

Impossibility of Quantum Virtual Black-Box Obfuscation of Classical Circuits

Gorjan Alagic^{1,2}, Zvika Brakerski³, Yfke Dulek^{4,6}, and Christian Schaffner^{5,6}

¹ Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD, United States

² National Institute of Standards and Technology, Gaithersburg, MD, United States
`galagic@gmail.com`

³ Weizmann Institute of Science, Rehovot, Israel
`zvika.brakerski@weizmann.ac.il`

⁴ Centrum Wiskunde & Informatica, Amsterdam, the Netherlands
`dulek@cw.nl`

⁵ University of Amsterdam, Amsterdam, the Netherlands
`c.schaffner@uva.nl`

⁶ QuSoft, Amsterdam, the Netherlands

Abstract. Virtual black-box obfuscation is a strong cryptographic primitive: it encrypts a circuit while maintaining its full input/output functionality. A remarkable result by Barak et al. (Crypto 2001) shows that a general obfuscator that obfuscates classical circuits into classical circuits cannot exist. A promising direction that circumvents this impossibility result is to obfuscate classical circuits into quantum states, which would potentially be better capable of hiding information about the obfuscated circuit. We show that, under the assumption that Learning With Errors (LWE) is hard for quantum computers, this quantum variant of virtual black-box obfuscation of classical circuits is generally impossible. On the way, we show that under the presence of dependent classical auxiliary input, even the small class of classical point functions cannot be quantum virtual black-box obfuscated.

1 Introduction

The obfuscation of a circuit is an object, typically another circuit, that allows a user to evaluate the functionality of the original circuit without learning any additional information about the structure of the circuit. Obfuscation is useful for publishing software without revealing the code, but it also has more fundamental applications in cryptography. For example, the strongest notion called *virtual black-box* obfuscation can transform any private-key encryption scheme into a public-key scheme, and transform public-key schemes into fully-homomorphic schemes. Unfortunately, this notion turns out to be impossible for general circuits [BGI⁺01] – at least, if we require the obfuscation of a circuit to be a circuit itself.

The impossibility result from [BGI⁺01] leaves open an intriguing possibility: what if the obfuscation of a (classical) circuit is allowed to be a *quantum state*?

Could a quantum state contain all the information about a functionality, allowing a user to produce correct outputs, without revealing all that information? This possibility seems hopeful, due to the unrevealing nature of quantum states, but the extent to which no-cloning affects cryptography is not fully understood. The question whether quantum obfuscation of classical circuits is possible appeared in Aaronson’s list of semi-grand challenges in quantum computing [Aar05]. At least one candidate construction was proposed [Chr14], but subsequently broken. An attempt to prove impossibility [AF16] also encountered several obstacles (discussed below). As a result, the question remained far from settled.

In this work, we answer the question by showing that virtual-black-box obfuscating classical circuits into quantum states is not possible. We adopt ideas from the classical proof [BGI⁺01], but while using a similar proof outline is natural, the actual execution is non-trivial: our work requires tools such as quantum fully homomorphic encryption [Mah18] and compute-and-compare obfuscation [WZ17,GKW17] that were only developed in recent years.

1.1 Related work

Barak et al. defined the obfuscating property of virtual black-box (vbb) obfuscators as follows: any information that an adversary can learn about a circuit from its obfuscation can also be learned by a simulator that does not have access to the obfuscation, but only to an oracle for the circuit’s functionality [BGI⁺01]. In this definition, the crucial difference between the adversary and the simulator is that the adversary has access to a short representation of the circuit (namely, the obfuscation), whereas the simulator only has access to an input/output interface that implements the functionality. Some circuit classes allow the adversary to exploit this difference by using the obfuscation as an input value to the circuit itself. Those circuit classes are unobfuscatable in the vbb sense, rendering vbb obfuscation impossible for the general class of circuits in \mathbf{P} [BGI⁺01].

In more detail, the impossibility proof in [BGI⁺01] relies on point functions, which output zero everywhere except at a single input value α , where they output a string β . The circuits in the unobfuscatable class can, depending on the input, do all of the following: (1) apply that point function, (2) return an encryption of α , (3) homomorphically evaluate a gate, or (4) check whether a ciphertext decrypts to β . An adversary holding the obfuscation is able to divide it into single gates, and can use those to homomorphically evaluate option (1), thereby converting a ciphertext for α into a ciphertext for β . That way, the adversary can tell whether he is holding an obfuscation with a point function from α to β , or one with the all-zero function. (In the second case, the homomorphic evaluation would yield a ciphertext for zero, rather than one for β .) A simulator, only having access to the input/output behavior, cannot perform the homomorphic evaluation, because it cannot divide the functionality into single gates.

The above construction rules out the existence of an obfuscator that maps classical circuits to classical circuits. It leaves open the possibility of an obfuscator that maps classical circuits to *quantum states*: such a quantum state,

together with a fixed public ‘interpreter map’, could be used to evaluate the obfuscated circuit. The possibility of quantum obfuscation was the object of study for Alagic and Fefferman [AF16], who attempted to port the impossibility proof from [BGI⁺01] to the quantum setting. In doing so, they encountered two issues:

Homomorphic evaluation. The interpreter map, that runs the obfuscation state on a chosen input, is a quantum map. It will likely have quantum states as intermediate states of the computation, so in order to homomorphically run the point function, one needs the ability to evaluate quantum gates on quantum ciphertexts. The unobfuscatable circuit class will thus need to contain quantum circuits to perform homomorphic evaluation steps.

Reusability. In the construction from [BGI⁺01], the obfuscated circuit needs to be used multiple times: for example, each homomorphic gate evaluation requires a separate call to the obfuscated circuit. If the obfuscation is a (classical or quantum) circuit, this poses no problem, but if it is a quantum *state*, multiple uses are not guaranteed.

These two issues limit the extent of the impossibility results in [AF16]: they show that it is impossible to vbb obfuscate *quantum* circuits into *reusable* obfuscated states (e.g., quantum circuits).

After it became clear [BGI⁺01] that obfuscating all classical circuits is impossible, efforts were made to construct obfuscators for smaller, but still non-trivial, classes of circuits. Successful constructions have been found for several classes of evasive functions, such as point functions [Wee05,CD08] and compute-and-compare functions [WZ17,GKW17]. Currently, no quantum obfuscators are known for circuit classes that cannot be classically obfuscated.

1.2 Our contributions

We strengthen the impossibility of virtual-black-box obfuscation of classical circuits by showing that classical circuits cannot be obfuscated into quantum states. We assume the existence of classical-client quantum fully homomorphic encryption and classical obfuscation of compute-and-compare functions. Both of these can be constructed from the learning-with-errors (LWE) assumption [Mah18,Bra18,WZ17,GKW17]. The compute-and-compare construction requires the strongest assumption in terms of the LWE parameters.

Theorem (informal). *If LWE is hard for quantum algorithms, then it is impossible to quantum vbb obfuscate the class of polynomial-size classical circuits (even with non-negligible correctness and security error, and even if the obfuscation procedure is inefficient).*

Our result uses the same proof strategy as in [BGI⁺01] and [AF16], overcoming the two main issues described above as follows:

Homomorphic evaluation. The constructions in [BGI⁺01] and [AF16] rely on the obfuscator to implement the homomorphic evaluations, by obfuscating

the functionality “decrypt, then apply a gate, then re-encrypt”. However, by now, we know how to build quantum fully-homomorphic encryption schemes directly [Mah18,Bra18], based on the learning-with-errors (LWE) assumption. Thus, in our construction, we can remove the homomorphic gate evaluation from the obfuscated circuits: the adversary can do the homomorphic evaluation of the point function herself, using a quantum fully-homomorphic encryption scheme. With the homomorphic evaluation removed from it, the class of circuits that we prove impossible to obfuscate can remain classical. This solution introduces a slight complication: part of the functionality of the circuit we construct is now to return the public evaluation key. However, unless one is willing to make an assumption on the circular security of the homomorphic encryption, the size of this key (and therefore the size of the circuit) scales with the size of the circuit that needs to be homomorphically evaluated. To get rid of this inconvenient dependence, our unobfuscatable circuit returns the public key in small, individual blocks that can be independently computed. We argue that any classical-key quantum fully-homomorphic encryption scheme has public keys that can be decomposed in this way.

Reusability. The circuits that we consider are classical and deterministic. Therefore, if the interpreter map is run on an obfuscation state ρ for a circuit C , plus a classical input x , then by correctness, the result is (close to) a computational-basis state $|C(x)\rangle$. This output can be copied out to a separate wire without disturbing the state, and the interpreter map can be reversed, recovering the obfuscation ρ to be used again. If the interpreter map is not unitary, then it can be run coherently (i.e., keeping purification registers around instead of measuring wires), and this coherent version can be reversed as long as the purification registers are not measured.

At one point in our proof, we will need to run the interpreter map homomorphically on (an encryption of) ρ and x . This may result in a superposition of different ciphertexts for $C(x)$, which cannot cleanly be copied out to a separate wire without entangling that wire with the output. Thus, recovering ρ is not necessarily possible after the homomorphic-evaluation step.

We circumvent this problem by making sure that the homomorphic evaluation occurs last, so that ρ is not needed anymore afterwards. This reordering is achieved by classically obfuscating the part of the circuit that checks whether a ciphertext decrypts to the value β . That way, this functionality becomes a constant output value that a user can request and store before performing the homomorphic evaluation, and use afterwards. To obfuscate the decryption check, we use a classical vbb obfuscator for compute-and-compare functions, which relies on a variant of the LWE assumption [WZ17,GKW17]. These vbb obfuscators have previously been successfully applied in a similar way, in the context of quantum extraction [AL20a,BS20].

Our impossibility result compares to the classical impossibility result from Barak et al. [BGI⁺01] as follows. First, as mentioned, we extend the realm of impossible obfuscators to include obfuscators that produce a quantum state, rather than

a classical circuit. Second, the impossibility result from [BGI⁺01] is unconditional, whereas we require the (standard) assumption that learning-with-errors is hard for quantum adversaries. It may be possible to relax this assumption if ρ can be recovered after the homomorphic evaluation, see Section 1.3 below. Third, the class of classical circuits that cannot be obfuscated is slightly different: in our work, it does not have the homomorphic-evaluation functionality built into it, and is therefore arguably simpler, strengthening the impossibility result. However, we stress that in both works, the unobfuscatable circuit class itself is somewhat contrived: the main implication is that its superclass \mathbf{P} is unobfuscatable.

As an intermediate result, we show that it is impossible to vbb obfuscate even just the class of classical multi-bit-output point functions into a quantum state, if the adversary and simulator have access to auxiliary classical information that contains an encryption of the non-zero input value α and a vbb obfuscation of a function depending on the secret key for that encryption.

Theorem (informal). *If LWE is hard for quantum algorithms, then it is impossible to quantum vbb obfuscate multi-bit-output point functions and the all-zero function under the presence of classical dependent auxiliary information (even with non-negligible soundness and security error).*

At first glance, that may seem to contradict constructions in [WZ17,GKW17], where vbb obfuscation for point functions is constructed, even in the presence of dependent auxiliary information. The crucial difference is that the constructions in [WZ17,GKW17] only allow a limited dependency of the auxiliary information, whereas in our impossibility proof, the dependence is slightly stronger. This subtle difference seems to indicate that the gap between possibility and impossibility of vbb obfuscation is closing.

Comparison with Concurrent Work. Independently of this work, Ananth and La Placa [AL20b] have concurrently shown the general impossibility of quantum copy-protection, thereby also ruling out quantum obfuscation of classical circuits. Their techniques are very similar to ours, but their adversary allows to completely de-obfuscate the program given non-black-box access (a property that we did not attempt to achieve). They also present some positive results in their work in the context of software protection. Their result requires an additional assumption compared to ours. Specifically, in addition to LWE being quantum-secure, they also require that the underlying homomorphic-encryption scheme is circularly secure. We avoid circularity by introducing a notion of decomposable public keys for homomorphic encryption. Our technique could be used to remove the circularity assumption from the copy-protection impossibility result [AL20b] as well.

1.3 Open questions

The strongest assumption in our work is the existence of the classical vbb obfuscator for compute-and-compare functions, which relies on a variant of LWE .

It is necessary because the QFHE evaluation may destroy the obfuscation state when the superposition of output ciphertexts is measured. However, it is not clear if this measurement actually destroys any information on the *plaintext* level, since the plaintext value is deterministic. Thus, it may be possible to recover the (plaintext) obfuscation state after the QFHE evaluation. In that case, it is not necessary to classically obfuscate the compute-and-compare function: it can simply be part of the quantum-obfuscated functionality.

Other open questions are about possibilities rather than impossibilities. What circuit classes *can* be vbb obfuscated into quantum states? Is quantum vbb obfuscation stronger than classical vbb obfuscation, in the sense that it can obfuscate circuit classes that classical vbb cannot? Also, the weaker notion of indistinguishability obfuscation (iO) (also introduced in [BGI⁺01]) is not affected by our impossibility result: it may still be possible to classically or quantumly iO obfuscate classical functionalities. Could such a construction be lifted into the quantum realm, so that we can (quantum) iO obfuscate [BK20] quantum functionalities?

1.4 Structure of this work

In Section 2, we give preliminary definitions of the relevant concepts for this work: (classical and quantum) obfuscation, quantum fully homomorphic encryption, and compute-and-compare functions. We also describe how the input of an (almost) deterministic quantum circuit can be recovered. In Section 4, we prove impossibility of quantum obfuscation of point functions under dependent auxiliary input. Building on the concepts in that section, Section 5 proves our main result, impossibility of quantum obfuscation of classical circuits without any auxiliary input.

2 Preliminaries

2.1 Notation

PPT stands for probabilistic polynomial-time algorithm, and QPT stands for quantum polynomial-time algorithm. If a classical or quantum algorithm A has oracle access to a classical function f , we write A^f . If A has access to multiple oracles with separate input/output interfaces, we write, e.g., $A^{f,g}$. Since our oracles will model or emulate evaluation of a known circuit, our quantum algorithms will always have superposition access to a classical oracle f . This amounts to oracle access to the unitary map $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus f(x)\rangle$ where \oplus is the bit-wise XOR operation.

Let $\text{poly}(x)$ denote an unspecified polynomial $p(x)$. Similarly, let $\text{negl}(x)$ denote an unspecified negligible function $\mu(x)$, i.e., for all constants $c \in \mathbb{N}$ there exists an $x_0 \in \mathbb{R}$ such that for all $x > x_0$, $|\mu(x)| < x^{-c}$. Let $\mathbf{Z}_n : \{0, 1\}^n \rightarrow \{0^n\}$ denote the all-zero function on n input bits: $\mathbf{Z}_n(x) = 0^n$ for all x .

If D is a distribution, we write $x \leftarrow D$ to signify that x is sampled according to D . For a finite set S , we write $x \leftarrow_R S$ to signify that x is sampled uniformly at

random from the set S . Two distribution ensembles $\{D_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{D'_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable, written $D_\lambda \stackrel{c}{\approx} D'_\lambda$, if no poly-time algorithm can distinguish between a sample from one distribution or the other, i.e., for all PPT A ,

$$\left| \Pr_{x \leftarrow D_\lambda} [A(x) = 1] - \Pr_{y \leftarrow D'_\lambda} [A(y) = 1] \right| \leq \text{negl}(\lambda).$$

We sometimes write $x \stackrel{c}{\approx} y$ if it is clear from which distributions x and y are sampled. If not even a QPT algorithm can distinguish them, the distributions are quantum-computationally indistinguishable.

A pure quantum state is written $|\psi\rangle$ or $|\varphi\rangle$, and a mixed quantum state is usually denoted by ρ or σ . As a special case, a computational-basis state is written $|x\rangle$ for some classical string $x \in \{0, 1\}^*$. We sometimes abuse notation and give a classical input x to a quantum algorithm A , writing $A(x)$: in that case, the algorithm A is actually given $|x\rangle$ as input.

X and Z denote the bit-flip gate and phase-flip gate, respectively. If we write X^a for some $a \in \{0, 1\}$, we mean that the gate X is applied if $a = 1$; otherwise, identity is applied.

Finally, for a mixed state ρ , let $\|\rho\|_{\text{tr}} := \text{Tr}(\sqrt{\rho^\dagger \rho})$ denote the trace norm. The trace distance $\frac{1}{2} \|\rho - \sigma\|_{\text{tr}}$ is a measure for how different two mixed states ρ and σ are.

2.2 Classical and Quantum Virtual-Black-Box Obfuscation

In this work we consider so-called *circuit* obfuscators: the functionalities to be hidden are represented by circuits. A virtual-black-box circuit obfuscator hides the functionality in such a way that the obfuscation looks like a “black box”: the only way to get information about its functionality is to evaluate it on an input and observe the output.

Definition 2.1 ([BGI⁺01, Definition 2.2]). *A classical virtual black-box obfuscator for the circuit class \mathfrak{F} is a probabilistic algorithm \mathcal{O} such that*

1. (*polynomial slowdown*) For every circuit $C \in \mathfrak{F}$, $|\mathcal{O}(C)| = \text{poly}(|C|)$;
2. (*functional equivalence*) For every circuit $C \in \mathfrak{F}$, the string $\mathcal{O}(C)$ describes a circuit that computes the same function as C ;
3. (*virtual black-box*) For any PPT adversary A , there exists a PPT simulator \mathcal{S} such that for all circuits $C \in \mathfrak{F}$,

$$\left| \Pr [A(\mathcal{O}(C)) = 1] - \Pr [A(\mathcal{S}^C(1^{|C|})) = 1] \right| \leq \text{negl}(|C|).$$

As a variation on the third requirement, one may assume that some auxiliary information (which may depend on the circuit C) is present alongside the obfuscation $\mathcal{O}(C)$. In that case, a simulator with access to that auxiliary information should still be able to simulate the adversary’s output distribution:

Definition 2.2 ([GK05, Definition 3]). *A classical virtual black-box obfuscator w.r.t. dependent auxiliary input for a circuit class \mathfrak{F} is a probabilistic algorithm \mathcal{O} that satisfies Definition 2.1, with the “virtual black-box” property redefined as follows:*

3. (virtual black-box) *For any PPT adversary A , there exists a PPT simulator \mathcal{S} such that for all circuits $C \in \mathfrak{F}$ and all strings $\mathbf{aux} \in \{0, 1\}^{\text{poly}(|C|)}$ (which may depend on C),*

$$\left| \Pr[A(\mathcal{O}(C), \mathbf{aux}) = 1] - \Pr[A(\mathcal{S}^C(1^{|C|}, \mathbf{aux})) = 1] \right| \leq \text{negl}(|C|).$$

In the quantum setting, we consider quantum obfuscators for classical circuit classes: that is, the obfuscation $\mathcal{O}(C)$ may be a quantum state. We adapt Definition 5 from [AF16], which defines quantum obfuscators for quantum circuits.

Definition 2.3. *A quantum virtual black-box obfuscator for the classical circuit class \mathfrak{F} is a quantum algorithm \mathcal{O} and a QPT \mathcal{J} such that*

1. (polynomial expansion) *For every circuit $C \in F$, $\mathcal{O}(C)$ is an m -qubit quantum state with $m = \text{poly}(n)$;*
2. (functional equivalence) *For every circuit $C \in F$ and every input x ,*

$$\frac{1}{2} \left\| \mathcal{J}(\mathcal{O}(C) \otimes |x\rangle\langle x|) - |C(x)\rangle\langle C(x)| \right\|_{\text{tr}} \leq \text{negl}(|C|);$$

3. (virtual black-box) *For every QPT adversary \mathcal{A} , there exists a QPT simulator \mathcal{S} (with superposition access to its oracle) such that for all circuits $C \in \mathfrak{F}$,*

$$\left| \Pr[\mathcal{A}(\mathcal{O}(C)) = 1] - \Pr[\mathcal{S}^C(1^{|C|}) = 1] \right| \leq \text{negl}(|C|).$$

There are a few differences with the classical definition. First, the obfuscation is a quantum state, and not a (classical or quantum) circuit. Second, due to the probabilistic nature of quantum computation, we allow a negligible error in the functional equivalence. Third, the simulator is slightly more powerful because of its superposition access to the functionality of C : a query performs the unitary operation specified by $|x\rangle|z\rangle \mapsto |x\rangle|z \oplus C(x)\rangle$. Note that a quantum adversary can always use a (classical or quantum) obfuscation to compute the obfuscated functionality on a superposition of inputs, obtaining a superposition of outputs. For this reason, the simulator gets superposition access to its oracle in the quantum setting. Throughout this work, all oracles supplied to quantum algorithms allow for superposition access.

We can again strengthen the virtual black-box property to include (classical or quantum) dependent auxiliary information: this auxiliary string or state would be provided to both the adversary and the simulator, in the same way as in Definition 2.2.

2.3 Quantum Fully Homomorphic Encryption

A fully homomorphic encryption (FHE) of a message m provides privacy by hiding the message, but allows ciphertexts to be transformed in a meaningful way. Given a ciphertext for m , some party that only knows the public key can produce a ciphertext for $f(m)$ for any efficiently computable function f . Any information that is necessary for this transformation is contained in the public key (in particular, we do not make a distinction between the public key and the evaluation key).

A *quantum* fully homomorphic encryption (QFHE) scheme allows quantum computations on encrypted quantum data. From the Learning with Errors assumption, it is possible to construct secure QFHE schemes where all client-side operations (key generation, encryption, and decryption) are classical [Mah18, Bra18].

Definition 2.4. *A quantum fully homomorphic encryption scheme QFHE consists of four algorithms, as follows:*

- **Key Generation:** $(pk, sk) \leftarrow \text{QFHE.KeyGen}(1^\lambda)$ produces a public key pk and a secret key sk , given a security parameter λ . This is a classical PPT algorithm.
- **Encryption:** $c \leftarrow \text{QFHE.Enc}_{pk}(m)$ encrypts a single-bit message $m \in \{0, 1\}$. For multi-bit messages $m \in \{0, 1\}^\ell$, we write $\text{QFHE.Enc}_{pk}(m)$ to denote the bit-by-bit encryption

$$(\text{QFHE.Enc}_{pk}(m_1), \text{QFHE.Enc}_{pk}(m_2), \dots, \text{QFHE.Enc}_{pk}(m_\ell)).$$

This algorithm is in general QPT but it only uses a classical random tape, and furthermore whenever m is classical, so is the encryption algorithm.

- **Decryption:** $m' = \text{QFHE.Dec}_{sk}(c)$ decrypts a ciphertext c into a single-bit message m' , using the secret key sk . If c is a ciphertext for a multi-bit message, we write $\text{QFHE.Dec}_{sk}(c)$ for the bit-by-bit decryption. Again this is QPT in general, but can be classical if c is classical.
- **Homomorphic evaluation:** $c' \leftarrow \text{QFHE.Eval}_{pk}(C, c)$ takes as input the public key, a classical description of a BQP circuit C with ℓ input wires and ℓ' output wires, and a bit-by-bit encrypted ciphertext c encrypting ℓ bits. It produces a c' , a sequence of ℓ' output ciphertexts. This is a QPT algorithm.

We say that a (Q)FHE scheme is (perfectly) correct if the homomorphic evaluation of any BQP circuit C on a ciphertext has the effect of applying C to the plaintext, i.e.,

$$\text{QFHE.Dec}_{sk}(\text{QFHE.Eval}_{pk}(C, \text{QFHE.Enc}_{pk}(m))) = C(m)$$

for all m, C , and $(pk, sk) \leftarrow \text{QFHE.KeyGen}(1^\lambda)$. A (Q)FHE scheme is secure if its encryption function is secure. We usually require quantum indistinguishability under chosen plaintext attacks (q-IND-CPA) [BJ15].

The QFHE schemes from [Mah18,Bra18] encrypt a message m using a quantum one-time-pad with random keys $a, b \in \{0, 1\}$, attaching classical FHE ciphertexts of the one-time pad keys:

$$\text{QFHE.Enc}_{pk}(m) = X^a Z^b |m\rangle \otimes |\text{FHE.Enc}_{pk}(a), \text{FHE.Enc}_{pk}(b)\rangle.$$

Note that this ciphertext can be classically represented as the tuple

$$(m \oplus a, \text{FHE.Enc}_{pk}(a), \text{FHE.Enc}_{pk}(b)),$$

so that encryption may be seen as a classical procedure. Conversely, a classical homomorphic encryption $\tilde{m} \leftarrow \text{FHE.Enc}_{pk}(m)$ can easily be turned into a valid quantum homomorphic encryption by preparing the state $|0\rangle \otimes |\tilde{m}, \text{FHE.Enc}_{pk}(0)\rangle$, which decrypts to m . Thus, it is possible to freely switch back and forth between quantum ciphertexts and classical ciphertexts, as long as the message is known to be classical.

The quantum one-time pad encryption also straightforwardly extends to encrypting general quantum states $|\psi\rangle$, rather than only computational-basis states $|m\rangle$: the quantum one-time pad is simply applied to the state $|\psi\rangle$, and the one-time-pad keys encrypted into a computational-basis state as above. Of course, encryption becomes a quantum procedure in this setting. We will use encryption of quantum states in our work, where we supply the encryption of a quantum-state obfuscation as the input to a homomorphic evaluation.

Leveled FHE and Bootstrappable FHE. In many cases in the literature, we wish to consider FHE schemes which require an a priori upper bound (polynomial in the security parameter) on the depth of circuits to be homomorphically evaluated. In the current state of the art, such schemes (referred to as *leveled FHE*) can be constructed under milder assumptions than unleveled schemes: in particular, they do not require circular-security-type assumptions. There are a few variants of leveled FHE defined in the literature, but for the purpose of this work we use the following.

Definition 2.5 (Leveled FHE). *A leveled (Q)FHE scheme is a scheme where the key generation takes an additional parameter $\text{KeyGen}(1^\lambda, 1^d)$ and outputs (sk, pk) for a (Q)FHE scheme. Correctness holds only for evaluating circuits of total depth at most d . Furthermore, the length of sk and the complexity of decryption are independent of d .*

We assume w.l.o.g. that the random tape used by KeyGen is always of length λ and does not depend on d (this is w.l.o.g. since it is always possible to use a PRG to stretch the random tape into the desired length).

One way to construct leveled FHE is via the bootstrapping technique as suggested by Gentry [Gen09]. Gentry showed that given a base scheme with homomorphic capacity greater than its decryption depth, it is possible to create a leveled scheme with the following properties.

Definition 2.6 (Leveled Bootstrapped FHE). A leveled bootstrapped (Q) FHE is a scheme where there exists a base-scheme with a key-generation algorithm $\text{SubKeyGen}(1^\lambda)$ and encryption, decryption and evaluation algorithms, such that the key generation algorithm $\text{KeyGen}(1^\lambda, 1^d)$ takes the following form.

1. Run $\text{SubKeyGen}(1^\lambda)$ with fresh randomness $(d+1)$ times to generate sub-keys (sk_i, pk_i) for $i = 0, \dots, d$.
2. Encrypt $c_i^* = \text{Enc}_{pk_i}(sk_{i-1})$ for all $i = 1, \dots, d$.
3. Output $pk = (pk_0, (pk_1, c_1^*), \dots, (pk_d, c_d^*))$ and $sk = sk_d$.

For our purposes, we will assume that the random tape that is being used for the SubKeyGen executions is generated using a pseudorandom function. That is, the random tape of KeyGen is used as a seed for a PRF, and for the i th execution of SubKeyGen we use a random tape that is derived by applying a PRF on i .

For the sake of completeness we note that the decryption algorithm of the bootstrapped scheme is the same as that of the base scheme, and that for the sake of encryption only pk_0 is needed.

2.4 Point Functions and Compute-and-Compare Functions

The class of compute-and-compare functions, as well as its subclass of point functions, plays an important role in this work. In this section we define these function classes.

Definition 2.7 (Point function). Let $y \in \{0, 1\}^n$. The point function \mathbf{P}_y is defined by

$$\mathbf{P}_y(x) := \begin{cases} 1 & \text{if } x = y \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The value y is called the *target value*. Point functions are a special type of compute-and-compare function, where the function f is the identity:

Definition 2.8 (Compute-and-compare function). Let $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ and $y \in \{0, 1\}^n$. The compute-and-compare function $\mathbf{CC}_{f,y}$ is defined by

$$\mathbf{CC}_{f,y}(x) := \begin{cases} 1 & \text{if } f(x) = y \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

One can also consider point functions or compute-and-compare functions with *multi-bit output*: in that case, the function outputs either some string z (instead of 1), or the all-zero string (instead of 0). We denote such functions with $\mathbf{P}_{y,z}$ and $\mathbf{CC}_{f,y,z}$.

2.5 Recovering the Input of a Quantum Circuit

We will consider (efficient) quantum operations as (polynomial-size) circuits, consisting of the following set of basic operations: unitary gates from some fixed constant-size gate set, measurements in the computational basis, and initialization of auxiliary wires in the $|0\rangle$ state.

While unitary gates are always reversible by applying their transpose ($U^\dagger U = I$ for any unitary U), measurement gates may not be, as they can possibly collapse a state. However, we can effectively delay all measurements in a circuit C until the very end, as follows. Define U_C as the unitary that computes C *coherently*: that is, for every computational-basis measurement in C on some wire w , U_C performs a CNOT operation from w onto a fresh auxiliary target wire initialized in the state $|0\rangle$. The circuit C is now equivalent to the following operation: initialize all auxiliary target wires in the $|0\rangle$ state⁷, apply the unitary U_C , and measure all auxiliary target wires in the computational basis.

In this work, we will encounter circuits C which, for specific inputs, yield a specific state in the computational basis with very high probability. In the proof of the following lemma, we specify how to use coherent computation in order to learn the output value while preserving the input quantum state.

Lemma 2.9. *Let C be a quantum circuit. There exists an input-recovering circuit C_{rec} such that for all inputs ρ_{in} , the following holds: if $\frac{1}{2} \|C(\rho_{\text{in}}) - |x\rangle\langle x|\|_{\text{tr}} \leq \varepsilon$ for some classical string x and some $\varepsilon > 0$, then*

$$\frac{1}{2} \|C_{\text{rec}}(\rho_{\text{in}}) - (\rho_{\text{in}} \otimes |x\rangle\langle x|)\|_{\text{tr}} \leq 2\sqrt{\varepsilon}.$$

The specification of C_{rec} is independent of the specific input state ρ_{in} . However, C_{rec} cannot necessarily recover *all* possible inputs ρ_{in} , only those that lead to an almost-classical output.

The input-recovering circuit consists of running C coherently, copying out the output register, and reverting the coherent computation of C . We formally prove Lemma 2.9 in Appendix A.

3 FHE with Decomposable Public Keys

For the purpose of our result in Section 5, we will need to obfuscate a class of circuits that allow to (quantumly) homomorphically evaluate operations of arbitrary polynomial depth. We nevertheless wish to rely only on leveled FHE for the sake of minimizing our assumptions. We therefore would like to define a class of circuits that are a priori polynomially bounded in size, but which are capable of encapsulating public-key generation of a leveled scheme for some

⁷ If, apart from the targets of the aforementioned CNOTs, the circuit C contains any other wires that are initialized in the $|0\rangle$ state inside the circuit, those wires are also considered part of the input of the unitary U_C . They should be initialized to $|0\rangle$ here as well.

depth d that is not fixed a priori. Note that in a leveled scheme even the length of pk depends on d .

To this end, we define the notion of a scheme with *decomposable* public key, which is defined below. Intuitively, in such a scheme, the public key can be generated by first generating a sequence of blocks, each of some size independent of d . These blocks can then be combined into the actual pk of the scheme. Crucially, the generation of the blocks can be done in parallel, and the complexity of generating each block (given the security parameter and the random tape) is independent of d . In other words, a decomposable public key can be generated on the fly, involving small “chunks” of computation that are independent of d . Formally, we recall Definition 2.5 and define decomposability as follows.

Definition 3.1 (Decomposable public key). *A leveled (Q)FHE scheme has a decomposable public key if there exists a polynomial $K = K(\lambda, d)$ and a polynomial-time deterministic function $\text{BlockGen}(1^\lambda, i, r, r')$ (where $r, r' \in \{0, 1\}^\lambda$) that generates classical strings (“blocks”) c_i such that the following holds:*

1. **Correctness:** *there exists a QPT Assemble such that for all λ, d, r , and r' , letting $K = K(\lambda, d)$, it holds that*

$$\text{Assemble}(c_0, c_1, c_2, \dots, c_K) = pk,$$

where $(pk, sk) = \text{KeyGen}(1^\lambda, 1^d, r)$, and $c_i = \text{BlockGen}(1^\lambda, i, r, r')$ for all i .

2. **Simulatability:** *there exists a QPT simulator \mathcal{S} such that for all d and r ,*

$$\mathcal{S}(1^\lambda, pk) \stackrel{c}{\approx} (c_0, c_1, c_2, \dots, c_K),$$

where $(pk, sk) = \text{KeyGen}(1^\lambda, d, r)$, and the distribution on (c_1, c_2, \dots, c_K) on the right-hand side is generated by selecting a uniformly random r' , and then for all i , setting $c_i = \text{BlockGen}(1^\lambda, i, r, r')$.

We emphasize that in Definition 3.1, the randomness strings r and r' are the same for every run of BlockGen . The reason for this choice is twofold. First, with our final goal in mind of obfuscating the BlockGen functionality, we want to avoid having to specify K independent randomness strings (as that would considerably increase the size of the circuit to obfuscate). Second, most schemes require some form of correlation to exist between the different blocks. Thinking of r and r' as short random seeds for a PRF, this correlation can be realized by running the PRF on the same inputs (see, for example Section 3.1).

3.1 Instantiation from Bootstrapped Schemes

For bootstrapped schemes (see Definition 2.6), decomposability follows immediately by definition. In this case, we do not even need the extra randomness r' and can simply set $\text{BlockGen}(1^\lambda, i, r, r')$ to be the process that evaluates $\text{PRF}_r(i-1)$ and $\text{PRF}_r(i)$ to generate random tapes for SubKeyGen , uses this randomness to generate (sk_{i-1}, pk_{i-1}) and (sk_i, pk_i) , generates c_i^* based on these values, and

outputs (pk_i, c_i^*) . In addition, for $i = 0$, it simply computes $PRF_r(0)$, and uses the resulting randomness to generate pk_0 .

Existing QFHE schemes are based on bootstrapping [Mah18, Bra18]. Without affecting security, we can assume that their randomness is sampled using a PRF as just described.

Lemma 3.2. *Bootstrapping-based leveled QFHE schemes with keys generated from a PRF have decomposable public keys.*

Proof. Define $K(\lambda, d) := d$, and $c_0 := pk_0$. For $i > 0$, define the blocks c_i , which are generated by $\text{BlockGen}(1^\lambda, i, r, r')$, as follows:

$$c_i := (pk_i, c_i^* = \text{Enc}_{pk_i}(sk_{i-1})), \text{ where } (pk_i, sk_i) \leftarrow \text{SubKeyGen}(1^\lambda; PRF_r(i)), \\ (pk_{i-1}, sk_{i-1}) \leftarrow \text{SubKeyGen}(1^\lambda; PRF_r(i-1)). \quad (3)$$

Note that for public keys of this form, BlockGen does not make use of the additional randomness r' .

The assembly function $\text{Assemble}(c_0, c_1, \dots, c_d)$ is a straightforward concatenation of all the blocks: $\text{Assemble}(c_0, c_1, \dots, c_d) := (c_0, c_1, \dots, c_d)$.

Simulatability as in Definition 3.1 is also easily satisfied: a simulator \mathcal{S} , for a public key pk and index i , reads out the pair (pk_i, c_i^*) . It can thereby exactly produce the list (c_1, \dots, c_d) .

3.2 Instantiation from Any Leveled (Q)FHE

We now observe that we can instantiate the a (Q)FHE with decomposable public keys from any leveled scheme, even ones that are not bootstrapped. Decomposing the public key of a general QFHE scheme is done via garbled circuits [Yao86, App17], as we will briefly outline here. A block c_i corresponds to a single garbled gate of the circuit for KeyGen . That is, $\text{BlockGen}(1^\lambda, i, r, r')$ returns a garbling of the i th gate⁸ of $\text{KeyGen}(1^\lambda, d, r)$, using r' as a PRF seed to generate sufficient randomness for the garbling. A separate block (e.g., c_0) contains the required encoding/decoding information to use the garbled circuit. To assemble the public key, a user concatenates all garbled gates, and evaluates the garbled circuit to obtain the output pk . Conversely, by the privacy property of garbled circuits [BHR12], a simulator given only the security parameter λ and the output pk of the garbled circuit, can reproduce a garbled circuit that is indistinguishable from the actual garbled circuit. It can then return the gates of that simulated garbled circuit as the blocks c_i .

Any result relying on the decomposability of the public key of a non-bootstrapping based QFHE scheme of course also relies on any computational assumptions required for the security of the garbled-circuit construction.

⁸ The total number of blocks, $K(\lambda, d)$, will be the number of gates in $\text{KeyGen}(1^\lambda, 1^d, r)$. Since the number of gates is polynomial in λ , it suffices for the length of the PRF seed r' to be linear in λ .

4 Impossibility with respect to Dependent Auxiliary Information

In this section, we show impossibility of virtual-black-box quantum obfuscation of classical point functions under dependent auxiliary information. It sets the stage for our main result, Theorem 5.1, where we incorporate the auxiliary information into the circuit, constructing a circuit class which is unobfuscatable even without the presence of any auxiliary information. Although the result in the current section is perhaps less surprising, the proof contains the most important technical details of our work.

The impossibility result requires two cryptographic primitives, both of which can be built from the hardness of LWE [Mah18,WZ17,GKW17]: (1) quantum fully homomorphic encryption with classical client-side operations (see Section 2.3), and (2) classical vbb obfuscation of compute-and-compare functions. Our result therefore holds under the assumption that LWE is hard. The least favorable LWE parameters are required for the obfuscation of compute-and-compare functionalities, and are discussed in Section 4.1.

In Section 4.1, we describe the classical obfuscator for compute-and-compare functions that we use. We will apply the construction from [WZ17,GKW17] to a specific class of compute-and-compare functions with a specific type of auxiliary information. In Section 4.2, we use this specific application to define a class of circuits and auxiliary-information strings that is unobfuscatable in the quantum vbb sense. The impossibility proof follows in Section 4.3.

4.1 Classical obfuscation of compute-and-compare functions

The works of [WZ17,GKW17] showed that under the assumption that LWE (with polynomial dimension and exponential modulus in the security parameter λ) is hard, it is possible to classically obfuscate compute-and-compare functions [WZ17,GKW17]. We will write “LWE*” to denote their specific variant of the LWE assumption. We note that LWE is known to be at least as hard as worst-case lattice problems [Reg05,PRS17]. In particular, the aforementioned parameter regime LWE* translates to the worst-case hardness of the Gap Shortest Vector Problem (GapSVP) with sub-exponential approximation factor (in the dimension of the lattice). There is currently no known super-polynomial quantum speedup for GapSVP, and the best known quantum (and classical) algorithms require sub-exponential running time.

The works of [WZ17,GKW17] achieve so-called distributional virtual-black-box obfuscation of functions $\mathbf{CC}_{f,y}$, assuming that the target value y has sufficient pseudo-entropy given a description of the function f . The obfuscation is even secure in the presence of (dependent) auxiliary information, so long as the pseudo-entropy of the target value remains high, even conditioned on this auxiliary information.

In our construction, we provide a classically-obfuscated compute-and-compare function as auxiliary information to a quantum obfuscation. We will require

that the target value of the compute-and-compare function is sufficiently random, even given the rest of the auxiliary information (including the quantum obfuscation).

More specifically, for any IND-secure public-key encryption scheme (KeyGen , Enc , Dec), fixed bit string α , and a classical obfuscation procedure $\mathcal{O}(\cdot)$, define a distribution ensemble $\{D_\lambda^{\alpha,d}\}_{\lambda \in \mathbb{N}}$ that samples

$$\begin{aligned} (pk, \tilde{\alpha}, o_{sk,\beta}) \leftarrow D_\lambda^{\alpha,d} \quad \text{as} \quad & (pk, sk) \leftarrow \text{KeyGen}(1^\lambda, 1^d), \\ & \tilde{\alpha} \leftarrow \text{Enc}_{pk}(\alpha), \\ & \beta \leftarrow_R \{0, 1\}^\lambda, \\ & o_{sk,\beta} \leftarrow \mathcal{O}(\text{CC}_{\text{Dec}_{sk},\beta}), \end{aligned} \quad (4)$$

where $\text{CC}_{\text{Dec}_{sk},\beta}$ is a compute-and-compare function as in Definition 2.8. For each α and λ , the target value β is chosen independently from all other information: its pseudo-entropy is λ , even conditioned on pk , $\tilde{\alpha}$ and Dec_{sk} . Therefore, there exists an obfuscation procedure for this class of compute-and-compare programs that has distributional indistinguishability in the following sense:

Lemma 4.1 (Application of [WZ17, Theorem 5.2]). *Under the LWE^* assumption, there exists a classical obfuscation procedure $O_{\text{CC}}(\cdot)$ and a (non-uniform) simulator \mathcal{S} such that for all α and d ,*

$$(pk, \tilde{\alpha}, o_{sk,\beta}) \stackrel{c}{\approx} (pk, \tilde{\alpha}, \mathcal{S}(1^\lambda, \text{params})), \quad (5)$$

where $(pk, \tilde{\alpha}, o_{sk,\beta}) \leftarrow D_\lambda^{\alpha,d}$ using $O_{\text{CC}}(\cdot)$ as the obfuscation procedure $\mathcal{O}(\cdot)$, and params is some information that is independent of sk and β (e.g., it may contain the size of the circuit and/or λ).

In the rest of this work, $O_{\text{CC}}(\cdot)$ will implicitly be the obfuscation procedure used in the distributions $D_\lambda^{\alpha,d}$.

We note that the proofs in [WZ17,GKW17] showed a classical reduction from the hardness of distinguishing the aforementioned distributions to the hardness of solving LWE . We note that proofs by (either Karp or Turing) classical polynomial-time reduction from A to B implies that any solver for A can be translated into a solver for B with comparable complexity, in particular if the solver for A runs in quantum polynomial time then so will the resulting solver for B .

As a consequence of Lemma 4.1, we show that it is hard to guess the value of α , given only a ciphertext $\tilde{\alpha}$ for α , and an obfuscation of the compute-and-compare function. Intuitively, since the information α is completely independent of the target value β , the obfuscation effectively hides the secret key sk that would be necessary to learn α .

Lemma 4.2. *Under the LWE^* assumption, there exists a negligible function $\text{negl}(\cdot)$ such that for any QPT algorithm A and any d ,*

$$\Pr[A(pk, \tilde{\alpha}, o_{sk,\beta}) = \alpha] \leq \text{negl}(\lambda). \quad (6)$$

Here, the probability is over $\alpha \leftarrow_R \{0, 1\}^\lambda$, $(pk, \tilde{\alpha}, o_{sk, \beta}) \leftarrow D_\lambda^{\alpha, d}$, and the execution of A .

Proof. The result follows almost directly from Lemma 4.1, except that we want to bound the probability that A outputs the multi-bit string α , whereas Lemma 4.1 only deals with algorithms with a single-bit output.

To bridge the gap, define an algorithm A'_α that runs A on its input, and compares the output of A to α : if they are equal, A'_α outputs 1; otherwise, it outputs 0.

For any *fixed* value of α , we have

$$\Pr[A(pk, \tilde{\alpha}, o_{sk, \beta}) = \alpha] = \Pr[A'_\alpha(pk, \tilde{\alpha}, o_{sk, \beta}) = 1] \quad (7)$$

$$\stackrel{(*)}{\approx} \Pr[A'_\alpha(pk, \tilde{\alpha}, \mathcal{S}(1^\lambda, \text{params})) = 1] \quad (8)$$

$$= \Pr[A(pk, \tilde{\alpha}, \mathcal{S}(1^\lambda, \text{params})) = \alpha]. \quad (9)$$

The approximation (*) follows from Lemma 4.1, and holds up to a difference of $\text{negl}(\lambda)$.

To complete the proof, note that $\mathcal{S}(1^\lambda, \text{params})$ depends neither on α nor on sk . Thus, randomizing over α again, and invoking privacy of the encryption, we get

$$\Pr[A(pk, \tilde{\alpha}, o_{sk, \beta}) = \alpha] \approx \Pr[A(pk, \tilde{\alpha}, \mathcal{S}(1^\lambda, \text{params})) = \alpha] \leq \text{negl}(|\alpha|) = \text{negl}(\lambda). \quad (10)$$

We have thus established that, even in the presence of an obfuscated compute-and-compare function that depends on the secret key, encryptions remain secure (in the one-way sense). For this security to hold, it is important that the target value β is sufficiently independent of the plaintext α .

4.2 An unobfuscatable circuit class

In this subsection, we define the class of circuits and auxiliary-information strings that we will prove unobfuscatable. Like in [BGI⁺01], we will exploit the idea that access to an object (circuit or quantum state) that allows the evaluation of a function is more powerful than mere black-box access to the functionality: in particular, it allows to evaluate the function homomorphically. For this argument to work, it is important that the function is not easily learnable through black-box access. We will use point functions, as in [BGI⁺01]: with black-box access only, it is hard to tell the difference between a point function and the all-zero function \mathbf{Z}_λ , that accepts inputs of length λ , and always returns 0^λ .

Consider the class $\mathcal{C}_{\lambda, d}^{\text{point}} \cup \mathcal{C}_{\lambda, d}^{\text{zero}}$ of circuits plus auxiliary information, where

$$\mathcal{C}_{\lambda, d}^{\text{point}} := \{(\mathbf{P}_{\alpha, \beta}, (pk, \tilde{\alpha}, o_{sk, \beta})) \mid \alpha \in \{0, 1\}^\lambda, (pk, \tilde{\alpha}, o_{sk, \beta}) \in \text{supp}(D_\lambda^{\alpha, d})\}, \quad (11)$$

$$\mathcal{C}_{\lambda, d}^{\text{zero}} := \{(\mathbf{Z}_\lambda, (pk, \tilde{\alpha}, o_{sk, \beta})) \mid \alpha \in \{0, 1\}^\lambda, (pk, \tilde{\alpha}, o_{sk, \beta}) \in \text{supp}(D_\lambda^{\alpha, d})\}. \quad (12)$$

The class $\mathcal{C}_{\lambda,d}^{\text{point}}$ contains all λ -bit point functions, together with an encryption of the point input α , a public key that enables evaluation of circuits up to depth d , and a function that checks whether a ciphertext decrypts to the target value β . $\mathcal{C}_{\lambda,d}^{\text{zero}}$ contains the all-zero function \mathbf{Z}_λ (which is itself a point function), but still with auxiliary information for the possible values of α and β .

Suppose that some quantum obfuscation ($O_Q(\cdot), \mathcal{J}$) exists. We define a QPT algorithm A , which expects an obfuscation $\rho = O_Q(\mathbf{P}_{\alpha,\beta})$ (or $O_Q(\mathbf{Z}_\lambda)$), together with the classical auxiliary information $\text{aux} = (pk, \tilde{\alpha}, o_{sk,\beta})$. On general inputs ρ and $\text{aux} = (\text{key}, \text{ctxt}, \text{obf})$ of this form, let A do as follows:

1. Run $\text{QFHE.Eval}_{\text{key}}(\mathcal{J}, \text{Enc}_{\text{key}}(\rho) \otimes |\text{ctxt}\rangle\langle\text{ctxt}|)$ to homomorphically evaluate the interpreter algorithm \mathcal{J} . If $\rho = O_Q(\mathbf{P}_{\alpha,\beta})$, $\text{key} = pk$, and $\text{ctxt} = \tilde{\alpha}$, then this step results in an encryption of β with high probability. If $\rho = O_Q(\mathbf{Z}_\lambda)$, $\text{key} = pk$, $\text{ctxt} = \tilde{\alpha}$, and d is at least the depth of \mathcal{J} , then it results in an encryption of 0^λ . Note that we use classical and quantum ciphertexts for the QFHE scheme interchangeably here: see Section 2.3 for a justification.
2. Run obf on the output of the previous step. If $\text{obf} = o_{sk,\beta}$, this will indicate whether the previous step resulted in a ciphertext for β or not.

The above algorithm A will almost certainly output 1 when given an element from $\mathcal{C}_{\lambda,d}^{\text{point}}$ for a sufficiently high value of d , because of the functional equivalence of the two obfuscations and the correctness of the homomorphic evaluation. Similarly, when given an element from $\mathcal{C}_{\lambda,d}^{\text{zero}} - \mathcal{C}_{\lambda,d}^{\text{point}}$, it will almost certainly output 0. Formally, for all $\alpha, \beta \in \{0, 1\}^\lambda - \{0^\lambda\}$, and d at least the depth of \mathcal{J} ,

$$\Pr[A(O_Q(\mathbf{P}_{\alpha,\beta}), pk, \tilde{\alpha}, o_{sk,\beta}) = 1] \geq 1 - \text{negl}(\lambda), \quad (13)$$

$$\Pr[A(O_Q(\mathbf{Z}_\lambda), pk, \tilde{\alpha}, o_{sk,\beta}) = 1] \leq \text{negl}(\lambda). \quad (14)$$

The vastly different output distribution of A when given an obfuscation of a point function versus the zero function are due the fact that A has an actual representation, ρ , of the function to feed into the interpreter \mathcal{J} . In the proof in the next subsection, we will see that a simulator, with only black-box access to these functionality, will not be able to make that distinction.

4.3 Impossibility proof

We are now ready to state and prove the impossibility theorem for quantum obfuscation of classical circuits with dependent auxiliary input. We reiterate that the two assumptions (quantum FHE and compute-and-compare obfuscation) can be realized under the LWE* assumption.

Define $\mathcal{C}_\lambda^{\text{point}} := \bigcup_{d \in [2^\lambda]} \mathcal{C}_{\lambda,d}^{\text{point}}$, and similarly $\mathcal{C}_\lambda^{\text{zero}} := \bigcup_{d \in [2^\lambda]} \mathcal{C}_{\lambda,d}^{\text{zero}}$.

Theorem 4.3 (Impossibility of quantum obfuscation w.r.t. auxiliary input). *Suppose that a classical-client quantum fully homomorphic encryption scheme QFHE exists that satisfies Definition 2.4, and a classical obfuscation procedure $O_{\text{CC}}(\cdot)$ for compute-and-compare functionalities exists that satisfies*

Lemma 4.1. Then any (not necessarily efficient) quantum obfuscator $(O_Q(\cdot), \mathcal{J})$ for the class $\mathcal{C}_\lambda^{\text{point}} \cup \mathcal{C}_\lambda^{\text{zero}}$ satisfying conditions 1 (polynomial expansion) and 2 (functional equivalence) from Definition 2.3 cannot be virtual black-box under the presence of classical dependent auxiliary input, i.e., cannot satisfy condition 3 from Definition 2.3 where both \mathcal{A} and \mathcal{S} get access to a classical string \mathbf{aux} (which may depend on C).

It may seem that the class $\mathcal{C}_\lambda^{\text{point}} \cup \mathcal{C}_\lambda^{\text{zero}}$, consisting of point functions, is classically obfuscatable using $O_{\text{CC}}(\cdot)$ from [WZ17,GKW17]. That obfuscation is secure if α (which is the target value if we view $\mathbf{P}_{\alpha,\beta}$ as the multi-bit output compute-and-compare function $\text{CC}_{\text{id},\alpha,\beta}$) is unpredictable given the auxiliary information $\mathbf{aux} = (pk, \tilde{\alpha}, o_{sk,\beta})$. On the surface, that seems to be the case: only an encryption of α is available in the auxiliary information. However, the secret key sk is present as part of the compute-and-compare function $\text{CC}_{\text{Dec}_{sk},\beta}$. That function is obfuscated, but the obfuscation is not secure in the presence of (an obfuscation of) $\mathbf{P}_{\alpha,\beta}$. Thus, the obfuscation result from [WZ17,GKW17] *almost* applies to the class $\mathcal{C}_\lambda^{\text{point}} \cup \mathcal{C}_\lambda^{\text{zero}}$, but not quite. Hence we are able to prove impossibility of obfuscating it, which we do below.

Proof. The proof structure is similar to [BGI⁺01], and is by contradiction: assume that a quantum obfuscation $(O_Q(\cdot), \mathcal{J})$ for the class $\mathcal{C}_\lambda^{\text{point}} \cup \mathcal{C}_\lambda^{\text{zero}}$ does exist that satisfies all three conditions. We will show that the output distribution of the algorithm A defined in Section 4.2 is approximately the same for every element of the class, contradicting Equations (13) and (14).

By the assumption of the existence of a secure quantum obfuscation $(O_Q(\cdot), \mathcal{J})$, there exists a simulator \mathcal{S} such that

$$|\Pr[A(O_Q(\mathbf{P}_{\alpha,\beta}), \mathbf{aux}) = 1] - \Pr[\mathcal{S}^{\mathbf{P}_{\alpha,\beta}}(1^\lambda, \mathbf{aux}) = 1]| \leq \text{negl}(\lambda), \text{ and} \quad (15)$$

$$|\Pr[A(O_Q(\mathbf{Z}_\lambda), \mathbf{aux}) = 1] - \Pr[\mathcal{S}^{\mathbf{Z}_\lambda}(1^\lambda, \mathbf{aux}) = 1]| \leq \text{negl}(\lambda). \quad (16)$$

The probability is taken over $\alpha \leftarrow_R \{0, 1\}^\lambda$ and $\mathbf{aux} = (pk, \tilde{\alpha}, O_{\text{CC}}(\text{CC}_{\text{Dec}_{sk},\beta})) \leftarrow D_\lambda^{\alpha,q}$ for q the depth of the interpreter circuit \mathcal{J} . Note that \mathcal{S} does not depend on α , β , sk , or pk .

In the remainder of this proof we show that for any \mathcal{S} (independent of α , β , sk , pk),

$$|\Pr[\mathcal{S}^{\mathbf{P}_{\alpha,\beta}}(1^\lambda, \mathbf{aux}) = 1] - \Pr[\mathcal{S}^{\mathbf{Z}_\lambda}(1^\lambda, \mathbf{aux}) = 1]| \leq \text{negl}(\lambda), \quad (17)$$

from which it can be concluded that

$$|\Pr[A(O_Q(\mathbf{P}_{\alpha,\beta}), \mathbf{aux}) = 1] - \Pr[A(O_Q(\mathbf{Z}_\lambda), \mathbf{aux}) = 1]| \leq \text{negl}(\lambda). \quad (18)$$

Since Equations (13) and (14) imply that this difference must be at least $1 - \text{negl}(\lambda)$, Equation (18) yields a contradiction.

To show that Equation (17) holds, i.e., to bound the difference in output probabilities of \mathcal{S} when given an oracle for $\mathbf{P}_{\alpha,\beta}$ versus an oracle for \mathbf{Z}_λ , we

employ the one-way to hiding theorem as it is stated in [AHU19, Theorem 3]. It says that there exists a QPT algorithm B such that

$$|\Pr[\mathcal{S}^{\mathbf{P}^{\alpha,\beta}}(1^\lambda, \mathbf{aux}) = 1] - \Pr[\mathcal{S}^{\mathbf{Z}^\lambda}(1^\lambda, \mathbf{aux}) = 1]| \leq 2d' \cdot \sqrt{\Pr[B^{\mathbf{Z}^\lambda}(1^\lambda, \mathbf{aux}) = \alpha]}, \quad (19)$$

where $d' = \text{poly}(\lambda)$ is the query depth of \mathcal{S} . However, by Lemma 4.2, the probability that B outputs α when given the auxiliary information $\mathbf{aux} = (pk, \tilde{\alpha}, o_{sk,\beta})$ is negligible in λ . Granting B access to the zero-oracle and the additional input 1^λ does not increase this probability, since the value of λ can already be deduced from \mathbf{aux} .

We can thus conclude that the difference in Equation (19) is negligible, and Equation (17) holds, as desired.

We end this section with a few remarks: we describe some variants and generalizations of Theorem 4.3 which almost immediately follow from the presented proof.

Remark 4.4. The proof for Theorem 4.3 also works if we replace $O_{\text{CC}}(\text{CC}_{\text{Dec}_{sk,\beta}})$ inside the distributions $D_\lambda^{\alpha,d}$ with $O_Q(\text{CC}_{\text{Dec}_{sk,\beta}})$, the quantum obfuscation we get from the assumption. This adaptation renders a quantum obfuscator for point functions impossible with respect to dependent auxiliary *quantum* input: a slightly weaker statement, but it does not require the existence of a classical obfuscator for compute-and-compare programs. In particular, the required LWE parameters are better, because we only need the assumption of quantum fully homomorphic encryption.

Remark 4.5. Even a quantum obfuscator $(O_Q(\cdot), \mathcal{J})$ for $\mathcal{C}_\lambda^{\text{point}} \cup \mathcal{C}_\lambda^{\text{zero}}$ with non-negligible errors in the functional equivalence and/or the virtual-black-box property would lead to a contradiction in the proof of Theorem 4.3. Concretely, let ε_f denote the error for functional equivalence, and ε_s denote the error for security in the virtual-black-box sense (they are both $\text{negl}(|C|) = \text{negl}(\lambda)$ in Definition 2.3). The impossibility proof works for any values of $\varepsilon_f, \varepsilon_s$ such that $\varepsilon_f + \varepsilon_s \leq \frac{1}{2} - \frac{1}{\text{poly}(\lambda)}$. So in particular, even a quantum obfuscator with small constant (instead of negligible) errors in both conditions cannot exist.

5 Impossibility without Auxiliary Information

In this section, we will show that quantum virtual-black-box obfuscation of classical circuits is impossible even when no auxiliary information is present. We will rely heavily on the class constructed in Section 4, essentially showing how the auxiliary information can be absorbed into the obfuscated circuit. As a result, the unobfuscatable circuit class itself becomes perhaps less natural, but still consists of classical polynomial-size circuits. Thus, our theorem implies impossibility of quantum vbb obfuscation of the class of all efficient classical circuits.

We would like to consider circuits of the following form:

$$C_{\alpha,\beta,\text{aux}}(b, x) := \begin{cases} \text{aux} = (pk, \tilde{\alpha}, o_{sk,\beta}) & \text{if } b = 0 \\ \mathbf{P}_{\alpha,\beta}(x) & \text{if } b = 1, \end{cases} \quad (20)$$

where $(pk, \tilde{\alpha}, o_{sk,\beta})$ is generated from $D_{\lambda}^{\alpha,d}$, as in Section 4. The input bit b is a choice bit: if it is set to 1, the function $\mathbf{P}_{\alpha,\beta}$ (or \mathbf{Z}_{λ}) is evaluated on the actual input x , whereas if it is set to 0, the auxiliary information is retrieved.

The idea would then be to retrieve the auxiliary information, followed by a homomorphic evaluation of the branch for $b = 1$. There is a problem with this approach, however: since the auxiliary information aux contains the public evaluation key pk , the circuit C grows with d , which affects the length of pk . But as the circuit grows, a (non-circularly-secure) QFHE scheme may require a larger pk to perform all evaluation steps.

To get around this issue, the unobfuscatable circuit will generate the public key step-by-step, in a construction inspired by [CLTV15]. We will assume that the public key of the leveled QFHE scheme is decomposable in the sense of Definition 3.1.

Given a scheme with a decomposable public key, we redefine the unobfuscatable circuit class as follows. Instead of returning the entire public key at once, the circuit allows the user to request individual blocks c_i , up to some depth d . An honest user can run the circuit $K + 1 = K(d, \lambda) + 1$ times to obtain $pk = \text{Assemble}(c_0, c_1, \dots, c_K)$. The depth d will not be fixed a priori, although it will be (exponentially) upper bounded: the circuit will only be able to handle inputs i where $|i| \leq \lambda$. Thus, only up to 2^λ blocks c_i can be retrieved.

The circuit class we consider in this section consists of circuits of the following form:

$$\hat{C}_{\alpha,\beta,d,r,r',\tilde{\alpha},o_{sk,\beta}}(b, x) := \begin{cases} (\tilde{\alpha}, o_{sk,\beta}) & \text{if } b = 0, \\ \text{BlockGen}(1^\lambda, x, r, r') & \text{if } b = 1 \text{ and } x \leq K(d, \lambda), \\ \perp & \text{if } b = 1 \text{ and } x > K(d, \lambda), \\ \mathbf{P}_{\alpha,\beta}(x) & \text{if } b = 2. \end{cases} \quad (21)$$

or

$$\hat{C}'_{\alpha,\beta,d,r,r',\tilde{\alpha},o_{sk,\beta}}(b, x) := \begin{cases} (\tilde{\alpha}, o_{sk,\beta}) & \text{if } b = 0, \\ \text{BlockGen}(1^\lambda, x, r, r') & \text{if } b = 1 \text{ and } x \leq K(d, \lambda), \\ \perp & \text{if } b = 1 \text{ and } x > K(d, \lambda), \\ \mathbf{Z}_{\lambda}(x) & \text{if } b = 2. \end{cases} \quad (22)$$

The first input b is now a choice trit: depending on its value, a different branch of the circuit is executed.

We alter the distribution $D_{\lambda}^{\alpha,d}$ from Equation (4), so that it does not explicitly generate the public key anymore. That information is now generated on-the-fly

by setting $b = 1$. The public and secret key are deterministically computed using r to generate the auxiliary information $(\tilde{\alpha}, o_{sk,\beta})$ for $b = 0$. Consider the distribution ensemble $\{D_\lambda^{\alpha,d,r}\}_{\lambda \in \mathbb{N}}$, where

$$\begin{aligned} (\tilde{\alpha}, o_{sk,\beta}) &\leftarrow D_\lambda^{\alpha,d,r} \quad \text{as} \quad (pk, sk) = \text{KeyGen}(1^\lambda, 1^d, r), \\ &\tilde{\alpha} \leftarrow \text{Enc}_{pk}(\alpha), \\ &\beta \leftarrow_R \{0, 1\}^\lambda \\ &o_{sk,\beta} \leftarrow \text{CC}_{\text{Dec}_{sk},\beta}. \end{aligned} \tag{23}$$

Note that the value of d does not influence the size of $\tilde{\alpha}$ or $o_{sk,\beta}$ (and thereby the circuit size of \hat{C} and \hat{C}').

We can then define the following parametrized circuit classes:

$$\hat{C}_{\lambda,d}^{\text{point}} := \{\hat{C}_{\alpha,\beta,d,r,r',\tilde{\alpha},o_{sk,\beta}} \mid \alpha \in \{0, 1\}^\lambda, r, r' \in \{0, 1\}^\lambda, (\tilde{\alpha}, o_{sk,\beta}) \in \text{supp}(D_\lambda^{\alpha,d,r})\}, \tag{24}$$

$$\hat{C}_{\lambda,d}^{\text{zero}} := \{\hat{C}'_{\alpha,\beta,d,r,r',\tilde{\alpha},o_{sk,\beta}} \mid \alpha \in \{0, 1\}^\lambda, r, r' \in \{0, 1\}^\lambda, (\tilde{\alpha}, o_{sk,\beta}) \in \text{supp}(D_\lambda^{\alpha,d,r})\}. \tag{25}$$

Define the circuit class $\hat{C}_\lambda^{\text{point}} \cup \hat{C}_\lambda^{\text{zero}}$, where $\hat{C}_\lambda^{\text{point}} := \bigcup_{d \in [2^\lambda]} \hat{C}_{\lambda,d}^{\text{point}}$ and similarly $\hat{C}_\lambda^{\text{zero}} := \bigcup_{d \in [2^\lambda]} \hat{C}_{\lambda,d}^{\text{zero}}$. Note that in all circuits in this class, the ‘‘auxiliary information’’ $(\tilde{\alpha}, o_{sk,\beta})$ is fixed. Hence, when the obfuscation of the compute-and-compare function is requested by setting $b = 0$, the circuit always returns the same obfuscation that depends on the same secret key sk .

Similarly to the setting with auxiliary input, there exists a QPT algorithm A' that has significantly different output distributions when given a circuit from $\hat{C}_{\lambda,d}^{\text{point}}$ versus a circuit from $\hat{C}_{\lambda,d}^{\text{zero}}$. Here, we define the algorithm A' that is able to distinguish *only* if it receives a circuit for $d = q$, where q is the depth of the interpreter circuit. If $d < q$, then A' will not be able to retrieve a long enough evaluation key, and will always output zero. However, for our impossibility result, a single value of d on which A' succeeds in distinguishing is sufficient. Note that we cannot define our circuit class to contain only circuits with $d = q$, since q depends on the specific obfuscator/interpreter pair.

On an input state ρ , we define A' as follows:

1. Run $\mathcal{J}_{\text{rec}}(\rho, |b = 0\rangle\langle b = 0| \otimes |0^\lambda\rangle\langle 0^\lambda|)$, where \mathcal{J}_{rec} is the input-recovering version of the interpreter circuit (see Lemma 2.9). If ρ is an obfuscation of a circuit in $\hat{C}^{\text{point}} \cup \hat{C}^{\text{zero}}$, this will result in a state (negligibly close to) $\rho \otimes |\tilde{\alpha}\rangle\langle \tilde{\alpha}| \otimes |o_{sk,\beta}\rangle\langle o_{sk,\beta}|$. Measure the second and third registers to obtain $\tilde{\alpha}$ and $o_{sk,\beta}$.
2. Let q be the depth of the interpreter \mathcal{J} . Because the interpreter is efficient, $q = \text{poly}(\lambda)$. Sequentially run $\mathcal{J}_{\text{rec}}(\rho, |b = 1\rangle\langle b = 1| \otimes |i\rangle\langle i|)$ for all $0 \leq i \leq K = K(q, \lambda)$ to obtain (c_0, c_1, \dots, c_K) , and compute the public evaluation key $pk = \text{Assemble}(c_0, c_1, \dots, c_K)$, suitable for homomorphic evaluations of up to depth q . Note that the key pk is only revealed in its entirety if the given circuit has parameter $d = q$. If $d < q$, A' will notice that \perp is returned for some queries, and outputs 0 at this point.

3. Run $\text{QFHE.Eval}_{pk}(\mathcal{J}, \text{Enc}_{pk}(\rho) \otimes |\text{Enc}_{pk}(b=2)\rangle\langle \text{Enc}_{pk}(b=2)| \otimes |\tilde{\alpha}\rangle\langle \tilde{\alpha}|)$. Similarly to Section 4.2, this will result in a ciphertext for β (if ρ was an obfuscation of a circuit in $\hat{\mathcal{C}}_{\lambda}^{\text{point}}$) or a ciphertext for 0^{λ} (if ρ was an obfuscation of a circuit in $\hat{\mathcal{C}}_{\lambda}^{\text{zero}}$), provided that $d = q$.
4. Run $o_{sk,\beta}$ on the output of the previous step. Doing so will indicate whether the previous step resulted in a ciphertext for β or not. If yes, output 1; otherwise output 0.

Let $(O_Q(\cdot), \mathcal{J})$ be an obfuscator. The algorithm A' , when given a random obfuscated circuit from $\hat{\mathcal{C}}_{\lambda,q}^{\text{point}}$, will almost certainly output 1, where q is the depth of \mathcal{J} . At the same time, an element from $\hat{\mathcal{C}}_{\lambda,q}^{\text{zero}} - \hat{\mathcal{C}}_{\lambda,q}^{\text{point}}$ will almost certainly result in the output 0. More formally, for all $\alpha, r \in \{0, 1\}^{\lambda}$ and $d = q$,

$$\Pr \left[A'(O_Q(\hat{\mathcal{C}}_{\alpha,\beta,d,r,r',\tilde{\alpha},o_{sk,\beta}})) = 1 \right] \geq 1 - \text{negl}(\lambda), \quad (26)$$

$$\Pr \left[A'(O_Q(\hat{\mathcal{C}}'_{\alpha,\beta,d,r,r',\tilde{\alpha},o_{sk,\beta}})) = 1 \right] \leq \text{negl}(\lambda). \quad (27)$$

The probability is taken over $D_{\lambda}^{\alpha,d,r}$, r' , and the internal randomness of A' . Compare these inequalities to Equations (13) and (14).

We are now ready to state our main theorem.

Theorem 5.1 (Impossibility of quantum obfuscation). *Suppose that a classical-client quantum fully homomorphic encryption scheme QFHE exists that satisfies Definitions 3.1 and 2.4, and a classical obfuscation procedure $O_{\text{CC}}(\cdot)$ for compute-and-compare functionalities exists that satisfies Lemma 4.1. Then any (not necessarily efficient) quantum obfuscator $(O_Q(\cdot), \mathcal{J})$ for the class $\hat{\mathcal{C}}_{\lambda}^{\text{point}} \cup \hat{\mathcal{C}}_{\lambda}^{\text{zero}}$ satisfying conditions 1 (polynomial expansion) and 2 (functional equivalence) from Definition 2.3 cannot be virtual black-box, i.e., cannot satisfy condition 3 from Definition 2.3.*

Corollary 5.2. *If the LWE^* assumption holds, the class of classical polynomial-size circuits cannot be quantum virtual-black-box obfuscated in the sense of Definition 2.3.*

Proof (Proof of Theorem 5.1). We again prove the statement by contradiction, assuming that there does exist an obfuscator $(O_Q(\cdot), \mathcal{J})$ that securely obfuscates $\hat{\mathcal{C}}_{\lambda}^{\text{point}} \cup \hat{\mathcal{C}}_{\lambda}^{\text{zero}}$. Let q be the depth of \mathcal{J} , so that $K(q, \lambda)$ is the number of blocks c_i of the evaluation key required by A' to successfully distinguish between an element of $\hat{\mathcal{C}}_{\lambda}^{\text{point}}$ and of $\hat{\mathcal{C}}_{\lambda}^{\text{zero}}$.

By the assumption that $(O_Q(\cdot), \mathcal{J})$ is secure, there must exist a simulator \mathcal{S}_0 such that for all $\alpha, r \in \{0, 1\}^\lambda$ (and setting $d = q$),

$$\left| \Pr[A'(O_Q(\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}})) = 1] - \Pr[\mathcal{S}_0^{\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}}}(1^\lambda) = 1] \right| \leq \text{negl}(\lambda), \quad (28)$$

$$\left| \Pr[A'(O_Q(\hat{C}'_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}})) = 1] - \Pr[\mathcal{S}_0^{\hat{C}'_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}}}(1^\lambda) = 1] \right| \leq \text{negl}(\lambda). \quad (29)$$

The probabilities are taken over $(\tilde{\alpha}, o_{sk, \beta}) \leftarrow D_\lambda^{\alpha, d, r}$ and $r' \leftarrow_R \{0, 1\}^\lambda$, and the internal randomness of A' and \mathcal{S}_0 .

The output distribution of \mathcal{S}_0 can be exactly simulated by another simulator, \mathcal{S}_1 , that has access only to an oracle for $\mathbf{P}_{\alpha, \beta}$ or \mathbf{Z}_λ , and gets the auxiliary information pk , $\tilde{\alpha}$, and $o_{sk, \beta}$ as input. \mathcal{S}_1 can simply run \mathcal{S}_0 , simulating each oracle query using its own oracle, auxiliary input, or a combination thereof. If (part of) the query of \mathcal{S}_0 is for some block c_i , \mathcal{S}_1 can use the decomposability of pk to compute the individual blocks. We formally show the existence of such an \mathcal{S}_1 in Corollary B.2.

We can thus conclude that for all $\alpha, r \in \{0, 1\}^\lambda$,

$$\left| \Pr[A'(O_Q(\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}})) = 1] - \Pr[\mathcal{S}_1^{\mathbf{P}_{\alpha, \beta}}(1^\lambda, \tilde{\alpha}, o_{sk, \beta}, pk) = 1] \right| \leq \text{negl}(\lambda), \quad (30)$$

$$\left| \Pr[A'(O_Q(\hat{C}'_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}})) = 1] - \Pr[\mathcal{S}_1^{\mathbf{Z}_\lambda}(1^\lambda, \tilde{\alpha}, o_{sk, \beta}, pk) = 1] \right| \leq \text{negl}(\lambda). \quad (31)$$

Again, the probabilities are over $D_\lambda^{\alpha, d, r}$ and r' , A' , and \mathcal{S}_1 .

However, by Equation (17) in the proof of Theorem 4.3, the output distribution of \mathcal{S}_1 can only differ negligibly between the two different oracles. Thus, we have

$$\left| \Pr[A'(O_Q(\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}})) = 1] - \Pr[A'(O_Q(\hat{C}'_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}})) = 1] \right| \leq \text{negl}(\lambda). \quad (32)$$

This contradicts our observation in Equations (26) and (27) that on input $\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}}$, A' will almost always output 1, whereas on input $\hat{C}'_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}}$, it will almost always output 0.

Acknowledgements

We thank Andrea Coladangelo, Urmila Mahadev and Alexander Poremba for useful discussions, and Serge Fehr for pointing out an error in the proof of Lemma 2.9. GA acknowledges support from the NSF under grant CCF-1763736, from the U.S. Army Research Office under Grant Number W911NF-20-1-0015,

and from the U.S. Department of Energy under Award Number DE-SC0020312. ZB is supported by the Binational Science Foundation (Grant No. 2016726), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701). YD is supported by the Dutch Research Council (NWO/OCW), as part of the Quantum Software Consortium programme (project number 024.003.037). CS is supported by a NWO VIDI grant (Project No. 639.022.519). Part of this work was done while the authors were attending The Quantum Wave in Computing at the Simons Institute for the Theory of Computing.

References

- Aar05. S. Aaronson. Ten semi-grand challenges for quantum computing theory. Blog post, 2005.
Online: <https://www.scottaaronson.com/writings/qchallenge.html>.
- ABC⁺19. D. Aharonov, Z. Brakerski, K.-M. Chung, A. Green, C.-Y. Lai, and O. Satath. On quantum advantage in information theoretic single-server PIR. In *Advances in Cryptology – EUROCRYPT*, pages 219–246. Springer, 2019.
DOI: 10.1007/978-3-030-17659-4_8.
- AF16. G. Alagic and B. Fefferman. On quantum obfuscation. 2016.
arXiv: 1602.01771.
- AHU19. A. Ambainis, M. Hamburg, and D. Unruh. Quantum security proofs using semi-classical oracles. In *Advances in Cryptology – CRYPTO*, pages 269–295. Springer, 2019.
DOI: 10.1007/978-3-030-26951-7_10.
- AL20a. P. Ananth and R. L. La Placa. Secure quantum extraction protocols. In *Theory of Cryptography Conference*, pages 123–152. Springer, 2020.
- AL20b. P. Ananth and R. L. La Placa. Secure software leasing. 2020.
arXiv: 2005.05289.
- App17. B. Applebaum. Garbled circuits as randomized encodings of functions: a primer. *Tutorials on the Foundations of Cryptography*, pages 1–44, 2017.
DOI: 10.1007/978-3-319-57048-8_1.
- BGI⁺01. B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology – CRYPTO*, pages 1–18. Springer, 2001.
DOI: 10.1007/3-540-44647-8_1.
- BHR12. M. Bellare, V. T. Hoang, and P. Rogaway. Foundations of garbled circuits. In *Proceedings of the 19th ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 784–796. ACM, 2012.
DOI: 10.1145/2382196.2382279.
- BJ15. A. Broadbent and S. Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity. In *Advances in Cryptology – CRYPTO*, pages 609–629. Springer, 2015.
DOI: 10.1007/978-3-662-48000-7_30.
- BK20. A. Broadbent and R. A. Kazmi. Indistinguishability obfuscation for quantum circuits of low t-count. 2020.
arXiv: 2005.14699.

- Bra18. Z. Brakerski. Quantum FHE (almost) as secure as classical. In *Advances in Cryptology – CRYPTO*, pages 67–95. Springer, 2018.
DOI: 10.1007/978-3-319-96878-0_3.
- BS20. N. Bitansky and O. Shmueli. Post-quantum zero knowledge in constant rounds. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 269–279, 2020.
- CD08. R. Canetti and R. R. Dakdouk. Obfuscating point functions with multi-bit output. In *Advances in Cryptology – EUROCRYPT*, pages 489–508. Springer, 2008.
DOI: 10.1007/978-3-540-78967-3_28.
- Chr14. P. Christiano. Quantum obfuscation of classical circuits. Seminar talk in the Quantum Games and Protocols program, Simons Institute for the Theory of Computing, 2014.
Online: <https://simons.berkeley.edu/talks/paul-christiano-2014-02-25>.
- CLTV15. R. Canetti, H. Lin, S. Tessaro, and V. Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In *Proceedings of the 12th International Conference on Theory of Cryptography (TCC)*, pages 468–497. Springer, 2015.
DOI: 10.1007/978-3-662-46497-7_19.
- Gen09. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 169–178. ACM, 2009.
DOI: 10.1145/1536414.1536440.
- GK05. S. Goldwasser and Y. T. Kalai. On the impossibility of obfuscation with auxiliary input. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 553–562. IEEE, 2005.
DOI: 10.1109/SFCS.2005.60.
- GKW17. R. Goyal, V. Koppula, and B. Waters. Lockable obfuscation. In *Proceedings of the 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–621. IEEE, 2017.
DOI: 10.1109/FOCS.2017.62.
- Mah18. U. Mahadev. Classical homomorphic encryption for quantum circuits. In *Proceedings of the 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 332–338. IEEE, 2018.
DOI: 10.1109/FOCS.2018.00039.
- PRS17. C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, pages 461–473. ACM, 2017.
DOI: 10.1145/3055399.3055489.
- Reg05. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 84–93. ACM, 2005.
DOI: 10.1145/1060590.1060603.
- Wee05. H. Wee. On obfuscating point functions. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 523–532. ACM, 2005.
DOI: 10.1145/1060590.1060669.
- WZ17. D. Wichs and G. Zirdelis. Obfuscating compute-and-compare programs under LWE. In *Proceedings of the 58th Annual Symposium on Foundations of*

Computer Science (FOCS), pages 600–611. IEEE, 2017.

DOI: 10.1109/FOCS.2017.61.

Yao86. A. C. Yao. How to generate and exchange secrets. In *Proceedings of the 27th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 162–167. IEEE, 1986.

DOI: 10.1109/SFCS.1986.25.

A Proof of Lemma 2.9

Proof. The input-recovering circuit C_{rec} will consist of running C coherently, copying out the output register, and reverting the coherent computation of C . Suppose the circuit C contains k measurement gates, ℓ initializations of wires in the $|0\rangle$ state, and outputs of length n . Define C_{rec} as:

1. Run U_C on input $\rho_{\text{in}}^{A_1} \otimes |0^\ell\rangle\langle 0^\ell|^{A_2} \otimes |0^k\rangle\langle 0^k|^M$, where U_C is the unitary that coherently executes C , $A = (A_1, A_2)$ is a register that contains the actual input and the auxiliary input $|0\rangle$ states for C , and M is the register that contains the auxiliary wires for the coherent measurements.
2. Copy the wires that are supposed to contain the output $C(\rho_{\text{in}})$ into a register Y , initialized to $|0^n\rangle\langle 0^n|$, using CNOTs. The source of the CNOTs is a register O , the subregister of A containing those output wires. Write \bar{O} for the registers in A that are not in O (these wires are normally discarded after the execution of C).
3. Run U_C^\dagger to recover the original input, and discard the registers A_2 and M .

The behavior of C_{rec} can be summarized as

$$C_{\text{rec}}(\rho_{\text{in}}) = \text{Tr}_{A_2 M} \left[U_C^\dagger \text{CNOT}_{O,Y}^{\otimes n} \left(U_C \left(\rho_{\text{in}}^{A_1} \otimes |0^{\ell+k}\rangle\langle 0^{\ell+k}|^{A_2 M} \right) U_C^\dagger \otimes |0\rangle\langle 0|^Y \right) \left(\text{CNOT}_{O,Y}^{\otimes n} \right)^\dagger U_C \right]. \quad (33)$$

To see that C_{rec} acts as promised, let ρ_{in} , x , and ε be s.t. $\|C(\rho_{\text{in}}) - |x\rangle\langle x|\|_{\text{tr}} \leq \varepsilon$. If ε is small, the CNOT in Step 2 does not create a lot of entanglement, since the control wires are (close to) the computational-basis state $|x\rangle\langle x|$. The output is therefore (almost) perfectly copied out.

More formally, note that $C(\rho_{\text{in}}) = \text{Tr}_{\bar{O}M} \left[U_C(\rho_{\text{in}} \otimes |0^{\ell+k}\rangle\langle 0^{\ell+k}|) U_C^\dagger \right]$. By Lemma A.1 in [ABC⁺19], the closeness of $C(\rho_{\text{in}})$ and $|x\rangle\langle x|$ implies that there exists a density matrix $\chi^{\bar{O}M}$ such that

$$\frac{1}{2} \left\| U_C(\rho_{\text{in}} \otimes |0^{\ell+k}\rangle\langle 0^{\ell+k}|) U_C^\dagger - |x\rangle\langle x| \otimes \chi^{\bar{O}M} \right\|_{\text{tr}} \leq \sqrt{\varepsilon}. \quad (34)$$

Next, we use the fact that a quantum map cannot increase the trace distance between two states to derive two inequalities from Equation (34).

For the first inequality, we append $|x\rangle\langle x|$ on both sides (into a separate Y register):

$$\frac{1}{2} \left\| U_C(\rho_{\text{in}} \otimes |0^{\ell+k}\rangle\langle 0^{\ell+k}|) U_C^\dagger \otimes |x\rangle\langle x|^Y - |x\rangle\langle x| \otimes \chi^{\overline{0}M} \otimes |x\rangle\langle x|^Y \right\|_{\text{tr}} \leq \sqrt{\varepsilon}. \quad (35)$$

For the second inequality, we instead append $|0\rangle\langle 0|$ into the Y register, followed by CNOTs from O onto Y . Note that on the second term inside the trace norm, the effect is the same as before:

$$\begin{aligned} & \frac{1}{2} \left\| \text{CNOT}_{O,Y}^{\otimes n} \left(U_C(\rho_{\text{in}} \otimes |0^{\ell+k}\rangle\langle 0^{\ell+k}|) U_C^\dagger \otimes |0\rangle\langle 0|^Y \right) (\text{CNOT}_{O,Y}^{\otimes n})^\dagger \right. \\ & \quad \left. - |x\rangle\langle x| \otimes \chi^{\overline{0}M} \otimes |x\rangle\langle x|^Y \right\|_{\text{tr}} \leq \sqrt{\varepsilon}. \end{aligned} \quad (36)$$

Thus, by the triangle inequality, the left-hand terms inside the trace norms in Equations (35) and (36) are $2\sqrt{\varepsilon}$ -close. Applying the map $\text{Tr}_{A_2M} [U_C^\dagger(\cdot)U_C]$ to both terms, which again does not increase the trace difference, we arrive at the desired statement:

$$\frac{1}{2} \left\| (\rho_{\text{in}} \otimes |x\rangle\langle x|) - C_{\text{rec}}(\rho_{\text{in}}) \right\|_{\text{tr}} \leq 2\sqrt{\varepsilon}. \quad (37)$$

B Auxiliary Lemmas for Theorem 5.1


Lemma B.1. *Let $g : \{0,1\}^m \rightarrow \{0,1\}^n$ for $m, n \in \mathbb{N}$, and let $c \in \{0,1\}^n$. Let $f : \{0,1\} \times \{0,1\}^m \rightarrow \{0,1\}^n$ be defined by*

$$f(b, x) := \begin{cases} c & \text{if } b = 0 \\ g(x) & \text{if } b = 1. \end{cases} \quad (38)$$

Then for every QPT A , there exists a simulator S such that for all f, g of the form described above, and all input states ρ :

$$\Pr[A^f(\rho) = 1] = \Pr[S^g(\rho, c) = 1]. \quad (39)$$

Proof. Recall that since A and S are quantum algorithms, they access their oracles in superposition: that is, A has access to the map defined by $|x\rangle|z\rangle \mapsto |x\rangle|z \oplus f(x)\rangle$, and S has access to the map defined by $|x\rangle|z\rangle \mapsto |x\rangle|z \oplus g(x)\rangle$. The simulator S runs A on input ρ , and simulates any oracle calls to f (on inputs registers BX and output register Z) using two oracle calls to g . It only needs to prepare an auxiliary register in the state $|0^n\rangle$, and run the following circuit:



To see that this circuit exactly simulates a query to f on BXZ , consider an arbitrary query state

$$\sum_i \alpha_i |b_i, x_i\rangle_{BX} |z_i\rangle_Z |\varphi_i\rangle_R, \quad (40)$$

where R is some purifying register. The state on $BXZR$ (plus the two auxiliary registers containing $|0^n\rangle$ and $|c\rangle$) after the above circuit is executed, is equal to

$$\sum_i \alpha_i |b_i, x_i\rangle_{XB} |0^n\rangle |z_i \oplus b \cdot g(x_i) \oplus (1-b) \cdot c\rangle_Z |c\rangle |\varphi_i\rangle_R \quad (41)$$

$$= \sum_i \alpha_i |b_i, x_i\rangle_{XB} |0^n\rangle |z_i \oplus f(x_i)\rangle_Z |c\rangle |\varphi_i\rangle_R, \quad (42)$$

which is exactly the state that would result from a direct query to f .

Corollary B.2. *Let $\hat{C}_\lambda^{\text{point}}$ and q be as in Section 5. Then for any QPT \mathcal{S}_0 , there exists a QPT simulator \mathcal{S}_1 such that for all $\alpha, r \in \{0, 1\}^\lambda$,*

$$\left| \Pr[\mathcal{S}_0^{\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}}}(1^\lambda) = 1] - \Pr[\mathcal{S}_1^{\mathbf{P}^{\alpha, \beta}}(1^\lambda, \tilde{\alpha}, o_{sk, \beta}, pk) = 1] \right| \leq \text{negl}(\lambda). \quad (43)$$

A similar statement holds for circuits from $\hat{C}_\lambda^{\text{zero}}$.

Proof. The statement is proven via an intermediate simulator \mathcal{S}_2 . This simulator is constructed by repeated application of Lemma B.1, so that for all α, r ,

$$\left| \Pr[\mathcal{S}_0^{\hat{C}_{\alpha, \beta, q, r, r', \tilde{\alpha}, o_{sk, \beta}}}(1^\lambda) = 1] - \Pr[\mathcal{S}_2^{\mathbf{P}^{\alpha, \beta}}(1^\lambda, \tilde{\alpha}, o_{sk, \beta}, c_0, c_1, c_2, \dots, c_K, \perp) = 1] \right| \quad (44)$$

is at most $\text{negl}(\lambda)$, where $K = K(q, \lambda)$ as in Definition 3.1. On the right-hand side, the probability is additionally over a random choice of r' (resulting in the sequence $(c_0, c_1, c_2, \dots, c_K)$).

Next, we apply the simulatability property of Definition 3.1. It states that there exists a simulator \mathcal{S}_3 that, given a public key, can generate the distribution over $(c_0, c_1, c_2, \dots, c_K)$ itself. Define

$$\mathcal{S}_1^{\mathbf{P}^{\alpha, \beta}}(1^\lambda, \tilde{\alpha}, o_{sk, \beta}, pk) := \mathcal{S}_2^{\mathbf{P}^{\alpha, \beta}}(1^\lambda, \tilde{\alpha}, o_{sk, \beta}, \mathcal{S}_3(pk), \perp), \quad (45)$$

and the corollary follows.