From Worst-Case Hardness of NP to Quantum Cryptography via Quantum Indistinguishability Obfuscation

Tomoyuki Morimae[1], Yuki Shirakawa[1], and Takashi Yamakawa[2,3,1]

[1]Yukawa Institute for Theoretical Physics, Kyoto University, Kyoto, Japan
tomoyuki.morimae@yukawa.kyoto-u.ac.jp
yuki.shirakawa@yukawa.kyoto-u.ac.jp
[2]NTT Social Informatics Laboratories, Tokyo, Japan
takashi.yamakawa@ntt.com
[3]NTT Research Center for Theoretical Quantum Information, Atsugi, Japan

Abstract

Indistinguishability obfuscation (iO) has emerged as a powerful cryptographic primitive with many implications. While classical iO, combined with the infinitely-often worst-case hardness of NP, is known to imply one-way functions (OWFs) and a range of advanced cryptographic primitives, the cryptographic implications of quantum iO remain poorly understood. In this work, we initiate a study of the power of quantum iO. We define several natural variants of quantum iO, distinguished by whether the obfuscation algorithm, evaluation algorithm, and description of obfuscated program are classical or quantum. For each variant, we identify quantum cryptographic primitives that can be constructed under the assumption of quantum iO and the infinitely-often quantum worst-case hardness of NP (i.e., $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$). In particular, we construct pseudorandom unitaries, QCCC quantum public-key encryption and (QCCC) quantum symmetric-key encryption, and several primitives implied by them such as one-way state generators, (efficiently-verifiable) one-way puzzles, and EFI pairs, etc. While our main focus is on quantum iO, even in the classical setting, our techniques yield a new and arguably simpler construction of OWFs from classical (imperfect) iO and the infinitely-often worst-case hardness of NP.

Contents

1 Introduction

Obfuscation enables us to transform a program into an unintelligible, "scrambled" form while preserving its functionality. The cryptographic study of obfuscation was initiated by Barak et al. [BGI+12], who formalized the notion and introduced the concept of virtual black-box (VBB) security. While they showed that VBB security is unachievable in general, they proposed a weaker notion called indistinguishability obfuscation (iO), leaving open the possibility of its realization.

About a decade later,[1] Garg et al. [GGH+16] proposed the first candidate construction of iO. This breakthrough led to a cascade of results: it was soon discovered that iO, when combined with one-way functions (OWFs), enables the construction of a wide array of powerful cryptographic primitives (e.g., [HSW13, GGG+14, SW14, BZ14, KNY14, CLP15, KLW15, CHN+16]), some of which had no known constructions prior to iO. As a result, iO has come to be viewed as a "central hub" [SW14] in cryptography.

Given its remarkable utility, iO has attracted extensive research interest, both in terms of proposing new constructions (e.g., [CLT13, BGK+14, BR14, CLT15, GGH15, BMSZ16, GMM+16, BDGM20], and developing attacks (e.g., [CHL+15, HJ16, MSZ16, CGH17, CHKL18, CCH+19]). This line of work culminated in a landmark result by Jain, Lin, and Sahai [JLS21], who gave a construction of iO based on long-studied and well-founded cryptographic assumptions.

Despite the tremendous power of iO when combined with OWFs, it is notable that iO alone does not yield any cryptographic primitives. For example, in a hypothetical world where $\mathsf{P} = \mathsf{NP}$ (or even $\mathsf{BPP} = \mathsf{NP}$), no cryptographic primitives can exist, yet iO still exists [KMN+14]. This highlights the fact that, for iO to be cryptographically useful, one must at least assume the worst-case hardness of NP.

In this context, Komargodski et al. [KMN+14] showed that, assuming only the (infinitely-often) worst-case hardness of NP (i.e., $\mathsf{NP} \not\subseteq \mathsf{i.o.BPP}$),[2] one can already construct OWFs from iO. Once OWFs are obtained, combining them with iO yields public-key encryption (PKE) and many other advanced cryptographic primitives.

The construction of a OWF based on iO in [KMN+14] is quite simple, at least assuming the perfect correctness of the obfuscator. Given an obfuscator Obf, one can define a function f by

$$f(r) \coloneqq \mathsf{Obf}(Z; r) \tag{1}$$

where Z denotes the zero-function that outputs 0 on all inputs, and the notation "$;r$" indicates that Obf uses randomness r. The authors showed that the one-wayness of f follows from the assumption that $\mathsf{NP} \not\subseteq \mathsf{i.o.BPP}$. They also demonstrated that even an imperfectly correct iO suffices to construct OWFs, although the construction in that case is more intricate.

As the above construction illustrates, a key requirement in [KMN+14] is that the obfuscator Obf is derandomizable, meaning that it behaves deterministically when given fixed randomness. In contrast, recent works [AF16, ABDS21, BK21, BM22, BKNY23, CG24, HT25] have considered quantum obfuscation, where Obf is a quantum algorithm that obfuscates either classical or quantum circuits. In this setting, the implicit assumption of derandomizability in [KMN+14] breaks down: quantum algorithms inherently involve randomness due to measurement, and thus cannot be derandomized. As a result, the classical approach of [KMN+14] does not extend to quantum obfuscation. In fact, it appears unlikely that quantum iO implies OWFs, since a quantum obfuscator is intrinsically a randomized object, making it intuitively useless for constructing

[1] A preliminary version of [BGI+12] was published at CRYPTO 2001 and a preliminary version of [GGH+16] was published at FOCS 2013.

[2] Here, a language L is in $\mathsf{i.o.BPP}$ if all $x \in L \cap \{0,1\}^n$ are correctly decided in probabilistic polynomial time for infinitely-many $n \in \mathbb{N}$.

deterministic primitives such as OWFs. Nonetheless, recent works [Kre21, KQST23, KQT24, LMW24] have identified several cryptographic primitives that may exist even in the absence of OWFs, including one-way state generators (OWSGs) [MY22, MY24], pseudorandom state generators (PRSGs) [JLS18], pseudorandom unitaries (PRUs) [JLS18, MH24], EFI pairs [BCQ23], efficiently-verifiable one-way puzzles (EV-OWPuzzs) [CGG24], one-way puzzles (OWPuzzs) [KT24], etc. These primitives lie below OWFs in known implications, and their existence under weaker assumptions raises the possibility that quantum iO, together with quantum worst-case hardness of NP, may suffice to construct them. This leads us to the central question of this work:

Does quantum iO imply any quantum cryptographic primitive, assuming only the quantum worst-case hardness of NP*?*

1.1 Our Result

In this work, we show that quantum iO for classical circuits, when combined with the infinitely-often quantum worst-case hardness of NP (i.e., $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$), implies a range of quantum cryptographic primitives. The specific primitives that can be constructed depend on which components of the assumed iO, such as the obfuscation algorithm, the evaluation algorithm, and the description of obfuscated circuit, are quantum and which remain classical. We elaborate on these distinctions and their implications below.

To capture the various flavors of quantum iO, we formalize it as a pair of quantum polynomial-time (QPT) algorithms: an obfuscation algorithm Obf and an evaluation algorithm Eval:

$\mathsf{Obf}(1^\lambda, C)$**:** The obfuscation algorithm takes the security parameter 1^λ and a classical circuit C as input and outputs an obfuscated encoding $\hat{C}$, which may be a quantum state.

$\mathsf{Eval}(\hat{C}, x)$**:** The evaluation algorithm takes an obfuscated encoding $\hat{C}$ and an input x, and outputs $C(x)$ (with overwhelming probability).

The security requirement is that, for any pair of functionally equivalent classical circuits C_0 and C_1, the obfuscations $\mathsf{Obf}(C_0)$ and $\mathsf{Obf}(C_1)$ must be computationally indistinguishable to any QPT distinguisher.

We consider several variants of quantum iO, distinguished by whether each component, namely, the obfuscator Obf, the evaluator Eval, and the obfuscated encoding $\hat{C}$, is quantum or classical. Specifically, for each $(\mathtt{X}, \mathtt{Y}, \mathtt{Z}) \in \{\mathtt{Q}, \mathtt{C}\}^3$, we define $(\mathtt{X}, \mathtt{Y}, \mathtt{Z})$-iO as follows:

- If $\mathtt{X} = \mathtt{Q}$, then Obf is a quantum algorithm; if $\mathtt{X} = \mathtt{C}$, then Obf is classical.
- If $\mathtt{Y} = \mathtt{Q}$, then Eval is a quantum algorithm; if $\mathtt{Y} = \mathtt{C}$, then Eval is classical.
- If $\mathtt{Z} = \mathtt{Q}$, then the obfuscated encoding $\hat{C}$ is a quantum state; if $\mathtt{Z} = \mathtt{C}$, then $\hat{C}$ is a classical string.

While there are eight possible combinations of $(\mathtt{X}, \mathtt{Y}, \mathtt{Z})$, not all are meaningful. In particular, the case $\mathtt{Z} = \mathtt{Q}$ only makes sense when both $\mathtt{X} = \mathtt{Q}$ and $\mathtt{Y} = \mathtt{Q}$, since a classical obfuscator cannot generate a quantum state and a classical evaluator cannot take a quantum state as input. Accordingly, we focus on the five meaningful variants: $(\mathtt{Q}, \mathtt{Q}, \mathtt{Q})$, $(\mathtt{Q}, \mathtt{Q}, \mathtt{C})$, $(\mathtt{Q}, \mathtt{C}, \mathtt{C})$, $(\mathtt{C}, \mathtt{Q}, \mathtt{C})$, and $(\mathtt{C}, \mathtt{C}, \mathtt{C})$.

We emphasize that, throughout, the circuit being obfuscated is always a classical circuit. This modeling choice only strengthens our results since our goal is to identify lower bounds of quantum iO.

Assuming $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$, we show the following results (see Figure 1 for the summary of the results):

- $(\mathtt{Q}, \mathtt{Q}, \mathtt{Q})$-iO implies IND-CPA secure quantum symmetric key encryption (QSKE), where the secret key is classical but the ciphertext is quantum. In particular, it implies OWSGs and EFI pairs.

- $(\mathtt{Q},\mathtt{Q},\mathtt{C})$-iO implies IND-CPA secure symmetric key encryption (SKE) in the quantum-computation classical-communication (QCCC) model, referred to as QCCC SKE, where all communication is classical, but local computations, such as encryption and decryption, may be quantum. In particular, it implies EV-OWPuzz, OWPuzzs, OWSGs, QEFID pairs and EFI pairs.
- $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO implies IND-CPA secure public key encryption (PKE) in the QCCC model, referred to as QCCC PKE. In particular, it implies EV-OWPuzz, OWPuzzs, OWSGs, QEFID pairs and EFI pairs.
- $(\mathtt{C},\mathtt{Q},\mathtt{C})$-iO implies IND-CPA secure QCCC PKE and (post-quantum) OWFs. In particular, it implies all Microcrypt primitives implied by PRUs.
- $(\mathtt{C},\mathtt{C},\mathtt{C})$-iO implies IND-CPA secure PKE and (post-quantum) OWFs. In particular, it implies all Microcrypt primitives implied by PRUs.

We remark that the implication of $(\mathtt{C},\mathtt{C},\mathtt{C})$-iO can be obtained via a straightforward adaptation of the construction in [KMN+14], as all components involved are classical, except that we consider quantum adversaries. Nonetheless, we believe that our proof is arguably simpler than that of [KMN+14], which requires cascaded obfuscation, i.e., obfuscating an already obfuscated circuit, whereas our approach avoids it. In addition, an advantage of our approach is that it only requires obfuscation for 3CNF formulas, rather than for general classical circuits. This resolves the open problem left by [KMN+14], namely, constructing OWFs from imperfectly correct iO under the assumption of worst-case hardness of NP. We note, however, that this open problem has been recently resolved (in a stronger form that only requires witness encryption) by completely different techniques [HN24, LMP24].

1.2 Technical Overview

Throughout this technical overview, we assume the infinitely-often quantum worst-case hardness of NP (i.e., $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$). Our starting point is the Valiant-Vazirani theorem [VV86], which provides a randomized classical reduction from any NP instance to a UP instance.[3] This immediately implies that $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ implies $\mathsf{UP} \not\subseteq \mathsf{i.o.BQP}$ as well. Therefore, in proving cryptographic implications, we may assume $\mathsf{UP} \not\subseteq \mathsf{i.o.BQP}$ without loss of generality.

The case of $(\mathtt{C},\mathtt{C},\mathtt{C})$-iO. We begin by focusing on the case of $(\mathtt{C},\mathtt{C},\mathtt{C})$-iO, as the other cases build on similar ideas. In this setting, we construct a (post-quantum) OWF and a PKE scheme. Since a PKE scheme can be easily constructed from $(\mathtt{C},\mathtt{C},\mathtt{C})$-iO and OWFs using the same approach as in the classical construction of [SW14], it suffices to focus on constructing OWFs.

Our main technical result is the following simple yet powerful statement. Let λ be the security parameter. Let m be a polynomial in λ. For a string $k \in \{0,1\}^m$, let P_k denote a circuit computing the *point function* with target k, i.e., $P_k(k) = 1$ and $P_k(k') = 0$ for all $k' \neq k$. Let Z_m be a circuit computing the *zero function* on m-bit inputs, that is, $Z_m(k') = 0$ for all $k' \in \{0,1\}^m$. Then, for some polynomial m, assuming $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$, we show that

$$\mathsf{Obf}(1^\lambda, P_k) \approx_c \mathsf{Obf}(1^\lambda, Z_m), \tag{2}$$

where $k \leftarrow \{0,1\}^m$,[4] and $\approx_c$ means computational indistinguishability against QPT distinguishers.[5]

[3] UP is a subclass of NP consisting of problems where every yes-instance admits a unique witness.

[4] Here $k \leftarrow \{0,1\}^m$ means that k is sampled uniformly at random from $\{0,1\}^m$.

[5] In the actual theorem, we account for the circuit sizes of P_k and Z_m, but omit these details here for simplicity.

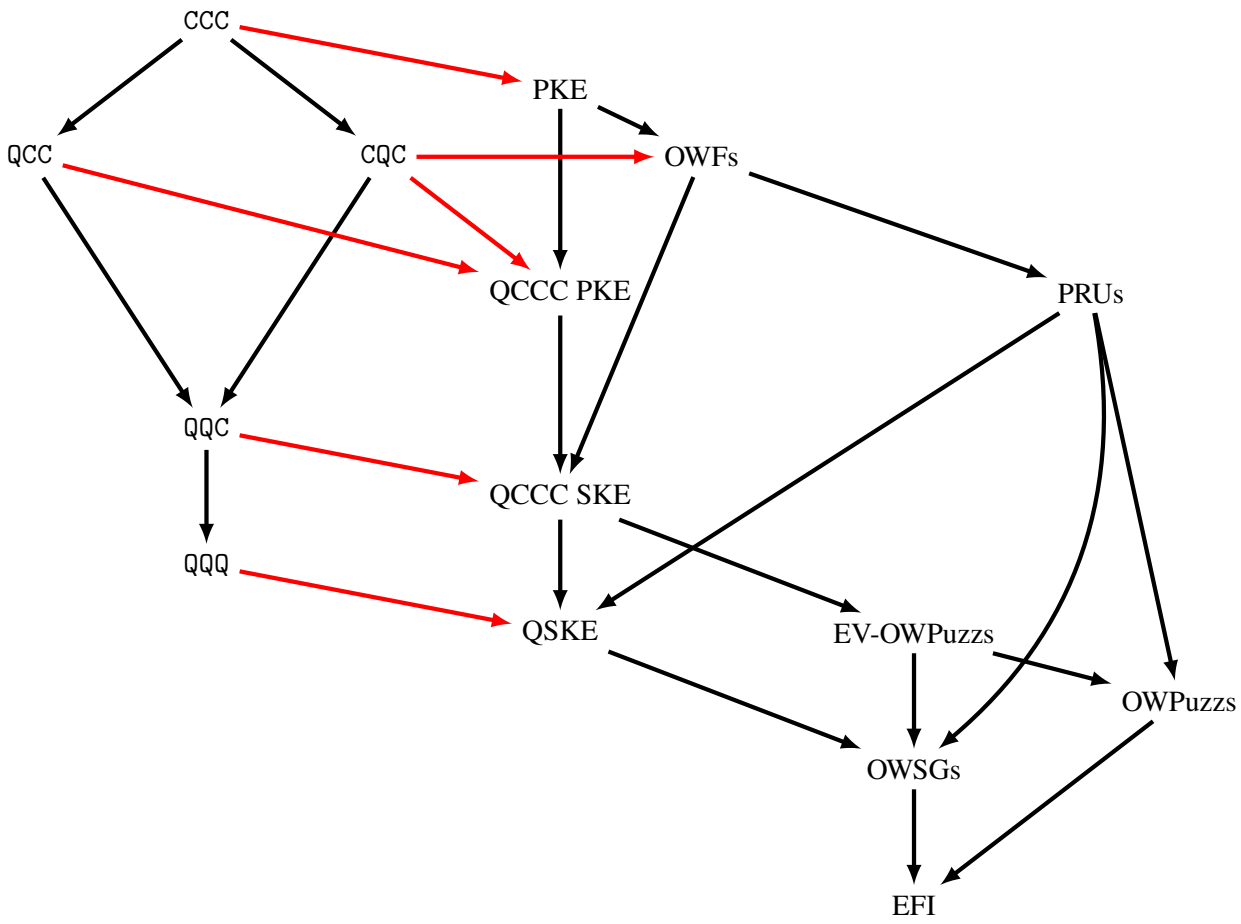

Figure 1: A summary of results (assuming $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$). Black arrows are known results or trivial implications. Red arrows are new results. $\mathtt{QQQ}$, for example, means the $(\mathtt{Q}, \mathtt{Q}, \mathtt{Q})$-iO.

This indistinguishability directly implies the existence of OWFs. Intuitively, the distributions $\mathsf{Obf}(1^\lambda, P_k)$ and $\mathsf{Obf}(1^\lambda, Z_m)$ should be *statistically far*, since the obfuscation of P_k is functionally equivalent to a point function, while the obfuscation of Z_m is functionally equivalent to a constant-zero function. Moreover, since Obf is assumed to be classical, both distributions are classically efficiently samplable. Such pairs of classically efficiently samplable, statistically far, but computationally indistinguishable distributions are known as *EFID pairs* [Gol90, BCQ23], and existentially equivalent to OWFs.[6]

We now describe the idea for proving the computational indistinguishability between $\mathsf{Obf}(1^\lambda, P_k)$ and $\mathsf{Obf}(1^\lambda, Z_m)$. As discussed above, we may assume $\mathsf{UP} \not\subseteq \mathsf{i.o.BQP}$. Therefore, to prove the above indistinguishability under the assumption, it suffices to show that any distinguisher between $\mathsf{Obf}(1^\lambda, P_k)$ and $\mathsf{Obf}(1^\lambda, Z_m)$ can be used to recover the unique witness of a yes-instance of a UP problem.

Let x be a yes-instance of a UP language with a unique witness $w \in \{0,1\}^m$, and let M be the corresponding verification algorithm. Our goal is to recover w given x and M. To do so, we take a uniformly random $v \in \{0,1\}^m$ and define a circuit $V[x,v](z)$ that outputs 1 if $M(x, v \oplus z) = 1$ and otherwise outputs 0. Note that $V[x,v](z)$ outputs 1 only when $z = w \oplus v$, so it is functionally equivalent to the point function $P_{w\oplus v}$. Hence, by the security of iO, for each $v \in \{0,1\}^m$, we have

$$\mathsf{Obf}(1^\lambda, V[x,v]) \approx_c \mathsf{Obf}(1^\lambda, P_{w\oplus v}). \tag{3}$$

Moreover, since v is chosen uniformly at random, $w \oplus v$ is also uniformly random, and so

$$\mathsf{Obf}(1^\lambda, P_{w\oplus v}) \equiv \mathsf{Obf}(1^\lambda, P_k), \tag{4}$$

where $\equiv$ denotes distributional equivalence and $v, k \leftarrow \{0,1\}^m$. Therefore, a distinguisher that distinguishes $\mathsf{Obf}(1^\lambda, P_k)$ with $k \leftarrow \{0,1\}^m$ from $\mathsf{Obf}(1^\lambda, Z_m)$ also distinguishes $\mathsf{Obf}(1^\lambda, V[x,v])$ with $v \leftarrow \{0,1\}^m$ from $\mathsf{Obf}(1^\lambda, Z_m)$.

While $V[x,v]$ and Z_m are not functionally equivalent, they differ on only a single input, namely $w \oplus v$. It is known that any iO also satisfies the notion of *differing-inputs obfuscation (diO)* [BGI+12, ABG+13, BCP14] in the single differing-input setting. This means that if one can distinguish two circuits that differ on only one input, then one can efficiently extract that input. Thus, a distinguisher between $\mathsf{Obf}(1^\lambda, V[x,v])$ with $v \leftarrow \{0,1\}^m$ and $\mathsf{Obf}(1^\lambda, Z_m)$ can be used to extract $w \oplus v$.

Since v is chosen independently of both x and w, we can construct a reduction algorithm that chooses v on its own, extracts $w \oplus v$, and thereby recovers w. This completes the proof of the computational indistinguishability.

The case of $(\mathtt{C}, \mathtt{Q}, \mathtt{C})$-iO. In this setting, we construct a (post-quantum) OWF and a QCCC PKE scheme.

To construct a OWF, we use the same argument as in the $(\mathtt{C}, \mathtt{C}, \mathtt{C})$ setting. Notably, that construction does not rely on the assumption that Eval is classical, so the proof remains valid even when Eval is quantum.[7]

For constructing a QCCC PKE scheme, we again follow the classical construction of [SW14]. The only difference is that evaluating the obfuscated program now involves quantum computation, making the encryption algorithm quantum and thus yielding a QCCC PKE scheme.

The case of $(\mathtt{Q}, \mathtt{C}, \mathtt{C})$-iO. In this setting, we construct a QCCC PKE scheme. Since Obf cannot be derandomized in this setting, it is unlikely that OWFs can be constructed. Therefore, we need a different approach from the classical construction of [SW14].

[6] The proof of this fact in [Gol90] only considers classical adversaries, but it extends to the post-quantum setting in a straightforward manner.

[7] Interestingly, the proof remains valid even if Eval is inefficient.

We begin by noting that the following indistinguishability still holds in this setting:

$$\mathsf{Obf}(1^\lambda, P_k) \approx_c \mathsf{Obf}(1^\lambda, Z_m), \tag{5}$$

where $k \leftarrow \{0,1\}^m$. Based on this, we construct a QCCC PKE scheme as follows:

- **Key generation:** Choose $k \leftarrow \{0,1\}^m$, and output the public key $\hat{P}_k \leftarrow \mathsf{Obf}(1^\lambda, P_k)$ and the secret key k.
- **Encryption:** On input a public key $\hat{P}_k$ and a message msg, define a classical circuit $C[\hat{P}_k, \mathsf{msg}]$ that takes $k' \in \{0,1\}^m$ as input and outputs msg if $\mathsf{Eval}(\hat{P}_k, k') = 1$, and outputs $\perp$ otherwise. Here, we assume for simplicity that Eval is deterministic (see the full proof in Section 5.3 for how to handle randomized Eval). The ciphertext is defined as $\hat{C}[\hat{P}_k, \mathsf{msg}] \leftarrow \mathsf{Obf}(1^\lambda, C[\hat{P}_k, \mathsf{msg}])$.
- **Decryption:** On input a ciphertext $\hat{C}[\hat{P}_k, \mathsf{msg}]$ and the secret key k, evaluate the obfuscated circuit on input k and output the result msg'.

The correctness of the scheme follows directly from the correctness of the iO.

We now argue IND-CPA security of the scheme. First, consider a hybrid where the public key is replaced with $\mathsf{Obf}(1^\lambda, Z_m)$. This is computationally indistinguishable from the real scheme by the indistinguishability between $\mathsf{Obf}(1^\lambda, P_k)$ and $\mathsf{Obf}(1^\lambda, Z_m)$, which has already been established.

Next, consider a second hybrid where the ciphertext is replaced with $\mathsf{Obf}(1^\lambda, Z_m)$. This is indistinguishable from the previous hybrid because, when the public key is $\mathsf{Obf}(1^\lambda, Z_m)$, the circuit $C[\mathsf{Obf}(1^\lambda, Z_m), \mathsf{msg}]$ is functionally equivalent to the zero function regardless of the message msg, and hence its obfuscation is indistinguishable from that of Z_m by the security of iO.

In the final hybrid, the ciphertext reveals no information about the message msg, so the scheme satisfies IND-CPA security.

The case of $(\mathtt{Q}, \mathtt{Q}, \mathtt{C})$-iO. In this setting, we construct a QCCC SKE scheme.

The idea is quite simple. We construct a QCCC SKE scheme for single-bit messages as follows: let $k \leftarrow \{0,1\}^m$ be the secret key. To encrypt the message 0, output the ciphertext $\mathsf{Obf}(1^\lambda, Z_m)$; to encrypt the message 1, output $\mathsf{Obf}(1^\lambda, P_k)$. Decryption is performed by evaluating the obfuscated program (i.e., the ciphertext) on input k and outputting the result, which will be either 0 or 1 accordingly.

While we have already established the indistinguishability between $\mathsf{Obf}(1^\lambda, Z_m)$ and $\mathsf{Obf}(1^\lambda, P_k)$, this alone is not sufficient to guarantee IND-CPA security of the above scheme. However, by carefully examining the proof of this indistinguishability, one can see that it extends to the case where the distinguisher is given multiple samples from the respective distributions, all generated using the same secret key k. This extension immediately implies the IND-CPA security of the scheme.

The case of $(\mathtt{Q}, \mathtt{Q}, \mathtt{Q})$-iO. In this setting, we construct a QSKE scheme, where the secret key is classical but the ciphertext is quantum. The construction is exactly the same as in the $(\mathtt{Q}, \mathtt{Q}, \mathtt{C})$ case described above. The only difference is that the output of Obf is a quantum state, which means the ciphertexts are quantum. As a result, the scheme realizes QSKE rather than QCCC SKE.

1.3 Related Work

Alagic and Fefferman [AF16] and Alagic, Brakerski, Dulek, and Schaffner [ABDS21] showed the impossibility of VBB obfuscation for classical circuits even when the obfuscator and the obfuscated programs are allowed to be quantum. Broadbent and Kazmi [BK21] constructed an iO for quantum circuits, whose efficiency depends exponentially on the number of T gates in the circuit being obfuscated. Bartusek and Malavolta [BM22] constructed iO for null quantum circuits in the classical oracle model.[8] Bartusek, Kitagawa, Nishimaki, and Yamakawa [BKNY23] extended it to iO for pseudo-deterministic quantum circuits in the classical oracle model.[9] Coladangelo and Gunn [CG24] introduced the notion of quantum state iO, which allows for obfuscating a quantum description of a classical function, and provided a construction in the quantum oracle model. This was later improved by Bartusek, Brakerski, and Vaikuntanathan [BBV24], who gave a construction in the classical oracle model. Huang and Tang [HT25] further improved it to support obfuscation of unitary quantum programs with quantum inputs and outputs.

Khurana and Tomer [KT25] presented a potential approach to constructing OWPuzzs based on the worst-case quantum hardness of $\#\mathsf{P}$. While the worst-case hardness of $\#\mathsf{P}$ is a significantly weaker assumption than that of NP, their current result relies on certain unproven conjectures related to quantum supremacy. While our work also relies on an additional assumption that quantum iO exists, our assumption is cryptographic in nature, whereas theirs pertains to quantum supremacy. Given the fundamental difference between the two, the approaches are not directly comparable.

Hirahara and Nanashima [HN24] proved that the infinitely-often worst-case (classical) hardness of NP and (classically-secure) zero-knowledge arguments for NP imply OWFs (see also a simplified exposition in [LMP24]). Since iO implies zero-knowledge arguments for NP, this result improves upon that of [KMN+14]. An interesting direction for future work is to establish a quantum analog of their result—for example, constructing quantum cryptographic primitives under the assumption of the infinitely-often worst-case quantum hardness of NP and the existence of quantum zero-knowledge arguments for NP. Note that such a result is known if we assume the *average-case* quantum hardness of NP instead of the worst-case quantum hardness [BCQ23].

A similar technique to ours—reducing an NP instance to a UP instance via the Valiant-Vazirani theorem, then rerandomizing the UP verification circuit through the obfuscation of a circuit in which the input is shifted by a random string—also appears in the work of Brakerski, Brzuska, and Fleischhacker [BBF16], albeit in a completely different context. Their motivation was to prove the impossibility of statistically secure iO with approximate correctness.

A concurrent work by Ilango and Lombardi [IL25] employs this technique in a context more closely related to ours. In particular, they independently present an alternative proof of the result in [KMN+14] using this idea. Nevertheless, the focus of their work is quite different: it is primarily set in the classical setting, aiming to establish fine-grained worst-case to average-case reductions using iO. While they do include one quantum result—providing proofs of quantumness from a worst-case assumption and iO—it relies on classical iO, and their work does not address quantum iO.

[8] A null quantum circuit is a quantum circuit with classical input and output that outputs 0 with high probability on all inputs.

[9] A pseudo-deterministic quantum circuit is a quantum circuit with classical input and output that computes a deterministic function with high probability.

2 Preliminaries

Notations. We use standard notations of quantum computing and cryptography. We use λ as the security parameter. $[n]$ means the set $\{1, 2, ..., n\}$. For a finite set S, $x \leftarrow S$ means that an element x is sampled uniformly at random from the set S. negl is a negligible function, and poly is a polynomial. PPT stands for (classical) probabilistic polynomial-time and QPT stands for quantum polynomial-time. We refer to a non-uniform QPT algorithm as a QPT algorithm with polynomial-size quantum advice. We stress that the running time of the algorithm can be polynomial in λ rather than in $\log \lambda$. For an algorithm A, $y \leftarrow A(x)$ means that the algorithm A outputs y on input x.

2.1 Cryptographic Primitives

Here, we give definitions of basic cryptographic primitives.

Definition 2.1 (One-Way Functions (OWFs)). *A function $f : \{0,1\}^* \to \{0,1\}^*$ is a (quantumly-secure) one-way function (OWF) if it is computable in classical deterministic polynomial-time, and for any QPT adversary $\mathcal{A}$, there exists a negligible function negl such that*

$$\Pr[f(x') = f(x) : x \leftarrow \{0,1\}^\lambda, x' \leftarrow \mathcal{A}(1^\lambda, f(x))] \leq \mathsf{negl}(\lambda). \tag{6}$$

Definition 2.2 (Quantum Symmetric Key Encryption (QSKE) [MY24]). *A QSKE scheme is a tuple $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ of QPT algorithms with the following syntax:*

- $\mathsf{Gen}(1^\lambda) \to \mathsf{sk}$*: A key generation algorithm takes the security parameter 1^λ as input and outputs a classical secret key sk.*
- $\mathsf{Enc}(\mathsf{sk}, \mathsf{msg}) \to \mathsf{ct}$*: An encryption algorithm takes a secret key sk and a message $\mathsf{msg} \in \{0,1\}^*$ as input and outputs a quantum ciphertext ct.*
- $\mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) \to \mathsf{msg}'$*: A decryption algorithm takes a secret key sk and a ciphertext ct as input and outputs a message $\mathsf{msg}' \in \{0,1\}^*$.*

We require the following correctness and IND-CPA security:

- ***Correctness:*** *For all $\mathsf{msg} \in \{0,1\}^*$ of polynomial length in λ,*

$$\Pr\left[\mathsf{msg}' = \mathsf{msg} : \begin{array}{r} \mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda) \\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{sk}, \mathsf{msg}) \\ \mathsf{msg}' \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) \end{array}\right] \geq 1 - \mathsf{negl}(\lambda). \tag{7}$$

- ***IND-CPA Security:*** *For a security parameter $\lambda \in \mathbb{N}$ and a bit $b \in \{0,1\}$, consider the following game between a challenger and an adversary $\mathcal{A}$:*
  1. *The challenger runs $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$.*
  2. *$\mathcal{A}$ can make arbitrarily many classical queries to the encryption oracle, which takes a message $\mathsf{msg} \in \{0,1\}^*$ as input and returns $\mathsf{Enc}(\mathsf{sk}, \mathsf{msg})$.*
  3. *$\mathcal{A}$ chooses $(\mathsf{msg}_0, \mathsf{msg}_1) \in (\{0,1\}^*)^2$ of the same length and sends them to the challenger.*
  4. *The challenger runs $\mathsf{ct}_b \leftarrow \mathsf{Enc}(\mathsf{sk}, \mathsf{msg}_b)$ and sends ct_b to $\mathcal{A}$.*

5. *Again, $\mathcal{A}$ can make arbitrarily many classical queries to the encryption oracle.*
6. *$\mathcal{A}$ outputs b'.*

We say that a QSKE scheme satisfies the IND-CPA security if for any QPT adversary $\mathcal{A}$,

$$|\Pr[b'=1|b=1]-\Pr[b'=1|b=0]|\leq \mathsf{negl}(\lambda). \tag{8}$$

We define SKE and PKE schemes in the quantum-computation classical-communication (QCCC) model, in which all local computations are quantum and all communication is classical.

Definition 2.3 (QCCC SKE [KT24]). *A QCCC SKE scheme is defined similarly to a QSKE scheme as defined in Definition 2.2 except that a ciphertext* ct *output by* Enc *is required to be classical.*

Definition 2.4 (QCCC Public Key Encryption (QCCC PKE) [KT24]). *A QCCC PKE scheme is a tuple* $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ *of QPT algorithms with the following syntax:*

- $\mathsf{Gen}(1^\lambda)\to(\mathsf{pk},\mathsf{sk})$*: A key generation algorithm takes the security parameter 1^λ as input and outputs a classical public key* pk *and classical secret key* sk*.*
- $\mathsf{Enc}(\mathsf{pk},\mathsf{msg})\to\mathsf{ct}$*: An encryption algorithm takes a public key* pk *and a message* $\mathsf{msg}\in\{0,1\}^*$ *as input and outputs a classical ciphertext* ct*.*
- $\mathsf{Dec}(\mathsf{sk},\mathsf{ct})\to\mathsf{msg}'$*: A decryption algorithm takes a secret key* sk *and a ciphertext* ct *as input and outputs a message* $\mathsf{msg}'\in\{0,1\}^*$*.*

We require the following correctness and IND-CPA security:

- ***Correctness:*** *For all* $\mathsf{msg}\in\{0,1\}^*$ *of polynomial length in λ,*

$$\Pr\left[\mathsf{msg}'=\mathsf{msg}:\begin{array}{r}(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^\lambda)\\ \mathsf{ct}\leftarrow\mathsf{Enc}(\mathsf{pk},\mathsf{msg})\\ \mathsf{msg}'\leftarrow\mathsf{Dec}(\mathsf{sk},\mathsf{ct})\end{array}\right]\geq 1-\mathsf{negl}(\lambda). \tag{9}$$

- ***IND-CPA Security:*** *For a security parameter $\lambda\in\mathbb{N}$ and a bit $b\in\{0,1\}$, consider the following game between a challenger and an adversary $\mathcal{A}$:*
  1. *The challenger runs $(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^\lambda)$ and sends pk to $\mathcal{A}$.*
  2. *$\mathcal{A}$ chooses $\mathsf{msg}_0,\mathsf{msg}_1\in(\{0,1\}^*)^2$ of the same length and sends them to the challenger.*
  3. *The challenger runs $\mathsf{ct}_b\leftarrow\mathsf{Enc}(\mathsf{sk},\mathsf{msg}_b)$ and sends ct_b to $\mathcal{A}$.*
  4. *$\mathcal{A}$ outputs b'.*

  We say that a QCCC PKE scheme satisfies the IND-CPA security if for any QPT adversary $\mathcal{A}$,

$$|\Pr[b'=1|b=1]-\Pr[b'=1|b=0]|\leq \mathsf{negl}(\lambda). \tag{10}$$

We omit definitions of other cryptographic primitives, such as EV-OWPuzzs [CGG24], OWPuzzs [KT24], OWSGs [MY24], QEFID pairs [CGG24], PRUs [JLS18], and EFI pairs [BCQ23], since we obtain them only as corollaries. For their definitions, we refer the reader to the respective cited works. One remark regarding the definition of OWSGs is that, unless stated otherwise, we refer to OWSGs with mixed-state outputs as defined in [MY24].

2.2 Complexity Theory

Here we explain basic complexity classes we use.

Definition 2.5 (i.o.BQP). *A promise problem $\Pi = (\Pi_{yes}, \Pi_{no})$ is in* i.o.BQP *if there exist a QPT algorithm Q and infinitely many $\lambda \in \mathbb{N}$ such that for all $x \in \Pi_{yes} \cup \Pi_{no}$,*

- *if $x \in \Pi_{yes} \cap \{0,1\}^\lambda$, then $\Pr[1 \leftarrow Q(x)] \geq 2/3$.*
- *if $x \in \Pi_{no} \cap \{0,1\}^\lambda$, then $\Pr[1 \leftarrow Q(x)] \leq 1/3$.*

Definition 2.6 (UP). *A promise problem $\Pi = (\Pi_{yes}, \Pi_{no})$ is in* UP *if there exist a classical polynomial-time (deterministic) Turing machine M and a polynomial m such that*

- *if $x \in \Pi_{yes}$, then there exists a unique $w \in \{0,1\}^{m(|x|)}$ such that $M(x,w) = 1$.*
- *if $x \in \Pi_{no}$, then for all $w \in \{0,1\}^{m(|x|)}$, $M(x,w) = 0$.*

The following lemma follows from the Valiant-Vazirani theorem [VV86].

Lemma 2.7. *If* NP $\not\subseteq$ i.o.BQP*, then* UP $\not\subseteq$ i.o.BQP.

3 Definitions of Quantum Obfuscation

We introduce definitions of quantum indistinguishability obfuscation for classical circuits and its variants.

Definition 3.1 (Quantum iO for Classical Circuits). *A quantum indistinguishability obfuscator (quantum iO) for classical circuits consists of two QPT algorithms $(\mathsf{Obf}, \mathsf{Eval})$ with the following syntax:*

- $\mathsf{Obf}(1^\lambda, C) \to \hat{C}$*: An obfuscation algorithm takes the security parameter 1^λ and a classical circuit C as input and outputs a quantum state $\hat{C}$, which we refer to as an obfuscated encoding of C.*
- $\mathsf{Eval}(\hat{C}, x) \to y$*: An evaluation algorithm takes an obfuscated encoding $\hat{C}$ and a classical input x as input and outputs a classical output y.*

We require the following correctness and security.

- ***Correctness:*** *For any family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$ of polynomial-size classical circuits of input length n_λ, and for any polynomial p, there exists $N \in \mathbb{N}$ such that*

$$\Pr_{\hat{C}_\lambda \leftarrow \mathsf{Obf}(1^\lambda, C_\lambda)}\left[\forall x \in \{0,1\}^{n_\lambda}, \Pr[\mathsf{Eval}(\hat{C}_\lambda, x) = C_\lambda(x)] \geq 1 - \frac{1}{p(\lambda)}\right] \geq 1 - \frac{1}{p(\lambda)} \tag{11}$$

  holds for all $\lambda \geq N$, where the inner probability is taken over the randomness of the execution of $\mathsf{Eval}(\hat{C}_\lambda, x)$.

- ***Security:*** *For any families $\{C_{0,\lambda}\}_{\lambda\in\mathbb{N}}$ and $\{C_{1,\lambda}\}_{\lambda\in\mathbb{N}}$ of polynomial-size classical circuits such that $C_{0,\lambda}$ and $C_{1,\lambda}$ are functionally equivalent and of the same size, and for any non-uniform QPT adversary $\mathcal{A}$,*

$$|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathsf{Obf}(1^\lambda, C_{0,\lambda}))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathsf{Obf}(1^\lambda, C_{1,\lambda}))]| \leq \mathsf{negl}(\lambda). \tag{12}$$

Remark 3.2. The correctness notion defined above may seem strong, as it requires that the evaluation returns the correct output for all inputs simultaneously with overwhelming probability. However, this stronger guarantee can be generically achieved assuming only a quantum iO with a weaker, input-wise correctness—that is, for each fixed input, the evaluation returns the correct output with overwhelming probability. To achieve the stronger correctness, we can simply repeat the evaluation algorithm multiple times and take a majority vote. By repeating sufficiently many times, the failure probability for each input can be reduced exponentially, and a union bound then ensures that the overall probability of failing on any input remains negligible. This argument closely parallels the classical case presented in [KMN+14, Appendix B], and thus we omit the details.

Remark 3.3. We require the security of iO to hold against non-uniform quantum QPT adversaries, even though our final goal is to construct cryptographic primitives with uniform security. This is because a uniform version of the security, where a uniform QPT adversary $\mathcal{A}$ chooses two functionally equivalent circuits C_0 and C_1 and then tries to distinguish obfuscations of them, would not suffice for our purpose, since we must consider a reduction algorithm that hardwires an arbitrary choice of an NP instance into the circuits. Although this point is not explicitly discussed, we believe the same applies even in the classical setting of [KMN+14]. In addition, we note that the non-uniform security notion aligns more closely with the standard formalization in the literature on iO.

Definition 3.4 (Variations of Quantum iO). *For* $(\mathtt{X},\mathtt{Y},\mathtt{Z}) \in \{\mathtt{Q},\mathtt{C}\}^3$, $(\mathtt{X},\mathtt{Y},\mathtt{Z})$*-iO for classical circuits is defined similarly to quantum iO for classical circuits as defined in Definition 3.1 except that:*

- *If* $\mathtt{X} = \mathtt{C}$, Obf *is a PPT algorithm whereas if* $\mathtt{X} = \mathtt{Q}$, Obf *is a QPT algorithm;*
- *If* $\mathtt{Y} = \mathtt{C}$, Eval *is a PPT algorithm whereas if* $\mathtt{Y} = \mathtt{Q}$, Eval *is a QPT algorithm;*
- *If* $\mathtt{Z} = \mathtt{C}$, *an encoding* $\hat{C}$ *output by* Obf *is a classical string whereas if* $\mathtt{Z} = \mathtt{Q}$, $\hat{C}$ *is a quantum state.*

Remark 3.5. While there are 8 possible choices for $(\mathtt{X},\mathtt{Y},\mathtt{Z})$, some of them are meaningless. In particular, it makes sense to have $\mathtt{Z} = \mathtt{Q}$ only if $\mathtt{X} = \mathtt{Y} = \mathtt{Q}$ since classical Obf cannot output quantum $\hat{C}$ and classical Eval cannot take quantum $\hat{C}$ as input. Thus, there are 5 meaningful choices: $(\mathtt{Q},\mathtt{Q},\mathtt{Q})$, $(\mathtt{Q},\mathtt{Q},\mathtt{C})$, $(\mathtt{Q},\mathtt{C},\mathtt{C})$, $(\mathtt{C},\mathtt{Q},\mathtt{C})$, and $(\mathtt{C},\mathtt{C},\mathtt{C})$ (see Figure 1 for their relationship). $(\mathtt{Q},\mathtt{Q},\mathtt{Q})$-iO corresponds to quantum iO as defined in Definition 3.1 and $(\mathtt{C},\mathtt{C},\mathtt{C})$-iO corresponds to (post-quantum) classical iO. We stress that we consider obfuscation of *classical* circuits and security against *quantum* adversaries in all the variant.

In the security definition of iO, the two circuits are required to be functionally equivalent, that is, they must agree on all inputs. The notion of differing-inputs obfuscation (diO) [BGI+12, ABG+13, BCP14] relaxes this requirement by allowing circuits that may differ on some inputs, as long as those inputs are hard to find. This results in a strictly stronger security notion than iO. It is known that in the classical setting, iO and diO are equivalent when the number of differing inputs is polynomial. We observe that this equivalence extends to the quantum setting as well. For our purposes, we present the definition of diO in the case where there is only a single differing input, which suffices for our applications.

Definition 3.6 (Single-Point Differing-Inputs Obfuscation (diO)). *For* $(\mathtt{X},\mathtt{Y},\mathtt{Z}) \in \{\mathtt{Q},\mathtt{C}\}^3$, $(\mathtt{X},\mathtt{Y},\mathtt{Z})$*-single-point diO for classical circuits is defined similarly to* $(\mathtt{X},\mathtt{Y},\mathtt{Z})$*-iO except that the security is replaced with extractability defined as follows:*

- ***Extractability (for single-differing-point):*** *For any QPT adversary* $\mathcal{A}$ *and any polynomial* p*, there exist a QPT algorithm* Ext *and a polynomial* q *for which the following holds. For any pair of families of*

polynomial-size classical circuits $\{C_{0,\lambda}\}_{\lambda\in\mathbb{N}}$ and $\{C_{1,\lambda}\}_{\lambda\in\mathbb{N}}$, such that for each λ, $C_{0,\lambda}$ and $C_{1,\lambda}$ have the same size and input length, and differ on at most a single input, and for any family of polynomial-size classical strings $\{z_\lambda\}_{\lambda\in\mathbb{N}}$, the following holds for all sufficiently large $\lambda\in\mathbb{N}$:

$$\Pr\left[b'=b:\begin{array}{c} b\leftarrow\{0,1\}\\ \hat{C}_\lambda\leftarrow\mathsf{Obf}(1^\lambda,C_{b,\lambda})\\ b'\leftarrow\mathcal{A}(1^\lambda,\hat{C}_\lambda,C_{0,\lambda},C_{1,\lambda},z_\lambda)\end{array}\right]\geq\frac{1}{2}+\frac{1}{p(\lambda)} \tag{13}$$

$$\implies\Pr\left[C_{0,\lambda}(x)\neq C_{1,\lambda}(x):x\leftarrow\mathsf{Ext}(1^\lambda,C_{0,\lambda},C_{1,\lambda},z_\lambda)\right]\geq\frac{1}{q(\lambda)}. \tag{14}$$

Remark 3.7. Zhandry [Zha23] observed that defining diO involves subtle challenges when considering security against quantum adversaries with quantum advice. In contrast, we restrict our attention to quantum adversaries with classical advice. As a result, these complications do not arise in our setting, allowing us to define diO in a manner that closely mirrors the classical definition from [BCP14].

In the classical advice setting considered above, the equivalence between iO and single-point diO can be proven using essentially the same argument as in the classical case, as shown in [BCP14].

Lemma 3.8. *For $(\mathtt{X},\mathtt{Y},\mathtt{Z})\in\{\mathtt{Q},\mathtt{C}\}^3$, if $(\mathsf{Obf},\mathsf{Eval})$ is an $(\mathtt{X},\mathtt{Y},\mathtt{Z})$-iO for classical circuits, then it is also $(\mathtt{X},\mathtt{Y},\mathtt{Z})$-single-point diO for classical circuits.*

4 Main Technical Theorem

We prove a technical theorem that is the basis of all our cryptographic implications.

Let s_{min} be a polynomial such that, for any $m\in\mathbb{N}$ and any $k\in\{0,1\}^m$, there exist classical circuits of size at most $s_{min}(m)$ that compute the following functions on m-bit inputs:

- the point function at target point k, which outputs 1 on input k and 0 on all other inputs; and
- the zero function, which outputs 0 on all m-bit inputs.

For $m\in\mathbb{N}$, $k\in\{0,1\}^m$, and $s\geq s_{min}(m)$, let $P_{k,s}$ denote a *canonical* classical circuit of size s that computes the point function on the target point k, and let $Z_{m,s}$ be a *canonical* classical circuit of size s that computes the zero-function on m-bit inputs. Here, "canonical" refers to a fixed but arbitrary choice of circuit construction, provided that the descriptions of $P_{k,s}$ and $Z_{m,s}$ are computable in classical polynomial time from $(k,1^s)$ and $(1^m,1^s)$, respectively. The specific choice of canonical circuits does not affect our results.

Then we prove the following theorem.

Theorem 4.1. *Suppose $\mathsf{NP}\not\subseteq\mathsf{i.o.BQP}$ and $(\mathsf{Obf},\mathsf{Eval})$ is an $(\mathtt{X},\mathtt{Y},\mathtt{Z})$-iO for classical circuits for $(\mathtt{X},\mathtt{Y},\mathtt{Z})\in\{\mathtt{Q},\mathtt{C}\}^3$. Then, there are classical-polynomial-time-computable polynomials m and s such that for any polynomial ℓ, the following two distributions (over classical bit strings if $\mathtt{Z}=\mathtt{C}$ and over quantum states if $\mathtt{Z}=\mathtt{Q}$) are computationally indistinguishable against uniform QPT adversaries:*

- *$\mathcal{D}_0(\lambda)$: Sample $k\leftarrow\{0,1\}^{m(\lambda)}$, run $\hat{P}^i_{k,s(\lambda)}\leftarrow\mathsf{Obf}(1^\lambda,P_{k,s(\lambda)})$ for $i\in[\ell(\lambda)]$, and output $(1^\lambda,\hat{P}^1_{k,s(\lambda)},\hat{P}^2_{k,s(\lambda)},\ldots,\hat{P}^{\ell(\lambda)}_{k,s(\lambda)})$.*
- *$\mathcal{D}_1(\lambda)$: Run $\hat{Z}^i_{m(\lambda),s(\lambda)}\leftarrow\mathsf{Obf}(1^\lambda,Z_{m(\lambda),s(\lambda)})$ for $i\in[\ell(\lambda)]$ and output $(1^\lambda,\hat{Z}^1_{m(\lambda),s(\lambda)},\hat{Z}^2_{m(\lambda),s(\lambda)},\ldots,\hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})$.*

Proof of Theorem 4.1. By Lemma 2.7, we have $\mathsf{UP} \not\subseteq \mathsf{i.o.BQP}$. Then, there exists a promise problem $\Pi = (\Pi_{yes}, \Pi_{no}) \in \mathsf{UP}$ such that $\Pi \notin \mathsf{i.o.BQP}$. By the definition of UP, there exist a polynomial m and a deterministic polynomial-time Turing machine M such that

- If $x \in \Pi_{yes}$, then there exists a unique $w \in \{0,1\}^{m(|x|)}$ such that $M(x,w) = 1$.
- If $x \in \Pi_{no}$, $M(x,w) = 0$ for any $w \in \{0,1\}^{m(|x|)}$.

Without loss of generality, we can assume that m is classical-polynomial-time-computable because we can pad w so that its length matches (an upper bound of) the running time of M.

For the sake of contradiction, let us assume that for any classical-polynomial-time-computable polynomial s, the following holds: There exist polynomials ℓ and p, and a uniform QPT algorithm $\mathcal{A}$ such that

$$\Pr_{k\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{k,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{k,s(\lambda)})] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})] \geq \frac{1}{p(\lambda)} \tag{15}$$

holds for infinitely many λ, where $\hat{P}^i_{k,s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, P_{k,s(\lambda)})$ and $\hat{Z}^i_{m(\lambda),s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ for each $i \in [\ell(\lambda)]$. Our goal is to construct a QPT algorithm that solves Π for infinitely many input lengths, thereby showing that $\Pi \in \mathsf{i.o.BQP}$. To do this, it suffices to show that there exist a QPT algorithm $\mathcal{B}$ and a polynomial r such that for infinitely many $\lambda \in \mathbb{N}$

$$\Pr[M(x,w) = 1 : w \leftarrow \mathcal{B}(x)] \geq \frac{1}{r(\lambda)} \tag{16}$$

is satisfied for all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$. Then, a QPT algorithm $\mathcal{C}$ that on input $x \in \{0,1\}^\lambda$, runs $w \leftarrow \mathcal{B}(x)$ and outputs $M(x,w)$ satisfies,

- if $x \in \Pi_{yes}$, $\Pr[1 \leftarrow \mathcal{C}(x)] = \Pr[M(x,w) = 1 : w \leftarrow \mathcal{B}(x)] \geq \frac{1}{r(\lambda)}$,
- if $x \in \Pi_{no}$, $\Pr[1 \leftarrow \mathcal{C}(x)] = \Pr[M(x,w) = 1 : w \leftarrow \mathcal{B}(x)] = 0$,

for infinitely many $\lambda \in \mathbb{N}$. The completeness-soundness gap is $1/r(\lambda) = 1/\mathrm{poly}(\lambda)$ and therefore $\Pi \in \mathsf{i.o.BQP}$.

In the remaining part, we show the existence of a polynomial r and a QPT algorithm $\mathcal{B}$ that satisfy Equation (16). There exists a family $\{V_\lambda[x,v]\}_{\lambda\in\mathbb{N}}$ of polynomial-size classical circuits such that for each λ, $V_\lambda[x,v]$ is parametrized by $x \in \{0,1\}^\lambda$ and $v \in \{0,1\}^{m(\lambda)}$, operates on $m(\lambda)$ bits, and computes the function

$$V_\lambda[x,v](z) := \begin{cases} 1 & \text{if } M(x, v\oplus z) = 1 \\ 0 & \text{otherwise,} \end{cases} \tag{17}$$

where the description of $V_\lambda[x,v]$ is computable in classical polynomial time from (x,v). Let $s(\lambda)$ be the size of $V_\lambda[x,v]$. Then, s is classical-polynomial-time-computable because there exists a Turing machine that on input 1^λ, computes $1^{m(\lambda)}$, computes the description of $V_\lambda[1^\lambda, 1^{m(\lambda)}]$ from $(1^\lambda, 1^{m(\lambda)})$, and outputs $s(\lambda) = |V_\lambda[1^\lambda, 1^{m(\lambda)}]|$. We can choose s such that $s(\lambda) \geq s_{min}(m(\lambda))$ by adding dummy gates that do not change the functionality of $V_\lambda[x,v]$.

For each $x \in \Pi_{yes}$ and $v \in \{0,1\}^{m(|x|)}$, $V_{|x|}[x,v]$ has the same functionality with $P_{w\oplus v, s(|x|)}$, where $w \in \{0,1\}^{m(|x|)}$ is the unique witness for x.

Then, by the security of iO, for any uniform QPT algorithm $\mathcal{A}'$ and for any polynomial b, there exists $N \in \mathbb{N}$ such that

$$\left| \Pr[1 \leftarrow \mathcal{A}'(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v})] - \Pr[1 \leftarrow \mathcal{A}'(1^\lambda, \hat{P}^1_{w\oplus v,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{w\oplus v,s(\lambda)})] \right| \leq \frac{1}{b(\lambda)} \tag{18}$$

for all $v \in \{0,1\}^{m(\lambda)}$, all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$ and all λ such that $\lambda \geq N$, where $\hat{V}^i_{x,v} \leftarrow \mathsf{Obf}(1^\lambda, V_\lambda[x,v])$ and $\hat{P}^i_{w\oplus v,s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, P_{w\oplus v,s(\lambda)})$ for each $i \in [\ell(\lambda)]$. Then, for the QPT algorithm $\mathcal{A}$ that satisfies Equation (15), there exist infinitely many $\lambda \in \mathbb{N}$ such that for all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$,

$$\Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v})] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})] \tag{19}$$

$$= \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{v,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{v,s(\lambda)})] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})] \tag{20}$$

$$+ \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v})] - \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{v,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{v,s(\lambda)})] \tag{21}$$

$$= \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{v,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{v,s(\lambda)})] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})] \tag{22}$$

$$+ \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v})] - \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{w\oplus v,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{w\oplus v,s(\lambda)})] \tag{23}$$

$$\geq \frac{1}{p(\lambda)} - \frac{1}{2p(\lambda)} \quad \text{(By Equations (15) and (18).)} \tag{24}$$

$$= \frac{1}{2p(\lambda)}, \tag{25}$$

where $\hat{V}^i_{x,v} \leftarrow \mathsf{Obf}(1^\lambda, V_\lambda[x,v])$, $\hat{P}^i_{k,s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, P_{k,s(\lambda)})$, and $\hat{Z}^i_{m(\lambda),s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ for each $i \in [\ell(\lambda)]$. Let us consider the following QPT algorithm $\mathcal{E}$:

1. Take $(1^\lambda, \hat{C}^1, ..., \hat{C}^{\ell(\lambda)}, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})$ as input, where $\hat{C}^i \in \{\hat{V}^i_{x,v}, \hat{Z}^i_{m(\lambda),s(\lambda)}\}$ for all $i \in [\ell(\lambda)]$.
2. Run $b \leftarrow \mathcal{A}(1^\lambda, \hat{C}^1, ..., \hat{C}^{\ell(\lambda)})$.
3. Output b.

By Equation (25), there exist infinitely many $\lambda \in \mathbb{N}$ such that for all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$,

$$\mathop{\mathbb{E}}_{v\leftarrow\{0,1\}^{m(\lambda)}}\Big[\Pr[1 \leftarrow \mathcal{E}(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v}, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})] \tag{26}$$

$$- \Pr[1 \leftarrow \mathcal{E}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)}, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})]\Big] \tag{27}$$

$$= \Pr_{v\leftarrow\{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v})] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})] \tag{28}$$

$$\geq \frac{1}{2p(\lambda)}. \tag{29}$$

For each $\lambda \in \mathbb{N}$ and each $x \in \{0,1\}^\lambda$, define a set

$$\mathsf{Good}_{\lambda,x} := \left\{ v \in \{0,1\}^{m(\lambda)} : \begin{array}{l} \Pr[1 \leftarrow \mathcal{E}(1^\lambda, \hat{V}^1_{x,v}, ..., \hat{V}^{\ell(\lambda)}_{x,v}, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})] \\ -\Pr[1 \leftarrow \mathcal{E}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)}, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})] \geq \frac{1}{4p(\lambda)} \end{array} \right\}. \tag{30}$$

By Equation (29),

$$\frac{1}{2p(\lambda)} \leq \Pr_{v \leftarrow \{0,1\}^{m(\lambda)}}[v \in \mathsf{Good}_{\lambda,x}] + \left(1 - \Pr_{v \leftarrow \{0,1\}^{m(\lambda)}}[v \in \mathsf{Good}_{\lambda,x}]\right) \frac{1}{4p(\lambda)}. \tag{31}$$

Thus, for infinitely many $\lambda \in \mathbb{N}$ and for all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$,

$$\Pr_{v \leftarrow \{0,1\}^{m(\lambda)}}[v \in \mathsf{Good}_{\lambda,x}] \geq \frac{1}{4p(\lambda) - 1}. \tag{32}$$

By Lemma 3.8, $(\mathsf{Obf}, \mathsf{Eval})$ is also a single-point diO for classical circuits. Therefore, there exist a QPT algorithm Ext and a polynomial q such that the following holds: There exist infinitely many $\lambda \in \mathbb{N}$ such that for all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$ and all $v \in \mathsf{Good}_{\lambda,x}$,

$$\Pr\left[V_\lambda[x,v](z) \neq Z_{m(\lambda),s(\lambda)}(z) : z \leftarrow \mathsf{Ext}(1^\lambda, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})\right] \geq \frac{1}{q(\lambda)}. \tag{33}$$

By using such Ext, we construct a QPT algorithm $\mathcal{B}$ as follows:

- Take $x \in \{0,1\}^*$ as input. Set $\lambda := |x|$.
- Compute $m(\lambda)$. Here $m(\lambda)$ is the classical-polynomial-time-computable polynomial that corresponds to the witness length for x.
- Sample $v \leftarrow \{0,1\}^{m(\lambda)}$.
- Compute $s(\lambda)$. Here $s(\lambda)$ is the classical-polynomial-time-computable polynomial that corresponds to the size of $V_\lambda[x,v]$.
- Run $z \leftarrow \mathsf{Ext}(1^\lambda, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)})$.
- Output $z \oplus v$.

Then, for infinitely many $\lambda \in \mathbb{N}$,

$$\Pr[M(x,w) = 1 : w \leftarrow \mathcal{B}(x)] \tag{34}$$

$$\geq \Pr\left[v \in \mathsf{Good}_{\lambda,x} \wedge V_\lambda[x,v](z) \neq Z_{m(\lambda),s(\lambda)}(z) : \begin{array}{c} v \leftarrow \{0,1\}^{m(\lambda)}; \\ z \leftarrow \mathsf{Ext}(1^\lambda, V_\lambda[x,v], Z_{m(\lambda),s(\lambda)}) \end{array}\right] \tag{35}$$

$$\geq \frac{1}{q(\lambda)(4p(\lambda) - 1)}, \tag{36}$$

holds for all $x \in \Pi_{yes} \cap \{0,1\}^\lambda$. We have the QPT algorithm $\mathcal{B}$ and a polynomial $r(\lambda) := q(\lambda)(4p(\lambda) - 1)$ that satisfy Equation (16), and therefore we complete the proof. □

5 Cryptographic Implications

In this section, we use Theorem 4.1 to demonstrate cryptographic implications of various quantum iO variants (as defined in Definition 3.4), when combined with the worst-case hardness of NP.

5.1 Q-Obf, Q-Eval, and Q-Encoding

Here, we study implications of $(\mathtt{Q},\mathtt{Q},\mathtt{Q})$-iO, where both Obf and Eval are quantum algorithms and an obfuscated encoding $\hat{C}$ is a quantum state. We prove the following theorem:

Theorem 5.1. *Suppose* $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{Q},\mathtt{Q},\mathtt{Q})$*-iO for classical circuits. Then there exists an IND-CPA secure QSKE scheme.*

Proof of Theorem 5.1. By Theorem 4.1, there exist classical-polynomial-time-computable polynomials m and s such that for any polynomial ℓ, the following two distributions over quantum states are computationally indistinguishable:

- $\mathcal{D}_0(\lambda)$: Sample $k \leftarrow \{0,1\}^{m(\lambda)}$, run $\hat{P}^i_{k,s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, P_{k,s(\lambda)})$ for $i \in [\ell(\lambda)]$, and output $(1^\lambda, \hat{P}^1_{k,s(\lambda)}, ..., \hat{P}^{\ell(\lambda)}_{k,s(\lambda)})$.
- $\mathcal{D}_1(\lambda)$: Run $\hat{Z}^i_{m(\lambda),s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ for $i \in [\ell(\lambda)]$ and output $(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{\ell(\lambda)}_{m(\lambda),s(\lambda)})$.

Without loss of generality, it suffices to construct an IND-CPA secure QSKE scheme $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ for single-bit message. We construct $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ as follows:

- $\mathsf{Gen}(1^\lambda) \to \mathsf{sk}$: Take the security parameter 1^λ as input, compute $m(\lambda)$, and sample $k \leftarrow \{0,1\}^{m(\lambda)}$. Output $\mathsf{sk} := k$.
- $\mathsf{Enc}(\mathsf{sk}, b) \to \mathsf{ct}$: Take the secret key sk and a message $b \in \{0,1\}$ as input. Let $C_0 := Z_{m(\lambda),s(\lambda)}$ and $C_1 := P_{\mathsf{sk},s(\lambda)}$. Sample $\hat{C}_b \leftarrow \mathsf{Obf}(1^\lambda, C_b)$ and output $\mathsf{ct} := \hat{C}_b$.
- $\mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) \to b'$: Take sk and ct as input. Run $b' \leftarrow \mathsf{Eval}(\mathsf{ct}, \mathsf{sk})$ and output b'.

The correctness of $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ follows from the correctness of iO: For all $b \in \{0,1\}$,

$$\Pr\left[b' = b : \begin{array}{r}\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda);\\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{sk}, b);\\ b' \leftarrow \mathsf{Dec}(\mathsf{sk}, \mathsf{ct})\end{array}\right] = \Pr\left[b' = b : \begin{array}{r}k \leftarrow \{0,1\}^{m(\lambda)};\\ \hat{C}_b \leftarrow \mathsf{Obf}(1^\lambda, C_b);\\ b' \leftarrow \mathsf{Eval}(\hat{C}_b, k)\end{array}\right] \tag{37}$$

$$= \Pr\left[\mathsf{Eval}(\hat{C}_b, k) = C_b(k) : \begin{array}{r}k \leftarrow \{0,1\}^{m(\lambda)};\\ \hat{C}_b \leftarrow \mathsf{Obf}(1^\lambda, C_b)\end{array}\right] \tag{38}$$

$$\geq 1 - \mathsf{negl}(\lambda). \tag{39}$$

We show the IND-CPA security of $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$. In the IND-CPA security game of $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$, it suffices to consider any QPT adversary that makes polynomially many queries only on message 1 to the encryption oracle because the ciphertext for message 0 can be computed without the secret key by simply running $\mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$. Thus, our goal is to show that for any polynomial t, the following two distributions are computationally indistinguishable:

- $\mathcal{D}'_0(\lambda)$: Sample $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$, run $\mathsf{ct}_1^i \leftarrow \mathsf{Enc}(\mathsf{sk}, 1)$ for $i \in [t(\lambda)]$ and $\mathsf{ct}_1^* \leftarrow \mathsf{Enc}(\mathsf{sk}, 1)$, and output $(\mathsf{ct}_1^1, ..., \mathsf{ct}_1^{t(\lambda)}, \mathsf{ct}_1^*)$.
- $\mathcal{D}'_1(\lambda)$: Sample $\mathsf{sk} \leftarrow \mathsf{Gen}(1^\lambda)$ run $\mathsf{ct}_1^i \leftarrow \mathsf{Enc}(\mathsf{sk}, 1)$ for $i \in [t(\lambda)]$ and $\mathsf{ct}_0^* \leftarrow \mathsf{Enc}(\mathsf{sk}, 0)$, and output $(\mathsf{ct}_1^1, ..., \mathsf{ct}_1^{t(\lambda)}, \mathsf{ct}_0^*)$.

When $\ell(\lambda) = t(\lambda) + 1$, the distribution $\mathcal{D}'_0(\lambda)$ is identical to the distribution $\mathcal{D}_0(\lambda)$. Thus, for any polynomial t and for any QPT adversary $\mathcal{A}$,

$$\left|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}'_0(\lambda))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}'_1(\lambda))]\right| \tag{40}$$

$$= \left|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}_0(\lambda))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}'_1(\lambda))]\right| \tag{41}$$

$$\leq \left|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}_0(\lambda))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}_1(\lambda))]\right| + \left|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}_1(\lambda))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}'_1(\lambda))]\right| \tag{42}$$

$$\leq \mathsf{negl}(\lambda) + \left|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}_1(\lambda))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}'_1(\lambda))]\right|. \quad \text{(By Theorem 4.1.)} \tag{43}$$

To show the computational indistinguishability between $\mathcal{D}'_0(\lambda)$ and $\mathcal{D}'_1(\lambda)$, it suffices to show that $\mathcal{D}_1(\lambda)$ and $\mathcal{D}'_1(\lambda)$ are computationally indistinguishable. For the sake of contradiction, assume that there exist polynomials t and p and a QPT algorithm $\mathcal{A}$ such that

$$\frac{1}{p(\lambda)} \leq \left|\Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}_1(\lambda))] - \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \mathcal{D}'_1(\lambda))]\right| \tag{44}$$

$$= \Bigg| \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{t(\lambda)}_{m(\lambda),s(\lambda)}, \hat{Z}^{t(\lambda)+1}_{m(\lambda),s(\lambda)})] \tag{45}$$

$$- \Pr_{k \leftarrow \{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{k,s(\lambda)}, ..., \hat{P}^{t(\lambda)}_{k,s(\lambda)}, \hat{Z}^{t(\lambda)+1}_{m(\lambda),s(\lambda)})]\Bigg| \tag{46}$$

for infinitely many $\lambda \in \mathbb{N}$, where $\hat{Z}^i_{m(\lambda),s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ and $\hat{P}^i_{k,s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, P_{k,s(\lambda)})$ for $i \in [t(\lambda) + 1]$. Let us consider a QPT algorithm $\mathcal{B}$ that on input $(1^\lambda, \hat{C}^1_b, ..., \hat{C}^{t(\lambda)}_b)$, runs $\hat{Z}_{m(\lambda),s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ and $b' \leftarrow \mathcal{A}(1^\lambda, \hat{C}^1_b, ..., \hat{C}^{t(\lambda)}_b, \hat{Z}_{m(\lambda),s(\lambda)})$, and outputs b'. Then,

$$\left| \Pr[1 \leftarrow \mathcal{B}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{t(\lambda)}_{m(\lambda),s(\lambda)})] - \Pr_{k \leftarrow \{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{B}(1^\lambda, \hat{P}^1_{k,s(\lambda)}, ..., \hat{P}^{t(\lambda)}_{k,s(\lambda)})]\right| \tag{47}$$

$$= \Bigg| \Pr[1 \leftarrow \mathcal{A}(1^\lambda, \hat{Z}^1_{m(\lambda),s(\lambda)}, ..., \hat{Z}^{t(\lambda)}_{m(\lambda),s(\lambda)}, \hat{Z}^{t(\lambda)+1}_{m(\lambda),s(\lambda)})] \tag{48}$$

$$- \Pr_{k \leftarrow \{0,1\}^{m(\lambda)}}[1 \leftarrow \mathcal{A}(1^\lambda, \hat{P}^1_{k,s(\lambda)}, ..., \hat{P}^{t(\lambda)}_{k,s(\lambda)}, \hat{Z}^{t(\lambda)+1}_{m(\lambda),s(\lambda)})]\Bigg| \tag{49}$$

$$\geq \frac{1}{p(\lambda)} \tag{50}$$

for infinitely many $\lambda \in \mathbb{N}$. This contradicts Theorem 4.1 and therefore the distributions $\mathcal{D}'_1(\lambda)$ and $\mathcal{D}_1(\lambda)$ are computationally indistinguishable. Hence we complete the proof of IND-CPA security. □

Since IND-CPA secure QSKE implies OWSGs and EFI pairs [MY24, KT24, BJ24], we have the following corollary.

Corollary 5.2. *Suppose* $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{Q},\mathtt{Q},\mathtt{Q})$*-iO for classical circuits. Then there exist OWSGs and EFI pairs.*

5.2 Q-Obf, Q-Eval, and C-Encoding

Here, we study implications of $(\mathtt{Q},\mathtt{Q},\mathtt{C})$-iO, where both Obf and Eval are quantum algorithms, but an obfuscated encoding $\hat{C}$ is classical. This can be regarded as an iO in the QCCC model. We prove the following theorem:

Theorem 5.3. *Suppose* $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{Q},\mathtt{Q},\mathtt{C})$*-iO for classical circuits. Then there exists an IND-CPA secure QCCC SKE scheme.*

Proof of Theorem 5.3. The proof of this theorem is quite similar to Theorem 5.1. The only difference is that the output of Obf is a classical string rather than a quantum state. The proof of Theorem 5.1 heavily relies on Theorem 4.1 that is also valid for $(\mathtt{Q},\mathtt{Q},\mathtt{C})$-iO. Thus, we can easily extend the proof of Theorem 5.1 to the $(\mathtt{Q},\mathtt{Q},\mathtt{C})$-iO. □

Since IND-CPA secure QCCC SKE implies EV-OWPuzz, which in turn implies OWPuzz, OWSGs, and QEFID pairs, and EFI pairs [KT24, CGG24], we have the following corollary.

Corollary 5.4. *Suppose* $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{Q},\mathtt{Q},\mathtt{C})$*-iO for classical circuits. Then there exist EV-OWPuzz, OWPuzz, OWSGs, QEFID pairs, and EFI pairs.*

5.3 Q-Obf, C-Eval, and C-Encoding

Here, we study implications of $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO, where Obf is a quantum algorithm, Eval is a classical algorithm, and an obfuscated encoding $\hat{C}$ is classical. We prove the following theorem:

Theorem 5.5. *Suppose* $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{Q},\mathtt{C},\mathtt{C})$*-iO for classical circuits. Then there exists an IND-CPA secure QCCC PKE scheme.*

To prove this theorem, we first show that $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO can be modified to satisfy a stronger notion of correctness that holds even for fixed randomness.

Lemma 5.6. *Suppose that there exists* $(\mathtt{Q},\mathtt{C},\mathtt{C})$*-iO for classical circuits. Then there exists* $(\mathtt{Q},\mathtt{C},\mathtt{C})$*-iO for classical circuits that satisfies the fixed randomness correctness, defined below:*

- ***Fixed randomness correctness:*** *For any family* $\{C_\lambda\}_{\lambda\in\mathbb{N}}$ *of polynomial-size classical circuits of input length* n_λ*,*

$$\Pr_{\substack{\hat{C}_\lambda\leftarrow\mathsf{Obf}(1^\lambda,C_\lambda)\\ r\leftarrow\mathcal{R}_\lambda}}\left[\forall x\in\{0,1\}^{n_\lambda},\mathsf{Eval}(\hat{C}_\lambda,x;r)=C_\lambda(x)\right]\geq 1-\mathsf{negl}(\lambda) \tag{51}$$

where $\mathcal{R}_\lambda$ *denotes the randomness space of* $\mathsf{Eval}(\hat{C}_\lambda,x)$*, and* $\mathsf{Eval}(\hat{C}_\lambda,x;r)$ *denotes the execution with the fixed randomness* r*.* [10]

[10] We assume without loss of generality that the randomness space of $\mathsf{Eval}(\hat{C}_\lambda,x)$ only depends on λ and does not depend on C_λ and x.

Proof of Lemma 5.6. The correctness (as in Definition 3.1) implies that, for any family $\{C_\lambda\}_{\lambda\in\mathbb{N}}$ of polynomial-size classical circuits of input length n_λ and for any $x \in \{0,1\}^{n_\lambda}$,

$$\Pr_{\substack{\hat{C}_\lambda\leftarrow\mathsf{Obf}(1^\lambda,C_\lambda)\\ r\leftarrow\mathcal{R}_\lambda}}\left[\mathsf{Eval}(\hat{C}_\lambda,x;r)=C_\lambda(x)\right]\geq 1-\mathsf{negl}(\lambda)\geq 2/3. \tag{52}$$

Thus, if we modify Obf to run $M=O(\lambda+n_\lambda)$ times to output M independently generated obfuscated encodings of C_λ, and Eval to evaluate each of the M obfuscated encodings and take the majority result, we can ensure that it satisfies[11]

$$\Pr_{\substack{\hat{C}_\lambda\leftarrow\mathsf{Obf}(1^\lambda,C_\lambda)\\ r\leftarrow\mathcal{R}_\lambda}}\left[\mathsf{Eval}(\hat{C}_\lambda,x;r)=C_\lambda(x)\right]\geq 1-2^{-(\lambda+n_\lambda)}. \tag{53}$$

By taking the union bound over $x\in\{0,1\}^{n_\lambda}$ it implies Equation (51). Moreover, the modification of Obf and Eval does not affect the security. Thus, this completes the proof of Lemma 5.6. □

Then we prove Theorem 5.5.

Proof of Theorem 5.5. Let $(\mathsf{Obf},\mathsf{Eval})$ be a $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO for classical circuits. By Lemma 5.6, we can assume that it satisfies fixed randomness correctness without loss of generality. Let m and s be polynomials as in Theorem 4.1. Then we construct a QCCC PKE scheme $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ as follows:

- $\mathsf{Gen}(1^\lambda)$: Choose $k\leftarrow\{0,1\}^{m(\lambda)}$, and compute $\hat{P}_{k,s(\lambda)}\leftarrow\mathsf{Obf}(1^\lambda,P_{k,s(\lambda)})$, where we recall that $P_{k,s(\lambda)}$ denotes the canonical circuit of size $s(\lambda)$ for the point function with target k. Output the classical public key $\mathsf{pk}:=(1^\lambda,\hat{P}_{k,s(\lambda)})$ and the classical secret key $\mathsf{sk}:=k$.
- $\mathsf{Enc}(\mathsf{pk}=(1^\lambda,\hat{P}_{k,s(\lambda)}),\mathsf{msg})$: Choose $r\leftarrow\mathcal{R}_\lambda$ where $\mathcal{R}_\lambda$ is the randomness space of $\mathsf{Eval}(\hat{P}_{k,s(\lambda)},k')$ for $k'\in\{0,1\}^{m(\lambda)}$. Let $C[\hat{P}_{k,s(\lambda)},\mathsf{msg},r]$ be a classical circuit that takes $k'\in\{0,1\}^{m(\lambda)}$ as input and outputs msg if $\mathsf{Eval}(\hat{P}_{k,s(\lambda)},k';r)=1$ and 0 otherwise. Compute $\hat{C}[\hat{P}_{k,s(\lambda)},\mathsf{msg},r]\leftarrow\mathsf{Obf}(1^\lambda,C[\hat{P}_{k,s(\lambda)},\mathsf{msg},r])$. Output the ciphertext $\mathsf{ct}=\hat{C}[\hat{P}_{k,s(\lambda)},\mathsf{msg},r]$.
- $\mathsf{Dec}(\mathsf{sk}=k,\mathsf{ct}=\hat{C}[\hat{P}_{k,s(\lambda)},\mathsf{msg},r])$: Compute $\mathsf{msg}'\leftarrow\mathsf{Eval}(\hat{C}[\hat{P}_{k,s(\lambda)},\mathsf{msg},r],k)$ and output msg'.

By the correctness of $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO, the above scheme clearly satisfies the correctness of PKE.

Below, we prove that it satisfies the IND-CPA security. For any QPT adversary $\mathcal{A}$ and $b\in\{0,1\}$, we consider the following hybrid experiments:

$H_{1,b}$**:** This is the original IND-CPA security experiment. That is, it works as follows:

1. The challenger generates $\hat{P}_{k,s(\lambda)}\leftarrow\mathsf{Obf}(1^\lambda,P_{k,s(\lambda)})$ for $k\leftarrow\{0,1\}^{m(\lambda)}$, and sends $\mathsf{pk}=(1^\lambda,\hat{P}_{k,s(\lambda)})$ to $\mathcal{A}$.
2. $\mathcal{A}$ chooses $\mathsf{msg}_0,\mathsf{msg}_1\in(\{0,1\}^*)^2$ of the same length and sends them to the challenger.
3. The challenger generates $\hat{C}[\hat{P}_{k,s(\lambda)},\mathsf{msg}_b,r]\leftarrow\mathsf{Obf}(1^\lambda,C[\hat{P}_{k,s(\lambda)},\mathsf{msg}_b,r])$ for $r\leftarrow\mathcal{R}_\lambda$, and sends $\mathsf{ct}_b=\hat{C}[\hat{P}_{k,s(\lambda)},\mathsf{msg}_b,r]$ to $\mathcal{A}$.
4. $\mathcal{A}$ outputs b', which is the output of the experiment.

[11] Due to the modifications to Eval, $\mathcal{R}_\lambda$ is also updated accordingly; it now consists of M-tuples of the randomness used in the original Eval.

Our goal is to prove that $|\Pr[H_{1,0}=1]-\Pr[H_{1,1}=1]|\leq \mathsf{negl}(\lambda)$.

$H_{2,b}$: This is identical to $H_{1,b}$ except that pk is set to be $\hat{Z}_{m(\lambda),s(\lambda)}\leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$, and consequently ct_b is set to be $\hat{C}[\hat{Z}_{m(\lambda),s(\lambda)},\mathsf{msg}_b,r]\leftarrow \mathsf{Obf}(1^\lambda, C[\hat{Z}_{m(\lambda),s(\lambda)},\mathsf{msg}_b,r])$ where we recall that $Z_{m(\lambda),s(\lambda)}$ denotes the canonical zero-function on $m(\lambda)$-bit inputs of size $s(\lambda)$. By a straightforward reduction to Theorem 4.1 for the case of $\ell=1$, we have $|\Pr[H_{1,b}=1]-\Pr[H_{2,b}=1]|\leq \mathsf{negl}(\lambda)$ for $b\in\{0,1\}$.

$H_{3,b}$: This is identical to $H_{2,b}$ except that ct_b is set to be $\hat{Z}_{m(\lambda),s'(\lambda)}\leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s'(\lambda)})$, where $s'(\lambda)$ is the size of $C[\hat{Z}_{m(\lambda),s(\lambda)},\mathsf{msg}_b,r]$. (We assume without loss of generality that the size only depends on λ by padding.) By the fixed randomness correctness of the $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO, $C[\hat{Z}_{m(\lambda),s(\lambda)},\mathsf{msg}_b,r]$ is functionally equivalent to $Z_{m(\lambda),s'(\lambda)}$ with overwhelming probability over the choice of r. Thus, by a straightforward reduction to the security of $(\mathtt{Q},\mathtt{C},\mathtt{C})$-iO, we have $|\Pr[H_{2,b}=1]-\Pr[H_{3,b}=1]|\leq \mathsf{negl}(\lambda)$ for $b\in\{0,1\}$.

Moreover, in $H_{3,b}$, no information of b is given to $\mathcal{A}$, and thus $\Pr[H_{3,0}=1]=\Pr[H_{3,1}=1]$.

Combining the above, we obtain $|\Pr[H_{1,0}=1]-\Pr[H_{1,1}=1]|\leq \mathsf{negl}(\lambda)$. This completes the proof of IND-CPA security. □

Since IND-CPA secure QCCC PKE implies EV-OWPuzz, which in turn implies OWPuzz, OWSGs, QEFID pairs, and EFI pairs [KT24, CGG24], we have the following corollary.

Corollary 5.7. *Suppose* $\mathsf{NP}\not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{Q},\mathtt{C},\mathtt{C})$*-iO for classical circuits. Then there exist EV-OWPuzz, OWPuzz, OWSGs, QEFID pairs, and EFI pairs.*

5.4 C-Obf, Q-Eval, and C-Encoding

Here, we study implications of $(\mathtt{C},\mathtt{Q},\mathtt{C})$-iO, where Obf is a classical algorithm, Eval is a quantum algorithms, and an obfuscated encoding $\hat{C}$ is classical. We prove the following theorem:

Theorem 5.8. *Suppose* $\mathsf{NP}\not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{C},\mathtt{Q},\mathtt{C})$*-iO for classical circuits. Then there exist OWFs and an IND-CPA secure QCCC PKE scheme.*

To show Theorem 5.8, we rely on the following lemma:

Lemma 5.9 ([Gol90]). *The following two conditions are equivalent:*

- *There exist OWFs.*
- *There exist pairs of classical-polynomial-time-samplable distributions that are statistically far but computationally indistinguishable.*

Lemma 5.10 ([SW14]). *If there exist* $(\mathtt{C},\mathtt{C},\mathtt{C})$*-iO for classical circuits and OWFs, then there exist IND-CPA secure PKE schemes.*

In the construction of [SW14], the encryption algorithm runs Eval. Thus, by adapting their construction to $(\mathtt{C},\mathtt{Q},\mathtt{C})$-iO in which only Eval is quantum algorithm but both of Obf and the obfuscated encoding are classical, we obtain QCCC PKE scheme in which only the encryption algorithm is quantum.

Corollary 5.11. *If there exist* $(\mathtt{C},\mathtt{Q},\mathtt{C})$*-iO for classical circuits and OWFs, then there exist IND-CPA secure QCCC PKE schemes.*

Now we are ready to prove Theorem 5.8.

Proof of Theorem 5.8. By Lemma 5.9 and Corollary 5.11, it suffices to construct a pair of classical-polynomial-time-samplable distributions that are statistically far but computationally indistinguishable. By applying Theorem 4.1 for $\ell = 1$, there exist classical-polynomial-time-computable polynomials m and s such that the following two distributions are computationally indistinguishable:

- $\mathcal{D}_0(\lambda)$: Sample $k \leftarrow \{0,1\}^{m(\lambda)}$, run $\hat{P}_{k,s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, P_{k,s(\lambda)})$, and output $\hat{P}_{k,s(\lambda)}$.
- $\mathcal{D}_1(\lambda)$: Run $\hat{Z}_{m(\lambda),s(\lambda)} \leftarrow \mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ and output $\hat{Z}_{m(\lambda),s(\lambda)}$.

Moreover both of $\mathcal{D}_0(\lambda)$ and $\mathcal{D}_1(\lambda)$ are classical-polynomial-time-samplable because m and s are classical-polynomial-time-computable and Obf is a PPT algorithm. Thus, to complete the proof, we show that $\mathcal{D}_0(\lambda)$ and $\mathcal{D}_1(\lambda)$ are statistically far. To show this, we construct an unbounded-time distinguisher $\mathcal{A}$ that distinguishes $\mathcal{D}_0(\lambda)$ and $\mathcal{D}_1(\lambda)$.

$\mathcal{A}(\hat{C})$**:** Upon receiving an obfuscated encoding $\hat{C}$, it computes $p_{k'} := \Pr[\mathsf{Eval}(\hat{C}, k') = 1]$ for all $k' \in \{0,1\}^{m(\lambda)}$. If there is $k' \in \{0,1\}^{m(\lambda)}$ such that $p_{k'} \geq 1/2$, it outputs 0, otherwise it outputs 1.

If $\hat{C} = \hat{P}_{k,s(\lambda)} \leftarrow \mathcal{D}_0(\lambda)$, the correctness of the iO implies that $\Pr[p_k \geq 1 - \mathsf{negl}(\lambda)] \geq 1 - \mathsf{negl}(\lambda)$ where the probability is taken over the randomness in the sampling procedure of $\mathcal{D}_0$. Thus, we have $\Pr[\mathcal{A}(\hat{C}) = 0] \geq 1 - \mathsf{negl}(\lambda)$. On the other hand, if $\hat{C} = \hat{Z}_{m(\lambda),s(\lambda)} \leftarrow \mathcal{D}_1(\lambda)$, the correctness of the iO implies that $\Pr[\forall k' \in \{0,1\}^{m(\lambda)}, p_{k'} \leq \mathsf{negl}(\lambda)] \geq 1 - \mathsf{negl}(\lambda)$ where the probability is taken over the randomness in the sampling procedure of $\mathcal{D}_1$. Thus, we have $\Pr[\mathcal{A}(\hat{C}) = 1] \geq 1 - \mathsf{negl}(\lambda)$. Therefore, $\mathcal{D}_0(\lambda)$ and $\mathcal{D}_1(\lambda)$ are statistically far and we complete the proof. □

Remark 5.12. One might think that Theorem 5.8 directly follows from an adaptation of the technique of [KMN+14] since we can derandomize Obf when it is classical. In fact, this is true in the perfectly correct case. On the other hand, this does not work in the imperfect case (as in Definition 3.1) since their proof in the imperfect setting involves obfuscation of an obfuscated circuit, but this is not possible in our setting since Eval is a quantum algorithm and thus cannot be obfuscated by iO for classical circuits as is considered in this paper.

Note that $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ and the existence of $(\mathtt{C}, \mathtt{Q}, \mathtt{C})$-iO for classical circuits imply all Microcrypt primitives since they imply the existence of OWFs (and therefore PRUs [MH24]).

5.5 C-Obf, C-Eval, and C-Encoding

Here, we study implications of $(\mathtt{C}, \mathtt{C}, \mathtt{C})$-iO, where both Obf and Eval are classical algorithms and an obfuscated encoding $\hat{C}$ is classical. This is oftend referred to as post-quantum iO. We prove the following theorem:

Theorem 5.13. *Suppose* $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ *and there exists* $(\mathtt{C}, \mathtt{C}, \mathtt{C})$*-iO for classical circuits. Then there exist OWFs and an IND-CPA secure PKE scheme.*

Proof of Theorem 5.13. This proof is essentially same as the proof of Theorem 5.8. By Lemmata 5.9 and 5.10, it suffices to construct a pair of classical-polynomial-time-samplable distributions that are statistically far but computationally indistinguishable. By applying Theorem 4.1 for $\ell = 1$, there exist classical-polynomial-time-computable polynomials m and s such that the distributions over $\mathsf{Obf}(1^\lambda, P_{k,s(\lambda)})$ and $\mathsf{Obf}(1^\lambda, Z_{m(\lambda),s(\lambda)})$ are computationally indistinguishable, where $k \leftarrow \{0,1\}^{m(\lambda)}$. We can show that they are classical-polynomial-time-samplable and statistically far by adapting the same argument used in the proof of Theorem 5.8. □

Remark 5.14. We could also prove Theorem 5.13 by a straightforward adaptation of [KMN+14]. On the other hand, an advantage of our approach is that it in fact only needs iO for 3CNF formulas rather than general classical circuits. Constructing OWFs from imperfect iO for 3CNF formulas was an open problem left by [KMN+14], and our alternative proof resolves this open problem, though the open problem itself was also recently resolved (in a stronger form) in [HN24, LMP24] by completely different techniques.

Note that $\mathsf{NP} \not\subseteq \mathsf{i.o.BQP}$ and the existence of $(\mathsf{C}, \mathsf{C}, \mathsf{C})$-iO for classical circuits imply all Microcrypt primitives since they imply the existence of OWFs.

Acknowledgements. We thank Minki Hhan, Giulio Malavolta, and Kabir Tomer for insightful discussions on the initial ideas of this work during QIP 2024. TM is supported by JST CREST JPMJCR23I3, JST Moonshot R&D JPMJMS2061-5-1-1, JST FOREST, MEXT QLEAP, the Grant-in Aid for Transformative Research Areas (A) 21H05183, and the Grant-in-Aid for Scientific Research (A) No.22H00522. YS is supported by JST SPRING, Grant Number JPMJSP2110.

References

[ABDS21] Gorjan Alagic, Zvika Brakerski, Yfke Dulek, and Christian Schaffner. Impossibility of quantum virtual black-box obfuscation of classical circuits. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 497–525, Virtual Event, August 2021. Springer, Cham. (Cited on page 1, 7.)

[ABG+13] Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. (Cited on page 5, 11.)

[AF16] Gorjan Alagic and Bill Fefferman. On quantum obfuscation, 2016. (Cited on page 1, 7.)

[BBF16] Zvika Brakerski, Christina Brzuska, and Nils Fleischhacker. On statistically secure obfuscation with approximate correctness. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 551–578. Springer, Berlin, Heidelberg, August 2016. (Cited on page 7.)

[BBV24] James Bartusek, Zvika Brakerski, and Vinod Vaikuntanathan. Quantum state obfuscation from classical oracles. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 1009–1017. ACM Press, June 2024. (Cited on page 7.)

[BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 52–73. Springer, Berlin, Heidelberg, February 2014. (Cited on page 5, 11, 12.)

[BCQ23] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. ITCS 2023, 2023. (Cited on page 2, 5, 7, 9.)

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate iO from homomorphic encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 79–109. Springer, Cham, May 2020. (Cited on page 1.)

[BGI+12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *Journal of the ACM*, 59(2):6:1–6:48, 2012. (Cited on page 1, 5, 11.)

[BGK+14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 221–238. Springer, Berlin, Heidelberg, May 2014. (Cited on page 1.)

[BJ24] Rishabh Batra and Rahul Jain. Commitments are equivalent to statistically-verifiable one-way state generators. In *65th FOCS*, pages 1178–1192. IEEE Computer Society Press, October 2024. (Cited on page 18.)

[BK21] Anne Broadbent and Raza Ali Kazmi. Constructions for quantum indistinguishability obfuscation. In Patrick Longa and Carla Ràfols, editors, *Progress in Cryptology - LATINCRYPT 2021 - 7th International Conference on Cryptology and Information Security in Latin America, Bogotá, Colombia, October 6-8, 2021, Proceedings*, volume 12912 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2021. (Cited on page 1, 7.)

[BKNY23] James Bartusek, Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. Obfuscation of pseudo-deterministic quantum circuits. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1567–1578. ACM Press, June 2023. (Cited on page 1, 7.)

[BM22] James Bartusek and Giulio Malavolta. Indistinguishability obfuscation of null quantum circuits and applications. In Mark Braverman, editor, *ITCS 2022*, volume 215, pages 15:1–15:13. LIPIcs, January / February 2022. (Cited on page 1, 7.)

[BMSZ16] Saikrishna Badrinarayanan, Eric Miles, Amit Sahai, and Mark Zhandry. Post-zeroizing obfuscation: New mathematical tools, and the case of evasive circuits. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 764–791. Springer, Berlin, Heidelberg, May 2016. (Cited on page 1.)

[BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 1–25. Springer, Berlin, Heidelberg, February 2014. (Cited on page 1.)

[BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Berlin, Heidelberg, August 2014. (Cited on page 1.)

[CCH+19] Jung Hee Cheon, Wonhee Cho, Minki Hhan, Jiseung Kim, and Changmin Lee. Statistical zeroizing attack: Cryptanalysis of candidates of BP obfuscation over GGH15 multilinear map. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 253–283. Springer, Cham, August 2019. (Cited on page 1.)

[CG24] Andrea Coladangelo and Sam Gunn. How to use quantum indistinguishability obfuscation. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 1003–1008. ACM Press, June 2024. (Cited on page 1, 7.)

[CGG24] Kai-Min Chung, Eli Goldin, and Matthew Gray. On central primitives for quantum cryptography with classical communication. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 215–248. Springer, Cham, August 2024. (Cited on page 2, 9, 18, 20.)

[CGH17] Yilei Chen, Craig Gentry, and Shai Halevi. Cryptanalyses of candidate branching program obfuscators. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 278–307. Springer, Cham, April / May 2017. (Cited on page 1.)

[CHKL18] Jung Hee Cheon, Minki Hhan, Jiseung Kim, and Changmin Lee. Cryptanalyses of branching program obfuscations over GGH13 multilinear map from the NTRU problem. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 184–210. Springer, Cham, August 2018. (Cited on page 1.)

[CHL+15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12. Springer, Berlin, Heidelberg, April 2015. (Cited on page 1.)

[CHN+16] Aloni Cohen, Justin Holmgren, Ryo Nishimaki, Vinod Vaikuntanathan, and Daniel Wichs. Watermarking cryptographic capabilities. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 1115–1127. ACM Press, June 2016. (Cited on page 1.)

[CLP15] Kai-Min Chung, Huijia Lin, and Rafael Pass. Constant-round concurrent zero-knowledge from indistinguishability obfuscation. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 287–307. Springer, Berlin, Heidelberg, August 2015. (Cited on page 1.)

[CLT13] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 476–493. Springer, Berlin, Heidelberg, August 2013. (Cited on page 1.)

[CLT15] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 267–286. Springer, Berlin, Heidelberg, August 2015. (Cited on page 1.)

[GGG+14] Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 578–602. Springer, Berlin, Heidelberg, May 2014. (Cited on page 1.)

[GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Berlin, Heidelberg, March 2015. (Cited on page 1.)

[GGH+16] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM J. Comput.*, 45(3):882–929, 2016. (Cited on page 1.)

[GMM+16] Sanjam Garg, Eric Miles, Pratyay Mukherjee, Amit Sahai, Akshayaram Srinivasan, and Mark Zhandry. Secure obfuscation in a weak multilinear map model. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 241–268. Springer, Berlin, Heidelberg, October / November 2016. (Cited on page 1.)

[Gol90] Oded Goldreich. A note on computational indistinguishability. Information Processing Letters 34.6 (1990), pp.277–281., 1990. (Cited on page 5, 20.)

[HJ16] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 537–565. Springer, Berlin, Heidelberg, May 2016. (Cited on page 1.)

[HN24] Shuichi Hirahara and Mikito Nanashima. One-way functions and zero knowledge. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 1731–1738. ACM Press, June 2024. (Cited on page 3, 7, 22.)

[HSW13] Susan Hohenberger, Amit Sahai, and Brent Waters. Full domain hash from (leveled) multilinear maps and identity-based aggregate signatures. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 494–512. Springer, Berlin, Heidelberg, August 2013. (Cited on page 1.)

[HT25] Mi-Ying (Miryam) Huang and Er-Cheng Tang. Obfuscation of unitary quantum programs. Cryptology ePrint Archive, Paper 2025/891, 2025. (Cited on page 1, 7.)

[IL25] Rahul Ilango and Alex Lombardi. Cryptography meets worst-case complexity: Optimal security and more from iO and worst-case assumptions. Cryptology ePrint Archive, Paper 2025/1087, 2025. (Cited on page 7.)

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Cham, August 2018. (Cited on page 2, 9.)

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 60–73. ACM Press, June 2021. (Cited on page 1.)

[KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. Indistinguishability obfuscation for Turing machines with unbounded memory. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 419–428. ACM Press, June 2015. (Cited on page 1.)

[KMN+14] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. In *55th FOCS*, pages 374–383. IEEE Computer Society Press, October 2014. (Cited on page 1, 3, 7, 11, 21, 22.)

[KNY14] Ilan Komargodski, Moni Naor, and Eylon Yogev. Secret-sharing for NP. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 254–273. Springer, Berlin, Heidelberg, December 2014. (Cited on page 1.)

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1589–1602. ACM Press, June 2023. (Cited on page 2.)

[KQT24] William Kretschmer, Luowen Qian, and Avishay Tal. Quantum-computable one-way functions without one-way functions, 2024. (Cited on page 2.)

[Kre21] W. Kretschmer. Quantum pseudorandomness and classical complexity. *TQC 2021*, 2021. (Cited on page 2.)

[KT24] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 968–978. ACM Press, June 2024. (Cited on page 2, 9, 18, 20.)

[KT25] Dakshita Khurana and Kabir Tomer. Founding quantum cryptography on quantum advantage, or, towards cryptography from #p hardness. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 178–188. ACM, 2025. (Cited on page 7.)

[LMP24] Yanyi Liu, Noam Mazor, and Rafael Pass. A note on zero-knowledge for NP and one-way functions. Cryptology ePrint Archive, Paper 2024/800, 2024. (Cited on page 3, 7, 22.)

[LMW24] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 979–990. ACM Press, June 2024. (Cited on page 2.)

[MH24] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. Cryptology ePrint Archive, Paper 2024/1652, 2024. (Cited on page 2, 21.)

[MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 629–658. Springer, Berlin, Heidelberg, August 2016. (Cited on page 1.)

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Cham, August 2022. (Cited on page 2.)

[MY24] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. In Frédéric Magniez and Alex Bredariol Grilo, editors, *19th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2024, September 9-13, 2024, Okinawa, Japan*, volume 310 of *LIPIcs*, pages 4:1–4:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. (Cited on page 2, 8, 9, 18.)

[SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014. (Cited on page 1, 3, 5, 20.)

[VV86] Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986. (Cited on page 3, 10.)

[Zha23] Mark Zhandry. Tracing quantum state distinguishers via backtracking. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 3–36. Springer, Cham, August 2023. (Cited on page 12.)