# Toward the Impossibility of Perfect Complete Quantum PKE from OWFs

**Longcheng Li**
State Key Lab of Processors, Institute of Computing Technology,
Chinese Academy of Sciences, Beijing, China

**Qian Li**
Shenzhen International Center for Industrial and Applied Mathematics,
Shenzhen Research Institute of Big Data, China

**Xingjian Li**
Tsinghua University, Beijing, China

**Qipeng Liu**
University of California, San Diego, La Jolla, CA, USA

## Abstract

In this paper, we study the impossibility of constructing perfect complete quantum public key encryption (QPKE) from quantumly secure one-way functions (OWFs) in a black-box manner. We show that this problem is connected to a fundamental conjecture about the roots of low-degree polynomials on the Boolean hypercube. Informally, the conjecture asserts that for every nonconstant low-degree polynomial, there exists a universal (randomized) way to modify a small number of input bits such that, for every input string, the polynomial evaluated on the modified input string avoids 0 with sufficiently large probability (over the choice of how the input string is modified). Assuming this conjecture, we demonstrate the impossibility of constructing QPKE from quantumly secure one-way functions in a black-box manner, by employing the information-theoretical approach recently developed by Li, Li, Li, and Liu (CRYPTO'24). Towards resolving this conjecture, we provide various pieces of evidence supporting it and prove some special cases. In particular, we fully rule out perfect QPKE from OWFs when the key generation algorithm only makes a logarithmic number of quantum queries, improving the previous work, which can only handle classical queries.

**2012 ACM Subject Classification** Theory of computation → Cryptographic primitives; Theory of computation → Quantum complexity theory

**Keywords and phrases** Qautnum public-key encryption, Boolean function analysis

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2025.71

**Funding** *Qian Li*: Qian Li's work was supported by Hetao Shenzhen-Hong Kong Science and Technology Innovation Cooperation Zone Project (No.HZQSWS-KCCYB-2024016).

## 1 Introduction

Public-key encryption (PKE) is a fundamental primitive in modern cryptography. It enables secure communication through an insecure channel. A server generates two keys, the secret key and the public key; the secret key is owned only by the server, whereas the public key is broadcasted to everyone else. PKE allows everyone to send messages to the server securely; even if a malicious party, who keeps listening to the channel, has no idea about the actual messages. Its counterpart, secret-key encryption, requires pre-shared keys to conduct secure

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).
Editor: Raghu Meka; Article No. 71; pp. 71:1–71:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

communication. Since the first proposal [1] by Diffie and Hellman [13], PKE has become one of the most important cryptographic primitives and concepts, with impacts ranging from theoretical computer science to real-world constructions.

Despite enjoying all the advancements in secret-key encryption (SKE), all PKE schemes have more structures. Therefore, PKE seems to require strictly stronger assumptions than one-way functions. This observation was later confirmed by Impagliazzo and Rudich [16]: there was no black-box construction of PKE solely from one-way functions (or in the random oracle model).

Quantum information has changed the landscape of cryptography, especially weakening the assumptions needed for various cryptographic primitives; for example, quantum key distribution [10], oblivious transfer/multi-party computation [7, 15], commitment [3, 23]. Talking about PKE, first introduced by Morimae and Yamakawa [22], then followed by [12, 17, 19, 6], QPKE with either quantum keys or quantum ciphertext can be constructed from OWFs. However, quantum public keys and quantum ciphertext are often more difficult to broadcast, use and authenticate than their classical counterparts. Thus, the possibility of basing quantum PKE with classical public key, secret key and ciphertext on one-way functions is still intriguing and largely unanswered. In the rest of the work, we will simply denote such QPKE with classical keys and ciphertext by "QPKE". Following the similar questions in [16], in this work, we are interested in the problem:

*Question 1: Whether QPKE can be constructed solely in the quantum random oracle model (QROM)?*

The direction was investigated first by Austrin, Chung, Chung, Fu, Lin and Mahmoody [5] and later by Li, Li, Li and Liu [18]. Austrin et al. proposed the so-called "polynomial compatible conjecture" (or PCC for short); basing on PCC, they fully rule out QPKE from OWFs. However, their conjecture is tailored to this separation problem; to which we will compare our conjecture later on. Besides, they are able to prove the non-existence of perfect complete QPKE in the QROM when (i) Enc or (ii) both Gen and Dec make only classical queries to a random oracle. Li et al. improved the result (ii) and showed that as long as the key generation algorithm Gen only makes classical queries, perfect correct QPKE can not exist in the QROM.

## 1.1 Our Results

In this work, we show that Question 1 is closely related to a fundamental conjecture about the roots of low-degree polynomials on the Boolean hypercube. More specifically, we propose the following conjecture, and show that it implies a full impossibility result of perfect complete QPKE in the QROM.

► **Conjecture 1.** *For any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, the following holds. There exists a distribution $\mathcal{D}$ on the set of partial assignments $\rho$ with $|\rho| \leq \mathrm{poly}(d)$ such that: for any $x \in \{-1,1\}^n$,*

$$\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/\mathrm{poly}(d).$$

[1] The UK Government Communications Headquarters proposed and developed the same scheme even earlier than Diffie and Hellman (1973 and 1974), and was later declassified by the British government in 1997.

A partial assignment is denoted by $\rho \in \{-1, 1, \star\}^n$ and $|\rho|$ is defined as the number of non-$\star$ entries. $x^\rho$ is defined as an input string whose $i$-th bit will be equal to $x_i$ if $\rho_i = \star$, and otherwise equal to $\rho_i$. In other words, the conjecture asserts that for every degree $d$ nonconstant $f$, there is a universal (randomized) way to modify poly$(d)$ bits such that, every input string $x$ has an inverse-polynomial probability of evaluating to non-zero under this randomized modification.

▶ Remark 2 (Comparing PCC in [5] with our conjecture). PCC (Conjecture 1.2) roughly asserts that, for two distributions $F, G$ over low degree polynomials with bounded influences, there always exists $f \in F, g \in G$ and $x$ such that $f(x)g(x) \neq 0$. To the best of our knowledge, there is no obvious connection between PCC and our conjecture. Nonetheless, we hold that our conjecture appears simpler and offers an alternative path towards the ultimate objective; since our conjecture only involves a single polynomial (instead of two distributions) and resembles many fundamental notions in the literature of boolean function analysis. Furthermore, PCC might be too strong to hold: as it will imply attacks with only a polynomial number of classical queries, even if both Alice and Bob make quantum queries; on the other hand, our Eve makes quantum queries.

With the conjecture, we now present our main theorem.

▶ **Theorem 3** (Main Theorem). *If Conjecture 1 is true, then perfect complete QPKE does not exist in the QROM.*

▶ Remark 4. In the proof of Theorem 3, $f$ is treated as a probability and ranges in $[0, 1]$. However, it would not weaken Conjecture 1 if we restrict the range of $f$ to $[0, 1]$. Precisely, in Conjecture 1, it does not lose generality to assume that $f$ is nonnegative, because we can consider $f^2$ instead; then we can scale the function to fit in the range $[0, 1]$.

Conjecture 1 is closely related to notions like certificate complexity, sensitivity in the literature of boolean function analysis (BFA), and also to the celebrated Combinatorial Nullstellensatz of Alon [2], making it a considerably natural conjecture. With the tools in BFA and the Combinatorial Nullstellensatz, we can provide many pieces of evidence and prove special cases. We mention some of the most important results below while leaving others in the main body.

## Special Case 1: Gen makes only classical queries

This is the case that already was analyzed in [18]; there, they did not go through BFA, but rather proved it directly. We demonstrate the versatility of our conjecture and framework, by showing:

▶ **Lemma 5.** *Conjecture 1 holds for all $f$ that can be expressed as the acceptance probability of a (randomized) classical decision tree.*

▶ **Corollary 6** (Recasting [18]). *Perfect complete QPKE does not exist in the QROM, if* Gen *only makes classical queries.*

## Evidence 2: The Combinatorial Nullstellensatz

The following relaxed version of Conjecture 1 directly follows from the Combinatorial Nullstellensatz of Alon (Lemma 22).

▶ **Lemma 7.** *For any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, the following holds. There exists a distribution $\mathcal{D}$ on the set of partial assignments $\rho$ with $|\rho| \le d$ such that: for any $x$,*

$$\Pr_{\rho \sim \mathcal{D}}[f(x^\rho) \neq 0] \ge 1/2^d.$$

Somewhat surprisingly, we find that if the probability of $f(x^\rho)$ being non-zero in Lemma 7 can be improved from $1/2^d$ to $1/2^{d^c}$ for an arbitrary $c < 1$, then Conjecture 1 would be affirmed. Formally, the following conjecture is equivalent to Conjecture 1:

▶ **Conjecture 8.** *There is a universal constant $c < 1$, such that for any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, the following holds. There exists a distribution $\mathcal{D}$ on the set of partial assignments $\rho$ with $|\rho| \le \mathrm{poly}(d)$ such that: for any $x$,*

$$\Pr_{\rho \sim \mathcal{D}}[f(x^\rho) \neq 0] \ge 1/2^{d^c}.$$

▶ **Lemma 9.** *Conjecture 8 holds if and only if Conjecture 1 holds.*

## Special Case 3: Gen makes $O(\log n)$ quantum queries

With Lemma 7, by following a similar argument in [18], we can directly conclude that when the generation algorithm only makes $O(\log n)$ quantum queries, such QPKE does not exist in the QROM.

▶ **Lemma 10.** *Perfect complete QPKE does not exist in the QROM, if* Gen *only makes $O(\log n)$ quantum queries.*

Thus, we make a step forward by improving the result (ii) in [5] and that in [18].

**Table 1** Provable separations in [5, 18] and this work. Q and C denote a polynomial number of quantum and classical queries, respectively.

| | [5] (i) | [5] (ii) | [18] | This work |
|---|---|---|---|---|
| Gen | Q | C | C | log Q |
| Enc | C | Q | Q | Q |
| Dec | Q | C | Q | Q |

## Special Case 4: Gen has “uniform” outputs

Here, by **“uniform”**, we mean that the output of $\mathsf{Gen}^H$ has support of the same size for every possible oracle $H$, and the output distribution is uniform over the support. In other words, $|\mathsf{supp}(\mathsf{Gen}^H)| = |\mathsf{supp}(\mathsf{Gen}^{H'})|$ for every pair of $H, H'$, and $\mathsf{Gen}^H$ is uniform over samples with a non-zero probability. To resolve this case, we prove the following special case of Conjecture 1:

▶ **Lemma 11.** *Conjecture 1 holds for all $f$ that takes at most $\mathrm{poly}(d)$ values.*

▶ **Corollary 12.** *Perfect complete QPKE does not exist in the QROM, if* Gen *is “uniform”.*

### Evidence 5: Other equivalent conjectures

By the min-max principle, Conjecture 1 has the following equivalent form.

▶ **Conjecture 13.** *For any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, and any distribution $X$ on $\{-1,1\}^n$, there exists a partial assignment with $|\rho| \leq \text{poly}(d)$ such that $\Pr_{x\sim X}[f(x^\rho) \neq 0] \geq 1/\text{poly}(d)$.*

▶ **Lemma 14.** *Conjecture 13 holds if and only if Conjecture 1 holds. Thus, if Conjecture 13 is true, then perfect complete QPKE does not exist in the QROM.*

In the main body, we provide additional evidence for Conjecture 13 (and consequently for Conjecture 1). Conjecture 13 has the advantage of requiring only a single partial assignment that works for the given input distribution, rather than a distribution of partial assignments. This simplifies the reasoning considerably, and much of our evidence supports Conjecture 13.

## 1.2 Techniques

Here, we give a brief overview of why Conjecture 1 implies a full impossibility result of QPKE in the QROM and the main differences between [18] and this work.

We will be mostly focusing on key agreement protocols in the QROM. In the QROM, all parties can quantumly query a random oracle. A key agreement protocol is an interactive protocol between two query-efficient quantum algorithms Alice and Bob, exchanging classical messages. Their goal is to agree on a key, whereas any query-efficient adversary can not learn the key, even observing all the communication on this channel. It is easy to see that QPKE implies a two-round key agreement protocol:

- Alice sends a single message $m_0$ to Bob; in this case, $m_0$ will be the public key.
- Bob sends a message $m_1$ to Alice; in this case, $m_1$ will be the encryption of a random key (which will be the agreed key in the two-round key agreement protocol).
- Finally, they both output their own keys $k_A, k_B$; in this case, Alice decrypts the ciphertext using her secret key, whereas Bob simply outputs the random key.

Thus, to rule out QPKE in the QROM, it is sufficient to rule out the two-round key agreement in the QROM.

Li, Li, Li and Liu [18] proposed the following framework that rules out QPKE with classical-query key generation in the QROM. Let $\mathsf{View}_A, \mathsf{View}_B$ denote the view of Alice and Bob, right after $m_1$ is received by Alice. They showed that, based on a novel approach called approximate quantum Markov chain, a quantum-query Eve can reconstruct Alice's view $\mathsf{View}'_A$ such that

$$\mathsf{View}'_A\mathsf{View}_B \approx \mathsf{View}_A\mathsf{View}_B.$$

More precisely, Eve can reconstruct Alice's register such that the margin of the fake Alice and the real Bob is arbitrarily close to the real Alice and the real Bob. However, this is not sufficient; since to compute the key $k_A$, Alice potentially needs to perform computation depending on its current state, the random oracle and all the messages $m_0, m_1$. As the marginal is defined by tracing out the random oracle, we do not know what oracle to work with. For example, a direct attempt is to run fake Alice under the real random oracle $H$; but it is possible that under real $H$, conditioned on the messages being $m_0, m_1$, the view of Alice and Bob will never equal to $\mathsf{View}'_A\mathsf{View}_B$, making the attempt meaningless.

Li et al. solved this by adapting Bob; while perfectly preserving the functionality of the key agreement protocol, the new Bob has one additional nice property (let us call it "stability"):

- if under a random oracle $H$, input message $m_0$, Bob outputs certain $m_1$ with non-zero probability
- then for every $H'$ that has a small Hamming distance to $H^2$, Bob on $H'$ and $m_0$ still has non-zero probability to output $m_1$.

With this property, they show that the two-round key agreement can not exist, when $m_0$ together with $\mathsf{View}_A$ can be computed by a classical decision tree (or $\mathsf{Gen}$ only makes classical queries). For any fake Alice view $\mathsf{View}'_A$, it depends on at most $d$ locations of a random oracle; here $d$ is the number of queries made by the classical decision tree. Thus, for any oracle $H_0$, as long as it is consistent on these $d$ locations, Alice on $H_0$ can output $\mathsf{View}'_A$ and $m_0$ with non-zero probability. They construct an oracle $H'$ such that: (i) it is almost $H$, except (ii) on those $d$ locations, $H'$ is reprogrammed to be consistent with these $d$ locations.

It is clear that Alice on $H'$ still outputs $m_0, \mathsf{View}'_A$ with a non-zero probability (consistency of these $d$ locations). Moreover, the "stability" ensures that the adapted Bob on $H'$ and $m_0$ outputs $\mathsf{View}_B$ also with a non-zero probability. Thus, there is a non-zero probability, under the oracle $H'$, Alice and Bob will end up with $\mathsf{View}'_A, \mathsf{View}_B$ and transcripts $(m_0, m_1)$. Since we assume the QPKE is perfect, after the whole execution, Alice and Bob should agree on a key. As for Eve, it simply runs Alice with $\mathsf{View}'_A, m_1$ on $H'$ and will recover the key.

The above reasoning in [18] does not work even if Alice only makes a single quantum query, and thus they can only rule out QPKE when $\mathsf{Gen}$ makes classical queries. Since a single quantum query can "store" information about exponentially many inputs, it seems that fixing a smaller number of input-output behaviors does not fix Alice's behavior.

We realize that, we do not need to keep Alice's behavior the same; we are only required to keep these probabilities non-zero. More precisely, our solution is to directly look at the polynomial $f$ describing the probability that a quantum Alice on some oracle, outputs $(m_0, \mathsf{View}'_A)$,

$$f(x) = \Pr[(m_0, \mathsf{View}'_A) \leftarrow A^x],$$

where we treat $x$ as the random oracle fed into $A$. As the random oracle will be the input to a polynomial, to be consistent with the literature of BFA, we will denote a random oracle by $x$. Since $A$ only makes $d$ quantum queries, such a polynomial $f$ has degree at most $2d$. The real random oracle is some $x$ of which we have no knowledge about; the goal is to find another $x'$ such that

- $f(x') > 0$, meaning Alice on the random oracle $x'$ outputs $\mathsf{View}'_A, m_0$ with a non-zero probability, just like the classical case; and,
- $x'$ and $x$ have a small Hamming distance, so that Bob still has non-zero probability to output $\mathsf{View}_B, m_1$.

As we do not know $x$, the only way we can obtain $x'$ is by locally modifying some locations of $x$. For example, reprogramming some locations of the real oracle, and running the rest of Alice under the reprogrammed oracle – this is where the partial assignment takes place. If we can find a partial assignment $\rho$ with $|\rho|$ is small, such that for every $x$, $f(x^\rho) > 0$, then the problem is resolved! This is the high-level intuition why our conjecture (Conjecture 1) implies a full impossibility result of QPKE in the QROM.

[2] We ignore a detail here: $H'$ should be consistent with $H$ on a small set of inputs, whereas on other inputs, their Hamming distance is small. This does not change the reasoning, thus we make the simplification in the introduction.

Towards Conjecture 1, we remark that the Combinatorial Nullstellensatz [2] directly implies the existence of such $\rho$. More specifically, the Combinatorial Nullstellensatz claims that: given any maximal monomial $x_S$ of $f$, for any $x \in \{-1,1\}^n$, there exists a $\rho$ with $\mathsf{supp}(\rho) = S$ such that $f(x^\rho) \neq 0$. Therefore, Conjecture 1 holds if we swap the order of quantifiers of $\rho$ and $x$. Consequently, by letting $\mathcal{D}$ be the uniform distribution on the set of partial assignments $\rho$ with $\mathsf{supp}(\rho) = S$, we have $\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/2^d$ for any $x$, as claimed by Lemma 7. In particular, when $\mathsf{Gen}$ makes only $O(\log n)$ queries and then $\deg(f) = O(\log n)$, we have $\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/\mathrm{poly}(n)$, which implies Lemma 10.

## 1.3 Related works

In Conjecture 1, we characterize $d$-query quantum algorithms as degree-$2d$ polynomials [8]. This characterization was further strengthened by Aaronson and Ambainis [1] in terms of bounded degree-$2d$ block-multilinear polynomials. Later, Arunachalam, Briet, and Palazuelos [4] provided an exact characterization of quantum query algorithms in terms of the so-called completely bounded norm.

# 2 Preliminaries

## 2.1 Basic notions in Boolean function analysis

Every function $f : \{-1,1\}^n \to \mathbb{R}$ on the hypercube can be uniquely expressed as a multilinear polynomial

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

where $x_S := \Pi_{i\in S} x_i$. Indeed, this expression is called the *Fourier expansion* of $f$, where $a_S = 2^{-n}\sum_x f(x)\cdot x_S$. The *degree* of $f$, denoted $\deg(f)$, is defined as the degree of its multilinear polynomial expression, i.e., $\max\{|S| \mid a_S \neq 0\}$. A monomial $x_S$ is called *maximal* if it has degree $\deg(f)$, i.e., $|S| = \deg(f)$ and $a_S \neq 0$. The *rank* of $f$, denoted $\mathrm{rank}(f)$, is the maximum number of disjoint maximal monomials.

A *partial assignment* is a function $\rho : [n] \to \{-1,1,\star\}$. We define the support of $\rho$ as $\mathsf{supp}(\rho) := \{i|\rho_i \neq \star\}$, and the size as $|\rho| := |\mathsf{supp}(\rho)|$. For $x \in \{-1,1\}^n$, we define the modification of $x$ with $\rho$, denoted $x^\rho$, as the string $x' \in \{-1,1\}^n$ where $x'_i = \rho_i$ for $i \in \mathsf{supp}(\rho)$ and $x'_i = x_i$ for any other $i$.

We use $\mathsf{Var}(f) := \mathbb{E}_x[f(x)^2] - \left(\mathbb{E}_x f(x)\right)^2$ to denote $f$'s variance, and define $\|f\|_\infty := \max_x |f(x)|$. A real polynomial $p$ $\epsilon$-approximates $f$ if $|f(x) - p(x)| \leq \epsilon$ for every $x \in \{-1,1\}^n$. The *approximate degree* of $f$, denoted by $\widetilde{\deg}(f)$, is defined to be the minimum degree needed to 1/3-approximate $f$.

The following lemma will be used.

▶ **Lemma 15** ([8]). *Suppose a quantum algorithm makes $d$ queries to a Boolean string $x \in \{-1,1\}^n$, 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*

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

## 2.2 Quantum public key encryption

In this paper, we will consider constructions in the quantum random oracle model. Given security parameter $\lambda$, the oracle $H$ is chosen from the uniformly random distribution over the family of functions $\mathcal{H}_\lambda\colon [2^{n_\lambda}] \to \{0,1\}$, where $n_\lambda$ is a polynomial of $\lambda$. The quantum circuit has access to the oracle unitary $U_H$ that maps $|i,y\rangle$ to $|i, y \oplus H(i)\rangle$. We can also view the oracle unitary in the phase basis $U'_H |i,y\rangle \to (-1)^{H(i)y}|i,y\rangle$.

We will consider quantum public key encryption with classical public key $\mathsf{pk}$ and ciphertext $\mathsf{ct}$ in this paper. Particularly, we will consider the quantum public key encryption (QPKE) scheme in the quantum random oracle model (QROM) defined as follows:

▶ **Definition 16** (Quantum public key encryption in QROM)**.** *A public key encryption scheme, relative to a random oracle $H \leftarrow \mathcal{H}_\lambda$ consists of the following three bounded quantum query algorithms:*

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

The algorithms should 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 QPT adversary $\mathcal{E}^H$, for every two plaintexts $m_0 \neq m_1$ chosen by $\mathcal{E}^H(\mathsf{pk})$ we have

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

where we call $\epsilon(\lambda)$ as the advantage of the adversary. The scheme is IND-CPA secure if $\epsilon(\lambda)$ is negligible for any adversary $\mathcal{E}^H$.

# 3 Consequences of our BFA conjecture

In this section, we will show the consequence in cryptography of Conjecture 1. If Conjecture 1 holds, we can obtain a black box separation between post-quantum secure one-way functions and quantum public key encryption schemes.

▶ **Theorem 17** (Restate of Theorem 3)**.** *Assuming Conjecture 1 is true, for any quantum public key encryption scheme in the quantum random oracle model, if* $\mathsf{Gen}^H$ *and* $\mathsf{Enc}^H$ *make $d$ queries to $H$ in total, and* $\mathsf{Dec}^H$ *makes $D$ queries, there exists an adversary Eve $\mathcal{E}^H$ which can break the IND-CPA security with advantage* $1/\mathrm{poly}(d)$ *by making* $O(\mathrm{poly}(d) + D)$ *queries.*

It is known that we can use a QPKE scheme for a two-message key agreement protocol, by setting the first message $m_0 = \mathsf{pk}$, and the second message $m_1 = \mathsf{Enc}(\mathsf{pk}, k)$, where $k$ is the key the two parties agree on. We can assume the key $k$ is chosen uniformly random from $\{0,1\}^n$. In the following section, we will refer to the two parties as $\mathcal{A}$ and $\mathcal{B}$, and call the algorithm of $\mathcal{A}$ before sending $m_0$ as $\mathcal{A}_0$, and the algorithm after receiving message $m_1$ as $\mathcal{A}_1$. If we can learn the key $k$ with probability $p$, we can also break the semantic security of the QPKE scheme with advantage $p$, which implies that we also break the IND-CPA security of the protocol. In the language of key agreement, we are considering the case where $\mathcal{A}_0, \mathcal{A}_1$, and $\mathcal{B}$ can both make quantum queries while only sending classical messages, and we will show how to learn the key $k$ with probability $1/\mathrm{poly}(d)$.

Let us recall the results from [18]. In their paper, they are considering the case where $\mathcal{A}_0$ can only make classical queries. The key part of their analysis is to show how to construct a register $\mathsf{A}'$, such that there exists some oracle $H'$, the content $\mathsf{View}'_A$ in $\mathsf{A}'$ is consistent with $H'$ and transcript $\pi = (m_0, m_1)$. We can think the algorithm $\mathcal{B}$ proceeds as follows: it first receives the message $m_0$ from $\mathcal{A}$, and after making some queries to the oracle $H$, it makes a measurement in the computational basis, and generates the key $k_B$ and the second message $m_1$. The state under consideration in [18] $\sigma_{AB} = \mathbb{E}_H[\sigma^H_{AB}]$ is the state before $\mathcal{B}$ makes the final measurement and sends the second message $m_1$, where $\sigma^H_{AB}$ is the state under oracle $H$, and the expectation is over the posterior distribution of $H$ given the first message $m_0$.

To construct the register $\mathsf{A}'$, they used tools from quantum information theory. They used the following theorem from [14].

▶ **Theorem 18.** *For any state $\rho_{\mathsf{AEB}}$ over systems $\mathsf{AEB}$, there exists a channel $\mathcal{T} : \mathsf{E} \to \mathsf{E} \otimes \mathsf{B}'$ such that the trace distance between the reconstructed state $\sigma_{\mathsf{A}'\mathsf{E}'\mathsf{B}'} = \mathcal{T}(\rho_{\mathsf{AE}})$ and the original state $\rho_{\mathsf{AEB}}$ is at most*

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

The theorem says that if the conditional mutual information $I(\mathsf{A} : \mathsf{B}|\mathsf{E})$ is small, we can generate a consistent copy of $\mathsf{A}'$ by a channel only acting on $\mathsf{E}$.

To decrease the conditional mutual information $I(\mathsf{A} : \mathsf{B}|\mathsf{E})$, we have the following lemma from [18].

▶ **Definition 19** (Permutation Invariance)**.** *Let $\mathsf{A}_1, \mathsf{A}_2, \mathsf{A}_3, \dots, \mathsf{A}_t, \mathsf{B}$ be $(t+1)$-partite quantum system. Given the joint state $\rho_{\mathsf{BA}_1\mathsf{A}_2\cdots\mathsf{A}_t}$, we say $A_1, \dots, A_t$ are permutation invariant, if for any permutation $\pi$ on $[t]$, we have*

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

▶ **Lemma 20.** *Let $\mathsf{A}_1, \mathsf{A}_2, \mathsf{A}_3, \dots, \mathsf{A}_t, \mathsf{B}, \mathsf{E}$ be $(t+2)$-partite quantum system. Suppose the state of the composite system $\rho_{\mathsf{BEA}_1\mathsf{A}_2\cdots\mathsf{A}_t}$ is fully separable. If $\mathsf{A}_1, \mathsf{A}_2, \mathsf{A}_3, \dots, \mathsf{A}_t$ are permutation invariant, then there is a $0 \le i \le t-1$ such that*

$$I(\mathsf{A}_t : \mathsf{B} \mid \mathsf{E}, \mathsf{A}_1, \dots, \mathsf{A}_i)_\rho \le S(\mathsf{B})/t.$$

The lemma shows that if Eve parallel repeats multiple copies of $\mathcal{B}$, the mutual information will be decreased.

The following lemma from [9] will also be used.

▶ **Lemma 21.** *Consider a quantum algorithm $\mathcal{B}$ that makes $d$ queries to an oracle $H$. Denote the quantum state immediately after $t$ queries to the oracle as*

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

*where $w$ is the content of the workspace register. Denote the query weight $q_x$ of input $x$ as*

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

*For any oracle $\tilde{H}$, denote $|\phi_d\rangle$ as the final state before measurement obtained by running $\mathcal{B}$ with oracle $\tilde{H}$, we have that*

$$\||\psi_d\rangle - |\phi_d\rangle\| \le 2\sqrt{d}\sqrt{\sum_{x:\, \tilde{H}(x) \neq H(x)} q_x}.$$

In [18], they defined the heavy query of Bob as $\{x\colon q_x \geq \epsilon^2/d^2\}$. We summarize their adversary Eve's algorithm $\mathcal{E}$ as follows:

1. Eve parallelly simulates Bob's algorithm $\mathcal{B}^H(m_0)$ for multiple times. In this process, Eve records heavy queries of $\mathcal{B}$ under $H$ classically and maintains an input-output pair register $R_E = \{(i_E, H(i_E))\}$.
2. By Lemma 20, if Eve repeat $\mathcal{B}^H(m_0)$ for $\text{poly}(d, 1/\epsilon)$ times, the conditional mutual information $I(\mathsf{A} : \mathsf{B} \mid \mathsf{E})$ of state $\sigma_{ABE}$ can be reduced to smaller than $O(\epsilon^2)$. By Theorem 18,the channel $\mathcal{T}\colon \mathsf{E} \to \mathsf{A'E}$ provides a state $\mathcal{T}(\sigma_{BE}) = \sigma_{A'BE}$ is $\epsilon$ close to the state $\sigma_{ABE}$. In the following section, we will choose $\epsilon = 1/\text{poly}(d)$.
3. Eve measures the secret key register as well as the query input-output history register of $\mathsf{A}'$. Note that in their case, we can assume $\mathcal{A}_0^H$ records its classical queries. The measurement will give us some secret key $\mathsf{sk}'$ and input-output pairs $R_{A'} = \{(i_{A'}, H'(i_{A'}))\}$, since $\sigma_{A'BE} \approx_\epsilon \sigma_{ABE}$, with high probability, the measurement result $\mathsf{sk}'$ and $R_{A'}$ are consistent with message $m_0$ and $R_E$. Specifically, every query that is both in $R_{A'}$ and $R_E$ should be consistent, meaning $R_{A'}$ is consistent with $\mathcal{B}$'s heavy queries under oracle $H$.
4. Finally, by the observation beyond, if the oracle $H$ is reprogrammed to $H'$ using the input-output pairs in $R_{A'}$, from $\mathcal{B}$'s view, only $\text{poly}(d)$ light query points under $H$ are modified. Using Lemma 21, it can be shown that w.h.p., $\mathcal{B}^{H'}(m_0)$ will output $m_1$ with nonzero probability, implying that $\pi = (m_0, m_1)$ is a valid transcript under oracle $H'$. Thus by perfect completeness, if we simulate $\mathcal{A}_1$ given input $(\mathsf{sk}', m_1)$ over oracle $H'$, it can still agree with $\mathcal{B}$.

We refer interested readers to their paper for the proof details.

Our observation is that in step 3, we do not need to generate the input-output pair $R_{A'}$ by measuring $\mathsf{A}'$. Given the algorithm of $\mathcal{A}_0$, if we can find some $\text{poly}(d)$-sized assignment $R_{A'}$ that is consistent with $R_E$ and guarantees the output probability of $(\mathsf{sk}', m_0)$ is still non-zero, by similar arguments, we can see that $\mathcal{B}$ is consistent with the reprogrammed oracle $H'$, and will output $m_1$ with non-zero probability. From Lemma 15, if we define the algorithm outputs $(\mathsf{sk}', m_0)$ as acceptance, we can characterize the probability with a $2d$-degree polynomial $f(x)$, viewing the truth table of random oracle $H$ as a $N_\lambda = 2^{n_\lambda}$ length vector $x$.

Now we show how to leverage Conjecture 1 to prove our main theorem. We use $x$ for the truth table of original oracle $H$, setting $x_i = (-1)^{H(i)}$, and $\rho_E$ for the restriction given by $R_E$, and apply Conjecture 1 to polynomial $g(x) = f(x^{\rho_E})$. In an operational meaning, considering $g(x)$ means that we first fix the input-output pairs given by $R_E$ when selecting the random oracle. We now argue $g(x)$ is a non-zero polynomial with high probability. We can obtain the joint view $\mathsf{View}_{A'E} = (\mathsf{sk}', m_0, R_E)$ of $\mathsf{A'E}$ by performing a computational basis measurement on corresponding registers. Since $\sigma_{A'E}$ is $\epsilon$ close to $\sigma_{AE}$, we can see that w.p. $1-\epsilon$, $\mathsf{View}_{A'E} = (\mathsf{sk}', m_0, R_E)$ is a possible valid view obtained by measuring $\sigma_{AE}$.

By Conjecture 1, there exists a distribution of polynomial-sized restrictions $\mathcal{D}$ such that for any $x$, $\Pr_{\rho\sim\mathcal{D}}[g(x^\rho) \neq 0] \geq 1/\text{poly}(d)$, with $|\rho| \leq \text{poly}(d)$. In our case, this implies that $\Pr_{\rho\sim\mathcal{D}}[f(x^{\rho\cup\rho_E}) \neq 0] \geq 1/\text{poly}(d)$. If we select a polynomial-sized restriction $\rho$, there is $1/\text{poly}(d)$ probability such that $(\mathsf{sk}', m_0)$ is still a valid view of $\mathcal{A}_0$ under the new oracle given by $x^{\rho\cup\rho_E}$. The other arguments follow from the original proof.

## 4 Toward our BFA conjecture: evidences and special cases

In this section, we explore Conjecture 1 by providing some evidence and proving some special cases.

## 4.1 Main evidence: Combinatorial Nullstellensatz

The main evidence for Conjecture 1 comes from the celebrated Combinatorial Nullstellensatz of Alon [2], which implies Conjecture 1, except with $\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/2^d$ instead of $\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/\mathrm{poly}(d)$. Let us state the special case of the Combinatorial Nullstellensatz for polynomials on the hypercube.

▶ **Lemma 22** ([2]). *Let $f : \{-1,1\}^n \to \mathbb{R}$ be any nonconstant function with degree $d$, and $x_S$ be any maximal monomial of $f$. For any $x \in \{-1,1\}^n$, there exists a $\rho$ with $\mathsf{supp}(\rho) = S$ such that $f(x^\rho) \neq 0$.*

*Consequently, by letting $\mathcal{D}$ be the uniform distribution on the set of partial assignments $\rho$ with $\mathsf{supp}(\rho) = S$, we have $\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/2^d$ for any $x$.*

We remark that Lemma 22 was also proven by Midrijanis [21]. Even though Lemma 22 has an exponential rather than polynomial dependence on $1/d$, we observe that it already has a nontrivial implication on the separation between QPKE and OWFs. Precisely, it implies that

▶ **Theorem 23** (Restate of Lemma 10). *For any quantum public key encryption scheme in the quantum random oracle model, if $\mathsf{Enc}^H$ makes $d$ queries to $H$, $\mathsf{Gen}^H$ makes $O(\log d)$ queries, and $\mathsf{Dec}^H$ makes $D$ queries, there exists an adversary Eve $\mathcal{E}^H$ which can break the IND-CPA security with advantage $1/\mathrm{poly}(d)$ by making $O(\mathrm{poly}(d) + D)$ queries.*

The proof of Theorem 23 is the same as Theorem 17 by replacing Conjecture 1 with Lemma 22.

## 4.2 Proof of Lemma 9

In this subsection, we present the proof of Lemma 9, which claims that Conjecture 1 is equivalent to a seemingly much weaker conjecture, namely Conjecture 8.

▶ **Conjecture 8.** *There is a universal constant $c < 1$, such that for any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, the following holds. There exists a distribution $\mathcal{D}$ on the set of partial assignments $\rho$ with $|\rho| \leq \mathrm{poly}(d)$ such that: for any $x$,*

$$\Pr_{\rho\sim\mathcal{D}}[f(x^\rho) \neq 0] \geq 1/2^{d^c}.$$

▶ **Lemma 9.** *Conjecture 8 holds if and only if Conjecture 1 holds.*

**Proof.** Conjecture 1 implying Conjecture 8 is trivial. Conversely, given any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, define

$$\tilde{f}(x_1, \ldots, x_t) := \prod_{i=1}^{t} f(x_i)$$

where $x_i \in \{-1,1\}^n$. Note that $\deg(\tilde{f}) = td$. Assuming Conjecture 8 holds for $\tilde{f}$, there exists a distribution $\tilde{\mathcal{D}}_0$ of partial assignment $\tilde{\rho} = (\rho_1, \ldots, \rho_t) \in (\{-1,1,\star\}^n)^t$ with $|\tilde{\rho}| \leq (dt)^{c_1}$ such that

$$\min_{\tilde{x}} \Pr_{\tilde{\rho}\sim\tilde{\mathcal{D}}_0}\left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] \geq \frac{1}{2^{(dt)^{c_2}}}$$

where $c_1, c_2$ are universal constants satisfying $c_1 \geq 0$ and $0 \leq c_2 < 1$. Observe that

$$\max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho}\sim\tilde{\mathcal{D}}}\left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] \geq \min_{\tilde{x}} \Pr_{\tilde{\rho}\sim\tilde{\mathcal{D}}_0}\left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] \geq \frac{1}{2^{(dt)^{c_2}}}$$

where $\tilde{\mathcal{D}}$ is a distribution of partial assignment $\tilde{\rho} = (\rho_1, \ldots, \rho_t) \in (\{-1, 1, \star\}^n)^t$ with $|\rho_i| \leq (dt)^{c_1}$ for all $i$. Then by Lemma 24, we have

$$\left(\max_{\mathcal{D}} \min_{x} \Pr_{\rho \sim \mathcal{D}} [f(x^\rho) \neq 0]\right)^t = \max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho} \sim \tilde{\mathcal{D}}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] \geq \frac{1}{2^{(dt)^{c_2}}}$$

where $\mathcal{D}$ is a distribution of partial assignment $\rho \in \{-1, 1, \star\}^n$ with $|\rho| \leq (dt)^{c_1}$. Thus there exists a distribution $\mathcal{D}^*$ such that

$$\left(\min_{x} \Pr_{\rho \sim \mathcal{D}^*} [f(x^\rho) \neq 0]\right)^t \geq \frac{1}{2^{(dt)^{c_2}}}.$$

By setting $t = d^{\frac{c_2}{1-c_2}} \log^{\frac{1}{c_2-1}} d$, we have

$$\min_{x} \Pr_{\rho \sim \mathcal{D}^*} [f(x^\rho) \neq 0] \geq 2^{-(dt)^{c_2}/t} = 2^{-\log d} = 1/d. \blacktriangleleft$$

▶ **Lemma 24.** *Given positive integers $t, K$, and function $f : \{-1, 1\}^n \to \mathbb{R}$, let $\tilde{f}(x_1, \ldots, x_t) := \prod_{i=1}^t f(x_i)$. Then*

$$\max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho} \sim \tilde{\mathcal{D}}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] = \left(\max_{\mathcal{D}} \min_{x} \Pr_{\rho \sim \mathcal{D}} [f(x^\rho) \neq 0]\right)^t$$

*where $\mathcal{D}$ is a distribution of partial assignments $\rho \in \{-1, 1, \star\}^n$ with $|\rho| \leq K$, $\tilde{\mathcal{D}}$ is a distribution of partial assignments $\tilde{\rho} = (\rho_1, \ldots, \rho_t) \in (\{-1, 1, \star\}^n)^t$ with $|\rho_i| \leq K$ for all $i$.*

**Proof.** By the min-max principle, we have

$$\max_{\mathcal{D}} \min_{x} \Pr_{\rho \sim \mathcal{D}} [f(x^\rho) \neq 0] = \min_{X} \max_{\rho} \Pr_{x \sim X} [f(x^\rho) \neq 0] := p.$$

Let $\mathcal{D}^*$ be the optimal distribution that reaches the maximum, $X^*$ the optimal distribution that reaches the minimum. Let $\tilde{p} := \max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho} \sim \tilde{\mathcal{D}}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right]$.

**Claim 1: $\tilde{p} \leq p^t$.** By min-max principle,

$$\begin{aligned}\max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho} \sim \tilde{\mathcal{D}}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] &= \min_{\tilde{X}} \max_{\tilde{\rho}} \Pr_{\tilde{x} \sim \tilde{X}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] \\ &= \min_{\tilde{X}} \max_{\rho_1, \ldots, \rho_t} \Pr_{(x_1, \ldots, x_t) \sim \tilde{X}} \left[f(x_1^{\rho_1}) \neq 0, \ldots, f(x_t^{\rho_t}) \neq 0\right].\end{aligned}$$

By setting $\tilde{X}$ to be $X^* \times X^* \times \cdots \times X^*$, we have

$$\begin{aligned}\max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho} \sim \tilde{\mathcal{D}}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] &\leq \max_{\rho_1, \ldots, \rho_t} \Pr_{x_i \sim X^*} \left[f(x_1^{\rho_1}) \neq 0, \ldots, f(x_t^{\rho_t}) \neq 0\right] \\ &= \max_{\rho_1, \ldots, \rho_t} \prod_{i} \Pr_{x_i \sim X^*} \left[f(x_i^{\rho_i}) \neq 0\right] = \prod_{t} \max_{\rho_i} \Pr_{x_i \sim X^*} \left[f(x_i^{\rho_i}) \neq 0\right] = p^t.\end{aligned}$$

**Claim 2: $\tilde{p} \geq p^t$.** By setting $\tilde{\mathcal{D}}$ to be $\mathcal{D}^* \times \mathcal{D}^* \times \cdots \times \mathcal{D}^*$, we have

$$\begin{aligned}\max_{\tilde{\mathcal{D}}} \min_{\tilde{x}} \Pr_{\tilde{\rho} \sim \tilde{\mathcal{D}}} \left[\tilde{f}(\tilde{x}^{\tilde{\rho}}) \neq 0\right] &\geq \min_{x_1, \ldots, x_t} \Pr_{\rho_i \sim \mathcal{D}^*} \left[f(x_1^{\rho_1}) \neq 0, \ldots, f(x_t^{\rho_t}) \neq 0\right] \\ &= \min_{x_1, \ldots, x_t} \prod_{i} \Pr_{\rho_i \sim \mathcal{D}^*} \left[f(x_i^{\rho_i}) \neq 0\right] = \prod_{i} \min_{x_i} \Pr_{\rho_i \sim \mathcal{D}^*} \left[f(x_i^{\rho_i}) \neq 0\right] = p^t.\end{aligned}$$

Thus $\tilde{p} = p^t$ and $\tilde{p}$ can be reached by i.i.d. distribution $\mathcal{D}^* \times \cdots \times \mathcal{D}^*$. ◀

▶ Remark 25. By the same technique, the non-zero probability in Conjecture 1 can be further improved from $1/\text{poly}(d)$ to $1 - 1/\text{poly}(d)$, while still keeping $|\rho| \leq \text{poly}(d)$.

## 4.3 Special cases of boolean functions

We also affirm Conjecture 1 for some special classes of functions.

▶ **Lemma 26.** *Conjecture 1 holds when $f$ satisfies one of the following conditions:*

**(a)** *$f$ can be expressed as the acceptance probability of a classical randomized decision tree with depth $\leq \mathrm{poly}(d)$ (Lemma 5).*

**(b)** *$f$ is Boolean-valued, more generally when $f$ takes at most $\mathrm{poly}(d)$ values (Lemma 11).*

**(c)** *$f$ is a symmetric function.*

**Proof.**

**Part (a).** Recall that a randomized decision tree can be viewed as a distribution $\mu$ over deterministic decision trees, such that the tree is evaluated by (i) first sampling a deterministic decision tree from $\mu$, and (ii) then evaluating this deterministic decision tree. Since $f(x)$ is the probability that the output of the evaluation is "yes", and $f$ is non-zero, there exists a deterministic decision tree $T$ with $\mu(T) > 0$ such that at least one of its leaves is labeled with "yes". We can set $\rho$ to be the partial assignment corresponding to the path from the root to the leaf. Then $f(x^\rho) \neq 0$ for any $x$.

**Part (b).** Without loss of generality, assume $\|f\|_\infty = 1$. Since $f$ takes at most $\mathrm{poly}(d)$ values, there exists a gap $b - a \geq 1/\mathrm{poly}(d)$ such that $f(x) \in [-1, a] \cup [b, 1]$ for all $x$. Define $f'(x) := \left(1 + \left|\frac{a+b}{2}\right|\right)^{-1} \left(f(x) - \frac{a+b}{2}\right)$. Then $f'(x) \in [-1, -1/\mathrm{poly}(d)] \cup [1/\mathrm{poly}(d), 1]$ for all $x$. Define a boolean function

$$\tilde{f}(x) := \begin{cases} -1 & \text{if } f'(x) < 0 \\ 1 & \text{if } f'(x) > 0 \end{cases}.$$

Then $\tilde{f}$ is $(1 - 1/\mathrm{poly}(d))$-approximated by $f'$. $f'$ can be amplified to a degree-$\mathrm{poly}(d)$ polynomial that 1/3-approximates $\tilde{f}$. Thus we have $\widetilde{\deg}(\tilde{f}) \leq \mathrm{poly}(d)$, which further implies $\tilde{f}$ can be computed by a decision tree with depth $\mathrm{poly}(d)$ [11].

(i) If $a < 0 < b$, then $f(x) \neq 0$ alway holds. (ii) If $a \geq 0$, by Part (a), there exists a partial assignment $|\rho| \leq \mathrm{poly}(d)$ such that $\tilde{f}(x) = 1$, which implies $f(x^\rho) \geq b > 0$ for any $x$. (iii) If $b \leq 0$, by Part (a), there exists a partial assignment $|\rho| \leq \mathrm{poly}(d)$ such that $\tilde{f}(x) = -1$, which implies $f(x^\rho) \leq a < 0$ for any $x$.

**Part (c).** Let $x_S = x_{i_1} x_{i_2} \cdots x_{i_d}$ be a maximal monomial of $f$, and $\rho_k$ be the partial assignment such that $\rho_k(i_1) = \cdots = \rho_k(i_k) = 1$ and $\rho_k(i_{k+1}) = \cdots = \rho_k(i_d) = -1$. By Lemma 22, for any $x$, there exists a $\rho$ with $\mathsf{supp}(\rho) = S$ such that $f(x^\rho) \neq 0$. Since $f$ is symmetric, we have $f(x^{\rho_j}) = f(x^\rho) \neq 0$ where $j$ is the number of 1's in $\rho$. Let $\mathcal{D}$ be the uniform distribution over $\{\rho_0, \rho_1, \ldots, \rho_d\}$. Then $\Pr_{\rho \sim \mathcal{D}}[f(x^\rho) \neq 0] \geq 1/(d+1)$ for any $x$. ◀

The above lemma implies the non-existence of perfect complete QPKE in QROM which the key generation takes special forms, e.g., the special cases 1 and 3 mentioned in the introduction.

▶ **Corollary 27.** *Perfect complete QPKE does not exist in the QROM, if one of the following holds:*

**(1)** $\mathsf{Gen}$ *only makes classical queries. (Corollary 6)*

**(2)** $\mathsf{Gen}$ *is uniform, i.e., it satisfies (i) $|\mathsf{supp}(\mathsf{Gen}^H)| = |\mathsf{supp}(\mathsf{Gen}^{H'})|$ for every pair of $H, H'$, and (ii) $\mathsf{Gen}^H$ is uniform over $\mathsf{supp}(\mathsf{Gen}^H)$ with a non-zero probability. (Corollary 12)*

**Proof.** Consider the polynomial $g(H)$ defined in the proof of Theorem 17, which is defined to be the probability of $\mathsf{Gen}^{H^{\rho_E}}$ outputting $(\mathsf{sk}', m_0)$.

**(1)** $g$ can be expressed as the acceptance probability of a classical randomized decision tree with depth $d$. Then the proposition follows from Part (a) of Lemma 26.

**(2)** $g$ takes only two values 0 and $1/|\mathsf{supp}(\mathsf{Gen}^H)|$. Then the proposition follows from Part (b) of Lemma 26. ◀

## 4.4 More evidences for Conjecture 13

Additionally, we provide evidence for the equivalent form Conjecture 13 (and consequently for Conjecture 1). The first piece of evidence is from the anti-concentration result for low-degree polynomials by Meka, Nguyen and Vu [20], which implies Conjecture 13 when $X$ is uniform distribution on $\{-1,1\}^n$, except that $|\rho|$ is allowed to be as large as $d^{8d+4}$ instead of $|\rho| \leq \text{poly}(d)$.

▶ **Theorem 28** ([20]). *Let $f : \{-1,1\}^n \to \mathbb{R}$ be a nonconstant polynomial with degree $d$, and $U$ denote the uniform distribution on $\{-1,1\}^n$. If $\text{rank}(f) \geq d^{8d+2}$, then $\Pr_{x\sim U}[f(x) \neq 0] = 1 - o(1)$.*

▶ **Lemma 29.** *For any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$, there exists a partial assignment with $|\rho| \leq d^{8d+4}$ such that $\Pr_{x\sim U}[f(x^\rho) \neq 0] \geq 1 - o(1)$.*

**Proof.** We construct a partial assignment $\rho$ by the following algorithm, which contains at most $d$ rounds and each round reduces $\deg(f)$ by at least 1. Denote the function at round $t$ by $f^t$. Initia,lly $\rho$ is empty. At round $t$,

1. If $\text{rank}(f^t) \geq d^{8d+2}$, the algorithm stops. By Theorem 28, we have
$$\Pr_{x\sim U}[f(x^\rho) \neq 0] = \Pr_{x\sim U}[f^t(x) \neq 0] = 1 - o(1).$$
2. If $\text{rank}(f^t) < d^{8d+2}$, let $T$ be the variables contained a maximum disjoint set of maximal monomials of $f^t$. The algorithm assigns all the variables in $T$ such that the function after the assignment is still nonconstant and adds them to $\rho$. This step increases $|\rho|$ by $|T| = \deg(f^t)\text{rank}(f^t) < d^{8d+3}$. For any maximal monomial $S$ of $f^t$, we have $S \cap T \neq \emptyset$ because otherwise $S \cup T$ is a disjoint set of maximal monomials larger than $T$. Thus this step reduces $\deg(f^t)$ by at least one.

Because there are at most $d$ rounds, we have $|\rho| \leq d^{8d+4}$. ◀

The second evidence is that Conjecture 13 holds for any function $f$ with large variance when $X$ is uniform.

▶ **Lemma 30.** *For any nonconstant function $f : \{-1,1\}^n \to \mathbb{R}$ with degree $d$ and $\mathsf{Var}(f) \geq \|f\|_\infty^2/\text{poly}(d)$, there exists a partial assignment with $|\rho| \leq \text{poly}(d)$ such that $\Pr_{x\sim U}[f(x^\rho) \neq 0] \geq 1/\text{poly}(d)$.*

**Proof.** $\mathsf{Var}(f) = \mathbb{E}_x[f^2(x)] - (\mathbb{E}_x f(x))^2 \leq \mathbb{E}_x[f^2(x)] \leq \Pr_x[f(x) \neq 0]\|f\|_\infty^2$. Thus $\Pr_x[f(x) \neq 0] \geq \frac{\mathsf{Var}(f)}{\|f\|_\infty^2} \geq 1/\text{poly}(d)$. ◀

## References

1 Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. *SIAM J. Comput.*, 47(3):982–1038, 2018. `doi:10.1137/15M1050902`.

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

3 Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *Annual International Cryptology Conference*, pages 208–236. Springer, 2022. doi:10.1007/978-3-031-15802-5_8.

4 Srinivasan Arunachalam, Jop Briët, and Carlos Palazuelos. Quantum query algorithms are completely bounded forms. *SIAM J. Comput.*, 48(3):903–925, 2019. doi:10.1137/18M117563X.

5 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. doi:10.1007/978-3-031-15979-4_6.

6 Khashayar Barooti, Alex B. Grilo, Loïs Huguenin-Dumittan, Giulio Malavolta, Or Sattath, Quoc-Huy Vu, and Michael Walter. Public-key encryption with quantum keys, 2023. doi:10.48550/arXiv.2306.07698.

7 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. doi:10.1007/978-3-030-84242-0_15.

8 Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. doi:10.1145/502090.502097.

9 Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. doi:10.1137/S0097539796300933.

10 Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, December 2014. doi:10.1016/j.tcs.2014.05.025.

11 Harry Buhrman and Ronald De Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288(1):21–43, 2002. doi:10.1016/S0304-3975(01)00144-X.

12 Andrea Coladangelo. Quantum trapdoor functions from classical one-way functions. *arXiv preprint arXiv:2302.12821*, 2023. URL: https://eprint.iacr.org/2023/282.

13 Whitfield Diffie and Martin E Hellman. New directions in cryptography. In *Democratizing Cryptography: The Work of Whitfield Diffie and Martin Hellman*, pages 365–390. 2022. doi:10.1145/3549993.3550007.

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

15 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. doi:10.1007/978-3-030-77886-6_18.

16 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. doi:10.1145/73007.73012.

17 Fuyuki Kitagawa, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum public-key encryption with tamper-resilient public keys from one-way functions. *arXiv preprint arXiv:2304.01800*, 2023. URL: https://eprint.iacr.org/2023/490.

18 Longcheng Li, Qian Li, Xingjian Li, and Qipeng Liu. How (not) to build qpke in minicrypt. In *Annual International Cryptology Conference*. Springer, 2024. doi:10.1007/978-3-031-68394-7_6.

19 Giulio Malavolta and Michael Walter. Robust quantum public-key encryption with applications to quantum key distribution. Cryptology ePrint Archive, Paper 2023/500, 2023. doi:10.1007/978-3-031-68394-7_5.

20 Raghu Meka, Oanh Nguyen, and Van Vu. Anti-concentration for polynomials of independent random variables. *Theory of Computing*, 12:11, 2016. doi:10.4086/toc.2016.v012a011.

21 Gatis Midrijanis. Exact quantum query complexity for total boolean functions, 2004. arXiv:quant-ph/0403168.

22 Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *arXiv preprint arXiv:2210.03394*, 2022. URL: https://eprint.iacr.org/2022/1336.

23 Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Annual International Cryptology Conference*, pages 269–295. Springer, 2022. doi:10.1007/978-3-031-15802-5_10.