Thesis Proposal

How to Use Quantum States for Cryptography

Alper Çakan

December 3, 2026

Computer Science Department
Kenneth C. Griffin School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Vipul Goyal, Chair (Carnegie Mellon University and NTT Research)
Aayush Jain
Ryan O'Donnell
Takashi Yamakawa
(NTT Research Center for Theoretical Quantum Information
and Yukawa Institute for Theoretical Physics, Kyoto University)

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science.

Copyright © 2026 Alper Çakan

Abstract

Classical information can be copied easily, but quantum information behaves differently: quantum states cannot be copied/cloned. In 1968, Wiesner suggested that we could use this to our advantage to achieve cryptographic capabilities such as unclonable digital money, which are impossible with classical information no matter how strong cryptography is. Later, in 1984, Bennett and Brassard proposed quantum key distribution, showing the possibility of cryptography from quantum states without assuming classical cryptography.

These ideas have inspired decades of extensive research, leading to numerous advances in both directions. However, on the first front, much of the work has focused on a small set of primitives, such as encryption, and specifically on unclonability, leaving open whether we can achieve classically-impossible guarantees for broader classes of functionalities and also security guarantees beyond unclonability. On the second front, some works have shown that quantum computation can enable cryptography using classical information even if there is no cryptography using classical computation. However, the algorithms literature produced many examples of quantum algorithms later being dequantized, motivating a stronger question: can quantum states enable cryptography even when cryptography using (even quantumly computed) classical information does not exist?

In the first part of the thesis, we show how to construct unclonable keys for a general class of functionalities called puncturable functionalities. Then, we introduce a notion called unbounded (and LOCC) leakage-resilience that proposes storing secret information as quantum states to achieve security against ubiquitous leakage attacks. We then construct such encoding schemes for secret keys of puncturable functionalities, thus showing that quantum states can achieve cryptographic guarantees beyond unclonability that are impossible to achieve with classical information.

In the second part of the thesis, we show a classical oracle relative to which P equals QCMA, but various quantum cryptographic primitives such as public-key encryption with (copy-protected) quantum decryption keys exist. This shows that quantum states can help us achieve cryptographic capabilities even when there is no cryptography from classical information.

Finally, we propose to extend our techniques to show that even stronger primitives can exist when $\mathsf{P} = \mathsf{QCMA}$, such as functional encryption, collusion-resistant copy-protection, and LOCC leakage-resilient quantum keys.

Contents

1 Introduction **1**
- 1.1 Completed work . . . 2
- 1.2 Proposed work . . . 3

2 Completed work **4**
- 2.1 Quantumly protecting our secrets: puncturable functionalities . . . 4
  - 2.1.1 Background and previous work . . . 4
  - 2.1.2 Results . . . 4
  - 2.1.3 Technical overview . . . 6
- 2.2 MicroCrypt: public-key encryption when $\mathsf{P} = \mathsf{QCMA}$ . . . 9
  - 2.2.1 Background and previous work . . . 9
  - 2.2.2 Results . . . 9
  - 2.2.3 Technical overview . . . 10

3 Proposed work **13**
- 3.1 Functional encryption when $\mathsf{P} = \mathsf{QCMA}$ . . . 13
- 3.2 Fully collusion-resistant copy-protection when $\mathsf{P} = \mathsf{QCMA}$ . . . 14
- 3.3 Stretch direction: LOCC leakage-resilient quantum keys . . . 15
- 3.4 Timeline . . . 16

1 Introduction

Perhaps the most important aspect of classical information is the fact that it can be copied easily. This is what enabled the Internet, and what enabled the digital revolution we have experienced in the last century. Much of our civilization's knowledge can be stored in databases in large computers in a few locations, and anyone with an inexpensive Internet connection can search things in this vast amount of information and download what they like to their own computer, without deducting anything from the shared database.

Quantum cryptography. Quantum states (i.e. quantum information) behave differently from classical information: as the well-known No-Cloning Principle [WZ82, Die82] says, quantum states cannot be copied/cloned in general. While this may initially seem annoying, Stephen Wiesner in 1968 suggested that we could use this property of quantum information to our advantage in *cryptography*: he came up with the idea of *quantum money*, where he proposed using quantum states (specifically, a class of states called *conjugate coding*) as banknotes so that bad actors cannot duplicate them [Wie83, ACM26].

While Wiesner's manuscript was initially rejected by journals and his ideas attracted little attention, physicist Charles H. Bennett (who was friends with Wiesner since their undergraduate days at Brandeis University) learned this idea from Wiesner and occasionally discussed it with various people. Eventually, in 1979, Bennett mentioned Wiesner's work to Gilles Brassard, and based on Wiesner's ideas, they devised a way called quantum key distribution for two parties to use quantum states to agree on a secret key over a public channel without using public-key encryption assumptions, which is required in the classical setting [BB84]. These works started the field of *quantum cryptography*, which has evolved into a very fruitful and active research area at the intersection of theoretical computer science, cryptography, and physics/quantum information theory.

Two directions in quantum cryptography. Interestingly, decades later, many works in modern quantum cryptography still roughly fall into two areas that follow these two works. The first area, akin to quantum money, asks whether quantum information can give us guarantees that are impossible with classical information, no matter how strong our cryptography is. Much of this area concerns *unclonable cryptography*, but, as we discuss below, the same question extends beyond unclonability, so we will refer to this direction more broadly as *uniquely quantum cryptography*. The second area, akin to quantum key distribution, asks whether quantum cryptography can exist without assuming classical cryptography. While some earlier works considered the possibility of unconditional cryptography based on quantum information, most primitives are impossible in that setting even with quantum information. In this thesis, we focus on the modern line of work, commonly called *MicroCrypt*, that asks whether quantum cryptographic primitives (with security against efficient adversaries) can exist from assumptions that are weaker than existence of classical cryptography.

Terminology. We now introduce some terminology. By *classical cryptography*, we will mean cryptography using classical algorithms and classical information. We say *cryptography using classical information* to mean the type of cryptography where all the keys, honest parties' memories and communication are classical, but the schemes' algorithms are allowed to be quantum (e.g. the secret keys are classical strings sampled by a quantum computer). Finally, when we say quantum cryptography, we allow all information to be quantum as well (e.g. honest parties have quantum memories).

Finally, we will also use the (informal) terms *quantum computational advantage* and *quantum state advantage*. By the former, we mean the existence of an efficient quantum prover and an efficient (quantum) verifier such that they only hold classical memory and use classical interaction, but the verifier can distinguish the prover from any efficient classical prover. In contrast, by quantum state advantage, informally, we mean the existence of a task for which quantum states allow us to accomplish it while no classical state succeeds, e.g. $\mathsf{BQP/qpoly}$ versus $\mathsf{BQP/poly}$ (quantum versus classical advice), or QMA versus QCMA (verifying membership in a language using quantum versus classical witnesses).

Broadening uniquely quantum cryptography. On the first front, constructions of copy-protection for primitives such as public-key encryption [GZ20, CLLZ21] show how quantum states can prevent a decryption key from being copied, but leave open how far this protection can extend to more general functionalities. Moreover, copying is not the only way in which a secret can be compromised: preventing an adversary from producing two usable copies of a key does not by itself guarantee security against classical leakage from that key. These limitations lead to the first questions considered in this thesis:

How can we use quantum states to protect our secrets, beyond simple primitives like encryption? Beyond unclonability, can quantum states offer solutions to other problems concerning the storage of secrets that classical information cannot solve?

Cryptography without quantum computational advantage. On the second front, work by Kretschmer, Qian, Sinha, and Tal showed that quantum cryptographic primitives can exist relative to a classical oracle even when $\mathsf{P} = \mathsf{NP}$ [KQST23]. Building on this, Kretschmer, Qian, and Tal constructed quantum-computable one-way functions and public-key encryption with classical keys and ciphertexts in such an oracle world [KQT25]. Thus, quantum computation can enable cryptography using classical information even when classical cryptography does not exist. However, their oracle necessarily satisfies $\mathsf{P} \neq \mathsf{BQP}$, so these constructions still rely on quantum computational advantage.

This reliance on quantum computational advantage is significant because several proposed quantum speedups have subsequently been *dequantized* under comparable data-access assumptions [Tan19, CGL+20]. These dequantization results motivate asking whether quantum states themselves can support cryptography when this computational advantage is absent:

Can quantum states enable cryptography even when cryptography using classical information does not exist?

We study this question under the hypothetical collapse $\mathsf{P} = \mathsf{QCMA}$, where QCMA consists of problems admitting efficient quantum verification with polynomial-length classical witnesses. Under this collapse, $\mathsf{P} = \mathsf{BQP}$, quantum-computable one-way functions do not exist, and there is no quantum computational advantage [Çak26a].

1.1 Completed work

Quantumly protecting our secrets. The first part of this thesis presents constructions of copy-protection and LOCC leakage-resilience for *puncturable functionalities*, developed in my joint work with Behera and Goyal [BÇG26] and with Goyal [ÇG25], respectively. Informally, puncturing a secret implementation C at an input $x^\star$ preserves its behavior at every other input while hiding the information needed to solve the functionality's security challenge at $x^\star$. Under post-quantum cryptographic assumptions, these works give the following guarantees:

(i) **Copy-protection.** We encode C as a reusable quantum state ρ_C that allows its holder to evaluate C, while preventing an efficient adversary from producing two independently usable implementations. The guarantee also covers correlated challenges, strengthening the independent-challenge security of my earlier joint work with Goyal [ÇG25, BÇG26].

(ii) **Unbounded and LOCC leakage-resilience.** We encode the secret keys of puncturable functionalities so that an efficient attack may extract any polynomial amount of classical information without enabling a receiver to break the functionality's security on a fresh challenge. In the stronger model of local operations and classical communication (LOCC), the party accessing the quantum key and the receiver may both perform quantum computations and exchange classical messages over multiple rounds, starting without shared entanglement; their communication must end before the challenge is revealed [ÇGLZR24, ÇG25].

The leakage guarantee is impossible for a classically stored key, since unrestricted classical leakage can reveal the entire key. A leakage attack may also consume the quantum key completely, so ruling out such an attack requires more than preventing the production of two working copies. Together, the copy-protection and leakage-resilience results extend uniquely quantum cryptography in both respects posed by the first research question: they protect a general class of functionalities and provide security beyond unclonability.

MicroCrypt: public-key encryption when $\mathsf{P} = \mathsf{QCMA}$. The second part of this thesis presents my construction of public-key encryption with classical public keys and ciphertexts and reusable quantum decryption keys [Çak26a]. The scheme is defined relative to a classical oracle $\mathcal{O}$ satisfying

$$\mathsf{P}^{\mathcal{O}} = \mathsf{QCMA}^{\mathcal{O}}.$$

All algorithms run in quantum polynomial time, and the decryption keys additionally satisfy copy-protection: although a key can be used repeatedly to decrypt ciphertexts, an efficient adversary cannot turn one issued key into two independently usable decryptors.

Because public-key encryption using classical information is impossible under this collapse, the quantum decryption key is essential to the security of the scheme. Thus, quantum state advantage can enable public-key encryption without quantum computational advantage, while copy-protection provides the additional guarantee studied in the first part of the thesis. The same oracle also supports copy-protected pseudorandom functions with quantum secret keys and public-key quantum money with classical verification keys [Çak26a].

1.2 Proposed work

The public-key encryption result leaves open how far cryptography from quantum state advantage can extend to more expressive functionalities and stronger forms of quantum protection when $\mathsf{P} = \mathsf{QCMA}$. I propose to study functional encryption and fully collusion-resistant copy-protection relative to classical oracles satisfying this collapse, with LOCC leakage-resilient quantum keys as a stretch goal.

Functional encryption. Whereas a decryption key in public-key encryption reveals the entire message m, a functional key for f should reveal only $f(m)$, while hiding information beyond the outputs authorized by the issued keys. I will try to extend the function-encoding and state-preparation techniques of my public-key encryption construction for $\mathsf{P} = \mathsf{QCMA}$ to provide this restricted access through quantum functional keys ρ_f, while keeping public keys and ciphertexts classical.

Fully collusion-resistant copy-protection. Functional encryption may also strengthen copy-protection: we showed how to combine FE with single-key quantum protection to obtain full collusion resistance, in joint work with Goyal, Kitagawa, Nishimaki, and Yamakawa [ÇGK+26]. I will try to adapt this transformation to the proposed FE and the copy-protected PKE for $\mathsf{P} = \mathsf{QCMA}$. The goal is to prevent an adversary that receives k independently issued keys from producing $k+1$ working decryptors, for any polynomially bounded k, even when the adversary processes all the issued keys jointly.

Going further: LOCC leakage-resilient quantum keys. Finally, I will investigate whether quantum decryption keys can resist interactive classical leakage when $\mathsf{P} = \mathsf{QCMA}$, extending the leakage protection from the first part of the thesis to the setting of the second. I will draw on the techniques for analyzing interactive leakage in my joint work with Goyal and Shmueli on quantum fire [ÇGS26] and my work on the limitations of quantum proof systems with classical interaction [Çak26b], and try to extend those techniques to account for the oracle access that enforces the collapse. Such a result would show that quantum states can protect secrets beyond unclonability even when cryptography using classical information is impossible.

2 Completed work

The completed work addresses the two questions introduced in Section 1. First, we present quantum encodings of puncturable functionalities that allow repeated evaluation while protecting the underlying secrets against copying and classical leakage. We then present public-key encryption with quantum decryption keys relative to a classical oracle satisfying $\mathsf{P} = \mathsf{QCMA}$, where public-key encryption using classical information is impossible even if its algorithms are quantum.

2.1 Quantumly protecting our secrets: puncturable functionalities

2.1.1 Background and previous work

Quantum copy-protection asks us to encode a program as a quantum state that allows repeated evaluation, while preventing its holder from producing two independently usable implementations [Aar09]. For decryption, for example, the holder should be able to decrypt ciphertexts but should be unable to produce two successful decryptors. Previous constructions achieve this guarantee using coset states [GZ20, CLLZ21], raising the question of whether the same quantum states can protect more general classes of functionalities.

Beyond individual primitives. Ananth and Behera [AB24] took a step toward this goal by giving a modular construction for puncturable circuits. However, their guarantee required uniformly distributed challenge inputs and a negligible baseline success probability, excluding encryption experiments in which guessing the challenge bit already succeeds with probability $1/2$. Their proof also relied on a conjecture about simultaneous extraction from entangled adversaries. This leaves the problem of obtaining copy-protection that accommodates the security experiments of different cryptographic functionalities without relying on the extraction conjecture.

2.1.2 Results

We construct reusable copy-protection for puncturable functionalities, in joint work with Behera and Goyal [BÇG26]. We also construct LOCC leakage-resilient quantum implementations of punc-

turable functionalities, in joint work with Goyal [ÇG25]. These results use puncturability to protect functionalities such as pseudorandom functions and decryption against both copying and classical leakage, thereby extending quantum protection beyond individual primitives and beyond unclonability.

To state the two guarantees in a common framework, we describe a functionality by a secret circuit C and a security experiment in which a challenger generates a challenge ch and private verification information τ. The adversary answers with y, and the predicate $V_\tau(y)$ determines whether it succeeds. In our running examples, the challenge is an input x to C, although the framework also allows additional challenge information; public information, such as an encryption scheme's public key, remains available to the adversary throughout. We write λ for the security parameter and $\mathsf{negl}(\lambda)$ for a negligible function, and consider quantum polynomial-time (QPT) adversaries unless specified otherwise.

The security bound is measured against a baseline success probability p_{triv} achievable without the secret. In an encryption indistinguishability experiment, for example, the challenge is a ciphertext encrypting one of two chosen messages, so guessing which message was encrypted gives $p_{\mathrm{triv}} = 1/2$. By comparison, predicting the output of a pseudorandom function with a long output has a negligible baseline.

Puncturability. Puncturability requires that the challenge remain secure even when the adversary receives a circuit that agrees with C away from the challenge point. We give an informal definition, omitting the setup and challenge-sampling algorithms of the full definitions [ÇG25, BÇG26].

Definition 1 (Puncturability, informal)**.** *A circuit family is puncturable if there is an efficient algorithm that, given a secret circuit C and a challenge point $x^\star$, produces a circuit $C_{\{x^\star\}}$ such that*

$$C_{\{x^\star\}}(x) = C(x) \qquad \text{for every } x \neq x^\star,$$

while revealing $C_{\{x^\star\}}$ and $x^\star$ leaves the challenge $\mathsf{ch}^\star$ generated at that point secure. In particular, in a prediction experiment this requires

$$\Pr\big[V_\tau\big(\mathcal{A}(C_{\{x^\star\}}, x^\star, \mathsf{ch}^\star)\big) = 1\big] \leq p_{\mathrm{triv}} + \mathsf{negl}(\lambda),$$

with the analogous indistinguishability requirement for an indistinguishability-based functionality.

For a PRF, puncturing at $x^\star$ preserves evaluation at all other inputs while leaving the value at $x^\star$ pseudorandom; for decryption, it preserves decryption away from the challenge ciphertext while hiding the encrypted message. To obtain copy-protection, we require puncturing at both recipients' challenge points and sufficiently unpredictable challenges, since otherwise the holder could evaluate the program on a small set of likely inputs in advance and distribute the answers.

A reusable quantum implementation. A protection scheme consists of algorithms $\mathsf{Protect}$ and Eval with syntax

$$\rho_C \leftarrow \mathsf{Protect}(1^\lambda, C), \qquad (y, \rho'_C) \leftarrow \mathsf{Eval}(\rho_C, x).$$

The encoding ρ_C may contain classical information as well as quantum registers. Correctness requires that $y = C(x)$ except with negligible probability and that the residual state ρ'_C support further evaluations, for any polynomial number of sequential uses. The adversary receives the entire encoding and can inspect and copy all classical descriptions supplied with it.

Copy-protection. In the copy-protection experiment, a pirate $\mathcal{A}$ receives ρ_C and produces two registers R_B, R_D, which may be entangled. The recipients $\mathcal{B}$ and $\mathcal{D}$ then receive their respective challenges $\mathsf{ch}_B, \mathsf{ch}_D$ and must answer without communicating. Security requires

$$\Pr[V_{\tau_B}(y_B) = 1 \ \wedge \ V_{\tau_D}(y_D) = 1] \leq p_{\text{triv}} + \mathsf{negl}(\lambda). \tag{1}$$

The baseline corresponds to passing the original encoding to one recipient and leaving the other to answer without the secret. For decryption, both recipients must identify their challenge bits, giving the bound $1/2 + \mathsf{negl}(\lambda)$. Since the experiment tests the recipients' answers, the adversary is free to change the representation of the key completely: the guarantee concerns its ability to produce two working implementations.

Theorem 1 (Copy-protection for puncturable functionalities, informal [BÇG26])**.** *Assuming post-quantum indistinguishability obfuscation and the quantum hardness of Learning with Errors, puncturable functionalities admit reusable copy-protection in the plain model. The guarantee also covers correlated challenges, extending the independent-challenge guarantee of my earlier joint work with Goyal [ÇG25].*

Beyond unclonability: LOCC leakage-resilience. A leakage attack may consume the original quantum key while transferring enough information to compromise its security, so preventing the production of two usable copies does not by itself address leakage. We introduced LOCC leakage-resilience to formalize security against interactive classical leakage from quantum encodings, in joint work with Goyal, Liu-Zhang, and Ribeiro [ÇGLZR24]. We subsequently extended this guarantee to puncturable functionalities, in joint work with Goyal [ÇG25].

The leakage experiment begins with a party $\mathcal{L}$ holding ρ_C and a receiver $\mathcal{R}$ holding only public information. Starting without shared entanglement, the two parties may perform local quantum computations and exchange classical messages; after their communication has ended, the receiver alone obtains a fresh challenge ch and produces an answer y. Security requires

$$\Pr[V_\tau(y) = 1] \leq p_{\text{triv}} + \mathsf{negl}(\lambda). \tag{2}$$

The amount of classical communication is unrestricted apart from the adversaries' polynomial running time, and the receiver may retain a quantum workspace throughout the interaction. In particular, it can process the received messages quantumly and choose its subsequent messages adaptively.

Theorem 2 (LOCC leakage-resilience for puncturable functionalities, informal [ÇG25])**.** *Assuming subexponentially secure post-quantum indistinguishability obfuscation and one-way functions, puncturable functionalities admit reusable quantum implementations satisfying LOCC leakage-resilience.*

Unlike a classical encoding, whose entire contents could be transmitted during the leakage phase, the quantum encoding remains secure against any such classical interaction, even if the interaction consumes the state. Requiring the communication to end before the challenge is essential, since otherwise the key holder could evaluate the challenge and send the answer. The theorem therefore guarantees security on a fresh challenge after the leakage, while correctness still allows an honest holder to evaluate the encoded functionality repeatedly.

2.1.3 Technical overview

We now describe the encoding of a puncturable circuit and explain how its security follows from protection properties of coset states [BÇG26, ÇG25]. The encoding combines a coset state with

two obfuscated programs whose outputs determine the circuit's value, and an honest holder uses the state in the computational and Hadamard bases to obtain both outputs without disturbing it. Puncturability then allows the security proof to isolate the challenge inputs and reduce an attack on the encoded circuit to an attack on the coset state.

Coset states. Let A be a uniformly random n-dimensional subspace of $\mathbb{F}_2^{2n}$, and let $s, t \in \mathbb{F}_2^{2n}$ be independent uniform shifts. The associated coset state is

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

Here inner products are over $\mathbb{F}_2$, $A^\perp$ is the subspace of vectors orthogonal to every vector in A, and H denotes the single-qubit Hadamard transform. Measuring the state in the computational basis gives a vector in $A + s$, whereas applying Hadamards first gives a state supported on $A^\perp + t$:

$$H^{\otimes 2n} |A_{s,t}\rangle = |(A^\perp)_{t,s}\rangle \quad \text{up to a global phase.}$$

Although measuring an individual vector generally destroys the ability to obtain a vector from the complementary coset, information that is constant on the entire coset can be obtained without disturbing the state. This observation allows the honest evaluator to use the same coset state in both bases, as we describe next.

Encoding a puncturable functionality. Let F_K be a puncturable PRF whose output length equals that of C. The encoding contains $|A_{s,t}\rangle$ and obfuscations of two classical programs:

$$Q_0(w, x) = \begin{cases} F_K(x) \oplus C(x), & w \in A + s, \\ \perp, & \text{otherwise,} \end{cases} \qquad Q_1(w, x) = \begin{cases} F_K(x), & w \in A^\perp + t, \\ \perp, & \text{otherwise.} \end{cases}$$

The membership tests are implemented by obfuscated programs, so the protected encoding takes the form

$$\rho_C = \big(|A_{s,t}\rangle, \widehat{Q}_0, \widehat{Q}_1\big),$$

where a hat denotes an obfuscation. The two outputs are additive shares of $C(x)$, masked using $F_K(x)$, and obtaining both shares allows the holder to recover $C(x)$. Although the obfuscated programs can be copied freely, each requires a vector from the corresponding coset before revealing its share.

To evaluate $C(x)$, we apply $\widehat{Q}_0(\cdot, x)$ coherently to the coset state. Every vector in its support passes the membership test and yields the same output $F_K(x) \oplus C(x)$, allowing us to copy this output into a fresh register and uncompute without disturbing the state. After applying $H^{\otimes 2n}$, we use $\widehat{Q}_1(\cdot, x)$ in the same way to obtain $F_K(x)$, and then apply Hadamards again to restore the original state. The XOR of the two outputs is $C(x)$, while the restored coset state remains available for subsequent evaluations.

Decisional monogamy of entanglement. To prove copy-protection, we must rule out an adversary that splits the coset state into possibly entangled registers and uses the obfuscated programs to answer both recipients' challenges. For this purpose, we establish a new *decisional* monogamy-of-entanglement property of coset states [BÇG26]. We first describe the underlying information-theoretic game, in which both recipients try to predict one bit:

(i) A pirate receives $|A_{s,t}\rangle$ and splits it into two registers.

(ii) After the split, sample nonzero $u \in A$ and $v \in A^\perp$ uniformly, and give the same pair (u, v) to both recipients.

(iii) The recipients, without communicating, each try to output $e = \langle u, t\rangle \oplus \langle v, s\rangle$.

We prove that even a computationally unbounded adversary wins this game with probability at most $1/2 + 2^{-\Omega(n)}$, whereas the holder of the unsplit state can obtain e without disturbing it. The latter follows because each inner product is constant on the corresponding coset, allowing the holder to compute one in the computational basis and the other after applying Hadamards.

Applying this property to copy-protection requires accounting for the classical programs distributed with the state. We therefore also prove computational monogamy in the presence of obfuscated membership tests and suitable simulatable auxiliary information: even with these additional inputs, every QPT adversary wins with probability at most $1/2 + \mathsf{negl}(\lambda)$. The membership programs allow the adversary to implement the verification projector $|A_{s,t}\rangle\langle A_{s,t}|$, so the state is *unlearnable* even with access to verification. This is the form of monogamy needed when the quantum state is accompanied by programs that help evaluate the protected functionality.

From decisional monogamy to copy-protection. The new monogamy game is particularly useful because its answer depends only on the cosets, rather than on the choice of vectors within them. Indeed, for every $w \in A + s$ and $w' \in A^\perp + t$,

$$e_0 := \langle v, s\rangle = \langle v, w\rangle, \qquad e_1 := \langle u, t\rangle = \langle u, w'\rangle.$$

The proof first punctures the PRF at the challenge points and replaces its values there by uniform masks, after which it rewrites the masked outputs in terms of these inner products. To illustrate the rewriting step with a single-bit decryption challenge, suppose that the challenge bit b is shared as $b \oplus r$ and r. For independent uniform bits q, ζ, we can equivalently sample the challenge bit and its mask as

$$b = e_0 \oplus e_1 \oplus \zeta, \qquad r = q \oplus e_1, \qquad b \oplus r = q \oplus \zeta \oplus e_0.$$

The resulting pair (b, r) is distributed as a fresh uniform challenge bit and mask, while the programs can compute the shares using $\langle v, w\rangle$ and $\langle u, w'\rangle$ in place of e_0, e_1. Since these inner products are constant on the accepted cosets, the rewritten programs have the same functionality, and indistinguishability obfuscation justifies replacing their obfuscations. A recipient's correct answer b, XORed with the known mask ζ, is then already the monogamy answer $e_0 \oplus e_1$. Puncturability allows us to make this change at the challenge points while preserving evaluation everywhere else.

Earlier approaches based on search monogamy required extracting a full coset vector from each recipient, even when the copy-protection task only asked it to predict a bit. With entangled recipients, extracting from one register could disturb the other and prevent the second extraction from succeeding. Our decisional game avoids this difficulty because each recipient's answer bit directly yields an answer to the monogamy challenge; the reduction also avoids the need to select canonical coset representatives. It therefore applies naturally to decryption and other functionalities whose security is expressed through indistinguishability.

Simulatable auxiliary information. For general functionalities, the reduction must also account for when the protected programs are distributed. The pirate receives them before the split, whereas the vectors (u, v) used in the rewritten programs arrive only afterward in the monogamy game. Our computational monogamy theorem allows the pirate to receive additional challenge-dependent data alongside the membership obfuscations, provided that the additional data can be jointly simulated statistically from public parameters without knowing A, u, v.

We obtain this simulation property by hiding (u, v) with a challenge point $x^\star$. For a freshly sampled universal hash function h, we set

$$z = (u\|v) \oplus h(x^\star).$$

The proof ensures that $x^\star$ retains enough entropy conditioned on the other auxiliary information for universal hashing to hide (u, v). Using independent hash functions for the two challenge points makes the data supplied before the split jointly statistically simulatable, while the later challenge inputs recover the vectors needed to evaluate the inner products.

For LOCC leakage-resilience, we use the same encoding of the functionality and the same approach to simulating the auxiliary information, replacing the two-recipient monogamy guarantee with the corresponding LOCC protection property of coset states [ÇG25]. The underlying experiment now consists of an interactive classical leakage phase followed by a fresh challenge. Thus, puncturability and the simulation of auxiliary information allow the protection provided by the coset state to extend to the encoded functionality, both for copy-protection and for LOCC leakage-resilience.

2.2 MicroCrypt: public-key encryption when $\mathsf{P} = \mathsf{QCMA}$

2.2.1 Background and previous work

Research in *MicroCrypt* asks whether quantum cryptography can exist without classical one-way functions [Kre21]. Kretschmer, Qian, and Tal [KQT25] showed that, relative to a classical oracle, quantum-computable one-way functions and public-key encryption can exist even when $\mathsf{P} = \mathsf{NP}$. Their encryption scheme uses classical keys and ciphertexts, so it demonstrates that quantum computation can enable cryptography using classical information even when classical cryptography does not exist.

However, their oracle necessarily satisfies $\mathsf{P} \neq \mathsf{BQP}$, leaving open whether quantum states can enable cryptography when quantum computational advantage is also absent. As discussed in the introduction, the dequantization of several proposed quantum speedups further motivates this question [Tan19, CGL+20].

We consider the stronger collapse $\mathsf{P} = \mathsf{QCMA}$, which rules out quantum-computable one-way functions and allows classical computers to solve every problem that admits efficient quantum verification with a classical witness. Public-key encryption using classical information is impossible in this setting, but the collapse does not settle what quantum decryption keys can provide. We therefore ask:

Can public-key encryption exist when $\mathsf{P} = \mathsf{QCMA}$*, using quantum decryption keys but classical public keys and ciphertexts?*

2.2.2 Results

I answer this question affirmatively by constructing a classical oracle relative to which $\mathsf{P} = \mathsf{QCMA}$, yet public-key encryption exists with classical public keys and ciphertexts and reusable quantum decryption keys [Çak26a]. Moreover, the decryption keys satisfy copy-protection, so quantum states both enable encryption in the absence of cryptography using classical information and protect the decryption keys against copying.

Encryption with a reusable quantum key. The scheme consists of quantum polynomial-time algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ for key generation, encryption, and decryption. Key generation produces a classical public key pk and a quantum decryption key ρ_{sk}, while encryption takes the public key and a message and produces a classical ciphertext. Decryption returns the message together with a residual key, which must remain usable for polynomially many honest decryptions, with negligible probability of any error.

Message security is the usual indistinguishability under chosen-plaintext attack: an adversary given pk chooses two messages m_0, m_1 and, upon receiving $\mathsf{Enc}(\mathsf{pk}, m_b)$ for a uniform bit b, must guess b with probability at most $1/2+\mathsf{negl}(\lambda)$. The honest parties and the adversary may all perform quantum computations and query the same classical oracle in superposition; however, encryption requires only the classical public key, whereas decryption requires the quantum key.

Copy-protection of the decryption capability. For bit encryption, copy-protection is defined by specializing the experiment in Equation (1) to ciphertext challenges. A pirate receives the quantum decryption key and produces two possibly entangled registers for two noncommunicating recipients, who then receive independently generated ciphertexts

$$c_i \leftarrow \mathsf{Enc}(\mathsf{pk}, m_{i,b_i}), \qquad b_i \leftarrow \{0,1\}, \quad i \in \{1,2\},$$

for message pairs chosen by the pirate. The probability that both recipients identify their challenge bits must be at most $1/2 + \mathsf{negl}(\lambda)$, even though the pirate was given a key that can correctly decrypt either challenge on its own.

Theorem 3 (Public-key encryption when $\mathsf{P} = \mathsf{QCMA}$, informal [Çak26a])**.** *There is a classical oracle $\mathcal{O}$ such that*

$$\mathsf{P}^{\mathcal{O}} = \mathsf{QCMA}^{\mathcal{O}},$$

and, relative to $\mathcal{O}$, there is public-key encryption with classical public keys and ciphertexts and reusable quantum decryption keys, satisfying message indistinguishability and copy-protection.

The oracle in this theorem is a classical Boolean function to which the honest parties and the adversary have the same access. Since public-key encryption using classical information is impossible when $\mathsf{P} = \mathsf{QCMA}$, the quantum decryption key is essential: the theorem gives cryptography from quantum state advantage without quantum computational advantage. The same oracle also supports copy-protected pseudorandom functions with quantum secret keys and public-key quantum money with classical verification keys [Çak26a].

2.2.3 Technical overview

The public-key encryption construction uses a new oracle problem, *Codeword-Hash pairs in Noise* (CHN), that admits efficient verification with a quantum witness but is hard for the QCMA hierarchy in the query complexity model. We first describe the problem and the main ideas behind its hardness, then explain how to encode a function into the oracle so that the quantum witness enables evaluation. The tools used to prove hardness will also allow us to preserve cryptographic security when we add the oracle enforcing $\mathsf{P} = \mathsf{QCMA}$ [Çak26a].

Codeword-Hash pairs in Noise. Let $\kappa = \mathsf{poly}(\lambda)$, let $\mathcal{C} \subseteq \Sigma^q$ be a suitable linear code with $q = \mathsf{poly}(\kappa)$, and sample the values of predicates $H_i : \Sigma \to \{0,1\}$ independently from $\mathsf{Bernoulli}(\kappa^{-4})$. Their symbol-wise hash defines the relation

$$H(v) = H_1(v_1) \| \cdots \| H_q(v_q), \qquad \mathcal{P}_H = \{(v,y) \in \mathcal{C} \times \{0,1\}^\kappa : H(v) = y \| 0^{q-\kappa}\}.$$

The CHN problem asks us to distinguish the following oracle distributions:

$$\mathcal{O}^{\mathsf{plant}} = B \vee \mathbf{1}_{\mathcal{P}_H}, \qquad \mathcal{O}^{\mathsf{noise}} = B, \qquad B(v,y) \leftarrow \mathsf{Bernoulli}(\beta), \quad \beta = 2^{-\lceil\sqrt{\kappa}\rceil},$$

where the noise bits are independent. The algorithms access only the sampled oracle, with no separate access to H or B.

The presence of noise in both distributions is essential because the codeword-hash relation alone can be recognized by a counting attack at a low level of the hierarchy: each codeword has at most one target y in $\mathcal{P}_H$, whereas a noise oracle accepts exponentially many targets for a fixed codeword. In CHN, both distributions contain the same background noise, so planting a codeword-hash pair changes this count by at most one and eliminates the noticeable relative gap used by the attack.

A quantum witness. A quantum witness can recognize the planted codeword-hash pairs even in the presence of this noise. The witness is the polynomial-size *predicate-preimage state* $|\Psi_H\rangle$, consisting of a uniform superposition over each of the sets $H_i^{-1}(0)$ and $H_i^{-1}(1)$, in separate registers, for every i. For suitable code parameters, the Yamakawa–Zhandry decoding algorithm [YZ24] uses this state to find a codeword satisfying $H(v) = y\|0^{q-\kappa}$ for any chosen y, without querying H. Since the same witness works for every target, a verifier can sample a fresh uniform y, use the witness to find a corresponding codeword, and check whether the oracle accepts the pair. This gives QMA verification of the planted distribution, with completeness and soundness holding outside negligible exceptional sets of instances.

Hardness for the QCMA hierarchy. I prove that every constant level of the hierarchy, defined by iterated oracle access to QCMA, has only negligible distinguishing advantage using polynomially many queries [Çak26a]. Eliminating the counting attack alone does not establish this result, since a hierarchy machine may make quantum queries and quantify over classical witnesses at every level. We therefore first simulate the quantum verifiers using deterministic classical queries, then remove their witnesses using a switching lemma, and repeat these steps to reduce the depth of the hierarchy.

(i) *Deterministic-classical-query simulation.* The codeword-hash relation and sparse noise satisfy *contractive simultaneous incompressibility*, which bounds how many distinct nonzero oracle locations can be recovered from a short classical description. This property allows us to identify a small database of nonzero locations containing almost all of the verifier's query weight on nonzero answers; replacing the remaining answers by zero then changes its acceptance probability only slightly. We thereby obtain a simulator that makes polynomially many deterministic classical queries, followed by randomized oracle-free computation. With overwhelming probability over the sampled oracles, the simulation approximates the verifier simultaneously for all classical witnesses and for both the planted and noise oracles.

(ii) *Removing the witness by switching.* Because the simulator makes deterministic classical queries, each possible query transcript specifies only polynomially many oracle bits. We can therefore express the existence of an accepting witness as a polynomial-width DNF by taking a disjunction over witnesses and transcripts for which the simulator's post-processing accepts with probability at least $1/2$. After fixing the common noise outside a smaller set of possible codeword-hash locations, a switching lemma converts this DNF into a deterministic classical-query algorithm that takes no witness. The fixed answers are shared by the planted and noise oracles, which is why the same replacement works in both cases.

(iii) *Repeated depth reduction.* Replacing the deepest verifier by an algorithm that takes no witness reduces the hierarchy depth by one, allowing us to repeat the simulation and switching steps. Since each replacement depends on some hash predicates, we partition the suffix predicates $H_{\kappa+1}, \ldots, H_q$ into groups and use a fresh group at each level, preserving the independence needed for the next application. At the final level, the switching guarantee yields an algorithm that makes no queries and has the same description for the planted and noise oracles. It must therefore output the same bit in both cases, proving indistinguishability.

The simulators are efficient in their number of queries, although their local computation need not be efficient. The resulting lower bound shows that a constant hierarchy of classical witnesses and polynomially many queries cannot distinguish the planted codeword-hash pairs from noise, even though a single polynomial-size quantum witness enables verification.

From CHN to function encoding. To use CHN for cryptography, we follow the function-encoding strategy of [KQT25] and replace the constant 1 planted at each codeword-hash pair by the value of a Boolean function $\mathfrak{F} : \{0,1\}^\kappa \to \{0,1\}$:

$$E(v,y) = \begin{cases} \mathfrak{F}(y), & v \in \mathcal{C} \text{ and } H(v) = y \| 0^{q-\kappa}, \\ U(v,y), & \text{otherwise}, \end{cases}$$

where the $U(v,y)$ are independent uniform bits, so the encoding uses uniform noise in place of the sparse noise of the decision problem. Given $|\Psi_H\rangle$ and an input y, we can find a codeword v satisfying $H(v) = y\|0^{q-\kappa}$ and obtain $\mathfrak{F}(y)$ by querying $E(v,y)$. Since every such codeword gives the same answer, we can copy the answer to an output register and uncompute, preserving the witness up to negligible error. The oracle E is public, while the quantum witness enables evaluation of the encoded function at chosen inputs.

A quantum decryption key. For public-key encryption, we apply the function encoding to a coset-state copy-protection construction [ALL+21, CLLZ21]. The public key is a serial number $s \in \{0,1\}^\lambda$, and each serial number indexes independent random Boolean functions $F_{s,0}, F_{s,1}$ on $\{0,1\}^\lambda$, with $F_s = F_{s,0} \oplus F_{s,1}$. For a random subspace A_s and shifts a_s, a'_s, let $|\chi_s\rangle$ be the coset state supported on $A_s + a_s$ in the computational basis and on $A_s^\perp + a'_s$ in the Hadamard basis. We use these two cosets to control access to the two shares of F_s through the function

$$C_s(x,c,\xi) = \begin{cases} F_{s,0}(x), & c = 0 \text{ and } \xi \in A_s + a_s, \\ F_{s,1}(x), & c = 1 \text{ and } \xi \in A_s^\perp + a'_s, \\ \perp, & \text{otherwise.} \end{cases}$$

As in the coset-state encoding from Section 2.1, the holder of $|\chi_s\rangle$ can evaluate C_s coherently in each basis to recover both shares of $F_s(x)$ without significantly disturbing the state. We now encode C_s into E_s, representing its validity and output by Boolean values, and use

$$|K_s\rangle = |\Psi_{H_s}\rangle \otimes |\chi_s\rangle$$

as the decryption key. The predicate-preimage state $|\Psi_{H_s}\rangle$ enables evaluation of the encoded function C_s, while the coset state $|\chi_s\rangle$ supplies valid inputs in both bases, allowing the holder to recover $F_s(x)$.

Public encryption by sampling. To encrypt a message, it suffices to sample a uniform input u together with $F_s(u)$, whereas decrypting a given ciphertext requires evaluating F_s at the input specified by that ciphertext. We enable public encryption by providing a way to prepare the graph state

$$|\Gamma_s\rangle = 2^{-\lambda/2} \sum_{u\in\{0,1\}^\lambda} |u\rangle|F_s(u)\rangle.$$

Measuring this state yields the required pair $(u, F_s(u))$, from which we encrypt a one-bit message $m \in \{0,1\}$ as

$$\mathsf{ct} = (u, m \oplus F_s(u)).$$

The holder of $|K_s\rangle$ can evaluate $F_s(u)$ and recover m, whereas polynomially many public samples are unlikely to include a previously specified input u.

Key generation likewise prepares $2^{-\lambda/2}\sum_s |s\rangle|K_s\rangle$ and measures the serial register, obtaining a key for a uniformly random serial number. To implement key generation and encryption using only classical oracles, I construct Boolean oracles from which a uniform quantum algorithm prepares the required states with negligible error. The adversary also has access to these preparation oracles, so their effect must be included in the security proof.

Security when $\mathsf{P} = \mathsf{QCMA}$. Finally, we enforce $\mathsf{P} = \mathsf{QCMA}$ by adding recursive classical oracles that answer existential questions and simulate quantum acceptance probabilities. To show that public-key encryption remains secure with this additional oracle access, we use the tools underlying the QCMA-hierarchy hardness of CHN: classical-query simulation and switching convert an attack with access to the function encoding and collapse oracles into a query-efficient attack with direct access to the encoded function. Combining this simulation with the security guarantees for the state-preparation oracles reduces message security to the unpredictability of F_s at the challenge input and copy-protection to coset-state monogamy [Çak26a, BÇG26]. In Section 3.1, I propose to extend the function encoding and state-preparation techniques to functional encryption.

3 Proposed work

The public-key encryption scheme in Section 2.2 shows that quantum decryption keys can enable encryption even when $\mathsf{P} = \mathsf{QCMA}$. I propose to extend this PKE construction to functional encryption and fully collusion-resistant copy-protection, using my earlier transformation from FE and single-key protection to full collusion resistance to guide the latter direction. As a stretch goal, I will investigate LOCC leakage-resilient decryption keys under the same collapse, bringing the protection against classical leakage from the first part of the thesis into a setting where cryptography using classical information is impossible. I will seek classical-oracle constructions, with quantum polynomial-time algorithms throughout.

3.1 Functional encryption when $\mathsf{P} = \mathsf{QCMA}$

In public-key encryption, a decryption key allows its holder to recover the entire encrypted message, whereas functional encryption gives more selective access: a key associated with a function f should reveal only $f(m)$ from an encryption of m. Constructing FE when $\mathsf{P} = \mathsf{QCMA}$ would therefore show that quantum states can enable this more general form of access to encrypted information even when cryptography using classical information is impossible.

Target guarantee. My goal is functional encryption with a classical master public key mpk and classical ciphertexts, but quantum functional keys ρ_f, relative to a classical oracle $\mathcal{O}$ satisfying $\mathsf{P}^{\mathcal{O}} = \mathsf{QCMA}^{\mathcal{O}}$. A functional key should support repeated evaluation, with correctness requiring

$$\mathsf{FDec}\big(\rho_f, \mathsf{FEnc}(\mathsf{mpk}, m)\big) = f(m)$$

except with negligible error. The corresponding security requirement is that encryptions of two messages m_0, m_1 remain computationally indistinguishable to an adversary holding functional keys whenever

$$f(m_0) = f(m_1) \qquad \text{for every function } f \text{ authorized by those keys.}$$

Proposed approach. I will try to adapt the function encoding and state-preparation techniques from my PKE construction [Çak26a]. Recall from Section 2.2 that a Boolean function is encoded into a classical oracle, and a quantum witness allows its holder to evaluate that function at chosen inputs. In the PKE scheme, the encoded coset function and the quantum key together allow the holder to recover the decryption mask and hence the entire message. For FE, the encoded functionality and the associated quantum key must instead permit only the designated computation $f(m)$, without giving the holder access to the full decryption mask.

I will therefore investigate how to make the encoded functionality and functional-key generation depend on f, while retaining the quantum sampling procedure that makes encryption public. Alongside this construction, I will try to extend the simulation argument used to prove PKE security in the presence of the collapse oracle, so that the simulation preserves the restriction to the functions authorized by the issued keys. The choice of functionality will also be guided by the application to collusion resistance below, where functional evaluation must produce a ciphertext for an independently issued protected key.

3.2 Fully collusion-resistant copy-protection when $\mathsf{P} = \mathsf{QCMA}$

The PKE scheme in Section 2.2 prevents the holder of one quantum decryption key from producing two successful decryptors. When several legitimate keys are issued, however, an adversary may process all of them jointly before distributing its outputs, and security for a single issued key does not by itself rule out such an attack. I will investigate whether copy-protection can be made fully collusion-resistant when $\mathsf{P} = \mathsf{QCMA}$, so that even a joint attack cannot increase the number of independently usable decryptors.

Target guarantee. For any polynomially bounded k, an adversary receiving k issued quantum decryption keys should be unable to produce $k+1$ successful decryptors. We can express this through the decryption experiment of Section 2.2: after receiving the keys, the adversary outputs $k+1$ possibly entangled registers, whose recipients are tested on fresh, independently generated ciphertexts with challenge bits b_i. If $\widehat{b}_i$ denotes the answer of recipient i, the desired bound is

$$\Pr\left[\bigwedge_{i=1}^{k+1} \widehat{b}_i = b_i\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

The baseline corresponds to giving each of k recipients an issued key and having the remaining recipient guess its challenge bit. Full collusion resistance requires one scheme to satisfy this bound for every polynomially bounded k, relative to a classical oracle satisfying $\mathsf{P} = \mathsf{QCMA}$.

Connection to functional encryption. We showed how functional encryption can be used to upgrade single-key protection of decryption keys to full collusion resistance, in joint work with Goyal, Kitagawa, Nishimaki, and Yamakawa [ÇGK+26]. Each issued key in our construction pairs a fresh protected decryption key with an FE functional key. The function associated with the FE key takes the message and auxiliary randomness as inputs and produces an encryption under the public key corresponding to the protected key. Thus, decryption first converts the outer FE ciphertext into a ciphertext for the holder's protected key, which the quantum key can then decrypt. This allows the same outer ciphertext to be decrypted using independently issued protected keys, even though those keys have different underlying public keys.

I propose to adapt this transformation by combining the proposed FE with the copy-protected PKE of Section 2.2. A central obstacle is that the earlier transformation evaluates the underlying scheme's *classical encryption circuit* through FE, whereas encryption in the $\mathsf{P} = \mathsf{QCMA}$ scheme is a quantum sampling procedure. Consequently, constructing FE under the collapse would not by itself suffice to apply the earlier transformation; the functional-evaluation step must also accommodate the sampling needed by the PKE scheme.

Proposed approach. I will investigate whether the quantum sampling needed to generate the inner PKE ciphertext can be carried out through functional evaluation while keeping the outer encrypted message hidden. This makes the FE and collusion-resistance directions closely related: the transformation specifies a concrete task for functional evaluation, while the FE construction must provide a way to perform that task without revealing a general decryption key. I will then try to adapt the earlier security reduction to turn a joint attack on the issued keys into an attack on a single protected decryption key, while accounting for the adversary's access to the collapse oracle.

3.3 Stretch direction: LOCC leakage-resilient quantum keys

Beyond preventing an adversary from producing additional decryptors, quantum states can also protect a secret against classical leakage, as shown by the completed work in Section 2.1. As a stretch goal, I will investigate whether LOCC leakage-resilient decryption keys can exist when $\mathsf{P} = \mathsf{QCMA}$. Such a construction would extend the first part of the thesis to a setting where quantum states are needed not only for leakage-resilience, but also for encryption itself, since cryptography using classical information is impossible under the collapse.

Target guarantee. The goal is public-key encryption with classical public keys and ciphertexts and reusable quantum decryption keys satisfying the LOCC leakage-resilience definition from Section 2.1. One party initially holds the quantum key and interacts with a receiver over a classical channel, while both parties may perform quantum computations and query the oracle. They start without shared entanglement and may exchange any polynomial amount of classical information, but their communication ends before the challenge. The receiver then chooses m_0, m_1, receives

$$\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b), \qquad b \leftarrow \{0,1\},$$

and should satisfy $\Pr\left[\widehat{b} = b\right] \leq 1/2 + \mathsf{negl}(\lambda)$. Thus, the receiver should remain unable to distinguish a fresh encryption even if the preceding leakage process consumes the original key completely.

Proposed approach. I will draw on the incompressibility and simulation techniques for interactive classical communication developed in my joint work with Goyal and Shmueli on quantum fire [ÇGS26], and in my separation of QMA from QCIP [Çak26b]. Incompressibility limits the useful

oracle information conveyed by classical messages, while the simulations account for the receiver's quantum workspace conditioned on the transcript. Accounting for this workspace is essential because the receiver chooses its messages using a private quantum state that evolves throughout the interaction, so a bound for a single classical message does not immediately extend to an interactive attack.

I will try to combine these techniques with an interactive version of the simulation used in the PKE security proof. The simulation must preserve the restriction to classical communication and the fact that only one party initially holds the quantum key, while controlling the information that the receiver obtains from its own oracle queries. Unlike the earlier leakage-resilience analyses, it must also account for both parties' access to the oracle that enforces $\mathsf{P} = \mathsf{QCMA}$. Finally, the quantum-fire construction shows that LOCC leakage-resilience is compatible with efficient cloning, so unclonability need not be a requirement for the leakage-resilient keys sought here.

3.4 Timeline

- December 2026: Work on proving the security of the proposed construction ideas for functional encryption (FE) and collusion-resistant copy-protection
- January 2027: If the above two goals succeed, start writing their proofs and move onto working on LOCC leakage-resilience question. Otherwise, keep working on those two goals.
- February 2027: Complete dissertation draft; advisor and committee feedback and revisions
- March 2027: Defense

References

[Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *24th Annual IEEE Conference on Computational Complexity*, pages 229–242. IEEE, 2009.

[AB24] Prabhanjan Ananth and Amit Behera. A modular approach to unclonable cryptography. In *Advances in Cryptology—CRYPTO 2024, Part VII*, volume 14926 of *Lecture Notes in Computer Science*, pages 3–37. Springer, 2024.

[ACM26] ACM. What led you to that insight? ACM, Advances in Computing, 2026. Interview with Charles H. Bennett and Gilles Brassard.

[ALL+21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. New approaches for quantum copy-protection. In *Advances in Cryptology—CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 526–555. Springer, 2021.

[BB84] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.

[BÇG26] Amit Behera, Alper Çakan, and Vipul Goyal. Decisional monogamy-of-entanglement for coset states and applications to unclonable cryptography with correlated challenges. arXiv:2608.18841, 2026. Version 2.

[Çak26a] Alper Çakan. Classical oracle separation of QMA from the QCMA hierarchy and quantum cryptography even when $P = QCMA$. arXiv:2611.26574, 2026. Version 1.

[Çak26b] Alper Çakan. Separating QMA from QCIP with a classical oracle, or, the power of quantum proofs over classical interaction for quantum verifiers. arXiv:2610.03648, 2026. Version 1.

[ÇG25] Alper Çakan and Vipul Goyal. How to copy-protect all puncturable functionalities without conjectures: A unified solution to quantum protection. Cryptology ePrint Archive, Paper 2025/1197, 2025.

[ÇGK+26] Alper Çakan, Vipul Goyal, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Multi-copy security in quantum cryptography and more. arXiv:2510.12626, 2026. Version 2.

[CGL+20] Nai-Hui Chia, András Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 387–400. ACM, 2020.

[ÇGLZR24] Alper Çakan, Vipul Goyal, Chen-Da Liu-Zhang, and João Ribeiro. Unbounded leakage-resilience and intrusion-detection in a quantum world. In *Theory of Cryptography—TCC 2024, Part II*, volume 15365 of *Lecture Notes in Computer Science*, pages 159–191. Springer, 2024.

[ÇGS26] Alper Çakan, Vipul Goyal, and Omri Shmueli. Public-key quantum fire and key-fire from classical oracles. In *Advances in Cryptology—CRYPTO 2026, Part V*, pages 341–371. Springer, 2026. Full version: arXiv:2504.16407v2.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In *Advances in Cryptology—CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 556–584. Springer, 2021.

[Die82] Dennis Dieks. Communication by EPR devices. *Physics Letters A*, 92(6):271–272, 1982.

[GZ20] Marios Georgiou and Mark Zhandry. Unclonable decryption keys. Cryptology ePrint Archive, Paper 2020/877, 2020.

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1589–1602. ACM, 2023.

[KQT25] William Kretschmer, Luowen Qian, and Avishay Tal. Quantum-computable one-way functions without one-way functions. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 189–200. ACM, 2025.

[Kre21] William Kretschmer. Quantum pseudorandomness and classical complexity. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics*, pages 2:1–2:20, 2021. Revised full version: arXiv:2103.09320v5.

[Tan19] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 217–228. ACM, 2019.

[Wie83] Stephen Wiesner. Conjugate coding. *ACM SIGACT News*, 15(1):78–88, 1983.

[WZ82] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.

[YZ24] Takashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. *Journal of the ACM*, 71(3), 2024.