in the support of $\left\{|\psi_i^1\rangle\,|\psi_j^2\rangle\right\}$ such that $|\lambda_i^1 - 1/2| > \gamma_1 - 2\epsilon$ and $|\lambda_j^2 - 1/2| > \gamma_2 - 2\epsilon$.

(3) The running time of the algorithm is polynomial in the running time of P_1, P_2, $1/\epsilon$ and $\log(1/\delta)$.

2.5 Unclonable Encryption

In this subsection, we provide the definition of unclonable encryption schemes. By unclonable encryption, we are refering to the security defined in [AK21]. This is a variant of the original security definition in [BL20], which forces one of m_0, m_1 to be uniformly random. We would remark that our security is stronger than the original one in [BL20], since in our definition m_0, m_1 can be arbitrarily chosen.

Definition 2.9. *An unclonable encryption scheme is a triple of efficient quantum algorithms* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *with the following interface:*

- $\mathsf{Gen}(1^\lambda) : \mathsf{sk}$ *on input a security parameter 1^λ, returns a classical key* sk.
- $\mathsf{Enc}(\mathsf{sk}, |m\rangle\langle m|) : \rho_{ct}$ *takes the key* sk *and the message $|m\rangle\langle m|$ for $m \in \{0,1\}^{\mathrm{poly}(\lambda)}$, outputs a quantum ciphertext ρ_{ct}.*
- $\mathsf{Dec}(\mathsf{sk}, \rho_{ct}) : \rho_m$ *takes the key* sk *and the quantum ciphertext ρ_{ct}, outputs a message in the form of quantum states ρ_m.*

Correctness. *The following must hold for the encryption scheme. For* $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$, *we must have* $\mathrm{Tr}[|m\rangle\langle m| \mathsf{Dec}(\mathsf{sk}, \mathsf{Enc}(\mathsf{sk}, |m\rangle\langle m|))] \geq 1 - \mathsf{negl}(\lambda)$.

Unclonability. *In the following sections, we focus on unclonable IND-CPA security. To define our unclonable security, we introduce the following security game.*

Definition 2.10 (Unclonable IND-CPA game). *Let $\lambda \in \mathbb{N}^+$. Given encryption scheme $\mathcal{S}$, consider the following game against the adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$.*

- *The adversary $\mathcal{A}$ generates $m_0, m_1 \in \{0,1\}^{n(\lambda)}$ and sends to the challenger as the chosen plaintext.*

- *The challenger randomly chooses a bit $b \in \{0,1\}$ and returns $\mathsf{Enc}(\mathsf{sk}, m_b)$ to $\mathcal{A}$. $\mathcal{A}$ produces a quantum state ρ_{BC} in register B and C, and sends corresponding registers to $\mathcal{B}$ and $\mathcal{C}$.*
- *$\mathcal{B}$ and $\mathcal{C}$ receive the key sk, and output bits $b_{\mathcal{B}}$ and $b_{\mathcal{C}}$ respectively*

and the adversary wins if $b_{\mathcal{B}} = b_{\mathcal{C}} = b$.

We denote the advantage (success probability) of above game by $\mathsf{adv}_{\mathcal{G},\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda)$. We say that scheme $\mathcal{S}$ is informational (computational) secure if for all(efficient) adversaries $(\mathcal{G}, \mathcal{A}, \mathcal{B}, \mathcal{C})$,

$$\mathsf{adv}_{\mathcal{G},\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda) \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

3 On the Impossibility of Deterministic Schemes

In this section, we provide an impossibility result for *deterministic information-theoretically secure* schemes. This result suggests that either computational assumptions or randomness is necessary for achieving unclonable encryption with optimal security. We also noticed that previously in [MST21], the authors have provided an impossibility result for more general schemes. Nevertheless, our result provides a better asymptotic lower bound for deterministic schemes and is based on observations on Haar random states.

To be precise, we define deterministic schemes as follows:

Definition 3.1 (Deterministic Scheme)**.** *We call an encryption scheme* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is a deterministic encryption scheme if it satisfies following:*

- *The encryption algorithm* Enc *can be realized as a unitary U_{sk} acting on the plaintext register $|m\rangle$ and ancillary bits initialized to 0, resulting in the ciphertext pure state in the form $|c_{\mathsf{sk}}\rangle$ of length λ.*
- *The decryption algorithm* Dec *acts the inverse $U_{\mathsf{sk}}^{\dagger}$ on received registers, then measures in computational basis to obtain the message.*

The correctness of deterministic schemes is satisfied. An example of deterministic scheme is the following: let sk encode two (arbitrary and) orthogonal states $|\phi_0\rangle, |\phi_1\rangle$; a message b is mapped to $|\phi_b\rangle$. Another example is the conjugate encryption defined in [BL20]:

1. $\mathsf{sk} = (r, \theta)$ where r, θ is independent random samples from $\{0,1\}^n$
2. $\mathsf{Enc}(\mathsf{sk}, m) = |(m \oplus r)^{\theta}\rangle \langle (m \oplus r)^{\theta}|$, where $|x^{\theta}\rangle = H^{\theta_1} \otimes H^{\theta_2} \otimes \cdots \otimes H^{\theta_n} |x\rangle$ is the BB84 state.
3. $\mathsf{Dec}(\mathsf{sk}, \rho)$: computes $\rho' = H^{\theta} \rho H^{\theta}$, measures ρ' in computational basis to obtain c, obtaining $m = c \oplus r$.

Though the authors in [BL20] have already proven this scheme does not satisfy the unclonable IND-CPA security, our attack scheme provided a no go theorem for a larger class of possible constructions.

For these schemes, we provide a universal adversary for the unclonable IND-CPA game.

Theorem 3.2. *For any deterministic encryption scheme, we have a universal information-theoretical adversary $(\mathcal{G}, \mathcal{A}, \mathcal{B}, \mathcal{C})$ that satisfies*

$$\mathsf{adv}_{\mathcal{G},\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda) \geq 0.568,$$

as $\lambda \to \infty$.

Since any deterministic encryption scheme can only suffice one-time security, we also considered whether our result can be extended to general encryption schemes that take randomness as input, such as the following scheme inspired by [GL89].

- $\mathsf{sk} = (\theta, u)$ for $\theta, u \leftarrow \{0,1\}^\lambda$.
- $Enc_k(m, r) = |r^\theta\rangle \, |\langle r, u\rangle \oplus m\rangle$ for $m \in \{0,1\}$, $r \leftarrow \{0,1\}^\lambda$.
- $Dec_k(\rho)$: Decode r by applying H^θ on first λ register, and measure ρ in computational basis to get ct. We can extract $m = \langle ct_{1\ldots\lambda}, u\rangle \oplus ct_{\lambda+1}$.

However, our impossibility result met some barriers in the generalization. We would try to characterize them as following:

- Since in quantum algorithms, randomness is generated intrinsically from measurements. Consider implementing a classical randomized algorithm by quantum circuits, the random bits in the classical algorithm would be replaced by measuring $|+\rangle$ states in the computational basis. Thus for general encryption algorithms, they should be modeled as quantum channels rather than unitaries, with cipher texts modeled as mixed states accordingly. However, the understanding on the actions of random unitaries on mixed states is a much less studied and more complicated problem.
- Our adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ also relies on *all* information of the cipher text states to decide its measurement. But if the encryption algorithm additionally takes some randomness, then the adversary $\mathcal{B}$ and $\mathcal{C}$ cannot decide the actual ciphertext state.

3.1 Preliminaries on Haar Measure

To prove our result, we provide a quick introduction to the theorems related to Haar measure in this subsection. For more information on Haar measure, readers can refer to [Wat18]. We denote the uniform spherical measure on unit sphere $\mathcal{S}((\mathbb{C}^2)^{\otimes n})$ as μ_n, the Haar measure on the unitary group $\mathcal{U}((\mathbb{C}^2)^{\otimes n})$ as η_n.

The following lemma relates the Haar measure on unitary operators to uniform spherical measure.

Lemma 3.3. *Let f be a function from $\mathcal{S}((\mathbb{C}^2)^{\otimes n}) \times \mathcal{S}((\mathbb{C}^2)^{\otimes n}) \to \mathbb{R}$. Then for any two fixed vectors $|\phi_0\rangle, |\phi_1\rangle \in \mathcal{S}((\mathbb{C}^2)^{\otimes n})$ such that $\langle\phi_0|\phi_1\rangle = 0$, we have that*

$$\mathop{\mathbb{E}}_{U \leftarrow \eta_n} f(U|\phi_0\rangle, U|\phi_1\rangle) = \mathop{\mathbb{E}}_{|\psi_0\rangle,|\psi_1\rangle \leftarrow \mu_n, \langle\psi_0|\psi_1\rangle = 0} f(|\psi_0\rangle, |\psi_1\rangle).$$

We introduce Lévy's lemma, which could be viewed as the counterpart of Chernoff bound on the uniform spherical measure.

Lemma 3.4 (Lévy's Lemma). *Let f be a function from $\mathcal{S}((\mathbb{C}^2)^{\otimes n}) \to \mathbb{R}$ that satisfies*

$$|f(|\phi\rangle) - f(|\psi\rangle)| \leq \kappa || \, |\phi\rangle - |\psi\rangle \, ||_2,$$

for some $\kappa > 0$. Then there exists a universal $\delta > 0$ for which the following holds. For every $\epsilon > 0$:

$$\Pr_{|\psi\rangle \leftarrow \mu_n} \left[\left| f(|\psi\rangle) - \mathop{\mathbb{E}}_{|\phi\rangle \leftarrow \mu_n} [f(|\phi\rangle)] \right| \geq \epsilon \right] \leq 3 \exp\left(-\frac{\delta \epsilon 2^n}{\kappa^2} \right).$$

The following simplified theorem from [MZB16] plays a crucial role in our proof.

Theorem 3.5. *Let $|\psi_1\rangle, |\psi_2\rangle \in \mathcal{S}((\mathbb{C}^2)^{\otimes 2n})$ be two states independently sampled from μ_{2n}. Then let ρ_1, ρ_2 be the corresponding reduced density matrix in the first n qubit register. As $n \to \infty$, the trace distance $\mathsf{TD}(\rho_1, \rho_2)$ almost surely converges to*

$$\mathsf{TD}(\rho_1, \rho_2) \xrightarrow{a.s.} \frac{1}{4} + \frac{1}{\pi} \approx 0.568.$$

For simplicity, in this section, $\mathbb{E}_{|\psi\rangle}$ stands for taking expectation over $|\psi\rangle$ sampled from uniform spherical measure on corresponding Hilbert space, $\mathbb{E}_V$ stands for V over Haar measure respectively.

3.2 Attack schemes

We are ready to present an attack for any deterministic information-theoretically secure schemes.

Attack.

- For the adversary $\mathcal{A}$, it first chooses $00\ldots00, 00\ldots01$ and sends to the challenger. After receiving the n qubit ciphertext state $|ct_k\rangle$, it applies a random Haar unitary V, then divides the output register into two parts, $R_{\mathcal{B}}$ for the qubits indexed $[1, \frac{\lambda}{2}]$, $R_{\mathcal{C}}$ for the qubits indexed $[\frac{\lambda}{2} + 1, \lambda]$.
- $\mathcal{A}$ then sends two registers respectively to $\mathcal{B}$ and $\mathcal{C}$, together with the description of V [9].
- With the given information, $\mathcal{B}$ and $\mathcal{C}$ can perform POVMs $\{\Pi_b^B\}_b$ and $\{\Pi_b^C\}_b$ to distinguish different messages. We will define the POVMs in detail in the following section.

The success probability of our attack scheme is equal to the success probability of the following game.

Definition 3.6. *Let $\lambda \in \mathbb{N}^+$. Consider the following game with a challenger and an (unbounded) adversary $(\mathcal{B}, \mathcal{C})$.*

- *The challenger generates two Haar random states $|\phi_0\rangle, |\phi_1\rangle$ with restriction $\langle\phi_0|\phi_1\rangle = 0$ and sends the description of two states[10] to $\mathcal{B}$ and $\mathcal{C}$.*

[9] Here we actually mean sending the corresponding minimal distance $\widetilde{V}$ in the ϵ-net of $\mathrm{U}(\mathcal{H}_2^{\otimes\lambda})$ to approximate the distribution of V. Thus we can sample from a finite set instead. Since we have a constant advantage in the end, we can take ϵ small enough such that it will only have a negligible effect on our result.

[10] Similarly, $\mathcal{A}$ sends an element in the ϵ-net of $\mathcal{S}(\mathcal{H}_2^{\otimes\lambda})$ in implementation.

- *The challenger randomly chooses $b_{ch} \in \{0,1\}$, and divides the state $|\phi_{b_{ch}}\rangle$ into two parts, $R_{\mathcal{B}}$ for the qubits indexed $[1, \frac{\lambda}{2}]$, $R_{\mathcal{C}}$ for the qubits indexed $[\frac{\lambda}{2}+1, \lambda]$.*
- *$\mathcal{B}$ and $\mathcal{C}$ perform POVM $\{\Pi^B_{b_B}\}_{b_B}$ and $\{\Pi^C_{b_C}\}_{b_C}$ on their received registers, outputs b_B and b_C from measurement results.*

The adversary wins the game if $b_B = b_C = b_{ch}$.

The success probability of our distinguishing game is given by the following optimization problem:

$$\begin{aligned}\max_{\Pi^B_0,\Pi^B_1,\Pi^C_0,\Pi^C_1} \frac{1}{2}\left(\langle\phi_0|\,\Pi^B_0\otimes\Pi^C_0\,|\phi_0\rangle + \langle\phi_1|\,\Pi^B_1\otimes\Pi^C_1\,|\phi_1\rangle\right)\\ s.t.\ \ \Pi^B_0+\Pi^B_1 = I_{\frac{\lambda}{2}}, \Pi^C_0+\Pi^C_1 = I_{\frac{\lambda}{2}},\\ 0\le\Pi^B_i\le I_{\frac{\lambda}{2}}, 0\le\Pi^C_i\le I_{\frac{\lambda}{2}}.\end{aligned}$$

The $\frac{1}{2}$ comes from the requirement that the challenger sends $|\phi_0\rangle, |\phi_1\rangle$ with equal probability. We denote this probability as $G(|\phi_0\rangle, |\phi_1\rangle)$. For simplicity, in following sections we will abbreviate $\{\Pi^B_{b_B}\}_{b_B}$ and $\{\Pi^C_{b_C}\}_{b_C}$ as $\{\Pi^B\}$ and $\{\Pi^C\}$ respectively.

In our attack scheme, our success probability is given by $\mathbb{E}_k\,\mathbb{E}_V[G(VU_k\,|0\ldots00\rangle, VU_k\,|0\ldots01\rangle)]$. By lemma 3.3, we have that

$$\mathop{\mathbb{E}}_{V}\left[G\left(VU_k\,|0\ldots00\rangle, VU_k\,|0\ldots01\rangle\right)\right] = \mathop{\mathbb{E}}_{|\phi_0\rangle,|\phi_1\rangle\langle\phi_0|\phi_1\rangle=0}\left[G\left(|\phi_0\rangle, |\phi_1\rangle\right)\right] = \Pr[(\mathcal{B},\mathcal{C})\text{ wins}]$$

Then we can provide a lower bound for the success probability via following inequalities:

$$\begin{aligned}&\Pr[(\mathcal{B},\mathcal{C})\text{ wins}|\,|\phi_0\rangle, |\phi_1\rangle]\\ =&\max_{\{\Pi^B\},\{\Pi^C\}}\Pr[(b_B=b_{ch})\wedge(b_C=b_{ch})|\{\Pi^B\},\{\Pi^C\},|\phi_0\rangle,|\phi_1\rangle]\\ \ge&1-\min_{\{\Pi^B\}}\Pr[b_B\ne b_{ch}|\{\Pi^B\},|\phi_0\rangle,|\phi_1\rangle]-\min_{\{\Pi^C\}}\Pr[b_C\ne b_{ch}|\{\Pi^C\},|\phi_0\rangle,|\phi_1\rangle]\\ =&1-\frac{1}{2}(1-\mathsf{TD}(\rho^B_0,\rho^B_1))-\frac{1}{2}(1-\mathsf{TD}(\rho^C_0,\rho^C_1))\\ =&\frac{1}{2}(\mathsf{TD}(\rho^B_0,\rho^B_1)+\mathsf{TD}(\rho^C_0,\rho^C_1)),\end{aligned}$$

where the first line is by definition, the second line follows from union bound, the third line is by the property of trace distance.

Then by taking expectation, we have that for large enough λ,

$$\begin{aligned}\Pr[(\mathcal{B},\mathcal{C})\text{ wins}] =& \mathop{\mathbb{E}}_{|\phi_0\rangle,|\phi_1\rangle,\langle\phi_0|\phi_1\rangle=0}\Pr[(\mathcal{B},\mathcal{C})\text{ wins}|\,|\phi_0\rangle,|\phi_1\rangle]\\ \ge&\mathop{\mathbb{E}}_{|\phi_0\rangle,|\phi_1\rangle,\langle\phi_0|\phi_1\rangle=0}\left[\frac{1}{2}(\mathsf{TD}(\rho^B_0,\rho^B_1)+\mathsf{TD}(\rho^C_0,\rho^C_1))\right]\\ \ge&\mathop{\mathbb{E}}_{|\phi_0\rangle,|\phi_1\rangle}\left[\frac{1}{2}(\mathsf{TD}(\rho^B_0,\rho^B_1)+\mathsf{TD}(\rho^C_0,\rho^C_1))\right]-\mathsf{negl}(\lambda)\\ \ge&\frac{1}{4}+\frac{1}{\pi}-\epsilon\ge0.568,\end{aligned}$$

where the first line is by definition, the second line is by the inequality before, the third line is by concentration property of the Haar measure, the last line is by theorem 3.5 as $\lambda \to \infty$. Thus we finished the proof of theorem 3.2

Here we provide rigorous proof of the third line. Note that for an arbitrary $|\phi_1\rangle$, given $|\phi_0\rangle$ it can be written as $|\phi_1\rangle = a\,|\phi_0\rangle + \sqrt{1-|a|^2}\,|\phi_0^\perp\rangle$, where $a = \langle\phi_0|\phi_1\rangle$ and $\langle\phi_0|\phi_0^\perp\rangle = 0$. By symmetry, we have that $\mathbb{E}_{|\phi_1\rangle}[|a|^2] = \frac{1}{2^\lambda}$. Taking $\epsilon = \lambda 2^{-\frac{\lambda}{2}}, \kappa = 2$ in lemma 3.4, we obtain that

$$\Pr\left[\left||a|^2 - \frac{1}{2^\lambda}\right| \geq \frac{\lambda}{2^{\frac{\lambda}{2}}}\right] \leq 3\exp\left(-\frac{\delta\lambda^2}{4}\right),$$

thus we can derive that

$$\mathop{\mathbb{E}}_{|\phi_1\rangle}[|a|] \leq 3\exp\left(-\frac{\delta\lambda^2}{4}\right)\cdot 1 + 1\cdot\frac{\sqrt{\lambda}+1}{2^{\frac{\lambda}{4}}} = \mathsf{negl}(\lambda).$$

Consider the trace distance $\mathsf{TD}(\rho_0^B, \rho_1^B)$ for two random states $|\phi_0\rangle, |\phi_1\rangle$. By definition it can be rewritten as $|\operatorname{Tr}_C[|\phi_0\rangle\langle\phi_0| - |\phi_1\rangle\langle\phi_1|]|_1$, then following the decomposition of $|\phi_1\rangle = a\,|\phi_0\rangle + \sqrt{1-|a|^2}\,|\phi_0^\perp\rangle$, we expand the expectation of the term as

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\substack{a\\|\phi_0\rangle,|\phi_0^\perp\rangle\\\langle\phi_0|\phi_0^\perp\rangle=0}}\left[\frac{1}{2}\left|\operatorname{Tr}_C[(1-|a|^2)(|\phi_0\rangle\langle\phi_0| - |\phi_0^\perp\rangle\langle\phi_0^\perp|) - \sqrt{1-|a|^2}(a\,|\phi_0\rangle\langle\phi_0^\perp| + a^*\,|\phi_0^\perp\rangle\langle\phi_0|)]\right|_1\right]\\
\leq &\mathop{\mathbb{E}}_{\substack{|\phi_0\rangle,|\phi_1\rangle\\\langle\phi_0|\phi_1\rangle=0}}\left[\frac{1}{2}\left|\operatorname{Tr}_C[|\phi_0\rangle\langle\phi_0| - |\phi_1\rangle\langle\phi_1|]\right|_1\right] + \mathop{\mathbb{E}}_a[|a|]\\
\leq &\mathop{\mathbb{E}}_{\substack{|\phi_0\rangle,|\phi_1\rangle\\\langle\phi_0|\phi_1\rangle=0}}\left[\mathsf{TD}(\rho_0^B,\rho_1^B)\right] + \mathsf{negl}(\lambda),
\end{aligned}$$

where the second line is by definition, the third line is from the decomposition of $|\phi_1\rangle$, the fourth line is by triangle inequality and renaming $|\phi_0^\perp\rangle$ to $|\phi_1\rangle$, the last line is by definition and previous bounds on $\mathbb{E}[|a|]$.

4 More on Coset States

In this section, we will recall the basic properties of coset states. We will then introduce a strengthened unclonable game in the quantum random oracle model (QROM), upon which we will build our unclonable encryption scheme. The last subsection is devoted to prove the security of this strengthened game.

4.1 Preliminaries

In this subsection, we recall the basic definitions and properties of coset states in [CLLZ21]. Let $A \subseteq \mathbb{F}_2^n$ be a subspace. Define its orthogonal complement of A as $A^\perp = \{b \in \mathbb{F}_2^n \,|\, \langle a, b\rangle \bmod 2 = 0\,,\ \forall a \in A\}$. It satisfies $\dim(A) + \dim(A^\perp) = n$. We also let $|A| = 2^{\dim(A)}$ denote the size of A.

Definition 4.1 (Coset States). *For any subspace $A \subseteq \mathbb{F}_2^n$ and vectors $s, s' \in \mathbb{F}_2^n$, the coset state $|A_{s,s'}\rangle$ is defined as:*

$$|A_{s,s'}\rangle = \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle s', a \rangle} |a + s\rangle \ .$$

By applying $H^{\otimes n}$ to the state $|A_{s,s'}\rangle$, one obtains exactly $|A^{\perp}_{s',s}\rangle$. Given A, s, s', the coset state is efficiently constructible.

For a subspace A and vectors s, s', we define $A + s = \{v + s : v \in A\}$, and $A^{\perp} + s' = \{v + s' : v \in A^{\perp}\}$. We define P_{A+s} and $P_{A^{\perp}+s'}$ as the membership checking oracle for both cosets.

It is also convenient for later sections to define a canonical representation of a coset $A + s$, with respect to subspace A,

Definition 4.2 (Canonical Representative of a Coset). *For a subspace A, we define the function $\mathsf{Can}_A(\cdot)$ such that $\mathsf{Can}_A(s)$ is the lexicographically smallest vector contained in $A + s$. We call this the canonical representative of coset $A + s$.*

If $\tilde{s} \in A + s$, then $\mathsf{Can}_A(s) = \mathsf{Can}_A(\tilde{s})$. We also note that $\mathsf{Can}_A(\cdot)$ is polynomial-time computable given the description of A. Accordingly, we can efficiently sample from $\mathsf{CS}(A) := \{\mathsf{Can}_A(s) : s \in \mathbb{F}_2^n\}$, which denotes the set of canonical representatives for A.

For a fixed subspace A, the coset states $\left\{|A_{s,s'}\rangle\right\}_{s \in \mathsf{CS}(A), s' \in \mathsf{CS}(A^{\perp})}$ form an orthonormal basis. (See Lemma C.2 in [CLLZ21])

Next, we recall the regular direct product and MOE properties of coset states. These properties will be used to prove the strengthened unclonable property.

Direct Product Hardness

Theorem 4.3 (Theorem 4.5,4.6 in [CLLZ21]). *Let $A \subseteq \mathbb{F}_2^{\lambda}$ be a uniformly random subspace of dimension $\frac{\lambda}{2}$, and s, s' be two uniformly random vectors from $\mathbb{F}_2^{\lambda}$. Let $\epsilon > 0$ such that $1/\epsilon = o(2^{n/2})$. Given one copy of $|A_{s,s'}\rangle$ and oracle access to P_{A+s} and $P_{A^{\perp}+s'}$, an adversary needs $\Omega(\sqrt{\epsilon} 2^{\lambda/2})$ queries to output a pair (v, w) that $v \in A + s$ and $w \in A^{\perp} + s'$ with probability at least ϵ.*

An important corollary immediately follows.

Corollary 4.4. *There exists an exponential function $\exp$ such that, for any query-bounded (polynomially many queries to $P_{A+s}, P_{A^{\perp}+s'}$) adversary, its probability to output a pair (v, w) that $v \in A + s$ and $w \in A^{\perp} + s'$ is smaller than $1/\exp(\lambda)$.*

Monogamy-of-Entanglement (with Membership Checking Oracles).

Definition 4.5. *Let $\lambda \in \mathbb{N}^+$. Consider the following game between a challenger and an adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$.*

- *The challenger picks a uniformly random subspace $A \subseteq \mathbb{F}_2^{\lambda}$ of dimension $\frac{\lambda}{2}$, and uniformly random vectors $(s, s') \in \mathsf{CS}(A) \times \mathsf{CS}(A^{\perp})$. It sends $|A_{s,s'}\rangle$ to $\mathcal{A}$.*
- *$\mathcal{A}, \mathcal{B}, \mathcal{C}$ get (quantum) oracle access to P_{A+s} and $P_{A^{\perp}+s'}$.*

- *$\mathcal{A}$ creates a bipartite state on registers B and C. Then, $\mathcal{A}$ sends register B to $\mathcal{B}$, and C to $\mathcal{C}$.*
- *The description of A is then sent to both $\mathcal{B}, \mathcal{C}$.*
- *$\mathcal{B}$ and $\mathcal{C}$ return respectively (s_1, s_1') and (s_2, s_2').*

$(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins if and only if for $i \in \{1, 2\}$, $s_i = s$ and $s_i' = s'$.

We denote the advantage (success probability) of the above game by $\mathsf{adv}_{\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda)$. We have the following theorem.

Theorem 4.6 (Theorem 4.14, 4.15 in [CLLZ21])**.** *There exists an exponential function* exp *such that, for every $\lambda \in \mathbb{N}^+$, for any query-bounded (polynomially many queries to $P_{A+s}, P_{A^\perp+s'}$) adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$,*

$$\mathsf{adv}_{\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda) \leq 1/\exp(\lambda)\,.$$

Note that in [CLLZ21], the authors only proved the above theorem for a sub-exponential function and membership checking oracles are given in the form of indistinguishability obfuscation (iO). The proof trivially holds if we replace iO with VBB obfuscation (quantum access to these oracles). Culf and Vidick [CV21] further proved the theorem holds for an exponential function.

4.2 Strengthened MOE Game in the QROM

In this subsection, we will introduce the strengthened MOE game in the QROM and state our main theorem. We present the proof in the next section.

Definition 4.7. *Let $\lambda \in \mathbb{N}^+$. Consider the following security game between a challenger and an adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ with a random oracle $H : \mathbb{F}_2^\lambda \times \mathbb{F}_2^\lambda \to \{0,1\}^{n(\lambda)}$.*

- *The adversary $\mathcal{A}$ generates $\Delta \in \{0,1\}^{n(\lambda)}$ and sends Δ to the challenger.*
- *The challenger samples a random subspace $A \subseteq \mathbb{F}_2^\lambda$ of dimension $\lambda/2$ and two random vectors $(s, s') \in \mathsf{CS}(A) \times \mathsf{CS}(A^\perp)$. The challenger also randomly chooses a bit $b \in \{0,1\}$ and calculates $w = H(s, s') \oplus (b \cdot \Delta)$.*

  It gives $|A_{s,s'}\rangle$ and w to $\mathcal{A}$.
- *$\mathcal{A}, \mathcal{B}, \mathcal{C}$ get (quantum) oracle access to P_{A+s} and $P_{A^\perp+s'}$.*
- *$\mathcal{A}$ produces a quantum state over registers BC and sends B to $\mathcal{B}$ and C to $\mathcal{C}$.*
- *$\mathcal{B}, \mathcal{C}$ are given the description of A, they try to produce bits $b_\mathcal{B}, b_\mathcal{C}$.*

$(\mathcal{A}, \mathcal{B}, \mathcal{C})$ win if and only if $b_\mathcal{B} = b_\mathcal{C} = b$.

We denote the advantage of the above game by $\mathsf{adv}_{\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda)$. Note that since s, s' is defined as the canonical vector of both cosets, they are uniquely defined; similarly, $H(s, s')$ is also uniquely defined.

We show the following theorem:

Theorem 4.8. *Let $n = \Omega(\lambda)$, then for every $\lambda \in \mathbb{N}^+$ and all query-bounded algorithms $(\mathcal{A}, \mathcal{B}, \mathcal{C})$, $\mathsf{adv}_{\mathcal{A},\mathcal{B},\mathcal{C}}(\lambda) \leq \frac{1}{2} + \mathsf{negl}(\lambda)$.*

4.3 Proof for Theorem 4.8

Proof. We prove the theorem by following hybrid arguments.

Hybrid 0. This hybrid is the original game.

Hybrid 1. This hybrid follows Hybrid 0, but the oracle of $\mathcal{A}$ will be reprogrammed as $H_{s,s'}$ defined as follows:

$$H_{s,s'}(z,z') = \begin{cases} u & \text{if } z = s, z' = s' \\ H(z,z') & \text{otherwise} \end{cases},$$

where $u \in \{0,1\}^n$ is chosen uniformly at random.

Hybrid 2. This hybrid will modify the access to random oracle of $\mathcal{B}$ and $\mathcal{C}$.

- The adversary $\mathcal{A}$ generates $\Delta \in \{0,1\}^{n(\lambda)}$ and sends Δ to the challenger.
- The challenger samples a random subspace $A \subseteq \mathbb{F}_2^\lambda$ of dimension $\lambda/2$ and two random vectors $(s,s') \in \mathsf{CS}(A) \times \mathsf{CS}(A^\perp)$. The challenger uniform randomly samples a bit $b \in \{0,1\}$ and $r \in \{0,1\}^{n(\lambda)}$, and defines the oracle $H^b_{s,s'}$ as follows:

$$H^b_{s,s'}(z,z') = \begin{cases} r \oplus (b \cdot \Delta) & \text{if } z = s, z' = s' \\ H(z,z') & \text{otherwise} \end{cases},$$

  It gives $|A_{s,s'}\rangle$ and r to $\mathcal{A}$.
- $\mathcal{A}, \mathcal{B}, \mathcal{C}$ get (quantum) oracle access to P_{A+s} and $P_{A^\perp+s'}$.
- With access to quantum random oracle $H_{s,s'}$, $\mathcal{A}$ produces a quantum state over registers BC and sends B to $\mathcal{B}$ and C to $\mathcal{C}$.
- With access to quantum random oracle $H^b_{s,s'}$, $\mathcal{B}, \mathcal{C}$ are given the description of A, they try to produce bits $b_\mathcal{B}, b_\mathcal{C}$.

$(\mathcal{A}, \mathcal{B}, \mathcal{C})$ win if and only if $b_\mathcal{B} = b_\mathcal{C} = b$.

We denote by p_i the optimal success probability of the game in **Hybrid i**. For the relations between different p_i, we have following lemmas:

Lemma 4.9. $|p_0 - p_1| \leq \mathsf{negl}(\lambda)$.

Lemma 4.10. $p_1 = p_2$.

Lemma 4.11. $p_2 \leq \frac{1}{2} + \mathsf{negl}(\lambda)$.

Combining the three lemmas, we have completed the proof of Theorem 4.8.

□

Now we provide proofs for lemmas beyond.

Proof for Lemma 4.9. We prove by contradiction. Suppose $p_0 \geq p_1 + 1/q(\lambda)$ for some polynomial $q(\lambda)$, then we can construct an adversary $\mathcal{A}'$ that violates the direct product hardness of coset states. $\mathcal{A}'$ will perform as follows:

- $\mathcal{A}'$ samples a random oracle $H : \mathbb{F}_2^\lambda \times \mathbb{F}_2^\lambda \to \{0,1\}^{n(\lambda)}$.
- $\mathcal{A}'$ simulates $\mathcal{A}$ using H and applies computational basis measurement on a random quantum query made by $\mathcal{A}$ to the random oracle.

By Theorem 2.1, assuming $\mathcal{A}$ makes at most T queries, then $\mathcal{A}'$ gets (s, s') with probability at least $4/(q^2T)$, a contradiction to Corollary 4.4. □

Proof of Lemma 4.10. Fixing Δ and b, the two games are identical by renaming the $w = H(s, s') \oplus (b \cdot \Delta)$ to r. Since $H(s, s')$ is uniformly random, its distribution is identical to r. □

Proof of Lemma 4.11. Fixing A, r, Δ, two canonical vectors s, s', let $H_{-s,s'}$ be a partial random oracle that is defined on every input except (s, s'). Fix any partial random oracle $H_{-s,s'}$, we define two *projectors* Π_0^B, Π_1^B over register B as:

- Π_0^B: runs $\mathcal{B}$ on input A with oracle access to $H^0_{s,s'}$ where $H^0_{s,s'}$ is the same as $H_{-s,s'}$ except on input (s, s') it outputs r; it measures if the outcome is r; then it undoes all the computation.
- Π_1^B: similar to Π_0^B except on input (s, s'), the random oracle $H^1_{s,s'}$ outputs $r \oplus \Delta$ and it checks if the outcome is $r \oplus \Delta$.

Let $\{|\phi_i\rangle\}_i$ be a set of the eigenvectors of $(\Pi_0^B + \Pi_1^B)/2$ with eigenvalues $\{\lambda_i\}_i$.

Fixing the same A, s, s', r and $H_{-s,s'}$, we can similarly define Π_0^C, Π_1^C for $\mathcal{C}$. Let $\{|\psi_j\rangle\}_j$ be a set of the eigenvectors of $(\Pi_0^C + \Pi_1^C)/2$ with eigenvalues $\{\mu_j\}_j$.

Let $|\phi_{\mathsf{BC}}\rangle$ be the state prepared by $\mathcal{A}$. Without loss of generality, we can assume the state is pure. We write the state under the basis $\{|\phi_i\rangle\}_i$ and $\{|\psi_j\rangle\}_j$:

$$|\phi_{\mathsf{BC}}\rangle = \sum_{i,j} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} \otimes |\psi_j\rangle_{\mathsf{C}} .$$

Lemma 4.12. *Taken the randomness of A, s, s' and $H_{-s,s'}$, for every polynomial $p(\cdot)$, there exists a negligible function* negl *such that with overwhelming probability the following weight is bounded:*

$$\sum_{\substack{i:\ |\lambda_i - 1/2| > 1/p \\ j:\ |\mu_j - 1/2| > 1/p}} |\alpha_{i,j}|^2 \leq \mathsf{negl}(n).$$

The proof for this lemma is given at the end of this section.

With the above lemma, we can claim that over the randomness of A, s, s' and $H_{-s,s'}$, for every polynomial $p(\cdot)$, $|\phi_{\mathsf{BC}}\rangle$ is negligibly close to the following state $|\phi'_{\mathsf{BC}}\rangle$:

$$\sum_{i:|\lambda_i - 1/2| \leq 1/p} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} \otimes |\psi_j\rangle_{\mathsf{C}} + \sum_{\substack{i:|\lambda_i - 1/2| > 1/p \\ j:|\mu_j - 1/2| \leq 1/p}} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} \otimes |\psi_j\rangle_{\mathsf{C}} .$$

For convenience, we name the left part as $|\phi'_{\mathcal{B}}\rangle$ (indicating $\mathcal{B}$ can not win) and the right part as $|\phi'_{\mathcal{C}}\rangle$ (indicating $\mathcal{C}$ can not win). Thus, for every polynomial $p(\cdot)$, there exists a negligible function

$\mathsf{negl}(\cdot)$, $|\,|\phi_{\mathsf{BC}}\rangle - (|\phi'_{\mathcal{B}}\rangle + |\phi'_{\mathcal{C}}\rangle)|_1$ is at most $\mathsf{negl}(\cdot)$ (in expectation, taken the randomness of A, s, s', r and $H_{-s,s'}$).

The probability that $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins is at most:

$$(\left|(\Pi_0^B \otimes \Pi_0^C)\,|\phi'_{\mathsf{BC}}\rangle\right|^2 + \left|(\Pi_1^B \otimes \Pi_1^C)\,|\phi'_{\mathsf{BC}}\rangle\right|^2)/2.$$

$\Pi_0^B \otimes \Pi_0^C$ is the case that they both get access to H_0 and $\Pi_1^B \otimes \Pi_1^C$ for H_1.

The probability is at most

$$\begin{aligned}
&(\left|(\Pi_0^B \otimes \Pi_0^C)(|\phi'_{\mathcal{B}}\rangle + |\phi'_{\mathcal{C}}\rangle)\right|^2 + \left|(\Pi_1^B \otimes \Pi_1^C)(|\phi'_{\mathcal{B}}\rangle + |\phi'_{\mathcal{C}}\rangle)\right|^2)/2 \\
=&\frac{1}{2}\cdot\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{B}}\rangle + \langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{B}}\rangle + \langle\phi'_{\mathcal{C}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + \langle\phi'_{\mathcal{C}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right) \\
+&\mathsf{Re}\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + \langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right) \\
\leq&\frac{1}{2}\cdot\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes I)|\phi'_{\mathcal{B}}\rangle + \langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes I)|\phi'_{\mathcal{B}}\rangle + \langle\phi'_{\mathcal{C}}|(I \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + \langle\phi'_{\mathcal{C}}|(I \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right) \\
+&\mathsf{Re}\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + \langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right).
\end{aligned}$$

We bound each term separately.

- $\frac{1}{2}\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes I)|\phi'_{\mathcal{B}}\rangle + \langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes I)|\phi'_{\mathcal{B}}\rangle\right)$. It is equal to $\langle\phi'_{\mathcal{B}}|(\Pi_0^B + \Pi_1^B)/2 \otimes I|\phi'_{\mathcal{B}}\rangle$; by the definition of $|\phi'_{\mathcal{B}}\rangle$, it will be at most $(\frac{1}{2} + \frac{1}{p})|\,|\phi'_{\mathcal{B}}\rangle\,|^2$.
- $\frac{1}{2}\left(\langle\phi'_{\mathcal{C}}|(I \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + \langle\phi'_{\mathcal{C}}|(I \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right)$. Similar to the above case, it is at most $(\frac{1}{2}+\frac{1}{p})|\,|\phi'_{\mathcal{C}}\rangle\,|^2$.
- $\mathsf{Re}\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle\right)$. By Corollary 2.4, the inner product will be 0:

  $$\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle = \sum_{i:|\lambda_i - 1/2| \leq 1/p} \sum_{\substack{i':|\lambda_{i'} - 1/2| > 1/p \\ j':|\mu_{j'} - 1/2| \leq 1/p}} \alpha_{i,j}^\dagger \alpha_{i',j'} \langle\phi_i|\Pi_0^B|\phi_{i'}\rangle\langle\psi_j|\Pi_0^C|\psi_{j'}\rangle;$$

  since every possible i, i' satisfy $\lambda_i + \lambda_{i'} \neq 1$, we have $\langle\phi_i|\Pi_0^B|\phi_{i'}\rangle = 0$.
- $\mathsf{Re}\left(\langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right)$. By Corollary 2.4, the inner product will be 0 as well.

Therefore, the total probability will be at most $\left(\frac{1}{2} + \frac{1}{p}\right)(|\,|\phi'_{\mathcal{B}}\rangle\,|^2 + |\,|\phi'_{\mathcal{C}}\rangle\,|^2) + \mathsf{negl}(n) \leq \frac{1}{2} + \frac{1}{p} + \mathsf{negl}(n)$.

Since the above statement holds for every polynomial $p(\cdot)$, it finishes the proof for Theorem 4.8. □

Finally, we give the proof for Lemma 4.12.

Proof of Lemma 4.12. We prove by contradiction: suppose there exists an adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ such that the weight, which we call W, is non-negligible, i.e. $W > 1/q(\lambda)$ for some polynomial $q(\cdot)$, with some non-negligible probability $\eta(\lambda)$. For convenience, we will omit λ in the proof when it is clear from the context.

We construct the following adversary $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ that breaks the regular MOE game in Definition 4.5:

1. $\mathcal{A}', \mathcal{B}', \mathcal{C}'$ get (quantum) oracle access to P_{A+s} and $P_{A^\perp+s'}$.

2. $\mathcal{A}'$ first receives Δ from simulated $\mathcal{A}$, it samples $r \in \{0,1\}^{n(\lambda)}$ and a random oracle H. Given $|A_{s,s'}\rangle$, r and two membership checking oracles, it simulates $\mathcal{A}$ via reprogrammed $H_{s,s'}$, and produces $|\phi_{\mathsf{BC}}\rangle$; it gives B to $\mathcal{B}'$ and C to $\mathcal{C}'$.

   Note that, although H is a total random oracle, we will later reprogram H at the input (s, s'). Thus, H will only serve as $H_{-s,s'}$. Since $\mathcal{A}'$ does not know (s, s'), it is hard for $\mathcal{A}'$ to only sample $H_{-s,s'}$.

3. Define two projectors Π_0^B, Π_1^B over register B as what we have described at the beginning of the proof, with the random oracle $H^0_{s,s'}$ and $H^1_{s,s'}$ is defined as:

$$H^0_{s,s'}(z,z') = \begin{cases} r & \text{if } z = s, z' = s' \\ H(z,z') & \text{otherwise} \end{cases},$$

   and

$$H^1_{s,s'}(z,z') = \begin{cases} r \oplus \Delta & \text{if } z = s, z' = s' \\ H(z,z') & \text{otherwise} \end{cases}.$$

   Given $P_{A+s}, P_{A^\perp+s'}$ and the description of A, one can efficiently implement point functions that check the canonical vectors s and s'; thus, additionally given H, $H^0_{s,s'}$ and $H^1_{s,s'}$ can also be efficiently simulated. Therefore, $\mathcal{B}'$ can implement both Π_0^B, Π_1^B efficiently.

   $\mathcal{B}'$ gets B, it applies the efficient approximate threshold measurement $\mathsf{SATI}^{\epsilon,\delta}_{(P,Q),\gamma}$ in Theorem 2.8 with $P = (\Pi_0^B + \Pi_1^B)/2$, $Q = I - P$, $\gamma = 3/4p$, $\epsilon = 1/4p$ and $\delta = 2^{-\lambda}$.

   If the outcome is 1, $\mathcal{B}'$ then runs $\mathcal{B}$ on the leftover state with H_0 or H_1 picked uniformly at random. It measures and outputs a random query $\mathcal{B}$ makes to the random oracle.

4. Similarly define Π_0^C, Π_1^C as above on register C. $\mathcal{C}'$ gets C, it applies the efficient approximated threshold measurement $\mathsf{SATI}^{\epsilon,\delta}_{(P,Q),\gamma}$ with $P = (\Pi_0^C + \Pi_1^C)/2$, $Q = I - P$, $\gamma = 3/4p$, $\epsilon = 1/2p$, and $\delta = 2^{-\lambda}$.

   When the outcome is 1, $\mathcal{C}'$ runs $\mathcal{C}$ on the leftover state with H_0 or H_1 picked uniformly at random. It measures and outputs a random query to the random oracle.

By Theorem 2.8 bullet (1), conditioned on $W \geq 1/q$, both $\mathcal{B}'$ and $\mathcal{C}'$ will get outcome 1 with probability $1/q - 2\delta = O(1/q)$. When both outcomes are 1, by bullet (2) of Theorem 2.8, the leftover state is 4δ-close to the the following state:

$$\sum_{\substack{i:|\lambda_i - 1/2| > 1/4p \\ j:|\mu_j - 1/2| > 1/4p}} \beta_{i,j} |\phi_i\rangle_{\mathsf{B}} \otimes |\psi_j\rangle_{\mathsf{C}}.$$

Observe that when $\mathcal{B}$ does not query (s, s'), it will succeed with probability exactly $1/2$. Therefore, by Theorem 2.1, the query weight of $\mathcal{B}$ on (s, s') is at least $1/4p^2T - \mathsf{negl}(\lambda)$, where T is an upper-bound on the number of queries made by $\mathcal{B}$. Arguing similarly for $\mathcal{C}$, we conclude that the adversary $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ wins with probability at least $O(\eta/(qp^4T^2))$, which is non-negligible.

□

5 Unclonable Encryption in the QROM

The following is the unclonable encryption scheme for a single bit:

1. $\mathsf{sk} = A$ where A is a random subspace $A \subseteq \mathbb{F}_2^n$ of dimension $n/2$;
2. $\mathsf{Enc}^H(\mathsf{sk}, m)$: it samples $s \leftarrow \mathsf{CS}(A)$ and $s' \leftarrow \mathsf{CS}(A^\perp)$ uniformly at random; it outputs $|A_{s,s'}\rangle$, $c = H(s, s') \oplus m$;
3. $\mathsf{Dec}^H(\mathsf{sk} = A, (|A_{s,s'}\rangle, c))$:
   - It first computes s in superposition. We know that there is a classical algorithm that on any vector in $A + s$ and the description of A, outputs the canonical vector of $A + s$ (which is s in this case). See [CLLZ21] Definition 4.3 for more references.
     We can run this classical algorithm coherently on $|A_{s,s'}\rangle$ to learn s.
   - Since the algorithm on any vector in $A + s$ outputs the same vector, the quantum state stays intact. We can run the same algorithms coherently on the Hadamard basis and the description of $A^\perp$ to learn s'.
   - Output $c \oplus H(s, s')$.

With Theorem 4.8, we can show the scheme satisfy the unclonable IND-CPA security.

Proof. If we have some adversary $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ for the scheme beyond, we can construct an adversary $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ for the strengthened MOE game with the same advantage.

- The adversary $\mathcal{A}'$ gets $(m_0, m_1) \leftarrow \mathcal{A}$ and sends $\Delta = m_0 \oplus m_1$ to the challenger.
- After receiving $|A_{s,s'}\rangle$ and w from the challenger, $\mathcal{A}'$ calculates $c = w \oplus m_0$, and sends $(|A_{s,s'}\rangle, c)$ to $\mathcal{A}$. The output registers B, C of $\mathcal{A}$ are sent to $\mathcal{B}', \mathcal{C}'$ respectively.
- $\mathcal{B}', \mathcal{C}'$ exactly run the algorithm of $\mathcal{B}, \mathcal{C}$, and output their output respectively.

Thus we have concluded the unclonable IND-CPA security of our game.

□

Remark 5.1. *Notice that compared to the strengthened MOE game, our construction does not provide additional membership checking oracles.*

6 Copy-Protection for Point Functions in QROM

6.1 Copy-Protection Preliminaries

Below we present the definition of a copy-protection scheme.

Definition 6.1 (Copy-Protection Scheme)**.** *Let $\mathcal{F} = \mathcal{F}(\lambda)$ be a class of efficiently computable functions of the form $f : X \to Y$. A copy protection scheme for $\mathcal{F}$ is a pair of QPT algorithms* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *such that:*

- ***Copy Protected State Generation:*** $\mathsf{CopyProtect}(1^\lambda, d_f)$ *takes as input the security parameter* 1^λ *and a classical description* d_f *of a function* $f \in \mathcal{F}$ *(that efficiently computes* f*). It outputs a mixed state* $\rho_f \in \mathcal{D}(\mathcal{H}_Z)$*, where* Z *is the output register.*
- ***Evaluation:*** $\mathsf{Eval}(1^\lambda, \rho, x)$ *takes as input the security parameter* 1^λ*, a mixed state* $\rho \in \mathcal{D}(\mathcal{H}_Z)$*, and an input value* $x \in X$*. It outputs a bipartite state* $\rho' \otimes |y\rangle\langle y| \in \mathcal{D}(\mathcal{H}_Z) \otimes \mathcal{D}(\mathcal{H}_Y)$*.*

We will sometimes abuse the notation and write $\mathsf{Eval}(1^\lambda, \rho, x)$ to denote the classical output $y \in Y$ when the residual state ρ' is not significant.

Definition 6.2 (Correctness)**.** *A copy-protection scheme* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *for* $\mathcal{F}$ *is* δ***-correct*** *if the following holds: for every* $x \in X$*,* $f \in \mathcal{F}$*,*

$$\Pr\left[f(x) \leftarrow \mathsf{Eval}(1^\lambda, \rho_f, x) \ : \ \rho_f \leftarrow \mathsf{CopyProtect}(1^\lambda, d_f)\right] \geq \delta.$$

If $\delta \geq 1 - \mathsf{negl}(\lambda)$*, we simply say that the scheme is* ***correct****.*

Remark 6.3. *When* δ *is negligibly close to 1, the evaluation algorithm* Eval *can be implemented so that it does not disturb the state* ρ_f*. This ensures that* ρ_f *can be reused polynomially many times with arbitrary inputs.*

We define security via a piracy experiment.

Definition 6.4 (Piracy Experiment)**.** *A* ***piracy experiment*** *is a security game defined by a copy-protection scheme* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *for a class of functions* $\mathcal{F}$ *of the form* $f : X \to Y$*, a distribution* $\mathcal{D}_\mathcal{F}$ *over* $\mathcal{F}$*, and a class of distributions* $\mathfrak{D}_X = \{\mathfrak{D}_X(f)\}_{f\in\mathcal{F}}$ *over* $X \times X$*. It is the following game between a challenger and an adversary, which is a triplet of algorithms* $(\mathcal{A}, \mathcal{B}, \mathcal{C})$*:*

- ***Setup Phase:*** *The challenger samples a function* $f \leftarrow \mathcal{D}_\mathcal{F}$ *and sends* $\rho_f \leftarrow \mathsf{CopyProtect}(1^\lambda, d_f)$ *to* $\mathcal{A}$*.*
- ***Splitting Phase:*** $\mathcal{A}$ *applies a CPTP map to split* ρ_f *into a bipartite state* ρ_{BC}*; it sends the* B *register to* $\mathcal{B}$ *and the* C *register to* $\mathcal{C}$*. No communication is allowed between* $\mathcal{B}$ *and* $\mathcal{C}$ *after this phase.*
- ***Challenge Phase:*** *The challenger samples* $(x_B, x_C) \leftarrow \mathfrak{D}_X(f)$ *and sends* x_B, x_C *to* $\mathcal{B}, \mathcal{C}$*, respectively.*
- ***Output Phase:*** $\mathcal{B}$ *and* $\mathcal{C}$ *output* $y_B \in Y$ *and* $y_C \in Y$*, respectively, and send to the challenger. The challenger outputs 1 if* $y_B = f(x_B)$ *and* $y_C = f(x_C)$*, indicating that the adversary has succeeded, and 0 otherwise.*

The bit output by the challenger is denoted by $\mathsf{PirExp}^{\mathsf{CopyProtect},\mathsf{Eval}}_{\mathcal{D}_\mathcal{F},\mathfrak{D}_X}(1^\lambda, (\mathcal{A}, \mathcal{B}, \mathcal{C}))$*.*

As noted by [CMP20], the adversary can always succeed in this game with probability negligibly close to

$$p^{\mathsf{triv}}(\mathcal{D}_\mathcal{F}, \mathfrak{D}_X) := \max_{E\in\{B,C\}} \mathop{\mathbb{E}}_{\substack{f\leftarrow\mathcal{D}_\mathcal{F} \\ (x_B,x_C)\leftarrow\mathfrak{D}_X(f)}} \max_{y\in Y} \Pr\left[y \mid x_E\right]$$

by sending ρ_f to $\mathcal{B}$ and have $\mathcal{C}$ guess the most likely output y given input x_C (or vice versa). In other words, p^{triv} is the success probability of optimal guessing strategy for one party $E \in \{B, C\}$ given only the test input x_E.

Bounding the success probability of the adversary is bounded by p^{triv} captures the intuition that ρ_f is no more helpful for simultaneous evaluation than a black-box program that could only be given to one party.

Definition 6.5 (Copy-Protection Security)**.** *Let* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *be a copy-protection scheme for a class* $\mathcal{F}$ *of functions* $f : X \to Y$*. Let* $\mathcal{D}_{\mathcal{F}}$ *be a distribution over* $\mathcal{F}$ *and* $\mathfrak{D}_X = \{\mathfrak{D}_X(f)\}_{f\in\mathcal{F}}$ *a class of distributions over* X*. Then,* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *is called* $(\mathcal{D}_{\mathcal{F}}, \mathfrak{D}_X)$-***secure*** *if there exists a negligible function* negl *such that any QPT adversary* $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ *satisfies*

$$\Pr\left[b = 1 \ : \ b \leftarrow \mathsf{PirExp}^{\mathsf{CopyProtect},\mathsf{Eval}}_{\mathcal{D}_{\mathcal{F}},\mathfrak{D}_X}\left(1^\lambda, (\mathcal{A}, \mathcal{B}, \mathcal{C})\right)\right] \leq p^{\mathsf{triv}}(\mathcal{D}_{\mathcal{F}}, \mathfrak{D}_X) + \mathsf{negl}(\lambda).$$

Copy Protection for Point Functions A *point function* $f_y : \{0,1\}^m \to \{0,1\}$ is of the form

$$f_y(x) = \begin{cases} 1, & x = y \\ 0, & x \neq y \end{cases}.$$

When dealing with point functions, the classical description of f_y will simply be y, and accordingly the distribution $\mathcal{D}_{\mathcal{F}}$ over point functions will be represented by a distribution $\mathcal{D} = \mathcal{D}_\lambda$ over $\{0,1\}^m$. Since copy protection is trivially impossible for a learnable distribution $\mathcal{D}$, we are going to restrict our attention to unlearnable distributions.

Definition 6.6. *A distribution* $\mathcal{D}_\lambda$ *over* $\{0,1\}^m$*, with* $m = \mathsf{poly}(\lambda)$*, is called* ***unlearnable*** *if for any query-bounded adversary* $\mathcal{A}^{f_y(\cdot)}$ *with oracle access to* $f_y(\cdot)$*, we have*

$$\Pr\left[y' = y : \substack{y\leftarrow\mathcal{D}_\lambda \\ y'\leftarrow\mathcal{A}^{f_y(\cdot)}(1^\lambda)}\right] \leq \mathsf{negl}(\lambda).$$

Definition 6.7 (Copy-Protection Security for Point Functions)**.** *Let* $m = \mathsf{poly}(\lambda)$ *and* $\mathcal{F}$ *be the class of point functions* $f_y : \{0,1\}^m \to \{0,1\}$*. Let* $\mathfrak{D}_X = \{\mathfrak{D}_X(f)\}_{f\in\mathcal{F}}$ *be a class of input distributions over* $\{0,1\}^m \times \{0,1\}^m$*. A copy protection scheme* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *for* $\mathcal{F}$ *is called* $\mathfrak{D}_X$-***secure*** *if there exists a negligible function* negl *such that* $(\mathsf{CopyProtect}, \mathsf{Eval})$ *is* $(\mathcal{D}_\lambda, \mathfrak{D}_X)$-***secure*** *for all unlearnable distributions* $\mathcal{D}_\lambda$ *over* $\{0,1\}^m$*.*

6.2 Construction

In this section, we design copy-protection for a class of point functions. We set $n = 2\lambda$ and $d = \lambda$ throughout the section. Our construction will use two hash functions: (a) $G : \{0,1\}^\lambda \to \{0,1\}^{n\cdot d}$ and (b) $H : \mathbb{F}_2^n \times \mathbb{F}_2^n \to \{0,1\}^{4n+\lambda}$. In the security proof, we will treat G and H as random oracles. We will use $\mathbb{F}_2^n$ and $\{0,1\}^n$ interchangeably.

We denote the set of all d-dimensional subspaces of $\mathbb{F}_2^n$ by $\mathcal{S}_d$. We will need the following lemma for correctness.

Lemma 6.8. *There exists a set of efficient unitaries* $\{U_{A'}\}_{A'\in\mathcal{S}_d} \subseteq \mathcal{U}(\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Z}} \otimes \mathcal{H}_{\mathbf{anc}})$*, where* $\mathbf{X}, \mathbf{Z}, \mathbf{anc}$ *are registers of length* $n, 2n, \mathsf{poly}(\lambda)$*, such that the following holds for any* $A \in \mathcal{S}_d$*:*

- *For any* $s \in \mathsf{CS}(A)$*,* $s' \in \mathsf{CS}(A^\perp)$*, we have* $U_A \left|A_{s,s'}\right\rangle \left|0^{2n}\right\rangle_{\mathbf{Z}} \left|0^{\mathsf{poly}(\lambda)}\right\rangle_{\mathbf{anc}} = \left|A_{s,s'}\right\rangle \left|s, s'\right\rangle_{\mathbf{Z}} \left|0^{\mathsf{poly}(\lambda)}\right\rangle_{\mathbf{anc}}$.

- *For any $A' \in \mathcal{S}_d$ such that $\frac{|A' \cap A|}{2^d} \leq \nu(\lambda)$, for some negligible function $\nu(\cdot)$, there exists a negligible function $\nu'(\lambda)$ such that the following holds for all $s \in \mathsf{CS}(A), s' \in \mathsf{CS}(A^\perp)$:*

$$\left\| (I_{\mathbf{X}} \otimes |s,s'\rangle\langle s,s'|_{\mathbf{Z}} \otimes I_{\mathbf{anc}}) \left(U_{A'} |A_{s,s'}\rangle |0^{2n}\rangle_{\mathbf{Z}} |0^{\mathsf{poly}(\lambda)}\rangle_{\mathbf{anc}} \right) \right\|^2 \leq \nu'(\lambda).$$

Proof. To get unitaries satisfying the first bullet, recall that there exists an efficient procedure which computes $\mathsf{Can}_A(\cdot)$ given the description of A. We can represent this procedure by a unitary U followed by measurement of s, s'. We describe U_A as follows:

1. Apply U to the $\mathbf{X}, \mathbf{anc}$ registers. Copy the answer to the first half of the $\mathbf{Z}$ register. Note that the answer is always s since $|A_{s,s'}\rangle$ is a superposition of vectors in $A + s$.
2. Apply $U^\dagger$ to the $\mathbf{X}, \mathbf{anc}$ registers.
3. Apply QFT on the $\mathbf{X}$ register to obtain $|A^\perp_{s',s}\rangle$.
4. Repeat the first two steps and copy the answer s' to the second half of the $\mathbf{Z}$ register.
5. Appy QFT again to recover $|A_{s,s'}\rangle_{\mathbf{X}}$

We will show that the second bullet follows from the first bullet. We first observe that the inner product between the coset states $|A_{s,s'}\rangle$ and $|A'_{s,s'}\rangle$ is small. Indeed, since $|(A+s) \cap (A'+s)| = |A \cap A'| \leq 2^d \nu(\lambda)$, we have

$$\begin{aligned} \left|\langle A_{s,s'}|A'_{s,s'}\rangle\right|^2 &= \left| \frac{1}{\sqrt{|A|}} \sum_{a \in A} (-1)^{\langle s',a\rangle} \langle a+s| \frac{1}{\sqrt{|A|}} \sum_{a' \in A'} (-1)^{\langle s',a'\rangle} |a'+s\rangle \right|^2 \\ &\leq \left| \frac{1}{|A|} 2^d \nu(\lambda) \right|^2 = \nu(\lambda)^2. \end{aligned}$$

Fix $U_{A'}$. Recall that the coset states $\left\{ |A'_{t,t'}\rangle \right\}_{t \in \mathsf{CS}(A), t' \in \mathsf{CS}(A^\perp)}$ form an orthonormal basis. By the first bullet, we have

$$\begin{aligned} &\left\| (I_{\mathbf{X}} \otimes |s,s'\rangle\langle s,s'|_{\mathbf{Z}} \otimes I_{\mathbf{anc}}) \left(U_{A'} |A_{t,t'}\rangle_{\mathbf{X}} |0^{2n}\rangle_{\mathbf{Z}} |0^{\mathsf{poly}(\lambda)}\rangle_{\mathbf{anc}} \right) \right\|^2 \\ = &\left\| (I_{\mathbf{X}} \otimes |s,s'\rangle\langle s,s'|_{\mathbf{Z}} \otimes I_{\mathbf{anc}}) \left(|A_{t,t'}\rangle_{\mathbf{X}} |t,t'\rangle_{\mathbf{Z}} |0^{\mathsf{poly}(\lambda)}\rangle \right) \right\|^2 = 0 \end{aligned}$$

for any $(t,t') \neq (s,s')$. Therefore, we have

$$\begin{aligned} &\left\| (I_{\mathbf{X}} \otimes |s,s'\rangle\langle s,s'|_{\mathbf{Z}} \otimes I_{\mathbf{anc}}) \left(U_{A'} |A_{s,s'}\rangle_{\mathbf{X}} |0^{2n}\rangle_{\mathbf{Z}} |0^{\mathsf{poly}(\lambda)}\rangle_{\mathbf{anc}} \right) \right\|^2 \\ = &\left\| \sum_{\substack{t \in \mathsf{CS}(A) \\ t' \in \mathsf{CS}(A^\perp)}} (I_{\mathbf{X}} \otimes |s,s'\rangle\langle s,s'|_{\mathbf{Z}} \otimes I_{\mathbf{anc}}) \left(U_{A'} |A'_{t,t'}\rangle \langle A'_{t,t'}|A_{s,s'}\rangle_{\mathbf{X}} |0^{2n}\rangle_{\mathbf{Z}} |0^{\mathsf{poly}(\lambda)}\rangle_{\mathbf{anc}} \right) \right\|^2 \\ = &\left|\langle A_{s,s'}|A'_{s,s'}\rangle\right|^2 \leq \nu(\lambda)^2. \end{aligned}$$

as desired. □

Construction. We describe the copy-protection scheme $(\mathsf{CopyProtect}, \mathsf{Eval})$ for a class of point functions $\mathcal{F} = \{f_y(\cdot)\}_{y\in\{0,1\}^\lambda}$ as follows:

- $\mathsf{CopyProtect}\left(1^\lambda, y\right)$: it takes as input λ in unary notation, $y \in \{0,1\}^\lambda$ and does the following:
  1. Compute $\mathbf{v} = G(y)$. Parse $\mathbf{v}$ as a concatenation of d vectors $v_1, \ldots, v_d$, where each v_i has dimension n. Abort if the vectors $\{v_1, \ldots, v_d\}$ are not linearly independent.
  2. Let $A = \mathrm{Span}\,(v_1, \ldots, v_d)$.
  3. Sample $s \leftarrow \mathsf{CS}(A)$ and $s' \leftarrow \mathsf{CS}(A^\perp)$ uniformly at random.
  4. Output the copy-protected state $\sigma = |A_{s,s'}\rangle\langle A_{s,s'}|_{\mathbf{X}} \otimes |H(s,s')\rangle\langle H(s,s')|_{\mathbf{Y}}$.
- $\mathsf{Eval}(\sigma, x)$: on input the copy-protected state $\sigma \in \mathcal{D}(\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}})$, input $x \in \{0,1\}^\lambda$, it does the following:
  1. Measure the register $\mathbf{Y}$ of σ to obtain the value θ. Call the resulting state σ'.
  2. Compute $\mathbf{v} = G(x)$. Parse $\mathbf{v}$ as a concatenation of d vectors $v_1, \ldots, v_d$, where each v_i has dimension n. Abort if the vectors $\{v_1, \ldots, v_d\}$ are not linearly independent.
  3. Let $A = \mathrm{Span}\,(v_1, \ldots, v_d)$.
  4. Apply U_A (defined in Lemma 6.8) coherently on $\sigma' \otimes |0^{2n}\rangle\langle 0^{2n}|_{\mathbf{Z}} \otimes |0^{\mathsf{poly}(\lambda)}\rangle\langle 0^{\mathsf{poly}(\lambda)}|_{\mathbf{anc}}$ to obtain the state σ''.
  5. Query H on the register $\mathbf{Z}$ and store the answer in a new register $\mathbf{out}$.
  6. Measure the register $\mathbf{out}$ in the computational basis. Denote the post-measurement state by $\sigma_{\mathbf{out}}$ and the measurement outcome by θ'.
  7. If $\theta = \theta'$, output $\sigma_{\mathbf{out}} \otimes |1\rangle\langle 1|$. Otherwise, output $\sigma_{\mathbf{out}} \otimes |0\rangle\langle 0|$.

We first discuss at a high level why this construction works. Regarding correctness, we argue that Eval on input $x \neq y$ computes a random subspace A', such that $|A'_{s,s'}\rangle$ is nearly orthogonal to $|A_{s,s'}\rangle$. As a result, Eval recovers (s,s') incorrectly. Since as a sufficiently expanding hash function H is injective with high probability, Eval fails.

As for security, first we show that it is hard for $\mathcal{A}$ to query the oracles G, H on inputs $y, (s,s')$. Next, we argue that $\mathcal{B}$ and $\mathcal{C}$ cannot both recover (s,s'), otherwise they break the MOE game in Theorem 4.6.

Most meaningful input distributions $\mathfrak{D}_X(y)$ for a point function f_y can be parameterized by a triple (p,q,r):

- With probability p, output (y,y)
- With probability q, output (y, x_C), where $x_C \neq y$ is a random string.
- With probability r, output (x_B, y), where $x_B \neq y$ is a random string.
- With probability $1-p-q-r$, output (x_B, x_C), where $x_B, x_C \neq y$ are random strings.

We show that our scheme is secure with respect to product distributions, i.e. when $(p, q, r, 1-p-q-r)$ is of the form (pp', pq', qp', qq') with $p+q = p'+q' = 1$, in Lemma 6.13. We also show security for maximally correlated input distributions, i.e. when $q = r = 0$, in Corollary 6.20. The way the random strings x_B, x_C are sampled (uniformly or otherwise) turns out to be inconsequential in our security proof.

We give the formal statements below.

Lemma 6.9. $(\mathsf{CopyProtect}, \mathsf{Eval})$ *satisfies correctness.*

Proof of Lemma 6.9. We first argue that step 1 of CopyProtect aborts only with negligible probability:

Claim 6.10. *Let $n = 2d = 2\lambda$ and $v_1, v_2, \ldots, v_d \in \mathbb{F}_2^n$ be uniformly random independent vectors, then there exists a negligible function ν_0 such that $v_1, \ldots, v_d$ are linearly dependent with probability at most $\nu_0(\lambda)$.*

Proof. Let p_i be the probability that $\{v_1, \ldots, v_i\}$ is linearly independent given that $\{v_1, \ldots, v_{i-1}\}$ is linearly independent. Since v_i is uniformly random and the span of $\{v_1, \ldots, v_{i-1}\}$ has size 2^{i-1}, we have $p_i = 1 - 2^{i-1}/2^n$. Thus, the probability that $\{v_1, \ldots, v_d\}$ is linearly independent is given by

$$\prod_{i=1}^{d} p_i = \prod_{i=1}^{d} (1 - 2^{i-1-n}) \geq \left(1 - 2^{-\lambda}\right)^{\lambda} \geq 1 - \lambda 2^{-\lambda},$$

where we used the union bound in the last step. Hence, the claim holds for $\nu_0(\lambda) = \lambda 2^{-\lambda}$. □

We will condition on step 1 of CopyProtect not aborting henceforth. Let $y \in \{0,1\}^\lambda$ and $\sigma \leftarrow \mathsf{CopyProtect}\left(1^\lambda, y\right)$. Note that σ is of the form $|A_{s,s'}\rangle\langle A_{s,s'}|_{\mathbf{X}} \otimes |\theta\rangle\langle\theta|_{\mathbf{Y}}$, where the following holds:

1. $\mathbf{v} = G(y)$ and $\mathbf{v}$ is a concatenation of d linearly independent vectors $v_1, \ldots, v_d$
2. $A = \mathrm{Span}\,(v_1, \ldots, v_d)$
3. $s \in \mathsf{CS}(A)$ and $s' \in \mathsf{CS}(A^\perp)$ are selected uniformly at random.
4. $\theta = H(s, s')$

We now consider the two cases: $x = y$ and $x \neq y$.

Case 1. $\mathsf{Eval}(\sigma, y) = \sigma_{\mathbf{out}} \otimes |1\rangle\langle 1|$*, for some state* $\sigma_{\mathbf{out}}$. If we follow the first four steps of $\mathsf{Eval}(\sigma, y)$, we will end up with the subspace A (defined above). From Lemma 6.8, we have the following: after applying U_A on $|A_{s,s'}\rangle\,|0\rangle_{\mathbf{anc}}$, we obtain (s, s') in **anc** register. That is, **anc** register has the state $|s, s'\rangle$. After querying H on **anc**, the value stored in **out** is $H(s, s')$. Thus, measuring the register **out** yields the value $\theta' = H(s, s')$. Since $\theta' = \theta$, the output of $\mathsf{Eval}(\sigma, y)$ is $\sigma_{\mathbf{out}} \otimes |1\rangle\langle 1|$, where $\sigma_{\mathbf{out}}$ is the residual state.

Case 2. $\forall x \neq y$, $\mathsf{TD}(\mathsf{Eval}(\sigma, x), \sigma_{\mathbf{out}} \otimes |0\rangle\langle 0|) \leq \nu(\lambda)$*, for some negligible function* $\nu(\lambda)$ *and some state* $\sigma_{\mathbf{out}}$. To prove this, it suffices to show that the probability that $\mathsf{Eval}(\sigma, x)$ outputs 1 is negligible in λ. Consider the following claim:

Claim 6.11. *If $H : \{0,1\}^{2n} \to \{0,1\}^{4n+\lambda}$ is picked uniformly at random, the probability that H is not injective is at most $\nu_1(\lambda)$, for some negligible function $\nu_1(\cdot)$.*

Proof. For any $a \neq b \in \{0,1\}^{2n}$, the probability that $H(a) = H(b)$ is $\frac{1}{2^{4n+\lambda}}$. By a union bound argument, the probability that H is not injective is at most $\frac{\binom{2^{2n}}{2}}{2^{4n+\lambda}} \leq \frac{2^{4n}}{2^{4n+\lambda}} = \frac{1}{2^{\lambda}}$. □

Let us condition on the event that H is injective. We consider the first four steps of execution of $\mathsf{Eval}(\sigma, x)$:

- Measure the register $\mathbf{Y}$ of σ to obtain the value θ. Call the post-measurement state σ'.
- Compute $\mathbf{v} = G(x)$. Parse $\mathbf{v}$ as a concatenation of d vectors $v_1, \ldots, v_d$, where each v_i has dimension n. Abort if the vectors $\{v_1, \ldots, v_d\}$ are not linearly independent.
- Let $A' = \mathrm{Span}\,(v_1, \ldots, v_d)$.

Consider the following claim.

Claim 6.12. *If $x \neq y$, then there exists a negligible function $\nu_2(\lambda)$ such that the probability (over the coins of G) that $\frac{|A' \cap A|}{2^d} \leq \nu_2(\lambda)$ holds is at least $1 - \nu_2(\lambda)$.*

Proof. Since $x \neq y$ and G is a random oracle, A and A' are independently sampled. By Claim 6.10, A and A' are uniformly random independent subspaces of dimension d each with probability at least $1 - 2\nu_0(\lambda)$. Conditioned on this, we can bound the expected size of their intersection as

$$\mathbb{E}\left[|A \cap A'|\right] = \sum_{v \in \mathbb{F}_2^n} \Pr\left[v \in A \cap A'\right] = \sum_{v \in \mathbb{F}_2^n} \Pr\left[v \in A\right]^2 = 1 + (2^n - 1)\left(2^{d-n}\right)^2 < 2. \tag{4}$$

Let $\nu_2(\lambda) = 2^{-\lambda/5} + 2\nu_0(\lambda)$. Then, by Markov's Inequality and eq. (4) we have

$$\Pr\left[\frac{|A' \cap A|}{2^d} > \nu_2(\lambda)\right] \leq 2\nu_0(\lambda) + \frac{\mathbb{E}\left[|A' \cap A|\right]}{2^d \nu_2(\lambda)} < 2\nu_0(\lambda) + 2^{-\lambda/2} < \nu_2(\lambda).$$

□

We will condition on the event that $\frac{|A' \cap A|}{2^d} \leq \nu_2(\lambda)$. By Lemma 6.8, we have that

$$p := \left\| (I_{\mathbf{X}} \otimes |s, s'\rangle\langle s, s'|_{\mathbf{Z}} \otimes I_{\mathbf{anc}}) \left(U_{A'} |A_{s,s'}\rangle |0^{2n}\rangle_{\mathbf{Z}} |0^{\mathsf{poly}(\lambda)}\rangle_{\mathbf{anc}} \right) \right\|^2 \leq \nu_3(\lambda).$$

for some negligible function $\nu_3(\lambda)$. Since we have conditioned on the event that H is injective, the probability that $\mathsf{Eval}(1^{\lambda}, \sigma, x)$ outputs 1 is given by

$$\left\| (I_{\mathbf{X},\mathbf{Y},\mathbf{Z},\mathbf{anc}} \otimes |H(s,s')\rangle\langle H(s,s')|_{\mathbf{out}}) \left(I_{\mathbf{X},\mathbf{Z},\mathbf{anc}} \otimes O^H \right) U_{A'} |A_{s,s'}\rangle_{\mathbf{X}} |0^{\mathsf{poly}(\lambda)}\rangle_{\mathbf{anc}} |0\rangle_{\mathbf{out}} \right\|^2 = p,$$

where O^H is the unitary that computes H. Combining this with Claim 6.10, Claim 6.11 and Claim 6.12, we conclude that for any $x \neq y$,

$$\begin{aligned}
\Pr\left[1 \leftarrow \mathsf{Eval}(1^\lambda, \sigma, x)\right] &\leq \Pr\left[1 \leftarrow \mathsf{Eval}(1^\lambda, \sigma, x) \mid \substack{\frac{|A \cap A'|}{2^d} \leq \nu_2(\lambda),\ H_2 \text{ is injective},\\ \mathsf{CopyProtect}(1^\lambda, y) \text{ or } \mathsf{Eval}(1^\lambda, \sigma, x) \text{ doesn't abort}}\right] \\
&\quad + \Pr\left[\mathsf{CopyProtect}(1^\lambda, y) \text{ or } \mathsf{Eval}(1^\lambda, \sigma, x) \text{ aborts}\right] + \Pr\left[H_2 \text{ is not injective}\right] \\
&\quad + \Pr\left[\frac{|A \cap A'|}{2^d} > \nu_2(\lambda) \mid H_2 \text{ is injective}\right] \\
&\leq \nu_3(\lambda) + 2\nu_0(\lambda) + \nu_1(\lambda) + \nu_2(\lambda) \\
&\leq \mathsf{negl}(\lambda).
\end{aligned}$$

□

Lemma 6.13. $(\mathsf{CopyProtect}, \mathsf{Eval})$ *is a $\mathfrak{D}_X$-secure copy-protection scheme for point functions with input length λ, where $\mathfrak{D}_X(y) = \mathfrak{D}_y^B \times \mathfrak{D}_y^C$ is a product distribution.*

Proof of Lemma 6.13. Fix an unlearnable distribution $\mathcal{D} = \mathcal{D}_\lambda$. We will define a sequence of hybrids:

Hybrid 1. This is the real piracy experiment for $(\mathsf{CopyProtect}, \mathsf{Eval})$ defined in Definition 6.4, where $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ all have access to both random oracles G and H. The input to $\mathsf{CopyProtect}$ is denoted by $y \in \{0,1\}^\lambda$ as in the construction.

Hybrid 2. In this hybrid, we change, for $\mathcal{A}$ only, the oracle G to G_y, which is the punctured oracle defined as

$$G_y(x) = \begin{cases} u, & x = y \\ G(x), & x \neq y \end{cases},$$

where $u \in \{0,1\}^{nd}$ is a fresh uniformly random string.

Hybrid 3. In this hybrid, we have the challenger sample $A \subseteq \mathcal{S}_d$ uniformly at the start. Using this A, we change the oracle G for $\mathcal{B}$ and $\mathcal{C}$ both to G_y^A, which is the reprogrammed oracle defined as follows:

- Fix a random basis $(v_1, \ldots, v_d)$ of A.
- If $x = y$, then $G_y^A(x)$ outputs $(v_1, \ldots, v_d)$.
- If $x \neq y$, then $G_y^A(x)$ outputs $G(x)$.

Hybrid 4. In this hybrid, we change, for $\mathcal{A}$ only, the oracle H to the punctured oracle $H_{s,s'}$ defined as

$$H_{s,s'}(t,t') = \begin{cases} v, & (t,t') = (s,s') \\ H(t,t'), & (t,t') \neq (s,s') \end{cases},$$

where $v \in \{0,1\}^{4n+\lambda}$ is a fresh uniformly random string.

Let p_i be the probability that $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins in Hybrid i, and let $p^{\mathsf{triv}} = p^{\mathsf{triv}}(\mathcal{D}_\lambda, \mathfrak{D}_X)$. We will show the following lemmas about the hybrids:

Lemma 6.14. $|p_1 - p_2| \leq \mathsf{negl}(\lambda)$.

Lemma 6.15. $|p_2 - p_3| \leq \mathsf{negl}(\lambda)$.

Lemma 6.16. $|p_3 - p_4| \leq \mathsf{negl}(\lambda)$.

Lemma 6.17. $p_4 \leq p^{\mathsf{triv}} + \mathsf{negl}(\lambda)$.

Proof of Lemma 6.14. Let $\rho_{\mathsf{BC}}^{(i)}$ be the bipartite state sent by $\mathcal{A}$ to $\mathcal{B}$ and $\mathcal{C}$ in the ith Hybrid. We will show that $\mathsf{TD}\left(\rho_{\mathsf{BC}}^{(1)}, \rho_{\mathsf{BC}}^{(2)}\right) \leq \mathsf{negl}(\lambda)$. Since Hybrid 1 and Hybrid 2 are identical after the splitting phase and trace distance cannot increase by post-processing, this suffices to prove the lemma.

Suppose that $\mathsf{TD}\left(\rho_{\mathsf{BC}}^{(1)}, \rho_{\mathsf{BC}}^{(2)}\right)$ is non-negligible. Using $\mathcal{A}$ from Hybrid 2 we will construct an adversary $\mathcal{A}'$ which violates the unlearnability of $\mathcal{D}_\lambda$ (Definition 6.6) without using the oracle $f_y(\cdot)$.

- $\mathcal{A}'$ samples random oracles G, H, a random subspace $A \in S_d$, and random $(s, s') \in \mathsf{CS}(A) \times \mathsf{CS}(A^\perp)$.
- $\mathcal{A}'$ runs $\mathcal{A}$ on input $\left(G, H, |A_{s,s'}\rangle, H(s, s')\right)$. Then it measures a random query y' made by $\mathcal{A}$ to G, and outputs y'.

By Theorem 2.1, the probability $\Pr[y' = y]$ is non-negligible, thus $\mathcal{A}'$ breaks unlearnability.

□

Proof of Lemma 6.15. This easily follows by the fact that Hybrid 2 and Hybrid 3 are identical conditioned on the fact that $G(y)$ outputs a valid basis, which happens with overwhelming probability by Claim 6.10. □

Proof of Lemma 6.16. Similarly as before, it suffices to show $\mathsf{TD}\left(\rho_{\mathsf{BC}}^{(1)}, \rho_{\mathsf{BC}}^{(2)}\right) \leq \mathsf{negl}(\lambda)$. Suppose this is not the case, we will construct an adversary $\mathcal{A}'$ which breaks direct product hardness (Corollary 4.4) using $\mathcal{A}$ from Hybrid 4:

- $\mathcal{A}'$ receives $|A_{s,s'}\rangle$ from the challenger, where $(s, s') \in \mathsf{CS}(A) \times \mathsf{CS}(A^\perp)$. It samples random oracles G, H, and a random string $v \in \{0,1\}^{4n+\lambda}$.
- $\mathcal{A}'$ runs $\mathcal{A}$ on input $\left(G, H, |A_{s,s'}\rangle, v\right)$. It measures and outputs a random query (t, t') made to H during the execution.

By Theorem 2.1, the probability $\Pr[(t, t') = (s, s')]$ is non-negligible, thus $\mathcal{A}'$ breaks direct product hardness.

□

Proof of Lemma 6.17. We will use the same template as in the proof of Lemma 4.11. Let $\rho_{\mathsf{BC}} := \rho_{\mathsf{BC}}^{(4)}$ be the bipartite state created by $\mathcal{A}$. We can assume without loss of generality that $\rho_{\mathsf{BC}} := |\phi_{\mathsf{BC}}\rangle\langle\phi_{\mathsf{BC}}|$ is a pure state. Define POVM elements Π^B, Π^C as follows:

- Π^B: samples $x_B \leftarrow \mathfrak{D}_y^B$; it runs $\mathcal{B}$ on input oracles G_y^A, H and test input x_B; it measures if the output is $f_y(x_B)$; then it undoes all the computation.
- Π^C: defined similarly for $\mathcal{C}$.

Now we write the state in its spectral decomposition

$$|\phi_{\mathsf{BC}}\rangle = \sum_{i,j} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} |\psi_j\rangle_{\mathsf{C}},$$

where $|\phi_i\rangle_{\mathsf{B}}$ is an eigenvector of Π^B with eigenvalue λ_i and $|\psi_j\rangle_{\mathsf{C}}$ is an eigenvector of Π^B with eigenvalue μ_j. Let p_B^{triv} (p_C^{triv}) be the trivial guessing probability when $\mathcal{B}$ ($\mathcal{C}$) makes a blind guess, so that $p^{\mathsf{triv}} = \max(p_B^{\mathsf{triv}}, p_C^{\mathsf{triv}})$. We will need a lemma similar to Lemma 4.12.

Lemma 6.18. *Let $p(\cdot)$ be a polynomial. With overwhelming probability over $y, (s,s'), G_y, H_{s,s'}, H(s,s')$, and A, we have*

$$\sum_{\substack{i:\ \lambda_i > p_B^{\mathsf{triv}} + 1/p \\ j:\ \mu_j > p_C^{\mathsf{triv}} + 1/p}} |\alpha_{i,j}|^2 \leq \mathsf{negl}(\lambda).$$

Using this lemma, we can bound the success probability of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ as

$$\begin{aligned}
&\langle \phi_{\mathsf{BC}} | (\Pi^B \otimes \Pi^C) | \phi_{\mathsf{BC}} \rangle = \sum_{i,j} |\alpha_{i,j}|^2 \lambda_i \mu_j \\
&\leq \sum_{\substack{i:\ \lambda_i > p_B^{\mathsf{triv}} + 1/p \\ j:\ \mu_j > p_C^{\mathsf{triv}} + 1/p}} |\alpha_{i,j}|^2 \lambda_i \mu_j \\
&+ \left(p_B^{\mathsf{triv}} + \frac{1}{p}\right) \sum_{\substack{i:\ \lambda_i \leq p_B^{\mathsf{triv}} + 1/p \\ j:\ \mu_j > p_C^{\mathsf{triv}} + 1/p}} |\alpha_{i,j}|^2 + \left(p_C^{\mathsf{triv}} + \frac{1}{p}\right) \sum_{i,j:\ \mu_j \leq p_C^{\mathsf{triv}} + 1/p} |\alpha_{i,j}|^2 \\
&\leq p^{\mathsf{triv}} + \frac{1}{p} + \mathsf{negl}(\lambda).
\end{aligned}$$

Since $p(\cdot)$ was chosen as an arbitrary polynomial, this suffices for the proof.

Proof of Lemma 6.18. Suppose for the sake of contradiction that the sum of weights is non-negligible with non-negligible probability over the randomness of $y, (s,s'), G_y, H_{s,s'}, H(s,s')$. We will construct an adversary $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ that breaks the MOE game (Definition 4.5):

- $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ get oracle access to $P_{A+s}, P_{A^\perp + s'}$ and $\mathcal{A}'$ receives the state $|A_{s,s'}\rangle$ from the challenger.
- $\mathcal{A}'$ uniformly samples random oracles G', H', as well as random strings $y \leftarrow \mathcal{D}_\lambda, v \in \{0,1\}^{4n+\lambda}$. It runs $\mathcal{A}$ on input $\left(G', H', |A_{s,s'}\rangle, v\right)$ to obtain the bipartite state ρ_{BC}. It sends the B register to $\mathcal{B}'$ and the C register to $\mathcal{C}$. It also sends (G', H', y, v) to both $\mathcal{B}'$ and $\mathcal{C}'$.
- In the second phase, $\mathcal{B}'$ and $\mathcal{C}'$ learn the description of A. Define the binary POVM elements Π^B, Π^C over registers B, C as above. Note that Π^B and Π^C are mixtures of projections since one can sample from $\mathcal{D}_\lambda$ using classical randomness, so that they satisfy the condition of Theorem 2.7.
  We observe that $\mathcal{B}'$ can efficiently implement Π^B as follows:
  - Sample $x_B \leftarrow \mathfrak{D}_y^B$ and reprogram G' on input y to output A, obtaining $(G'_y)^A$.

- Sample uniformly $(s, s') \leftarrow \mathsf{CS}(A) \times \mathsf{CS}(A^\perp)$ and reprogram H' on input (s, s') to output v using the membership oracles $P_{A+s}, P_{A^\perp+s'}$, obtaining $(H'_{s,s'})^v$.
- Run $\mathcal{B}$ using the reprogrammed oracles $(G'_y)^A, (H'_{s,s'})^v$ and test input x_B.
- Measure $\mathcal{B}$'s output z. Undo all the computation.
- If $z = f_y(x_B)$, output 1; otherwise output 0.

Using this, $\mathcal{B}'$ applies the efficient approximated threshold measurement $\mathsf{ATI}^{\epsilon,\delta}_{(P,Q),\gamma_1}$ in Theorem 2.7 with $P = \Pi^B, Q = I - \Pi^B, \gamma_1 = p_B^{\mathsf{triv}} + 3/4p, \epsilon = 1/4p$, and $\delta = 2^{-\lambda}$, with outcome b_B.
If $b_B = 0$, $\mathcal{B}'$ aborts. If $b_B = 1$, then $\mathcal{B}'$ runs a *test execution* on $\mathcal{B}$, described as follows: $\mathcal{B}'$ runs the first three steps of Π^B above on $\mathcal{B}$, and measures a random query (t_B, t'_B) made by $\mathcal{B}$ during the third step to the oracle $(H'_{s,s'})^v$. Then, $\mathcal{B}'$ outputs (t_B, t'_B). We define $\mathcal{C}'$ symmetrically, so that it will measure b_C, and if $b_C = 1$ output a query (t_C, t'_C) made by $\mathcal{C}$ in the test execution.

By Theorem 2.7 bullet (1), $b_B = b_C = 1$ with non-negligible probability. We will finish the proof by showing that $\mathcal{B}'$ and $\mathcal{C}'$ both output (s, s') with non-negligible probability conditioned on $b_B = b_C = 1$. Note that we can intertwine the order of local operations between the two registers this way thanks to no-signalling.

If $b_B = b_C = 1$, then by Theorem 2.7 bullet (2) the post-measurement state is negligibly close to a state of the form

$$\sum_{\substack{i:\ \lambda_i > 1/2+1/4p \\ j:\ \mu_j > 1/2+1/4p}} \beta_{i,j} \left|\phi_i\right\rangle_{\mathsf{B}} \otimes \left|\psi_j\right\rangle_{\mathsf{C}}.$$

Therefore, in the *test execution*, if $\mathcal{B}'$ had not measured (t_B, t'_B) in the third step, $\mathcal{B}$ would correctly output $f_y(x_B)$ correctly with probability greater than $p_B^{\mathsf{triv}} + 1/4p$. Consider a modified adversary $\widetilde{\mathcal{B}}'$ which is identical to $\mathcal{B}'$ except it uses the oracle H' (without reprogramming) when running $\mathcal{B}$. We claim that if $\mathcal{B}'$ is replaced by $\widetilde{\mathcal{B}}'$, then $\mathcal{B}$ would output $f_y(x)$ correctly with probability at most p_B^{triv} at the end of the *test execution*, had $\mathcal{B}'$ not measured a query (t_B, t'_B). This claim and Theorem 2.1 imply that $(t_B, t'_B) = (s, s')$ with non-negligible probability.

To prove this claim, suppose the opposite. We will describe a sequence of games, starting with **Game 1**, between $(\mathcal{A}', \widetilde{\mathcal{B}}')$ acting as the challenger and $(\mathcal{A}, \mathcal{B})$ acting as the adversary:

- $\mathcal{A}$ gets oracle access to G', H' and gets input $\left|A_{s,s'}\right\rangle, v$, all of which are as sampled above by $\mathcal{A}'$ and the MOE challenger.
- $\mathcal{A}$ sends a quantum state ρ_{B} to $\mathcal{B}$
- $\mathsf{ATI}^{\epsilon,\delta}_{(P,Q),\gamma_1}$ as defined above, is applied to $(\mathcal{B}, \rho_{\mathsf{B}})$ by $\widetilde{\mathcal{B}}'$, which uses the oracle H' without reprogramming alongside $(G'_y)^A$, obtaining b_B. If $b_B = 0$, the game is aborted.
- A *test execution* by $\widetilde{\mathcal{B}}'$ is run on $\mathcal{B}$ with its leftover state and $\mathcal{B}$ outputs z.
- The adversary wins if the output is correct, i.e. $z = f_y(x)$.

Note that since H' is not reprogrammed, the value v is a random string independent from the rest of the game. Now we modify the game by replacing $|A_{s,s'}\rangle$ in the first step with the maximally mixed state, resulting in **Game 2**. The success probability of the adversary is unaffected due to the fact that the random strings (s, s') only occur in $|A_{s,s'}\rangle$ in **Game 1**, and

$$\sum_{\substack{s\in\mathsf{CS}(A)\\ s'\in\mathsf{CS}(A^\perp)}} |A_{s,s'}\rangle\langle A_{s,s'}| = I$$

for any subspace A.

Next, we replace the first oracle $(G'_y)^A$ with a random oracle, obtaining **Game 3**. The success probability of the adversary again is affected only negligibly since A is a random subspace independent of the rest of **Game 2**, which is statistically close to a random value by Claim 6.10. Now, y is an independent value from all of **Game 3** except for the test input x_B, hence the adversary is restricted to making a trivial guess, so that it cannot succeed with probability greater than p_B^{triv}.

Similarly, we argue that conditioned on $(t_B, t'_B) = (s, s')$, the probability that (t_C, t'_C) is non-negligible. This follows by a similar argument after observing that after $\mathcal{B}'$ measures a query, the post-measurement state is still negligibly close to a state of the form

$$\sum_{j:\ \mu_j > p_C^{\mathsf{triv}} + 1/4p} \theta_j \, |\sigma_j\rangle_{\mathsf{B}} \, |\psi_j\rangle_{\mathsf{C}} \,,$$

for some states $|\sigma_j\rangle$, so that $\mathcal{C}$ will output correctly with probability greater than $p_C^{\mathsf{triv}} + 1/4p$ during the final execution made by $\mathcal{C}'$.

□

□

Lemmas 6.14 to 6.17 together with triangle inequality imply that $p_1 \leq p^{\mathsf{triv}} + \mathsf{negl}(\lambda)$ as desired, finishing the proof of Lemma 6.13.

□

Remark 6.19. *In our security proof, the adversary can run in unbounded time as long as it is query-bounded.*

Following techniques from the proof of Theorem 4.8, we can show security for correlated input distributions as well.

Corollary 6.20. *Let $w \in [0, 1]$ and let $\mathfrak{D}_X^w(y)$ be the following input distribution:*

- *Sample $x_B, x_C \leftarrow \{0,1\}^\lambda \setminus \{y\}$ independently and uniformly at random.*
- *With probability w, output (x_B, x_C).*
- *With probability $1 - w$, output (y, y).*

Then, $(\mathsf{CopyProtect}, \mathsf{Eval})$ above is a $\mathfrak{D}_X^w$-secure copy-protection scheme for point functions with input length λ.

Proof. Fix an unlearnable distribution $\mathcal{D}_\lambda$ and define the following hybrids:

Hybrid 1. This is the real piracy experiment for $(\mathsf{CopyProtect}, \mathsf{Eval})$.

Hybrid 2. This hybrid matches Hybrid 4 in the proof of Lemma 6.13. In other words, we make the following changes:

- The oracles G, H for $\mathcal{A}$ are replaced with reprogrammed oracles $G_y, H_{s,s'}$, where $G_y(y)$ and $H_{s,s'}(s,s')$ are reprogrammed to freshly random values.
- In addition, the oracle G for $\mathcal{B}$ and $\mathcal{C}$ both is changed to G^A_y, where $G^A_y(y)$ is reprogrammed to output a random (fixed) basis $(v_1, \ldots, v_d)$ of A.

Let p_i be the probability that $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ wins in Hybrid i. Note that $p^{\mathsf{triv}} := p^{\mathsf{triv}}(\mathcal{D}_\lambda, \mathfrak{D}^w_X) = \max(w, 1-w)$. By Lemmas 6.14 to 6.16, we have $|p_1 - p_2| \leq \mathsf{negl}(\lambda)$, since these lemmas are proved irrespective of the input distribution. Thus, it suffices to show that $p_2 \leq \max(w, 1-w)$.

Let ρ_{BC} be the bipartite state created by $\mathcal{A}$ in Hybrid 2. Without loss of generality assume that $\rho_{\mathsf{BC}} = |\phi_{\mathsf{BC}}\rangle\langle\phi_{\mathsf{BC}}|$ is a pure state. Fix $y \leftarrow \mathcal{D}_\lambda$, fix $G_y, H_{s,s'}, H(s,s'), A$ which are randomly sampled, and fix random inputs $(x_B, x_C) \leftarrow \{0,1\}^\lambda \setminus \{y\}$. We define the following projectors:

- Π^B_0: runs $\mathcal{B}$ on input oracles G^A_y, H and test input x_B; it measures if the output is $f_y(x_B)$; then it undoes all the computation.
- Π^B_1: runs $\mathcal{B}$ on input oracles G^A_y, H and test input y; it measures if the output is $f_y(x_B)$; then it undoes all the computation.
- Π^C_0 and Π^C_1 are defined similarly for $\mathcal{C}$.

Now we write the state $|\phi_{\mathsf{BC}}\rangle$ in its spectral decomposition with respect to $(w\Pi^B_0 + (1-w)\Pi^B_1) \otimes (w\Pi^C_0 + (1-w)\Pi^C_1)$ as

$$|\phi_{\mathsf{BC}}\rangle = \sum_{i,j} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} |\psi_j\rangle_{\mathsf{C}},$$

where $|\phi_i\rangle_{\mathsf{B}}$ is an eigenvector of $(w\Pi^B_0 + (1-w)\Pi^B_1)$ with eigenvalue λ_i and $|\psi_j\rangle_{\mathsf{C}}$ is an eigenvector of $(w\Pi^C_0 + (1-w)\Pi^C_1)$ with eigenvalue μ_j.

□

We first make the following observation:

Lemma 6.21. *Let $p(\cdot)$ be a polynomial. With overwhelming probability over $y, (s,s'), G_y, H_{s,s'}, H(s,s'), A$, and (x_B, x_C), we have*

$$\sum_{\substack{i:\ |\lambda_i - 1/2| > |w-1/2| + 1/p \\ j:\ |\mu_j - 1/2| > |w-1/2| + 1/p}} |\alpha_{i,j}|^2 \leq \mathsf{negl}(\lambda).$$

Proof. Note that the condition $|\lambda_i - 1/2| > |w - 1/2| + 1/p$ is satisfied if and only if $\lambda_i > p^{\mathsf{triv}} + 1/p$ or $1 - \lambda_i > p^{\mathsf{triv}} + 1/p$. The proof is nearly identical to the proof of Lemma 6.18. To avoid repetition, we only mention a few notable differences:

- After sampling $y \leftarrow \mathcal{D}_y$, $\mathcal{A}'$ additionally samples random inputs $x_B \neq y$ and $x_C \neq y$.
- Instead of ATI, $\mathcal{B}'$ applies $\mathsf{SATI}^{\epsilon,\delta}_{P,Q,\gamma_1}$, with $P = w\Pi_0^B + (1-w)\Pi_1^B$, $Q = I - P$, $\gamma_1 = 3/4p$, and $\epsilon = 1/2p$. Similarly for $\mathcal{C}'$.
- When implementing Π_0^B, $\mathcal{B}'$ uses x_B as input, and it uses y when implementing Π_1^B. Similarly for $\mathcal{C}'$.
- In the end when we say that an adversary, with no knowledge of y other than the test input given, can succeed with probability at most p^{triv}, we instead argue the success probability of such an adversary, denoted by q, must satisfy $\max(q, 1-q) \leq p^{\mathsf{triv}}$. This is because the adversary can always flip its output bit to succeed with probability q instead of $1-q$.

□

By Lemma 6.21, with overwhelming probability $|\phi_{\mathsf{BC}}\rangle$ is negligibly close to the state $|\phi'_{\mathcal{B}}\rangle + |\phi'_{\mathcal{C}}\rangle$, where

$$|\phi'_{\mathcal{B}}\rangle = \sum_{i\,:\,|\lambda_i - 1/2| \leq |w-1/2|+1/p} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} |\psi_j\rangle_{\mathsf{C}},$$
$$|\phi'_{\mathcal{C}}\rangle = \sum_{\substack{i\,:\,|\lambda_i - 1/2| > |w-1/2|+1/p \\ j\,:\,|\mu_j - 1/2| \leq |w-1/2|+1/p}} \alpha_{i,j} |\phi_i\rangle_{\mathsf{B}} |\psi_j\rangle_{\mathsf{C}}.$$

The rest of the proof will imitate the analysis in the proof of Lemma 4.11:

$$\begin{aligned}
&w\left|(\Pi_0^B \otimes \Pi_0^C)(|\phi'_{\mathcal{B}}\rangle + |\phi'_{\mathcal{C}}\rangle)\right|^2 + (1-w)\left|(\Pi_1^B \otimes \Pi_1^C)(|\phi'_{\mathcal{B}}\rangle + |\phi'_{\mathcal{C}}\rangle)\right|^2 \\
=&\big(w\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{B}}\rangle + (1-w)\langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{B}}\rangle + w\langle\phi'_{\mathcal{C}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle \\
&+(1-w)\langle\phi'_{\mathcal{C}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\big) + 2\mathsf{Re}\left(w\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + (1-w)\langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right) \\
\leq&\big(w\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes I)|\phi'_{\mathcal{B}}\rangle + (1-w)\langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes I)|\phi'_{\mathcal{B}}\rangle + w\langle\phi'_{\mathcal{C}}|(I \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle \\
&+(1-w)\langle\phi'_{\mathcal{C}}|(I \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\big) + 2\mathsf{Re}\left(w\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle + (1-w)\langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes \Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right).
\end{aligned}$$

We bound each term separately.

- $\big(w\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes I)|\phi'_{\mathcal{B}}\rangle + (1-w)\langle\phi'_{\mathcal{B}}|(\Pi_1^B \otimes I)|\phi'_{\mathcal{B}}\rangle\big)$. It is equal to $\langle\phi'_{\mathcal{B}}|(w\Pi_0^B + (1-w)\Pi_1^B) \otimes I|\phi'_{\mathcal{B}}\rangle$; by the definition of $|\phi'_{\mathcal{B}}\rangle$, it will be at most
$$(1/2 + |w - 1/2| + 1/p)\,|\,|\phi'_{\mathcal{B}}\rangle\,|^2 = \max(w, 1-w)\,|\,|\phi'_{\mathcal{B}}\rangle\,|^2.$$
- $\big(w\langle\phi'_{\mathcal{C}}|(\Pi_0^C \otimes I)|\phi'_{\mathcal{C}}\rangle + (1-w)\langle\phi'_{\mathcal{C}}|(\Pi_1^C \otimes I)|\phi'_{\mathcal{C}}\rangle\big)$. Similar to the above case, it is at most $\max(w, 1-w)\,|\,|\phi'_{\mathcal{B}}\rangle\,|^2$.
- $\mathsf{Re}\left(\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle\right)$. By Corollary 2.4, this term will vanish:
$$\langle\phi'_{\mathcal{B}}|(\Pi_0^B \otimes \Pi_0^C)|\phi'_{\mathcal{C}}\rangle = \sum_{i\,:\,|\lambda_i - 1/2| \leq |w-1/2|+1/p} \sum_{\substack{i'\,:\,|\lambda_i - 1/2| > |w-1/2|+1/p \\ j'\,:\,|\mu_j - 1/2| \leq |w-1/2|+1/p}} \alpha^\dagger_{i,j}\alpha_{i',j'}\langle\phi_i|\Pi_0^B|\phi_{i'}\rangle\langle\psi_j|\Pi_0^C|\psi_{j'}\rangle;$$
since every possible i, i' satisfy $\lambda_i + \lambda_{i'} \neq 1$, we have $\langle\phi_i|\Pi_0^B|\phi_{i'}\rangle = 0$.

- $\mathsf{Re}\left(\langle\phi'_{\mathcal{B}}|(\Pi_1^B\otimes\Pi_1^C)|\phi'_{\mathcal{C}}\rangle\right)$. Similarly, this term vanishes as well.

Therefore, the total probability is at most

$$w\left|(\Pi_0^B\otimes\Pi_0^C)(|\phi'_{\mathcal{B}}\rangle+|\phi'_{\mathcal{C}}\rangle)\right|^2+(1-w)\left|(\Pi_1^B\otimes\Pi_1^C)(|\phi'_{\mathcal{B}}\rangle+|\phi'_{\mathcal{C}}\rangle)\right|^2+\mathsf{negl}(n)\leq\max(w,1-w)+\frac{1}{p}+\mathsf{negl}(n).$$

Since the polynomial $p(\cdot)$ is arbitrary, this suffices for the proof.

References

[Aar09] Scott Aaronson. "Quantum copy-protection and quantum money". In: *2009 24th Annual IEEE Conference on Computational Complexity*. IEEE. 2009, pp. 229–242 (cit. on pp. 1, 2, 12).

[AC12] Scott Aaronson and Paul Christiano. "Quantum money from hidden subspaces". In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. 2012, pp. 41–60 (cit. on p. 1).

[AGKZ20] Ryan Amos, Marios Georgiou, Aggelos Kiayias, and Mark Zhandry. "One-shot signatures and applications to hybrid quantum/classical authentication". In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 2020, pp. 255–268 (cit. on p. 1).

[AK21] Prabhanjan Ananth and Fatih Kaleoglu. "Unclonable Encryption, Revisited". In: *Theory of Cryptography Conference*. Springer. 2021, pp. 299–329 (cit. on pp. 2, 3, 5, 12, 18).

[AL21] Prabhanjan Ananth and Rolando L La Placa. "Secure Software Leasing". In: *Eurocrypt* (2021) (cit. on pp. 1, 2, 13).

[ALL+21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. "New approaches for quantum copy-protection". In: *Annual International Cryptology Conference*. Springer. 2021, pp. 526–555 (cit. on pp. 1, 12, 15, 17).

[BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. "Strengths and weaknesses of quantum computing". In: *SIAM journal on Computing* 26.5 (1997), pp. 1510–1523 (cit. on pp. 13, 14).

[BJL+21] Anne Broadbent, Stacey Jeffery, Sébastien Lord, Supartha Podder, and Aarthi Sundaram. "Secure software leasing without assumptions". In: *Theory of Cryptography Conference*. Springer. 2021, pp. 90–120 (cit. on p. 13).

[BL20] Anne Broadbent and Sébastien Lord. "Uncloneable Quantum Encryption via Oracles". In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*. Ed. by Steven T. Flammia. Vol. 158. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 4:1–4:22. DOI: 10.4230/LIPIcs.TQC.2020.4 (cit. on pp. 2, 4–8, 12, 18, 19).

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. "Hidden cosets and applications to unclonable cryptography". In: *Annual International Cryptology Conference*. Springer. 2021, pp. 556–584 (cit. on pp. 1, 5, 7, 8, 13, 15, 23–25, 30).

[CMP20] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. *Quantum copy-protection of compute-and-compare programs in the quantum random oracle model*. 2020. arXiv: 2009.13865 [quant-ph] (cit. on pp. 2, 3, 5, 12, 31).

[CV21] Eric Culf and Thomas Vidick. "A monogamy-of-entanglement game for subspace coset states". In: *arXiv preprint arXiv:2107.13324* (2021) (cit. on p. 25).

[GL89] O. Goldreich and L. A. Levin. "A Hard-Core Predicate for All One-Way Functions". In: *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*. STOC '89. Seattle, Washington, USA: Association for Computing Machinery, 1989, pp. 25–32. ISBN: 0897913078. DOI: 10.1145/73007.73010 (cit. on p. 20).

[Got02] Daniel Gottesman. "Uncloneable encryption". In: *arXiv preprint quant-ph/0210062* (2002) (cit. on p. 1).

[GZ20] Marios Georgiou and Mark Zhandry. "Unclonable decryption keys". In: *IACR Cryptol. ePrint Arch* 877.2020 (2020), p. 3 (cit. on p. 1).

[Lut10] Andrew Lutomirski. "An online attack against Wiesner's quantum money". In: *arXiv preprint arXiv:1010.0256* (2010) (cit. on p. 8).

[MST21] Christian Majenz, Christian Schaffner, and Mehrdad Tahmasbi. "Limitations on Uncloneable Encryption and Simultaneous One-Way-to-Hiding". In: (Nov. 2021). arXiv: 2103.14510 [quant-ph] (cit. on pp. 4, 5, 7, 12, 19).

[MW05] Chris Marriott and John Watrous. "Quantum arthur–merlin games". In: *computational complexity* 14.2 (2005), pp. 122–152 (cit. on pp. 15, 17).

[MZB16] Jose Mejia, Camilo Zapata, and Alonso Botero. "The difference between two random mixed quantum states: exact and asymptotic spectral analysis". In: *Journal of Physics A: Mathematical and Theoretical* 50.2 (Dec. 2016), p. 025301. DOI: 10.1088/1751-8121/50/2/025301 (cit. on p. 21).

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. DOI: 10.1017/CBO9780511976667 (cit. on p. 13).

[Reg05] Oded Regev. *Witness-preserving Amplification of QMA*. 2005. URL: https://cims.nyu.edu/~regev/teaching/quantum_fall_2005/ln/qma.pdf (cit. on p. 14).

[RS22] Roy Radian and Or Sattath. "Semi-quantum money". In: *Journal of Cryptology* 35.2 (2022), pp. 1–70 (cit. on p. 1).

[TFKW13] Marco Tomamichel, Serge Fehr, Jędrzej Kaniewski, and Stephanie Wehner. "A monogamy-of-entanglement game with applications to device-independent quantum cryptography". In: *New Journal of Physics* 15.10 (2013), p. 103002. DOI: 10.1088/1367-2630/15/10/103002 (cit. on p. 6).

[Unr15] Dominique Unruh. "Revocable Quantum Timed-Release Encryption". In: *J. ACM* 62.6 (Dec. 2015). ISSN: 0004-5411. DOI: 10.1145/2817206 (cit. on p. 7).

[Vid21] Thomas Vidick. *Lecture Notes on Interactive proofs with quantum devices*. 2021. URL: http://users.cms.caltech.edu/~vidick/teaching/fsmp/lecture1.pdf (cit. on p. 14).

[VZ21] Thomas Vidick and Tina Zhang. "Classical proofs of quantum knowledge". In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2021, pp. 630–660 (cit. on p. 7).

[Wat18] John Watrous. *The Theory of Quantum Information*. Cambridge: Cambridge University Press, 2018. ISBN: 978-1-107-18056-7. DOI: 10.1017/9781316848142 (cit. on p. 20).

[Wie83] Stephen Wiesner. "Conjugate coding". In: *ACM Sigact News* 15.1 (1983), pp. 78–88 (cit. on pp. 1, 6).

[Zha20] Mark Zhandry. "Schrödinger's pirate: How to trace a quantum decoder". In: *Theory of Cryptography Conference*. Springer. 2020, pp. 61–91 (cit. on pp. 9, 15, 17).

[Zha21] Mark Zhandry. "Quantum lightning never strikes the same state twice. or: quantum money from cryptographic assumptions". In: *Journal of Cryptology* 34.1 (2021), pp. 1–56 (cit. on pp. 1, 5).