Cryptomania v.s. Minicrypt in a Quantum World

Longcheng Li* Qian Li† Xingjian Li‡ Qipeng Liu§

Abstract

We prove that it is impossible to construct perfect-complete quantum public-key encryption (QPKE) with classical keys from quantumly secure one-way functions (OWFs) in a black-box manner, resolving a long-standing open question in quantum cryptography.

Specifically, in the quantum random oracle model (QROM), no perfect-complete QPKE scheme with classical keys, and classical/quantum ciphertext can be secure. This improves the previous works which require either unproven conjectures or imposed restrictions on key generation algorithms. This impossibility even extends to QPKE with quantum public key if the public key can be uniquely determined by the secret key, and thus is tight to all existing QPKE constructions.

1 Introduction

Quantum information and computation are topics with growing importance in cryptography. They reshape people's views on cryptography drastically, including breaking classical secure cryptosystems [Sho99], creating primitives that are impossible for classical [Wie83], and weakening assumptions [BB14]. However, quantum cryptography is not an all-powerful tool, as it also has its own limits. Therefore, characterizing the boundary of quantum cryptography under different assumptions has become a topic of great interest.

Boundaries between classical cryptographic primitives have already been studied extensively. In the seminal work by Impagliazzo and Rudich [IR89], they proposed the methodology of black-box separation. They showed that one-way functions are insufficient to build public key encryption (PKE) schemes in a black-box manner. In the famous work by Impagliazzo [Imp95], he characterized five possibilities on the hardness of NP problems and their complexity consequences. Among them, he used the word 'Minicrypt' to refer to a world where one-way functions exist, and the word 'Cryptomania' for a world with public key cryptography primitives. Thus, the separation result by Impagliazzo and Rudich can be viewed as a separation between Minicrypt and Cryptomania.

It turns out that the landscape of quantum cryptography varies depending on the definition of Minicrypt. For example, it is known that with quantum communication, many primitives in classical Cryptomania can be built from one-way functions, including key agreement [BB14], oblivious

*State Key Lab of Processors, Institute of Computing Technology, Chinese Academy of Sciences. Email: lilongcheng22s@ict.ac.cn

†Shenzhen International Center For Industrial And Applied Mathematics, Shenzhen Research Institute of Big Data. Email: liqian.ict@gmail.com

‡Tsinghua University. Email: lxj22@mails.tsinghua.edu.cn

§University of California San Diego. Email: qipengliu0@gmail.com

transfer [GLSV21, BCKM21], public key encryption (with quantum public keys) [Col23, MW24, KMNY24, BGH+23]. On the other hand, none of these constructions are known in the quantum computation classical communication (QCCC) setting. Given that quantum communication has various drawbacks, including the difficulties of authenticating, broadcasting, reusability, and potentially adding interactions; therefore, we will mainly focus on the following question:

Does there exist any separation between Minicrypt and Cryptomania in the QCCC setting?

Previous works. Several works have attempted to address this question, but classical proof techniques often fail due to fundamental differences between quantum and classical algorithms/information, including challenges related to cloning, rewinding, and the unique structure of quantum queries. As a result, all previous approaches have either relied on unproven conjectures or applied only to highly restricted quantum PKE schemes.

In the work by Austrin et al. [ACC+22], they initialized the study of separations between quantum key agreements and one-way functions in the quantum random oracle model (QROM). They showed that under some *conjecture named 'polynomial compatibility conjecture'*, quantum key agreements with perfect correctness don't exist. Since quantum PKE implies a two-round key agreement scheme, their result also implies a separation between quantum PKE and one-way functions. The same idea was also applied in separating PKE with quantum ciphertext and one-way functions [BGV+23], which we will discuss later.

However, it seems that some kind of conjecture may be necessary in their routine. To show a separation between key agreement and one-way function in the QROM, one needs to construct an eavesdropper that breaks the security of the key agreement by making polynomially many queries to the oracle. In the paper [ACC+22], the authors construct an eavesdropper who only makes classical queries to the oracle, while Alice and Bob can make quantum queries in general.

In another line of work, Li, Li, Li, and Liu [LLLL24] approach the problem from a different perspective. They introduce tools from quantum Markov chains [FR15] to construct an eavesdropper for quantum public key encryption schemes that have a *classical key generation process*.

To be specific, they consider the following scenario, where they construct a two-round key agreement from PKE: Alice first runs the key generation algorithm and sends the public key pk to Bob, Bob then encrypts a random key k by running the encryption algorithm $\mathsf{Enc}(\mathsf{pk}, k) \to \mathsf{ct}$, and sends back the ciphertext ct to Alice; Alice runs the decryption algorithm to retrieve the key k. Their idea is to create some Eve such that the conditional mutual information $I(\mathsf{A} : \mathsf{B}|\mathsf{E}) \leq \epsilon$ by making $\mathsf{poly}(1/\epsilon)$ queries, which implies the three systems form an approximate Markov chain. By the operational meaning of the quantum Markov chain, there exists some channel $\mathcal{T} : \mathsf{E} \to \mathsf{E} \otimes \mathsf{A}'$ that generates a copy of Alice system, while guaranteeing the joint state $\sigma_{\mathsf{A}'\mathsf{EB}} = \mathcal{T}(\rho_{\mathsf{EB}})$ is $O(\sqrt{\epsilon})$ close to the original joint state ρ_{AEB}. Thus, by using the information in the register A', Eve can simulate Alice and generate a copy of the key k. In their paper, they need the key generation algorithm to be classical and thus make classical queries. By measuring the register A', they can generate a classical query record $R_{\mathsf{A}'}$ of polynomial size, and Eve will simulate the run of A on the oracle reprogrammed by $R_{\mathsf{A}'}$.

The paper [LLLL24] relies on maintaining a polynomial-sized query record R_{A} to ensure that reprogramming does not significantly disturb Bob's state. Therefore, extending their result to a quantum key generation process seems difficult, as query transcripts are no longer well-defined in the quantum setting. In a following paper [LLLL25], the same authors introduce a view from

boolean function analysis in an attempt to attack PKE with quantum key generation. It is known that the probability of Alice outputting some pk can be written as a low-degree polynomial $f(x)$, where x is the truth table of the random oracle. The key observation is that if there is some polynomial size partial assignment μ such that $f(x^\mu) \neq 0$ for all x, the partial assignment μ can replace the $R_{\mathsf{A}'}$ in [LLLL24] and the algorithm can reprogram the oracle on the points defined by the partial assignment μ. However, such a short partial assignment does not exist as proved in the same paper, the authors make some conjecture on the existence of a distribution on such partial assignments: particularly, they make a conjecture on the zero point distribution of low-degree polynomial f under a distribution of partial assignments; they prove based on the conjecture that, there exists a separation between PKE and OWF in the QCCC model.

Our main result. In this work, we obtain a full separation between perfectly complete quantum PKE and one-way functions, i.e. the quantum version of [IR89] on public-key encryption, closing the gap in previous works.

Theorem 1.1. *Perfect-complete quantum public key encryption, with classical keys and classical ciphertext, does not exist in the quantum random oracle model.*

Remark. *Our result removes the conjecture used in [ACC+22] and the restriction of a classical key generation algorithm in [LLLL24, LLLL25], leaving perfect completeness as the only requirement. Perfect completeness is a natural requirement satisfied by many cryptographic schemes, both classical and quantum, including all known quantum PKE schemes even with quantum keys [Col23, MW24, KMNY24, BGH+23]. Therefore, focusing on the perfect-complete setting does not impose a strong restriction. We conjecture that allowing non-perfectness does not change the impossibility result.*

This result follows and adapts the idea in the previous two papers by Li, Li, Li, and Liu. While we do not prove the conjecture in [LLLL25], we find a win-win scenario for the partial assignment μ. We will explicitly construct some polynomial-size μ such that

- either $f(x^\mu) \neq 0$ for all x, or
- there exist many disjoint partial assignments μ' (possibly not known) such that $f(x^{\mu' \cdot \mu}) \neq 0$ for all x.

We will show that in both cases, Eve simulates Alice on the oracle reprogrammed by μ will break PKE with advantage $1 - O(\epsilon)$. The first case is exactly what was conjectured in [LLLL25] and the partial assignment μ can be used to reprogram the oracle and yield the correct key. In the second case, if there exist sufficiently many μ', by an averaging argument, there must be at least one μ' on which the decryption algorithm has small query weight. This implies that we can simulate the decryption algorithm with x^μ instead of $x^{\mu' \cdot \mu}$ (since we do not know μ'), while not changing the output distribution significantly.

Result 2: Extending to quantum ciphertext. Using a quantum public key can cause challenges in public-key distribution, authentication, and reusability. [BGV+23] raised the open question of whether quantum PKE from one-way functions is possible when using classical keys and a quantum ciphertext; since ciphertext does not require to be distributed and reused. They proved that it is impossible, as long as the conjecture in [ACC+22] was true. We extend Theorem 1.1 to the quantum ciphertext case, by removing the conjecture in [BGV+23].

Theorem 1.2. *Perfect-complete quantum public key encryption, with classical keys and quantum ciphertext, does not exist in the quantum random oracle model.*

Result 3: Extending to quantum public keys. As discussed earlier, we believe achieving reusability and non-interactivity are the core of public key encryption; thus we focused on classical keys. Still, we examine the possibility to achieve QPKE with quantum keys from one-way functions. Particularly, we show that:

Theorem 1.3. *Perfect-complete quantum public key encryption, with quantum public keys, classical secret keys and classical/quantum ciphertext, does not exist in the quantum random oracle model, if the public key is uniquely determined by the secret key.*

Here, a public key $|\mathsf{pk}\rangle$ is said to be uniquely determined by a secret key sk if it can be generated by a procedure that depends only on sk and not on the random oracle.

Notably, even when public keys depend solely on secret keys, it remains unclear how to dequantize the quantum state into a classical string. As a result, we cannot directly reduce Theorem 1.3 to Theorem 1.1 and must instead adopt a different approach. Finally, our impossibility result is tight for all known QPKE constructions with quantum keys [Col23, MW24, KMNY24, BGH+23], as they all require $|\mathsf{pk}\rangle$ to depend on both sk and the random oracle.

Discussions and future directions. We highlight several interesting future directions after this work. The current proof is tailored for perfectly correct quantum PKE, leaving the non-perfect case largely unexplored. While we believe that the approximate Markov chain framework used to establish impossibility will remain useful, handling the non-perfect case requires a new approach to reprogram the random oracle. We conjecture that even non-perfect QPKE with classical keys does not exist in the QROM.

Another direction is to extend the separation result of the two-round key agreement to the multi-round setting, which could further support the impossibility of non-perfect QPKE from one-way functions. Suppose it can be shown that no three-round perfect-complete key agreement protocol exists in the quantum random permutation model (with only forward queries). This would imply that no two-round non-perfect-complete key agreement protocol exists either. The intuition behind this reduction is that a non-perfect two-round key agreement can be transformed into a perfect three-round KA by leveraging a random permutation P. Specifically, the parties can use $P(k)$ – the evaluation of P on their shared key k – as a cross-check: if they agree, they proceed; otherwise, they both output 0. This transformation effectively shifts the completeness deficiency to the security.

Finally, although quantum keys are not good for authenticating/reusing, there are not formal arguments establishing the limitation. [LLLL24] showed that for restricted key generation algorithms, QPKE with quantum public keys is either not secure, or quantum public keys are unclonable. However, unclonable does not imply non-reusability as the scheme in [BGH+23] offers some forms of private reusability. Thus, formalizing and ruling out public reusability for all QPKE even with quantum public keys will further illustrate the fundamental challenges of constructing QPKE from one-way functions.

2 Preliminaries

We assume familiarity with the basics of quantum computing and quantum information. For a comprehensive background, we refer the reader to [NC10]. Below, we present some backgrounds that are heavily used in this work.

2.1 Distance Measures

We recall the definitions of total variation distance and trace distance.

Definition 2.1 (Total Variation Distance)**.** *Given two probability distributions D_X and D_Y over a finite domain $\mathcal{X}$, the total variation distance between them is defined as*

$$TV(D_X, D_Y) = \frac{1}{2} \sum_{x \in \mathcal{X}} |D_X(x) - D_Y(x)| .$$

Here, $D_X(x)$ and $D_Y(x)$ denote the probability of x drawn from D_X and D_Y respectively.

Definition 2.2 (Trace Distance)**.** *For any two quantum states ρ and σ, the trace distance is defined by*

$$TD(\rho, \sigma) = \frac{1}{2} \operatorname{Tr} \left[\sqrt{(\rho - \sigma)^\dagger (\rho - \sigma)} \right] = \sup_{0 \le \Lambda \le I} \operatorname{Tr} \left[\Lambda(\rho - \sigma) \right] .$$

The following lemma is standard (e.g., see [LLLL24]). We include the proof for completeness.

Lemma 2.3. *For two probability distributions D_X and D_Y over the same domain, if $TV(D_X, D_Y) \le \epsilon$, we have that*

$$\Pr_{x \leftarrow D_X} [x \notin \mathsf{supp}(D_Y)] \le 2\epsilon.$$

Proof. $\sum_{x \notin \mathsf{supp}(D_Y)} D_X(x) \le \sum_x \left| D_X(x) - D_Y(x) \right| = 2TV(D_X, D_Y) \le 2\epsilon.$ □

2.2 Quantum Oracle Model and Random Oracle

In the quantum oracle model, a quantum algorithm $\mathcal{A}$ can make quantum queries to an oracle function $H : [2^n] \to \{0,1\}$ via a unitary transformation U_H mapping $|i, b\rangle$ to $|i, b \oplus H(i)\rangle$. We denote such an algorithm by $\mathcal{A}^H$, which can be expressed as a sequence of unitaries: $U_1, U_H, U_2, U_H, \ldots, U_d, U_H, U_{d+1}$. Here $U_1, \cdots, U_{d+1}$ are local unitaries acting on $\mathcal{A}$'s internal register.

Definition 2.4 (Query Weight)**.** *Consider a quantum algorithm $\mathcal{A}$ that makes d queries to an oracle H. Denote the quantum state immediately after t queries to the oracle as*

$$|\psi_t\rangle = \sum_{i,w} \alpha_{i,w,t} |i, w\rangle ,$$

where w is the content of the workspace register. Define the query weight q_i of input i as

$$q_i = \sum_{t=1}^{d} \sum_{w} |\alpha_{i,w,t}|^2.$$

Lemma 2.5 ([BBBV97])**.** *Consider two oracles $H, \tilde{H}$, and a quantum query algorithm $\mathcal{A}$ which makes d queries. Let $|\psi_d\rangle$ and $|\phi_d\rangle$ denote the final state before measurement when running $\mathcal{A}$ on H and $\tilde{H}$ respectively, and q_i denote the query weight of input i when running $\mathcal{A}$ on H. Then we have that*

$$\||\psi_d\rangle - |\phi_d\rangle\| \leq 2\sqrt{d}\sqrt{\sum_{i:\, \tilde{H}(i)\neq H(i)} q_i}.$$

We will also consider the quantum random oracle model (QROM). In this setting, a quantum algorithm has access to a random oracle $H : [2^{n_\lambda}] \to \{0,1\}$, which is chosen from the uniformly random distribution over all functions mapping $[2^{n_\lambda}]$ to $\{0,1\}$.

2.3 Entropy and Information

Definition 2.6 (Von Neumann Entropy)**.** *Let $\rho \in \mathbb{C}^{2^n}$ be a quantum state describing the system A, and let $|\phi_1\rangle, |\phi_2\rangle, \ldots, |\phi_{2^n}\rangle$ be an eigenbasis for ρ, so that*

$$\rho = \sum_i \eta_i |\phi_i\rangle \langle\phi_i| .$$

The Von Neumann entropy of ρ, denoted by $S(\rho)$ or $S(\mathsf{A})_\rho$, is defined as

$$S(\mathsf{A})_\rho = S(\rho) = -\sum_i \eta_i \log(\eta_i).$$

For a composite system AB with joint state ρ_{AB}, the conditional Von Neumann entropy is defined by

$$S(\mathsf{A}|\mathsf{B})_\rho = S(\mathsf{AB})_\rho - S(\mathsf{B})_\rho.$$

In the following, we will omit the subscript ρ when the quantum state is clear from context. For example, we will write $S(\mathsf{A})$ instead of $S(\mathsf{A})_\rho$, and $I(\mathsf{A} : \mathsf{B})$ instead of $I(\mathsf{A} : \mathsf{B})_\rho$.

Definition 2.7 (Mutual Information)**.** *Given a quantum state ρ that describes the joint systems A and B, the mutual information between A and B is given by*

$$I(\mathsf{A} : \mathsf{B}) = S(\mathsf{A}) + S(\mathsf{B}) - S(\mathsf{AB}).$$

Definition 2.8 (Conditional Mutual Information, CMI)**.** *Let ρ be a quantum state describing the three joint systems A, B, and C. Then the conditional mutual information is defined as*

$$I(\mathsf{A} : \mathsf{B}|\mathsf{C}) = S(\mathsf{AC}) + S(\mathsf{BC}) - S(\mathsf{ABC}) - S(\mathsf{C}).$$

The strong subadditivity property states that both the mutual information and the conditional mutual information are always non-negative.

Lemma 2.9 (Strong Subadditivity, [AL70])**.** *For Hilbert spaces A, B, and C, it holds that*

$$S(\mathsf{AC}) + S(\mathsf{AB}) \geq S(\mathsf{ABC}) + S(\mathsf{C}).$$

In its conditional form, for Hilbert spaces A, B, C, *and* D, *we have*

$$S(\mathsf{AC}|\mathsf{D}) + S(\mathsf{AB}|\mathsf{D}) \geq S(\mathsf{ABC}|\mathsf{D}) + S(\mathsf{C}|\mathsf{D}).$$

Fawzi and Renner [FR15] provided an insightful characterization of quantum states when the conditional mutual information is nearly zero. Intuitively, a small value of $I(\mathsf{A}:\mathsf{B}|\mathsf{E})$ indicates that the system B can be approximately reconstructed from system E. Formally,

Theorem 2.10 (Approximate Quantum Markov Chain, [FR15])**.** *For any state* ρ_{AEB} *over systems* AEB, *there exists a channel* $\mathcal{T} : \mathsf{E} \to \mathsf{E} \otimes \mathsf{B}'$ *such that the trace distance between the reconstructed state* $\sigma_{\mathsf{A'E'B'}} = \mathcal{T}(\rho_{\mathsf{AE}})$ *and the original state* ρ_{AEB} *is at most*

$$\sqrt{\ln 2 \cdot I(\mathsf{A}:\mathsf{B}|\mathsf{E})_\rho}.$$

2.4 Notations in Boolean Function Analysis

Any function $f : \{-1,1\}^N \to \mathbb{R}$ has a unique expression as a multilinear polynomial

$$f(x) = \sum_{S \subseteq [N]} a_S \cdot x_S,$$

where $x_S := \Pi_{i \in S} x_i$, and the coefficient a_S is given by $a_S = 2^{-N} \sum_x f(x) \cdot x_S$. The *degree* of f, denoted $\deg(f)$, is defined as the degree of its multilinear polynomial expression, i.e., $\max\{|S| : a_S \neq 0\}$. A monomial x_S is called *maximum* if $a_S \neq 0$ and it has degree $\deg(f)$, i.e., $|S| = \deg(f)$. Two monomials x_S and x_T are called *disjoint* if $S \cap T = \emptyset$. We say that f is *not identically zero* if $f(x) \not\equiv 0$.

A *partial assignment* is a function $\mu : [N] \to \{-1, 1, \star\}$. We define the support of μ as $\mathsf{supp}(\mu) := \{i | \mu(i) \neq \star\}$, and the size as $|\mu| := |\mathsf{supp}(\mu)|$. μ is called *empty* if $|\mu| = 0$. For $x \in \{-1,1\}^N$, we define the modification of x with μ, denoted by x^μ, as the string $x' \in \{-1,1\}^N$ such that

$$x'_i := \begin{cases} \mu(i) & \text{if } i \in \mathsf{supp}(\mu), \\ x_i & \text{otherwise.} \end{cases}$$

Given two partial assignments μ and η, define their *product*, denoted by $\mu \cdot \eta$, to be the partial assignment satisfying that $x^{\mu \cdot \eta} = (x^\mu)^\eta$ for any $x \in \{-1,1\}^N$. Note that the product operator is associative but not commutative. We say that two partial assignments μ and η are *disjoint* if $\mathsf{supp}(\mu) \cap \mathsf{supp}(\eta) = \emptyset$.

Lemma 2.11 ([BBC+01])**.** *Suppose a quantum algorithm makes* d *queries to a Boolean string*[1] $x \in \{-1,1\}^N$, *and the acceptance probability is denoted by* $f(x)$. *Then the function* $f : \{-1,1\}^N \to \mathbb{R}$ *has degree at most* $2d$. *That is,* f *can be expressed as*

$$f(x) = \sum_{|S| \leq 2d} a_S \cdot x_S.$$

Lemma 2.12 ([Alo99, Mid04])**.** *Let* $f : \{-1,1\}^N \to \mathbb{R}$ *be any function that is not identically zero, and* x_S *be any maximum monomial of* f. *For any* $x \in \{-1,1\}^N$, *there exists a* μ *with* $\mathsf{supp}(\mu) = S$ *such that* $f(x^\mu) \neq 0$.

[1] We interpret the string $x \in \{-1,1\}^N$ as a function $x : [N] \to \{-1,1\}$, and model queries to x as oracle queries to this function.

2.5 Quantum Public-Key Encryption

This section provides the formal definition of Quantum Public-Key Encryption (QPKE) and Quantum Key Agreement (QKA) in QROM.

Definition 2.13. *Let $\lambda \in \mathbb{Z}_+$ be the security parameter and $H\colon [2^{n_\lambda}] \to \{0,1\}$ be a random oracle. A quantum public-key encryption scheme, relative to H, consists of the following three bounded-query quantum algorithms:*

- $\mathsf{Gen}^H(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$*: the key generation algorithm that generates a pair of classical public key* pk *and classical secret key* sk*.*
- $\mathsf{Enc}^H(\mathsf{pk}, m) \to \mathsf{ct}$*: the encryption algorithm that takes as input the public key* pk *and the plaintext m, and produces a classical ciphertext* ct*.*
- $\mathsf{Dec}^H(\mathsf{sk}, \mathsf{ct}) \to m'$*: the decryption algorithm that takes as input the secret key* sk *and the ciphertext* ct*, and outputs the plaintext m'.*

The algorithms need to satisfy the following requirements:

Perfect Completeness $\Pr\left[\mathsf{Dec}^H\left(\mathsf{sk}, \mathsf{Enc}^H(\mathsf{pk}, m)\right) = m\colon \mathsf{Gen}^H(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})\right] = 1.$

IND-CPA Security *For any adversary $\mathcal{E}^H$ that makes* $\mathsf{poly}(\lambda)$ *quantum queries, for every two plaintexts $m_0 \neq m_1$ chosen by $\mathcal{E}^H$ after seeing* pk*, we have*

$$\Pr\left[\mathcal{E}^H\left(\mathsf{pk}, \mathsf{Enc}^H(\mathsf{pk}, m_b)\right) = b\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

For simplicity, we use "QPKE" to refer to quantum public-key encryption schemes with classical secret key, public key, and ciphertext, unless specified otherwise. Besides, we will also consider QPKE schemes with quantum public keys, defined as follows.

Definition 2.14 (QPKE with quantum public key)**.** *Let $\lambda \in \mathbb{Z}_+$ be the security parameter and $H\colon [2^{n_\lambda}] \to \{0,1\}$ be a random oracle. A quantum public-key encryption scheme with quantum public key, relative to H, consists of the following four bounded-query quantum algorithms:*

- $\mathsf{SKGen}^H(1^\lambda) \to \mathsf{sk}$*: the secret key generation algorithm that generates a classical secret key* sk*.*
- $\mathsf{PKGen}^H(\mathsf{sk}) \to \rho_{\mathsf{pk}}$*: the public key generation algorithm that takes the secret key* sk *and generates a quantum state ρ_{pk} as the public key.*
- $\mathsf{Enc}^H(\rho_{\mathsf{pk}}, m) \to \mathsf{ct}$*: the quantum encryption algorithm that takes the public key ρ_{pk} and the plaintext m, and produces a classical or quantum ciphertext ρ_{ct}.*
- $\mathsf{Dec}^H(\mathsf{sk}, \rho_{\mathsf{ct}}) \to m'$*: the quantum decryption algorithm that takes the secret key* sk *and the ciphertext ρ_{ct}, and outputs the plaintext m'.*

The algorithms need to satisfy the following requirements:

Perfect Completeness

$$\Pr\left[\mathsf{Dec}^H\left(\mathsf{sk}, \mathsf{Enc}^H(\rho_{\mathsf{pk}}, m)\right) = m\colon \mathsf{SKGen}^H(1^\lambda) \to \mathsf{sk}, \mathsf{PKGen}^H(\mathsf{sk}) \to \rho_{\mathsf{pk}}\right] = 1.$$

IND-CPA Security *For any adversary $\mathcal{E}^H$ that receives $\mathsf{poly}(\lambda)$ copies of the public key and can make $\mathsf{poly}(\lambda)$ queries, and every two plaintexts $m_0 \neq m_1$ chosen by the adversary, we have*

$$\Pr\left[\mathcal{E}^H\left(\rho_{\mathsf{pk}}^{\otimes \mathsf{poly}(\lambda)}, \mathsf{Enc}^H(\rho_{\mathsf{pk}}, m_b)\right) = b\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

In this paper, we focus on the setting where the quantum public key is uniquely determined by the secret key sk; that is, the quantum algorithm $\mathsf{PKGen}(\mathsf{sk})$ makes no queries to the oracle H. This setting covers all possible QPKE schemes with classical public keys, as we may assume, without loss of generality, that sk contains a copy of pk. Furthermore, we may assume that ρ_{pk} is a pure state, since sk can be taken to include a purification of ρ_{pk}.

Lastly, we define Quantum Key Agreement in QROM.

Definition 2.15 (Quantum Key Agreement in the Oracle Model)**.** *Let $\lambda \in \mathbb{Z}_+$ be the security parameter and let $H : [2^{n_\lambda}] \to \{0,1\}$ be a random oracle. A Quantum Key Agreement (QKA) protocol involves two parties, Alice and Bob, who initially begin with all-zero states. They can perform any quantum operations, make $\mathsf{poly}(\lambda)$ quantum queries to the oracle H, and exchange classical messages. At the end of the protocol, Alice and Bob output classical strings k_A and k_B, respectively.*

The protocol needs to satisfy the following conditions:

Correctness $\Pr[k_A = k_B] \geq 1/\mathsf{poly}(\lambda)$*, where the probability is taken over the randomness of Alice and Bob's channels, and the random oracle H.*

Security *For any eavesdropper Eve that makes $\mathsf{poly}(\lambda)$ quantum queries to H, eavesdrops on classical communication between Alice and Bob and outputs k_E, we have $\Pr[k_A = k_E] = \mathsf{negl}(\lambda)$.*

Similar to the perfect completeness in QPKE, a QKA protocol is said to be *perfect complete* if it satisfies $\Pr[k_A = k_B] = 1$.

3 Helper Lemmas

This section presents some helper lemmas, which may be of independent interest.

3.1 Information-Theoretic Tools

The following two information-theoretic lemmas from [LLLL24] will be used to prove our main results. We provide their proofs in Appendix A to make this paper self-contained.

Lemma 3.1 upper bounds how much entropy a quantum algorithm can accumulate through oracle queries.

Lemma 3.1 ([LLLL24])**.** *Consider an algorithm $\mathcal{A}$ that starts with a pure state, and makes d quantum queries to a random oracle $H : [2^n] \to \{0,1\}$. Let A denote the whole register of $\mathcal{A}$ and let ρ be the quantum state right before the final measurement. Then, it holds that $S(\mathsf{A})_\rho \leq 2d(n+1)$.*

Lemma 3.3 claims that repetition decreases CMI.

Definition 3.2 (Permutation Invariance)**.** *Let $\mathsf{A}, \mathsf{B}_1, \mathsf{B}_2, \cdots, \mathsf{B}_t$ be a $(t+1)$-partite quantum system. Given the joint state $\rho_{\mathsf{AB}_1\mathsf{B}_2\cdots\mathsf{B}_t}$, we say that $\mathsf{B}_1, \ldots, \mathsf{B}_t$ are permutation invariant if, for any permutation π on $[t]$, it holds that*

$$\rho_{\mathsf{AB}_1\mathsf{B}_2\cdots\mathsf{B}_t} = \rho_{\mathsf{AB}_{\pi(1)}\mathsf{B}_{\pi(2)}\cdots\mathsf{B}_{\pi(t)}}.$$

Lemma 3.3 (Lemma 4.2, [LLLL24]). *Let* $\mathsf{A}, \mathsf{B}_1, \mathsf{B}_2, \ldots, \mathsf{B}_t, \mathsf{C}$ *be a* $(t+2)$*-partite quantum system. Suppose the joint state* $\rho_{\mathsf{AC}\mathsf{B}_1\mathsf{B}_2\cdots\mathsf{B}_t}$ *is fully separable. If* $\mathsf{B}_1, \mathsf{B}_2, \ldots, \mathsf{B}_t$ *are permutation invariant, then there exists some* $0 \le j \le t-1$ *such that*

$$I(\mathsf{B}_t : \mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \ldots, \mathsf{B}_j)_\rho \le \frac{S(\mathsf{A})}{t}.$$

3.2 A Structural Property of Low-Degree Polynomials

In the proof of our main results, a key step involves explicitly reprogramming an unknown $x \in \{-1,1\}^N$ (representing the oracle) by modifying at most $\mathsf{poly}(d)$ bits, in order to make a given degree-d polynomial f (representing the probability that Gen or SKGen outputs a particular sk) evaluate to non-zero.

The following lemma, which builds heavily on Lemma 2.12, establishes a win-win situation: by obliviously modifying a small number of bits of the unknown x, denoted by a partial assignment μ, either we can already guarantee that $f(x^\mu) \neq 0$, or there must exist many *disjoint* (albeit unknown) partial assignments to make f evaluate to non-zero.

Lemma 3.4. *Let* $m > 0$ *be an integer. For any degree-*d *function* $f : \{-1,1\}^N \to \mathbb{R}$ *that is not identically zero, we can explicitly construct a partial assignment* μ *of* $|\mu| \le md^2$ *such that: either*

(a) for any $x \in \{-1,1\}^N$, $f(x^\mu) \neq 0$*; or*

(b) for any $x \in \{-1,1\}^N$*, there must exist* m *pairwise disjoint partial assignments* $\mu_1, \ldots, \mu_m$ *of size at most* d *such that* $f(x^{\mu_\ell \cdot \mu}) \neq 0$ *for all* $\ell \in [m]$.

Proof. We propose an algorithm to construct such a partial assignment μ. The algorithm maintains a function $\tilde{f}$ and a partial assignment $\tilde{\mu}$. Initially, $\tilde{f} = f$ and $\tilde{\mu}$ is empty. The algorithm contains at most $\deg(f)$ rounds: in each but the last round, we extend $\tilde{\mu}$ by fixing at most md additional bits, and reduce the degree of $\tilde{f}$ by at least 1. Specifically, in each round, the algorithm first construct a maximal set $\mathcal{S}$ of disjoint maximum monomials of $\tilde{f}$ (so any maximum monomial of $\tilde{f}$ intersects with at least one monomial in $\mathcal{S}$); Then

1. If $|\mathcal{S}| > m$, then stop and return $\tilde{\mu}$;
2. Otherwise, fix all variables appearing in $\mathcal{S}$ while keeping the new $\tilde{f}$ not identically zero. To do this, process each variable x_j in $\mathcal{S}$ one by one. For each, choose a value $b \in \{1,-1\}$ such that $\tilde{f}$ remains not identically zero after setting x_j as b. Such a choice always exists since $\tilde{f}$ is not identically zero. Let η be the resulting partial assignment, and update $\tilde{\mu}$ as $\tilde{\mu} \cdot \eta$ and $\tilde{f}$ as $\tilde{f}^\eta$. Here, the function $\tilde{f}^\eta(x)$ is defined as $\tilde{f}(x^\eta)$. Now, if $\deg(\tilde{f}) = 0$, then stop and return $\tilde{\mu}$.

We now analyze the algorithm. First, we claim that the final $\tilde{\mu}$ satisfies either condition (a) or (b). This is because that:

- If the algorithm stops because $\deg(\tilde{f}) = 0$, then $\tilde{f}$ is a constant function that is not zero, say $\tilde{f}(x) \equiv c \neq 0$. Therefore, for any x, we have $f(x^{\tilde{\mu}}) = \tilde{f}(x) = c \neq 0$, and condition (a) is satisfied.
- If the algorithm stops because $|\mathcal{S}| > m$, then given any x, for each maximum monomial ℓ of $\tilde{f}$ in $\mathcal{S}$, by Lemma 2.12, there exists a partial assignment μ_ℓ on the $\le d$ variables of ℓ such that $f(x^{\mu_\ell \cdot \tilde{\mu}}) = \tilde{f}(x^{\mu_\ell}) \neq 0$. Recalling that the monomials in $\mathcal{S}$ are disjoint, we conclude that condition (b) is satisfied.

Next, we show that $|\mu| \leq md^2$, and therefore finish the proof. Since any maximum monomial of $\tilde{f}$ intersects with at least one monomial in $\mathcal{S}$, fixing all variables in $\mathcal{S}$ reduces $\deg(\tilde{f})$ by at least 1. Hence the number of round is at most d. Moreover, in each round, $|\mu|$ increases by at most md. So we have $|\mu| \leq d \cdot (md) = md^2$. □

4 Impossibility of Perfect-Complete Quantum PKE

This section will prove that perfect-complete QPKE schemes do not exist in QROM. More formally, we have the following theorem.

Theorem 4.1 (Restate of Theorem 1.1)**.** *For any perfect-complete QPKE in QROM, which makes d queries to a random oracle $H : [2^n] \to \{0,1\}$ during each of* $\mathsf{Gen}, \mathsf{Enc}$ *and* Dec*, there exists an adversary Eve that can break the scheme w.p. $1 - O(\epsilon)$ by making $O\left(d^7 \log(d/\epsilon)/\epsilon^4 + nd^2/\epsilon^2\right)$ queries to H.*

The remainder of this section presents the proof of Theorem 4.1. It is well-known that: given a perfect-complete QPKE scheme $(\mathsf{Gen}^H, \mathsf{Enc}^H, \mathsf{Dec}^H)$, one can construct a perfect-complete two-round key agreement protocol between two parties, Alice and Bob, as follows.

1. Alice computes $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}^H(1^\lambda)$ and sends $m_0 := \mathsf{pk}$ to Bob. Denote this stage by $\mathcal{A}_0$.
2. Bob randomly chooses $k_B \in \{0,1\}$, computes $\mathsf{ct} \leftarrow \mathsf{Enc}^H(m_0, k_B)$, sends $m_1 := \mathsf{ct}$ to Alice and outputs k_B. Denote this stage by $\mathcal{B}$.
3. Alice computes $k_A \leftarrow \mathsf{Dec}^H(\mathsf{sk}, m_1)$ and outputs k_A. Denote this stage by $\mathcal{A}_1$.

Each stage of the key agreement makes at most d queries. Thus breaking this QKA also breaks QPKE. We now construct an eavesdropper Eve that sees (m_0, m_1) and guesses the agreed key w.p. $1 - O(\epsilon)$ by making $O(\mathsf{poly}(n, d, 1/\epsilon))$ queries to H. Eve's attack algorithm consists of three steps.

4.1 Step 1: Identify $\mathcal{B}$'s heavy queries.

The first step is to identify Bob's heavy queries, i.e., inputs with large query weight. These heavy queries will be kept unchanged when reprogramming the oracle in later steps, in order to ensure that the reprogrammed oracle will be indistinguishable from the real oracle from Bob's perspective.

Specifically, in this step, Eve computes a query record $R_E := \{(i_E, H(i_E))\}$ by repeating the following process $\frac{d^6}{\epsilon^4} \log \frac{d^6}{\epsilon^5}$ times:

1. Randomly choose $t \leftarrow [d]$, simulate $\mathcal{B}^H(m_0)$ to its t-th query to the oracle, and measure the input register, obtain outcome $i \in [2^n]$;
2. Classically query i to the oracle and add $(i, H(i))$ to R_E.

We claim that, with high probability, Eve can identify all of Bob's heavy queries. Formally,

Lemma 4.2. *Let q_i be the query weight of input i when running $\mathcal{B}$ on H, and $W_B := \{i : q_i \geq \epsilon^4/d^5\}$. Then $\Pr[W_B \not\subseteq R_E] \leq \epsilon$.*

Proof. For each $i \in W_B$, it would be measured w.p. at least ϵ^4/d^6 at each repetition. Thus the probability that it is not measured is bounded by

$$\Pr[i \notin R_E] \leq \left(1 - \frac{\epsilon^4}{d^6}\right)^{\frac{d^6}{\epsilon^4}\log\frac{d^6}{\epsilon^5}} \leq \epsilon^5/d^6.$$

Since $\sum_i q_i = d$, we have $|W_B| \leq d^6/\epsilon^4$. Thus by a union bound, we have $\Pr[W_B \nsubseteq R_E] \leq \epsilon$. □

4.2 Step 2: Sample a fake secret key

The next step is to obtain a fake secret key that is indistinguishable from the real secret key from Bob's perspective. We take an information-theoretic approach: first reduce the mutual information between Alice and Bob's registers conditioned on Eve's registers to a small value, and then apply the reconstruction channel in Theorem 2.10 to sample a fake secret key.

Specifically, consider the time right before $\mathcal{B}$ performs the final measurement. Let A denote the registers of $\mathcal{A}_0$ and B denote the registers of $\mathcal{B}$. Eve repetitively runs $\mathcal{B}^H(m_0)$ and stops right before the final measurement for $4dn/\epsilon^2$ times, which yields $4dn/\epsilon^2$ copies $\mathsf{B}_1, \mathsf{B}_2, \ldots, \mathsf{B}_{4dn/\epsilon^2}$ of B. Observe that $\mathsf{B}, \mathsf{B}_1, \mathsf{B}_2, \ldots, \mathsf{B}_{4dn/\epsilon^2}$ are permutation invariant. By Lemma 3.1 and Lemma 3.3, there exists a $0 \leq j \leq 4dn/\epsilon^2$ such that

$$I(\mathsf{A} : \mathsf{B}|R_E, \mathsf{B}_1, \mathsf{B}_2, \cdots, \mathsf{B}_j) \leq \frac{S(\mathsf{A})}{4dn/\epsilon^2} \leq \frac{2d(n+1)}{4dn/\epsilon^2} \leq \frac{\epsilon^2}{\ln 2}.$$

Note that Eve can compute such a j without making queries to H because H is traced out in Lemma 3.3. Eve only keeps the j copies of B, so that $\mathsf{E} := (R_E, \mathsf{B}_1, \mathsf{B}_2, \ldots, \mathsf{B}_j)$.

Then Eve applies the quantum channel $\mathcal{T} : \mathsf{E} \to \mathsf{E} \otimes \mathsf{A}'$ in Theorem 2.10 and generates a fake copy A' of A such that

$$TD(\rho_{ABE}, \rho_{A'BE}) \leq \sqrt{\ln 2 \cdot I(\mathsf{A} : \mathsf{B}|\mathsf{E})} < \epsilon,$$

where ρ_{ABE} is the state of system ABE and $\rho_{A'BE}$ is state of system $\mathsf{A'BE}$. Since A' contains a register sk' storing the secret key, Eve will use sk' as the fake secret key. Also, the channel $\mathcal{T}$ can be implemented without making queries to H.

Lemma 4.3. *Let* $\mathsf{View}_{ABE} := (\mathsf{sk}, m_0, k_B, m_1, R_E)$, $\mathsf{View}_{A'BE} := (\mathsf{sk}', m_0, k_B, m_1, R_E)$, D_{ABE} *denote the distribution of* View_{ABE}, *and* $D_{A'BE}$ *denote the distribution of* $\mathsf{View}_{A'BE}$. *We have*

$$\Pr_{\mathsf{View}_{A'BE} \leftarrow D_{A'BE}}[\mathsf{View}_{A'BE} \notin \mathsf{supp}(D_{ABE})] \leq 2\epsilon.$$

Proof. View_{ABE} and $\mathsf{View}_{A'BE}$ are obtained from performing measurement in computational basis on the corresponding registers of state ρ_{ABE} and $\rho_{A'BE}$ respectively. Since $TD(\rho_{ABE}, \rho_{A'BE}) \leq \epsilon$, we have that $TV(D_{ABE}, D_{A'BE}) \leq \epsilon$ by the operational meaning of trace distance. By Lemma 2.3, we have $\Pr_{\mathsf{View}_{A'BE} \leftarrow D_{A'BE}}[\mathsf{View}_{A'BE} \notin \mathsf{supp}(D_{ABE})] \leq 2\epsilon$. □

The above lemma shows that, with high probability, the tuple $(\mathsf{sk}', m_0, k_B, m_1, R_E)$ corresponds to a valid execution and is therefore compatible with some oracle H'. Here, we say that a tuple $(\mathsf{sk}', m_0, k_B, m_1, R_E)$ and an oracle H' are compatible if (i) running the QKA protocol on H' generates the view $(\mathsf{sk}', m_0, k_B, m_1)$ with non-zero probability; and (ii) the list of input-output pairs R_E is

consistent with H'. If Eve had access to such an oracle H', it could just compute $k_E \leftarrow \mathsf{Dec}^{H'}(\mathsf{sk}', m_1)$ to break the QKA protocol, since the perfect completeness ensures that $k_E = k_B$.

4.3 Step 3: Reprogram the oracle and run $\mathcal{A}_1$

In the final step, Eve first obtains a reprogrammed oracle $\tilde{H}$ by modifying $O(d^4/\epsilon^2)$ entries of the original oracle H (the details of the reprogramming will be specified later), and then computes k_E by running the decryption algorithm $\mathcal{A}_1$ on $\tilde{H}$ using sk'. In some cases, $\tilde{H}$ is compatible with the tuple $(\mathsf{sk}', m_0, k_B, m_1, R_E)$, in which case we have $k_E = k_B$. However, this compatibility does not always hold. Nevertheless, we can argue that $(\mathsf{sk}', m_0, k_B, m_1, R_E)$ is compatible with another (possibly unknown) oracle H' that is very close to $\tilde{H}$ and agrees with it on the heavy queries made by $\mathcal{A}_1$. In this case, by Lemma 2.5, we can still conclude that $k_E = k_B$ with high probability.

Here are the details of reprogramming. Let $N := 2^n$ and fix sk' and m_0. Let $g(x)$ denote the probability $\Pr[\mathcal{A}_0^H \to (\mathsf{sk}', m_0)]$ where $x \in \{-1, 1\}^N$ is the truth table of H by setting $x_i = (-1)^{H(i)}$. Define $f(x) := g(x^{R_E})$ where we abuse R_E as a partial assignment that assigns the i-th bit of x as $(-1)^y$ for all $(i, y) \in R_E$. Since $\mathcal{A}_0$ makes at most d queries, we have $\deg(f) \le 2d$ by Lemma 2.11.

If g is a zero polynomial, then Eve aborts[2]. Otherwise, Eve applies Lemma 3.4 on polynomial f by setting $m = d^2/\epsilon^2$ and obtains a partial assignment $\mu : [N] \to \{-1, 1, \star\}$ of size at most $m \cdot (2d)^2 = d^4/\epsilon^2$. We can assume $\mathsf{supp}(\mu) \cap \mathsf{supp}(R_E) = \emptyset$ because changing x_i for $i \in \mathsf{supp}(R_E)$ has no effect on the value of $f(x) = g(x^{R_E})$. Then Eve reprograms the oracle as

$$\tilde{H}(i) := \begin{cases} (1 - \mu(i))/2 & \text{if } i \in \mathsf{supp}(\mu) \\ H(i) & \text{otherwise} \end{cases}.$$

We have the following lemma. Intuitively, by Lemma 2.5, $\tilde{H}$ is likely to be compatible with Bob's view (m_1, k_B), as it differs from H on only a small fraction of Bob's query weight.

Lemma 4.4. *For the reprogrammed oracle $\tilde{H}$ defined above and any quantum algorithm $\mathcal{B}$ making d queries to the oracle,*

$$\Pr_{(k_B, m_1) \leftarrow \mathcal{B}^H(m_0)} \left[(k_B, m_1) \in \mathsf{supp}\left(\mathcal{B}^{\tilde{H}}(m_0)\right) \middle| \mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE}) \right] \ge 1 - O(\epsilon),$$

where we slightly abuse the notation $\mathcal{B}^H(m_0)$ for the output distribution of the algorithm $\mathcal{B}$.

Proof. Combining Lemma 4.3 and Lemma 4.2 we have that

$$\Pr[\mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE}) \wedge W_B \subseteq R_E] \ge 1 - \Pr[\mathsf{View}_{A'BE} \notin \mathsf{supp}(D_{ABE})] - \Pr[W_B \not\subseteq R_E] \ge 1 - O(\epsilon).$$

Now consider when $\mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE})$ and $W_B \subseteq R_E$. Since $\mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE})$, (sk', m_0, R_E) is valid under some oracle, which implies $f(x)$ is not identically zero. Then Eve will

[2] As we will see, this will never happen unless $\mathsf{View}_{A'BE} \notin \mathsf{supp}(D_{ABE})$.

not abort and $\tilde{H}$ is well-defined. We have

$$\begin{aligned} TV\left(\mathcal{B}^{\tilde{H}}(m_0), \mathcal{B}^{H}(m_0)\right) &\leq 4||\,|\psi_d\rangle - |\phi_d\rangle\,|| \\ &\leq 8\sqrt{d}\sqrt{\sum_{i\colon \tilde{H}(i)\neq H(i)} q_i} \\ &\leq 8\sqrt{d}\sqrt{\frac{\epsilon^4}{d^5}|\{i\colon \tilde{H}(i)\neq H(i)\}|} \\ &\leq 8\sqrt{d}\sqrt{\frac{\epsilon^4}{d^5}\cdot\frac{d^4}{\epsilon^2}} = O(\epsilon), \end{aligned}$$

where $|\psi_d\rangle$ and $|\psi_d\rangle$ are the states of $\mathcal{B}^{\tilde{H}}(m_0)$ and $\mathcal{B}^{H}(m_0)$ respectively, and q_i is the query weight of input i when running $\mathcal{B}$ on H. The first inequality comes from [BBBV97, Theorem 3.1], the second inequality is Lemma 2.5, the third inequality is because H and $\tilde{H}$ only differ on inputs that are outside W_B, and the last inequality is because $|\{i\colon \tilde{H}(i)\neq H(i)\}| \leq |\mu| \leq d^4/\epsilon^2$.

By Lemma 2.3, we have that

$$\Pr_{(k_B,m_1)\leftarrow\mathcal{B}^{H}(m_0)}\left[(k_B,m_1)\in \mathsf{supp}\left(\mathcal{B}^{\tilde{H}}(m_0)\right)\middle|\mathsf{View}_{A'BE}\in\mathsf{supp}(D_{ABE})\wedge W_B\subseteq R_E\right]\geq 1-O(\epsilon).$$

The final statement follows from a conditional probability formula and Lemma 4.2. □

4.4 Putting things together

Now, we are ready to prove Theorem 4.1.

Proof of Theorem 4.1. We will prove that by the 3-step attack algorithm described above, Eve will output k_E such that $\Pr[k_E = k_B] = 1 - O(\epsilon)$. First of all, by Lemma 4.3, we have that

$$\Pr[\mathsf{View}_{A'BE}\in\mathsf{supp}(D_{ABE})]\geq 1-O(\epsilon). \tag{1}$$

Now consider the case when $\mathsf{View}_{A'BE}\in\mathsf{supp}(D_{ABE})$, which implies that $(\mathsf{sk}', m_0, k_B, m_1, R_E)$ will be a valid execution under some oracle. In this case, the function f is not identically zero, and Eve will obtain a partial assignment μ without aborting. By Lemma 3.4, one of the following cases must hold:

(a) $f(x^\mu)\neq 0$. Observe that x^μ can be viewed as the boolean string of the reprogrammed oracle $\tilde{H}$ since the i-th bit of x^μ equals $(-1)^{\tilde{H}(i)}$. Thus $\Pr[(\mathsf{sk}', m_0)\leftarrow\mathcal{A}_0^{\tilde{H}}] = f(x^\mu)$ is non-zero, which means $(\mathsf{sk}', m_0)\in\mathsf{supp}(\mathcal{A}_0^{\tilde{H}})$.

(b) There exist d^2/ϵ^2 pairwise disjoint partial assignments $\mu_1,\ldots,\mu_{d^2/\epsilon^2}$ of size $\leq 2d$ such that $f(x^{\mu_\ell\cdot\mu})\neq 0$ for all $\ell\in[d^2/\epsilon^2]$. We can assume that for all ℓ, $\mathsf{supp}(\mu_\ell)\cap\mathsf{supp}(R_E)=\emptyset$ as changing x_i for $i\in\mathsf{supp}(R_E)$ has no effect on the value of $f(x) = g(x^{R_E})$. Then observe that

$x^{\mu_\ell \cdot \mu}$ can be viewed as the boolean string of the following oracle

$$\tilde{H}_\ell(i) := \begin{cases} (1-\mu_\ell(i))/2 & \text{if } i \in \mathsf{supp}(\mu_\ell) \setminus \mathsf{supp}(\mu) \\ \tilde{H}(i) & \text{otherwise} \end{cases}.$$

Thus $\Pr[(\mathsf{sk}', m_0) \leftarrow \mathcal{A}_0^{\tilde{H}_\ell}] = f(x^{\mu_\ell \cdot \mu})$ is non-zero, which means $(\mathsf{sk}', m_0) \in \mathsf{supp}(\mathcal{A}_0^{\tilde{H}_\ell})$.

Next, we argue that in both cases, Eve will output $k_E = k_B$ with probability $1 - O(\epsilon)$.

Case (a) We argue that $\mathsf{View}_{A'BE} = (\mathsf{sk}', m_0, k_B, m_1, R_E)$ is compatible with $\tilde{H}$ with high probability. Obviously, $\tilde{H}$ is consistent with R_E. Moreover,

1. From perspective of A', $(\mathsf{sk}', m_0) \in \mathsf{supp}(\mathcal{A}_0^{\tilde{H}})$ implies that (sk', m_0) is compatible with $\tilde{H}$.
2. From perspective of Bob, $\mathcal{B}^H(m_0)$ represents a distribution over key-message pairs (k_B, m_1). Now suppose we run the algorithm $\mathcal{B}^{\tilde{H}}(m_0)$ instead. According to Lemma 4.4, with probability at least $1 - O(\epsilon)$, a pair (k_B, m_1) produced by $\mathcal{B}^H(m_0)$ will also lie within the support of $\mathcal{B}^{\tilde{H}}(m_0)$.

So, in particular, $\mathsf{View}_{A'B}$ is a valid execution under $\tilde{H}$ with probability $1 - O(\epsilon)$. Conditioned on that $\mathsf{View}_{A'B}$ is valid under $\tilde{H}$, the perfect completeness implies that $k_E = \mathcal{A}_1^{\tilde{H}}(\mathsf{sk}', m_1)$ must equal k_B. Thus we have $\Pr[k_E = k_B] = 1 - O(\epsilon)$ in Case (a).

Case (b) Let w_i be the query weight of input i when running $\mathcal{A}_1(\mathsf{sk}', m_1)$ on $\tilde{H}$. Note that $\sum_i w_i \le d$ because $\mathcal{A}_1$ makes at most d queries to H. Since $\mu_1, \ldots, \mu_{d^2/\epsilon^2}$ are disjoint partial assignments, there must exist a $\ell^* \in [d^2/\epsilon^2]$ such that

$$\sum_{i \in \mathsf{supp}(\mu_{\ell^*})} w_i \le \frac{d}{d^2/\epsilon^2} = \frac{\epsilon^2}{d}. \tag{2}$$

For simplicity, let H' denote $\tilde{H}_{\ell^*}$. First, imagine that Eve runs $\mathcal{A}_1(\mathsf{sk}', m_1)$ on H' and obtains a key k'_E. Observe that H' and H differ by at most $|\mu_{\ell^*} \cdot \mu| \le 2d + O(d^4/\epsilon^2) = O(d^4/\epsilon^2)$ positions and $\mathsf{supp}(\mu_{\ell^*} \cdot \mu) \cap \mathsf{supp}(R_E) = \emptyset$. By the same argument as in Lemma 4.4, we have

$$\Pr_{(k_B, m_1) \leftarrow \mathcal{B}^H(m_0)} \left[(k_B, m_1) \in \mathsf{supp}\left(\mathcal{B}^{H'}(m_0)\right) \middle| \mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE}) \right] = 1 - O(\epsilon). \tag{3}$$

Then, by the same argument as in Case (a), we have

$$\Pr_{k'_E \leftarrow \mathcal{A}_1^{H'}(\mathsf{sk}', m_1)} \left[k'_E = k_B \middle| (k_B, m_1) \in \mathsf{supp}\left(\mathcal{B}^{H'}(m_0)\right), \mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE}) \right] = 1. \tag{4}$$

Combining Eqs. (1), (3), and (4), we have $\Pr_{k'_E \leftarrow \mathcal{A}_1^{H'}(\mathsf{sk}', m_1)} \left[k'_E = k_B \right] = 1 - O(\epsilon)$, which means Eve will find the key with probability $1 - O(\epsilon)$ if it runs $\mathcal{A}_1(\mathsf{sk}', m_1)$ on oracle H'. However, Eve knows only the existence of H', but does not know how to access it. The next step is to argue that by running $\mathcal{A}_1(\mathsf{sk}', m_1)$ on oracle $\tilde{H}$ instead, as done in the attack algorithm, Eve can also find the key

with high probability. For any fixed real oracle H, we have that

$$\begin{aligned} TV\left(\mathcal{A}_1^{H'}(\mathsf{sk}', m_1), \mathcal{A}_1^{\tilde{H}}(\mathsf{sk}', m_1)\right) &\leq 4\| \, |\psi_d\rangle - |\phi_d\rangle \, \| \\ &\leq 8\sqrt{d}\sqrt{\sum_{i:H'(i)\neq\tilde{H}(i)} w_i} \\ &\leq 8\sqrt{d}\sqrt{\sum_{i\in\mu_{\ell^*}} w_i} \leq 8\sqrt{d}\sqrt{\frac{\epsilon^2}{d}} = O(\epsilon) \end{aligned}$$

where $|\psi_d\rangle$ and $|\psi_d\rangle$ are the states of $\mathcal{A}_1^{H'}(\mathsf{sk}', m_1)$ and $\mathcal{A}_1^{\tilde{H}}(\mathsf{sk}', m_1)$ right before the final measurement, respectively. The first inequality comes from [BBBV97, Theorem 3.1], the second inequality is Lemma 2.5, the third inequality is by definition of H', and the last inequality is Eq. (2).

Since the distribution of $k_E \leftarrow \mathcal{A}_1^{\tilde{H}}(\mathsf{sk}', m_1)$ and $k'_E \leftarrow \mathcal{A}_1^{H'}(\mathsf{sk}', m_1)$ are $O(\epsilon)$-close for any H, replacing k'_E with k_E in Eq. (4) will only cause $O(\epsilon)$ loss of the probability, i.e.,

$$\Pr_{k_E\leftarrow\mathcal{A}_1^{\tilde{H}}(\mathsf{sk}',m_1)}\left[k_E = k_B \middle| (k_B, m_1) \in \mathsf{supp}\left(\mathcal{B}^{H'}(m_0)\right), \mathsf{View}_{A'BE} \in \mathsf{supp}(D_{ABE})\right] = 1 - O(\epsilon). \tag{5}$$

Combining Eqs. (1), (3), and (5), we have $\Pr_{k_E\leftarrow\mathcal{A}_1^{\tilde{H}}(\mathsf{sk}',m_1)}[k_E = k_B] = 1 - O(\epsilon)$.

Finally, we analyze the query complexity of Eve's attack algorithm. Step 1 requires at most $(d+1)\cdot d^6\log(d^6/\epsilon^5)/\epsilon^4 = O(d^7\log(d/\epsilon)/\epsilon^4)$ queries. Step 2 requires at most $d\cdot 4dn/\epsilon^2 = O\left(d^2n/\epsilon^2\right)$ queries. Step 3 requires d queries. Thus Eve makes $O\left(d^7\log(d/\epsilon)/\epsilon^4 + nd^2/\epsilon^2\right)$ queries in total. □

4.5 Extending to quantum public key and ciphertext

By further reviewing our proof, we can extend the impossibility result to quantum m_0 and m_1. Specifically, as long as the public key $|m_0\rangle$ is a pure state that is uniquely determined by the secret key, and Eve can access polynomially many copies of $|m_0\rangle$, the attack algorithm still works, with a few minor modifications detailed below.

- Steps 1 and 2 of the attack require Eve to run $\mathcal{B}^H(|m_0\rangle)$ for polynomial times. As Eve can obtain polynomially many copies of $|m_0\rangle$, these two steps are doable and the related analysis still holds.
- In Step 3, since Eve does not have the full description of the quantum m_0, and thus is not capable of identifying the polynomial that represents the probability of outputting (sk', m_0). However, the key observation is that since $|m_0\rangle$ is uniquely determined by the secret key, we can instead define the polynomial f as the probability of the oracle outputting sk' alone, i.e.,

  $$g(x) := \Pr[\mathcal{A}_0^H \to \mathsf{sk}'], \quad f(x) := g(x^{R_E}).$$

  Then Eve applies Lemma 3.4 on this f, obtains a reprogrammed oracle $\tilde{H}$, and finally outputs $k_E \leftarrow \mathcal{A}_1^{\tilde{H}}(\mathsf{sk}', \rho_{m_1})$.

The proof is almost the same as before, and we only sketch the main ideas here.

1. From the perspective of A', the reprogramming the oracle using f will guarantee $\tilde{H}$ in Case (a) (and H' in Case (b)) to produce sk' with non-zero probability. We argue that the fake secret key sk' will produce the real public key $|m_0\rangle$ with high probability: Since B is not affected by channel $\mathcal{T}$, we uncompute $\mathcal{B}$ on state $\rho_{A'B}$. As the uncomputation does not increase trace distance, we can see that the state of $(\mathsf{sk}', |m_0\rangle)$ is ϵ-close to that of $(\mathsf{sk}, |m_0\rangle)$. Thus with probability $1-\epsilon$, sk' will produce $|m_0\rangle$.

2. From the perspective of Bob, the output state of $\mathcal{B}^H(|m_0\rangle)$ is a cq-state $\rho_B = \sum_{k_B} p_{k_B} |k_B\rangle\langle k_B| \otimes \rho_{k_B}$ where ρ_{k_B} is the state of m_1 conditioned on k_B. Similarly for $\mathcal{B}^{\tilde{H}}(|m_0\rangle)$ in Case (a) (and $\mathcal{B}^{H'}(|m_0\rangle)$ in Case (b)), we express the output state as $\sigma_B = \sum_{k_B} p'_{k_B} |k_B\rangle\langle k_B| \otimes \sigma_{k_B}$. By using Lemma 2.5 as in the proof of Theorem 4.1, we will get $TD(\rho_B, \sigma_B) \leq O(\epsilon)$. Let $\rho'_B := \sum_{k_B} p'_{k_B} |k_B\rangle\langle k_B| \otimes \rho_{k_B}$, $D_B := \{p_{k_B}\}$ and $D_{B'} := \{p_{k'_B}\}$. Then we have

$$TD(\rho'_B, \sigma_B) \leq TD(\rho_B, \rho'_B) + TD(\rho_B, \sigma_B) = TV(D_B, D_{B'}) + TD(\rho_B, \sigma_B) \leq 2TD(\rho_B, \sigma_B) \leq O(\epsilon),$$

where the first inequality is the triangle inequality, the second inequality is because partial trace does not increase trace distance, and the last inequality is $TD(\rho_B, \sigma_B) \leq O(\epsilon)$. Since $TD(\rho'_B, \sigma_B) = \mathbb{E}[TD(\rho_{k_B}, \sigma_{k_B})]$, by Markov inequality, we have

$$\Pr_{k_B \leftarrow D_B}[TD(\rho_{k_B}, \sigma_{k_B}) \leq \sqrt{\epsilon}] \geq 1 - O(\sqrt{\epsilon}).$$

By the above argument, with probability $1 - O(\sqrt{\epsilon})$, the tuple $(\mathsf{sk}', |m_0\rangle, k_B, \rho_{k_B}, R_E)$ is $O(\sqrt{\epsilon})$-close to a tuple $(\mathsf{sk}', |m_0\rangle, k_B, \sigma_{k_B}, R_E)$ which is compatible with oracle $\tilde{H}$ in Case (a) (and H' is Case (b)). In Case (a), by perfect completeness, it follows that $k_E = k_B$ with probability $1 - O(\sqrt{\epsilon})$. In Case (b), firstly for $k'_E \leftarrow \mathcal{A}_1^{H'}(\mathsf{sk}', \sigma_{k_B})$, we have $k'_E = k_B$ with probability $1 - O(\sqrt{\epsilon})$. We then apply Lemma 2.5 to the algorithm $\mathcal{A}_1(\mathsf{sk}', \sigma_{k_B})$ under oracles $\tilde{H}$ and H' to conclude that $k_E = k_B$ with probability $1 - O(\sqrt{\epsilon})$. We remark that although Lemma 2.5 is stated for pure-state inputs, it also applies to the mixed-state input σ_{k_B}, since we can always assume the input to be the purification of σ_{k_B}.

Recall the IND-CPA security notion from Definition 2.14. If the public key is a pure state, the adversary algorithm $\mathcal{E}$ can obtain polynomial number of copies of $|\mathsf{pk}\rangle$. Given any two plaintexts $m_0 \neq m_1$, we can create a one-bit key agreement by designating the ciphertext as the second message (choosing between ρ_{ct_0} and ρ_{ct_1}). Therefore, by executing our modified attack algorithm, we can break the IND-CPA security with an advantage of $1 - O(\sqrt{\epsilon})$.

Theorem 4.5 (Restate of Theorem 1.3)**.** *For any perfect-complete QPKE with quantum public key in QROM, which makes d queries to the random oracle $H : [2^n] \to \{0,1\}$ during each of* $\mathsf{SKGen}, \mathsf{Enc}$ *and* Dec*, and no queries during* PKGen*, there exists an adversary Eve that can break the scheme with probability $1 - O(\sqrt{\epsilon})$ by making $O\left(d^7 \log(d/\epsilon)/\epsilon^4 + nd^2/\epsilon^2\right)$ queries to H.*

Theorem 1.2 is immediately implied by Theorem 4.5.

References

[ACC+22] Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In *Annual International Cryptology Conference*, pages 165–194. Springer, 2022. 2, 3

[AL70] Huzihiro Araki and Elliott H Lieb. Entropy inequalities. *Communications in Mathematical Physics*, 18(2):160–170, 1970. 6

[Alo99] Noga Alon. Combinatorial nullstellensatz. *Comb. Probab. Comput.*, 8(1–2):7–29, jan 1999. 7

[BB14] Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical computer science*, 560:7–11, 2014. 1

[BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. 6, 14, 16

[BBC+01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. 7

[BCKM21] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. On the round complexity of secure quantum computation. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 406–435. Springer, 2021. 2

[BGH+23] Khashayar Barooti, Alex B. Grilo, Loïs Huguenin-Dumittan, Giulio Malavolta, Or Sattath, Quoc-Huy Vu, and Michael Walter. Public-Key Encryption with Quantum Keys. In Guy Rothblum and Hoeteck Wee, editors, *Theory of Cryptography*, pages 198–227, Cham, 2023. Springer Nature Switzerland. 2, 3, 4

[BGV+23] Samuel Bouaziz, Alex B Grilo, Damien Vergnaud, Quoc-Huy Vu, et al. Towards the impossibility of quantum public key encryption with classical keys from one-way functions. *Cryptology ePrint Archive*, 2023. 2, 3

[Col23] Andrea Coladangelo. Quantum trapdoor functions from classical one-way functions. *arXiv preprint arXiv:2302.12821*, 2023. 2, 3, 4

[FR15] Omar Fawzi and Renato Renner. Quantum conditional mutual information and approximate markov chains. *Communications in Mathematical Physics*, 340(2):575–611, 2015. 2, 7

[GLSV21] Alex B Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in miniqcrypt. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 531–561. Springer, 2021. 2

[Imp95] R. Impagliazzo. A personal view of average-case complexity. In *Proceedings of the 10th Annual Structure in Complexity Theory Conference (SCT'95)*, SCT '95, page 134, USA, June 1995. IEEE Computer Society. 1

[IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 44–61, 1989. 1, 3

[KMNY24] Fuyuki Kitagawa, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum public-key encryption with tamper-resilient public keys from one-way functions. In *Advances in Cryptology – CRYPTO 2024: 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2024, Proceedings, Part VII*, page 93–125, Berlin, Heidelberg, 2024. Springer-Verlag. 2, 3, 4

[LLLL24] Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. How (not) to build qpke in minicrypt. In *Annual International Cryptology Conference*. Springer, 2024. 2, 3, 4, 5, 9, 10, 19

[LLLL25] Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. Toward the Impossibility of Perfect Complete Quantum PKE from OWFs. In Raghu Meka, editor, *16th Innovations in Theoretical Computer Sciencshoe Conference (ITCS 2025)*, volume 325 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 71:1–71:16, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. 2, 3

[Mid04] Gatis Midrijanis. Exact quantum query complexity for total boolean functions, 2004. 7

[MW24] Giulio Malavolta and Michael Walter. Robust quantum public-key encryption with applications to quantum key distribution. In *Advances in Cryptology – CRYPTO 2024: 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2024, Proceedings, Part VII*, page 126–151, Berlin, Heidelberg, 2024. Springer-Verlag. 2, 3, 4

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition edition, 2010. 5

[Sho99] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999. 1

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

[Wil11] Mark M Wilde. From classical to quantum shannon theory. *arXiv preprint arXiv:1106.1445*, 2011. 20

A Missing Proofs in Section 3.1

For the convenience of readers, we provide the proofs of two lemmas from [LLLL24] that are used in the main text.

A.1 Proof of Lemma 3.1

We can realize the quantum query unitary U_H via a quantum communication protocol involving two parties: the algorithm $\mathcal{A}$ and Oracle. To execute U_H, the protocol proceeds as follows:

1. $\mathcal{A}$ sends both its input register and output register, $n+1$ qubits in total, to Oracle;
2. Oracle applies the unitary U_H on these $n+1$ qubits and then returns them to $\mathcal{A}$.

By the subadditivity of entropy, each such quantum communication can increase the entropy of $\mathcal{A}$'s whole register by at most $2(n+1)$. In addition, applying local unitary U_i does not change the entropy. Since $\mathcal{A}$'s register A initially contains a pure state (with zero entropy), it follows that $S(\mathsf{A})_\rho \leq 2d(n+1)$ after d such rounds.

A.2 Proof of Lemma 3.3

The following basic facts below will be used.

Fact A.1 ([Wil11]). *If ρ_{AB} is a separable state, then $S(\mathsf{A}|\mathsf{B}) \geq 0$.*

Fact A.2 (Chain rule). $I(\mathsf{B}_1, \mathsf{B}_2, \cdots, \mathsf{B}_t : \mathsf{A} \mid \mathsf{C}) = \sum_{i=1}^{t} I(\mathsf{B}_i : \mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \cdots, \mathsf{B}_{i-1})$.

Proof of Lemma 3.3. By the chain rule for conditional mutual information (Fact A.2), we have

$$\sum_{i=1}^{t} I(\mathsf{B}_i : \mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \ldots, \mathsf{B}_{i-1}) = I(\mathsf{B}_1, \ldots, \mathsf{B}_t : \mathsf{A} \mid \mathsf{C}). \tag{6}$$

Moreover, we have

$$I(\mathsf{B}_1, \ldots, \mathsf{B}_t : \mathsf{A} \mid \mathsf{C}) = S(\mathsf{A} \mid \mathsf{C}) - S(\mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \ldots, \mathsf{B}_t) \leq S(\mathsf{A} \mid \mathsf{C}) \leq S(\mathsf{A}), \tag{7}$$

where the inequalities follow from Fact A.1 and the non-negativity of $I(\mathsf{A} : \mathsf{C}) = S(\mathsf{A}) - S(\mathsf{A} \mid \mathsf{C})$. Combining (6) and (7), it follows that there exists some $i \in [t]$ for which

$$I(\mathsf{B}_i : \mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \ldots, \mathsf{B}_{i-1}) \leq \frac{S(\mathsf{A})}{t}.$$

Finally, by permutation invariance, we have

$$I(\mathsf{B}_i : \mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \ldots, \mathsf{B}_{i-1}) = I(\mathsf{B}_t : \mathsf{A} \mid \mathsf{C}, \mathsf{B}_1, \ldots, \mathsf{B}_{i-1}).$$

This completes the proof. □