

On the Computational Hardness Needed for Quantum Cryptography

Zvika Brakerski   

Weizmann Institute of Science, Rehovot, Israel

Ran Canetti   

Boston University, MA, US

Luowen Qian   

Boston University, MA, US

Abstract

In the classical model of computation, it is well established that *one-way functions* (OWF) are minimal for computational cryptography: They are essential for almost any cryptographic application that cannot be realized with respect to computationally unbounded adversaries. In the quantum setting, however, OWFs appear not to be essential (Kretschmer 2021; Ananth et al., Morimae and Yamakawa 2022), and the question of whether such a minimal primitive exists remains open.

We consider EFI pairs – efficiently samplable, statistically far but computationally indistinguishable pairs of (mixed) quantum states. Building on the work of Yan (2022), which shows equivalence between EFI pairs and statistical commitment schemes, we show that EFI pairs are necessary for a large class of quantum-cryptographic applications. Specifically, we construct EFI pairs from minimalistic versions of commitments schemes, oblivious transfer, and general secure multiparty computation, as well as from QCZK proofs from essentially any non-trivial language. We also construct quantum computational zero knowledge (QCZK) proofs for all of QIP from any EFI pair.

This suggests that, for much of quantum cryptography, EFI pairs play a similar role to that played by OWFs in the classical setting: they are simple to describe, essential, and also serve as a linchpin for demonstrating equivalence between primitives.

2012 ACM Subject Classification Theory of computation → Cryptographic primitives; Theory of computation → Cryptographic protocols; Theory of computation → Computational complexity and cryptography; Security and privacy → Mathematical foundations of cryptography; Theory of computation → Quantum complexity theory; Theory of computation → Quantum complexity theory

Keywords and phrases quantum cryptography, efi, commitment scheme, oblivious transfer, zero knowledge, secure multiparty computation

Digital Object Identifier 10.4230/LIPIcs.ITCS.2023.24

Related Version *Full Version:* <https://arxiv.org/abs/2209.04101>

Funding *Ran Canetti* and *Luowen Qian*: Supported by DARPA under Agreement No. HR00112020023.

Zvika Brakerski: Supported by the Israel Science Foundation (Grant No. 3426/21), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482).

Acknowledgements LQ thanks Prabhanjan Ananth and Henry Yuen for the motivating discussions prior to this work. LQ is also grateful to Scott Aaronson and William Kretschmer for their insightful discussions for Section 1.2. Some of these discussions occurred while LQ was visiting the Simons Institute for the Theory of Computing. We thank Fermi Ma for his helpful comments on a preliminary version of this work.



© Zvika Brakerski, Ran Canetti, and Luowen Qian;
licensed under Creative Commons License CC-BY 4.0

14th Innovations in Theoretical Computer Science Conference (ITCS 2023).

Editor: Yael Tauman Kalai; Article No. 24; pp. 24:1–24:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

One of the most fundamental achievements of cryptography has been the conceptualization and eventual formalization of the forms of computational hardness that are needed for obtaining prevalent cryptographic tasks. Notions such as one-way functions [29, 64] (capturing functions that can be computed efficiently but are hard to meaningfully invert) and pseudorandom generators [57, 64, 14] (capturing the ability to efficiently expand short random strings into longer strings that are hard to distinguish from fully random) became foundational pillars for the design and reduction-based analysis of cryptographic schemes that are only “computationally secure” (that is, secure only against computationally-bounded attacks). Furthermore, the celebrated equivalence between the two notions [14, 33, 37] has cemented the combined concept as the “foundational computational hardness for cryptography”: One that is essential for realizing almost any cryptographic task that requires computational hardness, and at the same time suffices for realizing a large class of cryptographic tasks.

However, in the quantum setting, where parties can generate, process, and communicate quantum information, the lay of the land of computational hardness turns out to be different. First, quintessential tasks such as key-exchange with only public communication, which classically can only be computationally secure, can be obtained without any need for computational hardness [11, 56]. Furthermore, quantum protocols can use (quantum-hard) one way functions to obtain tasks that are provably unobtainable from one way functions alone in the classical setting, at least in a relativizing manner. These include non-interactive commitments with either statistical hiding or statistical binding property [50, 63, 13], and oblivious transfer [41, 27, 8, 35].

Even further, it has been recently shown how to obtain commitments, oblivious transfer and general multiparty computation from a form of computational hardness that appears to be “purely quantum”, in the sense that it does not appear to imply one way functions (or any equivalent formulation of computational hardness) – not even ones against *classical* attackers [6, 51]. Superficially, this new form of computational hardness, called *Pseudorandom States (PRS)*, is a straightforward generalization of pseudorandom generators: it postulates the ability to efficiently generate quantum states that are hard to distinguish from a Haar-random state even when given multiple instances [44, 17, 18]. However this apparent similarity is deceiving; indeed, there are no relativizing constructions of one way functions from PRS [49]. Also, in spite of initial attempts [51], the celebrated classical equivalence between one-wayness and pseudorandomness does not appear to naturally generalize to the quantum setting, at least not with respect to PRS. Still, we do not know whether PRS are *essential* for realizing any of the above cryptographic primitives.

This leaves quantum cryptography devoid of a convenient form of “foundational computational hardness”, namely a form of computational hardness that is both necessary for any meaningful computational security, and sufficient for realizing a large class of tasks.

Our contributions. We formulate a relatively *simple and natural* primitive and show that its existence is both necessary and sufficient for a significant class of cryptographic applications in a quantum-enabled computational model. While many of these implications are either known or easily derived from known results, we hope that the proposed framing, along with the new implications, will help in understanding the computational foundations of quantum cryptography.

The proposed primitive draws from a classical primitive considered by Goldreich [32], as well as from the notion of canonical quantum commitments proposed by Yan [62]. Goldreich’s primitive is aimed at capturing non-trivial and “cryptographically useful” computational indistinguishability:

► **Definition 1** (EFID pairs [32]). *An EFID pair is a pair of efficient (classical) sampling algorithms such that their output distributions are statistically far but computationally indistinguishable.*

Goldreich’s work leverages the result of Impagliazzo, Levin, and Luby [40] to show that EFID pairs exist if and only if (classical) pseudorandom generators exist. This, together with what we know about pseudorandom generators, means the existence of a classical protocol for almost any cryptographic task that requires computational hardness implies existence of EFID pairs, and furthermore that EFID pairs suffice for realizing a large class of tasks. We consider a natural quantum analogue of EFID pairs:

► **Definition 2** (EFI pairs, informal). *An EFI pair is a pair of efficient quantum algorithms whose output states are statistically far but computationally indistinguishable.*

Clearly, any quantum-hard EFID pair is also an EFI pair. On the other hand, it is unknown whether existence of EFI pairs implies the existence of (classical) pseudorandom generators or one-way functions. Indeed, the implication is false in the relativizing setting (see [49], combined with [6, Theorem 4.1]).

A first indication that EFI pairs are central to quantum cryptography is the observation that they are essentially equivalent to the statically binding variant of canonical form quantum commitments [62]¹. Building on this initial connection, we demonstrate that the existence of EFI pairs is essential for the existence of *any* commitment scheme, oblivious transfer protocols, non-trivial multi-party computation protocols, and zero-knowledge proofs for non-trivial languages. Furthermore, for each one of these primitives, we use EFI pairs as a tool for demonstrating that existence of protocols for a minimal version of the primitive implies existence of protocols for a full-fledged version of that primitive. Informally:

- **Quantum commitment schemes.** A commitment scheme is a cryptographic protocol where a committer commits to a hidden bit so that it can be later revealed but not modified. As mentioned, EFI pairs can be readily used to build (canonical form) statistically binding non-interactive quantum commitments, which subsequently imply statistically hiding commitments [62].

We construct EFI pairs from any plain “semi-honest” interactive commitment scheme, namely an interactive commitment scheme that is (computationally) binding and hiding as long as both parties are honest during the commitment phase. (For commitment schemes that are either statistically hiding or statistically binding, this implication is essentially shown in [62]. We extend this results to any commitment.)

- **Quantum oblivious transfer.** An oblivious transfer scheme is a cryptographic protocol where a sender makes two bits available to a receiver in a way that enables the receiver to obtain exactly one of them, without the sender learning which bit it obtained. Fully secure (namely, simulation-secure against adversaries that deviate from the protocol) quantum oblivious transfer is known to be constructible from quantum statistically binding commitments [8, 6] (and so also from EFI pairs).

We show how to construct EFI pairs from any *semi-honest OT* protocol, namely any OT protocol that is only guaranteed to be secure when both parties follow the protocol instructions without abort and up to purifications (namely, without tracing out any register used by each party).

¹ Indeed, Yan [62] suggests studying the connections between statistically binding canonical-form commitment schemes and other cryptographic primitives. This work follows the same path, while distilling EFI as the notion of interest. See more details in Section 1.1.

- **Quantum secure multiparty computation protocols.** A multiparty secure computation protocol is a protocol where participating parties jointly compute the output of a function of their secret inputs without revealing anything but the function value. Known constructions of general MPC from any statistically binding quantum commitment [6] imply that EFI pairs can also be used to perform secure evaluation of any functionality. We show how any protocol for securely evaluating any non-trivial classical finite functionality (namely a function with an insecure minor as in [9]), even in the semi-honest model, implies the existence of EFI pairs.
- **Quantum computational zero-knowledge (QCZK) proofs.** Finally, a computational zero-knowledge proof is an interactive proof system where any malicious verifier cannot learn anything beyond the fact that the statement is true, in the sense that their view could be efficiently simulated given only the public instance. By observing that zero knowledge is a special case of secure two-party computation (2PC), we have that if EFI pairs exist then $\text{QCZK} = \text{QIP}$, and any language in QMA also admits QCZK proofs with negligible soundness and an efficient prover that uses only a single copy of the witness. Conversely, we build on results from [63] to construct EFI pairs from any *honest verifier* QCZK proof (QCZK_{HV}) for any language that is hard on average for BQP. (QCZK_{HV} is a relaxation of QCZK where zero knowledge property is guaranteed to hold only against purified verifiers with abort, rather than arbitrary polytime verifiers.)

Furthermore, all these equivalences relativize². Thus, Kretschmer’s oracle separation [49] immediately generalizes to show that none of the objects constructible from EFI pairs (or pseudorandom states) imply the existence of one-way functions (post-quantum or not) in a relativizing way.

1.1 Our techniques

This section presents an overview of the proofs for our results.

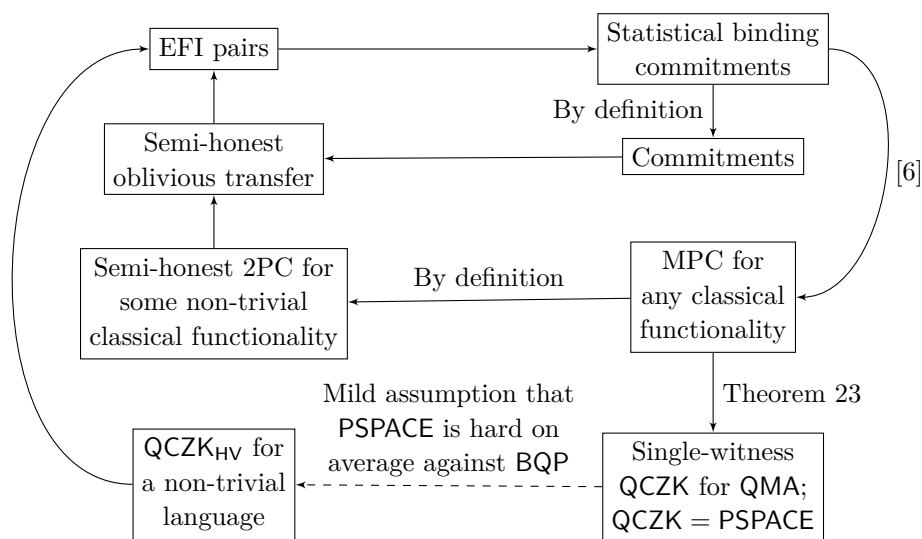
EFI pairs and statistical commitments. In the classical setting, building a commitment scheme from EFID pairs would naturally go via Goldreich’s transformation to a PRG, and then use, say, Naor’s commitment [52]. However, it is not clear how this transformation could be generalized to the quantum setting. In particular, Goldreich’s proof crucially relies on the fact that for a BPP (randomized) algorithm, it is possible to separate the randomness from the rest of the computation – or even arbitrarily *program* the randomness. Such techniques cannot work for quantum algorithms, as also observed by the recent work of Aaronson, Ingram, and Kretschmer [4] comparing the complexity classes BPP and BQP.

Still, as noticed several times in the literature, EFI pairs give quantum commitments in a rather direct way. To sketch this basic construction, we first recall the syntax of a canonical form quantum commitment scheme:³

² We use the term *relativizing* to denote that the construction works even in the presence of (quantum) oracles. In particular, this allows the construction to invoke the next message function of a protocol as a black box or even run its purification. An example of a relativizing implication is Watrous’s construction of a QSZK protocol from any honest-verifier QSZK_{HV} protocol [60].

We reserve the use of the term *black-box reductions* to denote reductions which are black box in a primitive, namely reductions that use oracle access to the ideal functionality describing the primitive, irrespective of any particular implementation (as e.g. in the work of Kilian [46]).

³ This (non-interactive) canonical form of quantum bit commitment schemes was first introduced by the work of Yan et al. [63], but the idea dates back to the work of Chailloux, Kerenidis, and Rosgen [24]. Subsequently, it has been shown that canonical form commitments are as useful as traditional bit commitments by Yan et al. [63, 31, 62] Its connection to (classical) EFID pairs was observed by Yan in a 2022 revision of their work [62].



■ **Figure 1** Our results for EFI pairs illustrated. We give a more detailed overview of these implications in Section 1.1.

- To commit to b , the committer efficiently generates a bipartite (i.e., two-register) state $|\psi_b\rangle_{CR}$, and sends the commitment register C .
- To open, simply reveal the other register R and b , and the receiver can perform a rank-1 projection onto the corresponding state $|\psi_b\rangle_{CR}$ (or equivalently, uncompute the state generation unitary for committing to b and check if we get back all zeroes) to check whether to accept the commitment.

It is possible to view the purified generation of the EFI pair as a canonical form commitment where the output corresponds to the commitment register C and the purification corresponds to the opening register R . When viewed this way, the statistical distance guarantee of the EFI pair translates to statistical binding property, whereas the computational indistinguishability of the EFI pair translates to the computational hiding property of the commitment. Thus an EFI pair is essentially a statistically binding canonical form commitment. This observation (which is implicit in [62]) is indeed the starting point of our work.

Furthermore, the round collapsing theorem in [62] shows that any quantum commitment can be compiled into the canonical form while preserving the hiding and binding properties. Since statistically binding commitments and statistically hiding commitments are equivalent [27, 62], we can construct EFI pairs from either one.

Showing how to construct EFI pairs from commitment schemes that are neither statistically binding nor statistically hiding appears more challenging. One may hope to somehow construct a candidate EFI pair of states, and prove computational indistinguishability from computational security of the commitment, and statistical distance via an inefficient attack on the commitment scheme. However, it is not clear how to transform an inefficient attack against the binding property into a distinguishing attack as needed for EFI pairs. (Classically, this part can be done using one-way functions, but, as argued above, these techniques do not have natural quantum analogues.)

We get around this difficulty by going through oblivious transfer, where the security for both ends can be naturally viewed as distinguishing tasks, and is thus more amenable to constructing EFIs even without statistical security. In fact, we observe that semi-honest OT suffices. Let us elaborate.

EFI pairs from semi-honest OT. We first define a new notion of quantum oblivious transfer, which considers only “purification attacks”. (This notion can be viewed as a quantum analogue of the classical “non-erasing, honest-but-curious” attacks.) Furthermore, the semi-honest adversaries must run the purified protocol till the end, i.e. are disallowed to abort. (Classically, this does not matter.)

We first describe constructing EFI from semi-honest OT, which is more straightforward. Consider the adversarial view in the following two executions:

- Here the sender is honest and the receiver is semi-honest. The sender chooses the bits uniformly at random and the semi-honest receiver chooses the choice bit uniformly at random. (By correctness, the receiver is always able to recover one bit specified by the choice bit with certainty. The task for the semi-honest receiver is to extract the sender’s other bit from his purified view.)
- Here the receiver is honest and the sender is semi-honest. The receiver chooses the choice bit uniformly at random and the semi-honest sender prepares two equal superposition states $|+\rangle^{\otimes 2}$ for her input bits. (Here the task for the semi-honest sender is to extract receiver’s choice bit from her purified view.)

It is easy to see that these views can be computed efficiently since the OT protocol is efficient; and by the semi-honest security of OT, both tasks should be impossible for efficient algorithms. At the same time, the impossibility of Chailloux, Gutoski, and Sikora [23] states that for every OT protocol, one of these two tasks can be accomplished with success probability $\geq \frac{2}{3}$ inefficiently.

Now the construction of EFI pairs follows by simply reinterpreting these bit extraction tasks as a distinguishing task. In particular, for $b = 0, 1$, the b -th state is simply the concatenation of the two above views, conditioned on the correct answers being b in both executions. In other words, the b -th state consists of the semi-honest receiver’s view of the first execution when the sender’s other bit is b , followed by the semi-honest sender’s view of the second execution when the receiver’s choice bit is b .

Semi-honest OT from commitments. Many prior works have already studied constructing quantum oblivious transfer from commitments [26, 12, 28, 27, 31, 62]. However, they all start with a commitment with some statistical security guarantee – either statistically binding or statistically hiding. On the other hand, we want to start from an arbitrary commitment scheme (which may be computationally binding and computationally hiding). While these constructions could probably still carry over, since here we are aiming for a much weaker security, we instead give a much simplified protocol with a self-contained description. Let us begin by considering the simplest (almost trivial) quantum oblivious transfer protocol inspired by Crépeau and Kilian [26], which is only secure if both parties are completely honest during the protocol:

1. The sender on input two message bits b_0, b_1 , sends two qubits $|b_0\rangle \otimes H|b_1\rangle$, where the first qubit encodes the first message bit in the standard basis and the second qubit encodes the second message bit in the Hadamard basis.
2. The receiver measures both qubits in the standard basis to recover b_0 and a random bit b'_1 , or in the Hadamard basis to recover a random bit b'_0 and b_1 .

Security is straightforward: the sender gets no information at all, and the receiver destroys the information about the other bit by measuring it in an incompatible basis.

However, this is obviously not semi-honest secure as a purified receiver could simply uncompute a purified measurement to recover the other bit as well. Indeed, a better way to “erase” information in the semi-honest model is to simply send it to the other party: in this

case, the receiver sends a copy of these two measured (classical) bits to the sender. Since the semi-honest receiver’s purified view at the end has the sender’s private registers traced out, doing this ensures that these measured bits do indeed collapse even for the purified view.

On the other hand, since these two bits contains information about the receiver’s choice bit, we cannot simply send it in the clear as now the sender can break. One simple fix is to have the receiver instead commit to the two measured bits to the sender. This strategy is also often employed when designing maliciously secure OT from commitments [12, 28, 27].

It is easy to see that this committed-measurement OT remains secure against semi-honest sender by hiding of the commitment, and security against semi-honest receiver remains to be seen. If the commitment is statistically binding, then it can be seen that the collapse still occurs; but the computational binding case seems less clear. Fortunately, this can be overcome through Yan’s computational collapse theorem [62], which on a high level states that a canonical form computationally binding commitment scheme computationally “collapses” the committed qubit (given that the commit phase was performed semi-honestly), even if the commitment is never opened. (Their theorem is established via reducing a collapsing distinguishing adversary to an adversary that breaks computational binding of the canonical form commitment.) This completes the argument.

Multiparty secure computations for classical functionalities. Using a known sequence of transformations outlined in existing works [6] (which builds on existing works including but not limited to [12, 8, 35]), it is already known how construct, given any statistically binding quantum commitment scheme (and hence also given any EFI pair), multi-party protocols that securely evaluate any classical function with any number of faults. Furthermore, these protocols provide statistical security guarantees against at least one of the parties. As an aside, via known results, these protocols can further be used to construct two-party secure computations for general quantum functions (or channels) where only one party obtains output [30], and in addition, reactive (meaning stateful and interactive) classical functionalities [25, 42].

For the converse direction, we can now use the powerful equivalence established for oblivious transfer above, and simply invoke the classical equivalence of Beimel, Malkin, and Micali [9] to complete the proof. While the [9] proof contains parts which do not naturally generalize to the quantum setting, the only thing we need from that proof is the reduction from semi-honest OT to semi-honest 2PC for any non-trivial classical functionality (i.e. if it contains an insecure minor), and this construction is black-box and hence extends to our setting. While the semi-honest models are slightly different, we verify that their semi-honest reduction also works for our model. Once we have semi-honest OT, we get EFI pairs by the equivalence above.

Zero knowledge proofs from EFI pairs. We now turn to establishing an equivalence between EFI pairs and non-trivial quantum computational zero knowledge (QCZK) proofs. We first consider the task of constructing QCZK protocols for QMA from EFI pairs. Here the commit-and-open QCZK protocol by Broadbent and Grilo [20] can be readily instantiated by any quantum commitment. However, this protocol uses sequential repetition, and as a consequence, requires multiple copies of the quantum witness to achieve negligible soundness. As also proposed by Broadbent et al. [21], this limitation can be avoided via performing 2PC for quantum (CPTP) functionalities, which can be constructed from OT [30] and thus EFI pairs.

We now move on to general QCZK proofs without any constraint on prover efficiency and show how to use EFI pairs to construct QCZK proofs for all of QIP. Before presenting our protocol, let us recall the celebrated construction of Ben-Or et al. [10] that transforms any

(without loss of generality, Arthur-Merlin or public-coin) IP protocol into a CZK protocol. In the transformed protocol, the parties first run the original public-coin protocol, where the prover only sends (statistically binding) commitments to its messages. Next, the parties engage in a zero-knowledge protocol where the instance consists of the transcript so far, and the language accepts a transcript if there exist valid openings to all the prover commitments that would have caused the original verifier to accept.

Since $\text{QIP} = \text{IP}$ [58, 43], it is natural to consider extending the [10] construction to our setting. However, direct extension hits a roadblock: the statement that needs to be proven in zero knowledge is now a *quantum statement* involving the commitment states when instantiated with quantum commitments. This is not a context that is traditionally considered by the zero knowledge literature. (Indeed, recall that even zero knowledge proofs for QMA still consider classical statements.) Even if we attempt to mimic a zero knowledge proof via statistical 2PC, we soon encounter another issue: how should the two parties agree on the quantum statement that is being proven? Sure, we could make the verifier send the quantum state in the statement to 2PC, but a malicious verifier could refuse to provide the correct state. This becomes an issue as the verifier might be able to manipulate the commitment message so that checking the validity of the commitments itself might reveal non-trivial information about the committed bit.

We thus take a different path: we have the prover and the verifier engage in a secure evaluation of the following reactive functionality (which also can be constructed from OT [25, 42]). The verifier inputs its random challenges in the underlying interactive proof, and the functionality uses these challenges to play the verifier role in an interactive proof with the external prover. Finally the functionality outputs the acceptance bit to the external verifier. Both soundness and zero knowledge follow from the security of the MPC.

EFI pairs from non-trivial QCZK_{HV}. Finally, we show how to construct EFI pairs from any QCZK proof for any language that is hard on average against BQP. Note that the computational indistinguishability given by the QCZK security does not give EFI pairs immediately as it might be possible to generate the honest view efficiently.

One possible approach might be to try to extend the classical result of Ostrovsky and Wigderson [55] to our setting. However, they use the non-existence of one way functions to build universal extrapolators that efficiently turn simulators into cheating provers, and it is not clear how to use the non-existence of EFI pairs to construct quantum universal extrapolators.

We instead turn to the works of Ong and Vadhan [59, 54], showing an equivalence between instance-dependent commitments and CZK, that is a language admits an instance-dependent commitment (a commitment, parameterized by an instance x , whose computational hiding and statistical binding properties only hold if x is in or not in the language, respectively) if and only if it admits a CZK proof. We note that if a hard-on-average language admits an instance-dependent commitment, then this commitment is essentially a full-fledged commitment, thus implying the existence of EFI pairs. Therefore, it remains to extend the equivalence to the quantum setting, i.e. we wish to establish that any language admits a QCZK proof if and only if it admits instance-dependent quantum commitments.

The bad news is that going from CZK to instance-dependent commitments again involve going through instance-dependent one-way (universal hash) functions. However, we note that the mixed states considered by Watrous [60] for handling QSZK readily gives an instance-dependent mixed state for QCZK protocols: a weak variant of EFI states that is only required to satisfy either statistical fairness or computational indistinguishability if x is in or not in the language, respectively. It can be seen that the transformations described before

also readily extends to the instance-dependent setting, and thus this gives an instance-dependent commitment. Indeed, this transformation from QCZK to instance-dependent quantum commitments has been observed by the work of Yan et al. [63]. We then conclude that if L is hard on average for BQP, then this instance-dependent mixed state averaged over the hard distribution immediately gives an EFI pair.

Upon the completion of this work, we discovered a result by Chailloux, Kerenidis, and Rosgen [24] that is similar to this part with very similar proof techniques. However, their separation is between $\text{QIP} = \text{PSPACE}$ and QMA , which is technically incomparable with our separation between QCZK and BQP. Furthermore, they consider worst case hardness instead of average case hardness here, and thus only getting quantum auxiliary-input EFI. In our case, this difference is rather minor and the results can translate back and forth; and in their case, it is not clear how to get standard EFI pairs from any notion of average case hardness of QIP against QMA.⁴

1.2 Discussions and open questions

We now give a few open questions in this direction, organized into three categories. To keep the discussion succinct, we point the readers to the references for details of the terminologies.

EFI and quantum complexity. One way functions (and equivalently pseudorandom generators and classical EFID pairs) have been one of the central objects in complexity theory [7]. Since EFI pairs are both essential and sufficient for much of quantum cryptography, and furthermore are very simple to describe, it is natural to ask whether EFI pairs could also be a useful object to study from the complexity point of view. Note that the computational hardness underlying EFI pairs, which is the quantum state distinguishability problem, seems especially relevant to the study of the complexity of quantum states and transformations [1].

One very important question, we think, is whether there is any barrier for proving the existence of EFI pairs. In the classical setting, existence of one-way functions implies $\text{P} \neq \text{NP}$, but is there any barrier for establishing the existence of quantum EFI pairs? EFI pairs would immediately imply a quantum circuit lower bound for an explicit two-outcome measurement, but is there any reason to believe that such a lower bound would be hard to establish?

For a more concrete example, is P vs PSPACE a (classical) barrier for the existence of EFI pairs? In other words, does the existence of EFI pairs separate BQP from PSPACE? We know that the existence of pseudorandom states do separate BQP from $\text{PP} = \text{PostBQP}$ [49], but nothing is known for EFI pairs. One way to achieve this could be to demonstrate a way to synthesize the Helstrom measurement given a PSPACE oracle, which is closely related to the unitary synthesis problem [5]: in particular, if the unitary synthesis problem could be done by a PSPACE oracle, then the existence of EFI pairs would separate PSPACE from BQP.

Hardness amplification for EFI. Another intriguing challenge is hardness amplification for EFI. Is it possible to construct full-fledged EFI pairs from weaker variants where either the computational distance is non-negligible, or the statistical distance is bounded away from 1, or both? What is the minimal initial gap between the statistical and computational distances that still allows amplification?

⁴ For readers that are familiar with [24], even if we assume QIP is hard on average against QMA for BQP samplable distributions, this still does not suffice for getting EFI pairs without quantum auxiliary input. The reason is that the quantum auxiliary input needs to specify the state that witnesses the diamond norm of the two channels, and it is not clear how this state could be prepared efficiently.

Candidate EFI? Given the oracle separation, are there any concrete candidate assumptions that imply quantum EFI, with formal evidence that it does not imply one-way functions?

Two natural candidates come from pseudorandom state candidates. One possible approach is to assume that “sufficiently large” quantum random circuits are pseudorandom unitaries. These random quantum circuits are already being investigated with motivations like quantum supremacy [2, 3] and the theory of black holes [22, 19, 36]. One could hope that pursuing this direction could ultimately lead to useful quantum cryptography that could be implemented on near-term quantum devices. The other possible approach is to consider the pseudorandom states proposed by Bouland, Fefferman, and Vazirani [15] from the physical description of wormholes. Intuitively, such a construction could be secure based on the physical belief that a wormhole is highly “scrambling”. In that paper, they also prove that their construction is a secure pseudorandom state generator if the evolution unitary is a black-box Haar random unitary.

Another candidate is proposed by Kawachi, Koshihara, Nishimura, and Yamakami [45]. On a high level, they consider computational indistinguishability between two types of random coset states, and show that it is at least as hard as the graph automorphism problem.

EFI and quantum cryptography. The importance of one way functions in classical cryptography cannot be overstated: virtually any non-trivial computational cryptography (those that cannot be realized with respect to computationally unbounded adversaries) implies the existence of one way functions classically. Yet it still remains to be seen how much of quantum cryptography is related to EFI. For instance:

Quantum pseudorandomness. The celebrated result of Goldreich [32] shows existence of classical EFI pairs imply pseudorandom generators and subsequently pseudorandom functions. While we know how to construct various quantum cryptography from quantum EFI pairs, the way we do it completely avoids the need to construct quantum pseudorandomness. Nevertheless, given the many applications of both classical and quantum pseudorandomness, an important question is whether it is possible to construct quantum pseudorandomness (pseudorandom states, unitaries, or any other meaningful pseudorandom objects) from EFI pairs.

Quantum unforgeability. Cryptographic primitives such as digital signatures, message authentication codes, or quantum money appear to inherently require some flavor of one-wayness, in that a break involves solving a computationally hard search problem where solutions exist and are efficiently verifiable given some additional secret information. (In the public key setting, where some classical verification key is made public, solutions are verifiable publicly.) Another related object already proposed previously is one-way state generators [51]. As discussed, this form of computational hardness appears very different than indistinguishability. Still, can we show that existence of any one of these primitives implies existence of EFI pairs? Can we construct any of these primitives from EFI pairs?

Quantum zero knowledge arguments. It is possible to extend our proof to show that if EFI pairs exist then we can give a QSZK argument for any QMA language with an efficient prover having a single copy of the witness. On the other hand, we are unable to show that if L admits an argument then the instance-dependent commitment is computationally binding for NO instances. On a high level, the difficulty is that known techniques in the statistical binding setting do not translate to the computational setting, since we do not have a hardness amplification procedure for computationally binding commitments: it is not clear whether parallel repetition of commitments decreases the computational binding error against malicious committers. We refer the readers to the related discussions in Yan’s work [62] for more details.

2 Preliminaries

2.1 Quantum information

We refer the reader to [53] for a comprehensive reference on the basics of quantum information and quantum computation. We use standard Dirac notation for quantum states.

We recall the notion of density matrices, which are PSD trace-1 matrices that represent the complete characterization of a state of a quantum system. The state of a system can be “pure”, i.e. in the form of a state $|\psi\rangle$, in which case the density matrix is $|\psi\rangle\langle\psi|$ (i.e. of rank 1), or “mixed” which corresponds to a distribution over pure states, and is represented by a density matrix of rank > 1 . Two quantum states are identical if and only if their density matrices are equal, and the distance between quantum states is also expressed as a function of their density matrices, as explained below.

We recall that quantum operations can always be expressed as unitary operators on *some* quantum system, and they act on the density matrix of this quantum system via conjugation. We may sometimes refer to a quantum operation that uses some auxiliary registers, or that removes (“traces out”) registers during the computation. Such a general quantum operation is known as a *quantum channel*. (Mathematically, a quantum channel can be expressed as a completely-positive trace-preserving (CPTP) map on the space of density matrices, but this formulation will not be required for our purposes.)

We use $\mathcal{D}(\mathcal{H})$ to denote the set of density matrices on a Hilbert space \mathcal{H} . Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be density matrices. We write $\text{TD}(\rho, \sigma)$ to denote the trace distance between them, i.e.,

$$\text{TD}(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\|_1$$

where $\|X\|_1 = \text{Tr}(\sqrt{X^\dagger X})$ denotes the trace norm. We also use $F(\rho, \sigma) = (\text{Tr}(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}))^2$ to denote the fidelity of ρ and σ .

► **Fact 3.** For any two mixed states ρ, σ , $(F(\rho, \sigma))^2 + (\text{TD}(\rho, \sigma))^2 \leq 1$.

► **Theorem 4** (Holevo–Helstrom [39, 38]). The best success probability to (inefficiently) distinguish two mixed states ρ, σ is given by $\frac{1}{2}(1 + \text{TD}(\rho, \sigma))$. The measurement that achieves this success probability is called the Helstrom measurement.

► **Corollary 5.** For any mixed states ρ, σ and integer $n > 0$,

$$\text{TD}(\rho^{\otimes n}, \sigma^{\otimes n}) \geq 1 - \exp(-n \text{TD}(\rho, \sigma)/2).$$

We recall that $\rho^{\otimes n}$ is the state containing n -copies of the state represented by ρ .

Proof. We prove this by building a majority-vote distinguisher. Let D be the distinguisher optimally distinguishing ρ from σ by Holevo–Helstrom theorem. We apply D on each copy, and take the majority vote. Corollary follows by applying Hoeffding’s inequality and Holevo–Helstrom again. ◀

2.2 Quantum algorithms

A quantum algorithm A is a family of generalized quantum circuits $\{A_n\}_{n \in \mathbb{N}}$ over a discrete universal gate set (such as $\{CNOT, H, T\}$). By generalized, we mean that such circuits can have a subset of input qubits that are designated to be initialized in the zero state, and a subset of output qubits that are designated to be traced out at the end of the computation.

24:12 On the Computational Hardness Needed for Quantum Cryptography

Thus a generalized quantum circuit A_n corresponds to a *quantum channel*. When we write $A_n(\rho)$ for some density matrix ρ , we mean the output of the generalized circuit A_n on input ρ . If we only take the quantum gates of A_n and ignore the subset of input/output qubits that are initialized to zeroes/traced out, then we get the *unitary part* of A_n , which corresponds to a unitary operator which we denote by \hat{A}_n . The *size* of a generalized quantum circuit is the number of gates in it, plus the number of input and output qubits.

We say that $A = \{A_n\}_n$ is a *quantum polynomial-time (QPT) algorithm* if there exists a polynomial p such that the size of each circuit A_n is at most $p(n)$. We furthermore say that A is *uniform* if there exists a deterministic polynomial-time Turing machine M that on input 1^n outputs the description of A_n .

We also define the notion of a *non-uniform* QPT algorithm A that consists of a family $\{(A_n, \rho_n)\}_n$ where $\{A_n\}_n$ is a polynomial-size family of circuits (not necessarily uniformly generated), and for each n there is additionally a subset of input qubits of A_n that are designated to be initialized with the density matrix ρ_n of polynomial length. This is intended to model nonuniform quantum adversaries who may receive quantum states as advice.

The notation we use to describe the inputs/outputs of quantum algorithms will largely mimic what is used in the classical cryptography literature. For example, for a state generator algorithm G , we write $G_n(k)$ to denote running the generalized quantum circuit G_n on input $|k\rangle\langle k|$, which outputs a state ρ_k .

Ultimately, all inputs to a quantum circuit are density matrices. However, we mix-and-match between classical, pure state, and density matrix notation; for example, we may write $A_n(k, |\theta\rangle, \rho)$ to denote running the circuit A_n on input $|k\rangle\langle k| \otimes |\theta\rangle\langle\theta| \otimes \rho$. In general, we will not explain all the input and output sizes of every quantum circuit in excruciating detail; we will implicitly assume that a quantum circuit in question has the appropriate number of input and output qubits as required by context.

We assume that all parties are quantum algorithms with (noiseless) quantum communication. Furthermore, all algorithms run with the same security parameter.

A function $f : \mathbb{N} \rightarrow \mathbb{R}^{\geq 0}$ is negligible, if for any polynomial p , $f(\lambda) \leq 1/p(\lambda)$ for all sufficiently large $\lambda \in \mathbb{N}$. Otherwise, we say it is noticeable, or equivalently when it is infinitely often at least $1/p(\lambda)$.

► **Definition 6** (Computational indistinguishability). *For two families of mixed states $\{\rho_\lambda\}_\lambda, \{\sigma_\lambda\}_\lambda$, we say that they are computationally indistinguishable (against BQP/qpoly), if for any QPT algorithm D , there exists a negligible function ε such that for any security parameter λ and advice state $\alpha = (\alpha_\lambda)_\lambda$,*

$$|\Pr[D(\alpha_\lambda, \rho_\lambda) = 1] - \Pr[D(\alpha_\lambda, \sigma_\lambda) = 1]| \leq \varepsilon(\lambda).$$

Extending the standard cryptographic convention, the above definition considers adversaries with non-uniform quantum advice. The definition can be adapted to the uniform setting by simply requiring fixing $\alpha = \perp$. We note that the reductions shown in this work are all uniform, or “advice preserving”: given an adversary with some advice, the generated adversary uses the same advice.

3 EFI pairs of states

We define the main object considered in this work, namely pairs of efficiently generatable mixed quantum states that are statistically far and yet computationally indistinguishable:

► **Definition 7.** We call $\xi = (\xi_{b,\lambda})$ a pair of EFI states if it satisfies the following criteria:

1. **Efficient generation:** There exists a uniform QPT quantum algorithm A that on input $(1^\lambda, b)$ for some integer λ and $b \in \{0, 1\}$, outputs the mixed state $\xi_{b,\lambda}$.
2. **Statistically Far:** $\text{TD}(\xi_{0,\lambda}, \xi_{1,\lambda})$ as a function of λ is at least inverse polynomial.
3. **Computational Indistinguishability:** $(\xi_{0,\lambda})_\lambda$ is computationally indistinguishable to $(\xi_{1,\lambda})_\lambda$.

Here we require exact generation of the mixed state. Since we only care about the existence of such object, this requirement does not make a difference. In particular, if we can approximately synthesize a certain family of states with inverse-exponential fidelity, then taking the output of the circuit directly would also satisfy the requirement.

4 Commitments and semi-honest oblivious transfer

In this work, we without loss of generality focus on the canonical form of quantum commitment schemes [62, Definition 5]. A commitment scheme consists of two phases. In the commitment phase of a canonical commitment scheme, Alice (the committer) chooses a bit b , and runs a uniform QPT circuit $Q_{\lambda,b}$ on all zeroes, which outputs two registers C, R ; she then proceeds to send the register C to Bob (the receiver). Later in the reveal phase, Alice sends the other register R and the bit b to Bob; Bob accepts the opening if he performs $Q_{\lambda,b}^\dagger$ on two registers and measures all zeroes in the computational basis. We now recall the requirements on the commitment schemes, specialized to canonical forms for convenience. It will be convenient to define the commitment message $\rho_{\lambda,b} := \text{Tr}_R(Q_{\lambda,b} |0\rangle\langle 0| Q_{\lambda,b}^\dagger)$.

► **Definition 8 (Computational hiding).** A commitment scheme satisfies computational hiding, if $\rho_{\lambda,0}$ is computationally indistinguishable to $\rho_{\lambda,1}$.

We are going to consider a specific more restricted variant of statistical binding called honest binding. We refer the readers to related works [63, 31, 51, 62] for a more thorough discussion on this variant (and how it is equivalent to statistical binding for canonical commitment schemes).

► **Definition 9 (Honest binding).** A canonical commitment scheme satisfies honest computational (resp. statistical) binding if for any auxiliary state $|\psi\rangle$ and any polynomial-time (resp. physically) realizable unitary U , we have that

$$\left\| \left(Q_{\lambda,1} |0\rangle\langle 0|_{CR} Q_{\lambda,1}^\dagger \otimes I_Z \right) (I_C \otimes U_{RZ}) (Q_{\lambda,0} |0\rangle\langle 0|_{CR} \otimes |\psi\rangle_Z) \right\|_2$$

is negligible.

In an oblivious transfer protocol, Bob (the sender) chooses two bits x_0, x_1 to send to Alice, and Alice (the receiver) chooses the bit b to receive. At the end of the protocol, Alice is able to recover x_b . Here, we assume the protocol is able to transmit x_b with probability 1.

Here, we say a (quantum) party is semi-honest (or secure against purified adversaries, analogous to the classical honest-but-curious security), if they follow the protocol (without abort) except that they can purify (without loss of generality) all measurements. We in addition also allow Bob (the sender) to purify his randomness for x_0, x_1 if he was to sample them randomly; on the other hand, we require Alice (the receiver) to specify a classical input to make it easier to define security. (Looking ahead, this is also needed to invoke the semi-honest inefficient attack [23].) At the end, they output their residual state as their view for the distinguisher as usual.

24:14 On the Computational Hardness Needed for Quantum Cryptography

► **Definition 10** (Statistical security against semi-honest Alice). *We say an oblivious transfer protocol is P_A^* -secure against semi-honest Alice, if for every bits b, c , at the end of the protocol with Alice's input being b , a semi-honest Alice's view when $x_b = c$ and $x_{1-b} = 0$ is at most P_A^* -close to that when $x_b = c$ and $x_{1-b} = 1$ in trace distance.*

► **Definition 11** (Statistical security against semi-honest Bob). *We say an oblivious transfer protocol is P_B^* -secure against semi-honest Bob, if for every possible (purified) Bob's inputs (meaning an arbitrary bipartite quantum state where the first part is a qubit indicating the input choice bit), at the end of the protocol, a semi-honest Bob's view when $b = 0$ is at most P_B^* -close to that when $b = 1$ in trace distance.*

Computational security against semi-honest Alice and Bob can be similarly defined, except considering these two views to be computationally indistinguishable instead of statistically indistinguishable.

We recall the impossibility due to Chailloux, Gutoski, and Sikora showing that oblivious transfer protocols that are statistically secure against both parties do not exist. While the “semi-honest” definition they have is different from here, we could open the proof and check that the cheating strategies constructed there are indeed semi-honest according to our definition.

► **Theorem 12** ([23, Theorem 1.1]). *For any oblivious transfer protocol, it holds that $2P_B^* + P_A^* \geq 2$, where Alice chooses the choice bit uniformly at random (classically) and Bob chooses the two bits as uniform superposition ($\frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle)$), and P_A^* is the best probability that a semi-honest Alice is able to predict Bob's choice correctly, and P_B^* is the best probability that a semi-honest Bob is able to predict both bits being sent by Alice correctly.*

We briefly recall the cheating strategies constructed in their proof. Alice's strategy is the following [23, Section 2.1]: she randomly chooses b and then follows the protocol semi-honestly according to our definition; at the end, she performs a gentle measurement to learn x_b (it is gentle since by completeness she is supposed to be able to learn x_b with almost certainty), and then performs the Helstrom measurement (Theorem 4) to learn x_{1-b} . Similarly, Bob's strategy is the following [23, Section 2.2]: he follows the protocol semi-honestly, purifying all measurements including the uniform sampling of x_0, x_1 ; at the end, he performs a post-processing to try to guess b . Let the success probability of Alice and Bob be p_a, p_b respectively. They establish that for these two strategies, $2p_b + p_a \geq 2$, and thus the theorem follows.

The following theorem is proven in [16]:

► **Theorem 13.** *The following assumptions are equivalent.*

1. *Existence of EFI states.*
2. *Existence of statistically binding (canonical-form) commitment schemes.*
3. *Existence of commitment schemes.*
4. *Existence of semi-honest oblivious transfer.*

5 Dichotomy for secure two party computations

In this section, we study secure two party computations with quantum parties for classical functionalities. A secure two-party computation protocol consists of two (interactive uniform quantum) algorithms A, B , where they receive (implicitly) the security parameter λ and their respective inputs a, b , take turns to run and exchange a message register back and

forth; in the end, we denote their joint state as $\langle A, B \rangle(a, b)$. We can also denote Alice's state to be $\langle A, B \rangle(a, b)_A$ and Bob's to be $\langle A, B \rangle(a, b)_B$. Without loss of generality, we consider the protocol so that only Bob gets the output and otherwise they do not learn any other information [9, Definition 1]. In this case, $\langle A, B \rangle(a, b)$ would simply be (\perp, z) as Alice outputs nothing and Bob outputs the evaluated result z . In particular, this evaluated result could be the output of any efficient quantum channel [30]. We can also define the output state $\langle A^*, B \rangle(a, b)$ for a malicious Alice (and analogously for Bob), where the malicious Alice can output anything she wants and not necessarily \perp .

We now describe the definition for malicious simulation security.

► **Definition 14** (Malicious simulation security). *Let $f = (f_\lambda)_\lambda$ be a quantum channel computable by a polynomial-size quantum circuit. A protocol computing f satisfies malicious simulation security for Alice, if the following holds. For any (malicious) QPT algorithm A^* , there exists a QPT simulator S such that for any QPT distinguisher D , there exists a negligible function ε such that for all security parameter λ , non-uniform bipartite advice state ρ_{AD} , and Bob's input b (permissible by f_λ),*

$$|\Pr[D(\langle A^*, B \rangle(\rho_A, b), \rho_D) = 1] - \Pr[D(S_f(\rho_A, b), \rho_D) = 1]| \leq \varepsilon(\lambda),$$

where $S_f(\rho_A, b)$ is the following algorithm:

- The two-stage algorithm $S(1^\lambda, \rho_A)$ is run, which outputs some a^* .
- Compute $(z_a, z_b) \leftarrow f_\lambda(a^*, b)$ to be the output of f . (In our setup, $z_a = \perp$ but z_b is the actual output.)
- Finish executing S with input z_a , which in the end outputs a certain state σ .
- Output (σ, z_b) .

Malicious simulation security for Bob can be defined in the same way as above, except exchanging the role of Alice and Bob.

We say the malicious simulation security is statistical if it holds even against any unbounded algorithms A^* and D , and in this case there need not be a running time bound on the simulator.

In this work, we focus on secure two-party computations although the consequences also generalize to secure multi-party computations where possibly more than two parties are involved and all of them could receive outputs. We refer the readers to the prior work [8] for related literature.

Combining our equivalence theorem from before and existing work constructing one-sided statistically secure 2PC from statistically binding (quantum) commitments [8, 6, 61], we immediately get the following corollary.

► **Corollary 15.** *Assuming EFI state pairs exist, then any P/poly functionalities can be computed with full malicious security and one-sided statistical security.*

For the rest of the section, we show EFI states are also implied by non-trivial 2PC protocols. For that purpose, we focus on 2PC protocols for finite functionalities. By “finite”, we mean that the function to be computed is a fixed-size function independent of the security parameter, say Yao's millionaires' problem.

► **Definition 16** (Insecure minor). *Let S_1, S_2, S_3 be finite sets and $f : S_1 \times S_2 \rightarrow S_3$ be a (finite) function. Then we say f contains an insecure minor, if there exists $x_0, x_1 \in S_1$ and $y_0, y_1 \in S_2$ such that $f(x_0, y_0) = f(x_1, y_0)$ and $f(x_0, y_1) \neq f(x_1, y_1)$.*

► **Lemma 17** ([9, Claim 1]). *If a function $f(\cdot, \cdot)$ does not contain an insecure minor, then there is a classical one-message perfectly secure computation protocol for f .*

► **Lemma 18.** *If a function $f(\cdot, \cdot)$ contains an insecure minor, then we can build a semi-honest OT protocol from an semi-honest secure computation protocol for f .*

Proof. This essentially follows from the work of Beimel, Malkin, and Micali [9, Claim 3]. Since the original proof is black-box, it also immediately generalizes to our setting when the parties are quantum. As the precise definition of semi-honest is different in our case, we give the proof for completeness.

Let x_0, x_1, y_0, y_1 be the values guaranteed by the insecure minor, and let Π_f be the semi-honest secure computation protocol for f . The (semi-honest) oblivious transfer protocol described in that proof works as follows:

- (Recall) Alice gets as input a_0, a_1 and Bob gets as input a choice bit b .
- Execute Π_f on input x_{a_0}, y_{1-b} , and Bob gets output z_0 .
- Execute Π_f on input x_{a_1}, y_b , and Bob gets output z_1 .
- Bob outputs 0 if $z_b = f(x_0, y_1)$, otherwise Bob outputs 1.

Correctness follows directly since the construction is black-box. A semi-honest Alice's view only consists of her semi-honest view from two protocol executions, and thus Alice does not learn anything about b . Similarly, a semi-honest Bob's view only consists of his semi-honest view from two protocol executions, and thus semi-honest security also follows from the property of insecure minor and Alice's privacy against semi-honest Bob for Π_f . ◀

Combining this with Theorem 12, we immediately get the following:

► **Corollary 19.** *If a function $f(\cdot, \cdot)$ contains an insecure minor, then f cannot be computed by statistically-secure semi-honest protocols.*

Combining our equivalence theorem with Lemmas 17 and 18, we obtain the following dichotomy theorem. We shall remark that this theorem, similar to the classical proof [9], is non-black-box in the use of the functionalities [26, 46, 47] due to the use of the equivalence theorem.

► **Theorem 20.** *If there is a semi-honest two-party secure computation protocol for a classical finite functionality $f(\cdot, \cdot)$, then either f can be computed perfectly securely in a single message or EFI states exist.*

6 Quantum computational zero knowledge proofs

The existence of one-way functions implies that all of PSPACE admit a computational zero knowledge proof [10, 58], and NP admits computational zero-knowledge proofs where proofs can be efficiently generated given a witness for membership [34]. Furthermore, the existence of computational zero-knowledge proofs for *any* non-trivial (i.e., average-case easy) language implies the existence of (infinitely-often) one-way functions [55]. In this section, we use our equivalence theorem to establish the quantum analogue of these classical results.

► **Definition 21.** *A language L is in QCZK if there is a (quantum) interactive protocol between an unbounded prover and a QPT verifier (specified by an interactive quantum Turing machine V) such that the following holds:*

- *Completeness: For any $x \in L$, there is an unbounded prover strategy P that would make V accept with probability at least $1 - 2^{-|x|}$.*

- *Soundness:* For any $x \notin L$ and any unbounded prover strategy, V accepts with probability at most $2^{|x|}$.
- *Computational zero knowledge:* For any malicious QPT verifier V^* , there exists a QPT simulator S such that for any QPT distinguisher D and non-uniform bipartite advice state ρ_{AD} , there exists a negligible function ε such that for any $x \in L$,

$$|\Pr[D(\langle P, V^*(\rho_A) \rangle(x, x)_{V^*}, \rho_D) = 1] - \Pr[D(S(x, \rho_A), \rho_D) = 1]| \leq \varepsilon(|x|).$$

We can also consider a (much) weaker variant of this zero knowledge requirement called computationally zero knowledge against purified verifiers *with abort* (or “honest verifier”), where we restrict the malicious V^* to only purifying his state and aborting after any fixed number of rounds⁵. We call the corresponding class QCZK_{HV} , similar to QSZK_{HV} with respect to QSZK for statistical zero knowledge [60].

Note that here, unlike semi-honest OT where we disallow a semi-honest party to prematurely abort, here we must allow a purified verifier to abort prematurely (this makes the complexity class smaller). This is because otherwise we can even show the corresponding class (for even quantum *perfect* zero knowledge) is trivially equal to IP , by simply asking the IP verifier at the end to destroy all the other information he has learned by measuring them in Hadamard basis and then returning them to the prover.

► **Fact 22** ([58, 60, 43, 48]). $\text{BQP} \subseteq \text{QSZK} = \text{QSZK}_{\text{HV}} \subseteq \text{QCZK} = \text{QCZK}_{\text{HV}} \subseteq \text{QIP} = \text{IP} = \text{PSPACE}$.

We first consider QCZK protocols with efficient provers (in which case the largest complexity class we can consider is QMA), and then move on to QCZK with inefficient provers. The work by Broadbent and Grilo on QCZK [20] show how to build a commit-and-open zero knowledge protocol for QCZK using a commitment scheme, and thus combining it with a quantum commitment scheme, we get a QCZK proof, but it requires multiple copies of the advice to boost the soundness to negligible. We strengthen this to show that we can achieve the same thing from the same assumption, but with a single copy of the witness. Overall, the following theorems are proven in [16]:

► **Theorem 23.** *If EFI states exist, then $\text{QCZK} = \text{QIP}$. Furthermore, any language in QMA has a quantum computational zero knowledge proof with negligible soundness and with an efficient prover that uses only a single copy of the witness.*

► **Theorem 24.** *If there is a language in QCZK_{HV} that is hard on average for BQP for some BQP -samplable distribution D , i.e. any BQP algorithm has negligible success probability in deciding L on average over D and the algorithm’s randomness, then EFI state pairs (secure against uniform BQP adversaries) exist.*

References

- 1 Scott Aaronson. The complexity of quantum states and transformations: From quantum money to black holes, 2016. [arXiv:1607.05256](https://arxiv.org/abs/1607.05256).
- 2 Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013. doi:10.4086/toc.2013.v009a004.

⁵ Formally, we require that for a k -round protocol, there exists a simulator S whose output is computationally indistinguishable to $\rho_1 \otimes \dots \otimes \rho_k$, where ρ_i for $i \in [k]$ is the purified verifier’s state immediately after receiving i -th message from the prover.

- 3 Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*, volume 79 of *LIPICs*, pages 22:1–22:67. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.CCC.2017.22.
- 4 Scott Aaronson, DeVon Ingram, and William Kretschmer. The acrobatics of BQP. In *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPICs*, pages 20:1–20:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.CCC.2022.20.
- 5 Scott Aaronson and Greg Kuperberg. Quantum versus classical proofs and advice. *Theory of Computing*, 3(7):129–157, 2007. doi:10.4086/toc.2007.v003a007.
- 6 Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *CRYPTO 2022*, volume 13507 of *LNCS*, pages 208–236. Springer, 2022. doi:10.1007/978-3-031-15802-5_8.
- 7 Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, USA, 1st edition, 2009.
- 8 James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In *CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*, volume 12825 of *LNCS*, pages 467–496. Springer, 2021. doi:10.1007/978-3-030-84242-0_17.
- 9 Amos Beimel, Tal Malkin, and Silvio Micali. The all-or-nothing nature of two-party secure computation. In *CRYPTO '99*, volume 1666 of *LNCS*, pages 80–97. Springer, 1999. doi:10.1007/3-540-48405-1_6.
- 10 Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, and Phillip Rogaway. Everything provable is provable in zero-knowledge. In *CRYPTO '88*, volume 403 of *LNCS*, pages 37–56. Springer, 1988. doi:10.1007/0-387-34799-2_4.
- 11 Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of International Conference on Computers, Systems & Signal Processing, Dec. 9-12, 1984, Bangalore, India*, pages 175–179, 1984. arXiv:2003.06557.
- 12 Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. Practical quantum oblivious transfer. In *CRYPTO '91*, volume 576 of *LNCS*, pages 351–366. Springer, 1991. doi:10.1007/3-540-46766-1_29.
- 13 Nir Bitansky and Zvika Brakerski. Classical binding for quantum commitments. In *TCC*, volume 13042 of *LNCS*, pages 273–298. Springer, 2021. doi:10.1007/978-3-030-90459-3_10.
- 14 Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudorandom bits. *SIAM Journal on Computing*, 13(4):850–864, 1984. doi:10.1137/0213053.
- 15 Adam Boulund, Bill Fefferman, and Umesh V. Vazirani. Computational pseudorandomness, the wormhole growth paradox, and constraints on the ads/cft duality (abstract). In *11th ITCS*, volume 151 of *LIPICs*, pages 63:1–63:2. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ITCS.2020.63.
- 16 Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography, 2022. arXiv:2209.04101.
- 17 Zvika Brakerski and Omri Shmueli. (pseudo) random quantum states with binary phase. In *TCC*, volume 11891 of *LNCS*, pages 229–250. Springer, 2019. doi:10.1007/978-3-030-36030-6_10.
- 18 Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In *CRYPTO 2020*, volume 12171 of *LNCS*, pages 417–440. Springer, 2020. doi:10.1007/978-3-030-56880-1_15.
- 19 Fernando G.S.L. Brandão, Wissam Chemissany, Nicholas Hunter-Jones, Richard Kueng, and John Preskill. Models of quantum complexity growth. *PRX Quantum*, 2:030316, July 2021. doi:10.1103/PRXQuantum.2.030316.
- 20 Anne Broadbent and Alex B. Grilo. Qma-hardness of consistency of local density matrices with applications to quantum zero-knowledge. In *61st IEEE FOCS, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 196–205. IEEE, 2020. doi:10.1109/FOCS46700.2020.00027.

- 21 Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for qma. *SIAM Journal on Computing*, 49(2):245–283, 2020. doi:10.1137/18M1193530.
- 22 Adam R. Brown and Leonard Susskind. Second law of quantum complexity. *Phys. Rev. D*, 97:086015, April 2018. doi:10.1103/PhysRevD.97.086015.
- 23 André Chailloux, Gus Gutoski, and Jamie Sikora. Optimal bounds for semi-honest quantum oblivious transfer. *Chicago Journal of Theoretical Computer Science*, 2016(13), September 2016. doi:10.4086/cjtcs.2016.013.
- 24 André Chailloux, Iordanis Kerenidis, and Bill Rosgen. Quantum commitments from complexity assumptions. In *ICALP*, volume 6755 of *LNCS*, pages 73–85. Springer, 2011. doi:10.1007/978-3-642-22006-7_7.
- 25 Claude Crépeau, Jeroen 5 Graaf, and Alain Tapp. Committed oblivious transfer and private multi-party computation. In *CRYPTO '95*, volume 963 of *LNCS*, pages 110–123. Springer, 1995. doi:10.1007/3-540-44750-4_9.
- 26 Claude Crépeau and Joe Kilian. Achieving oblivious transfer using weakened security assumptions (extended abstract). In *29th FOCS, White Plains, New York, USA, 24-26 October 1988*, pages 42–52. IEEE Computer Society, 1988. doi:10.1109/SFCS.1988.21920.
- 27 Claude Crépeau, Frédéric Légaré, and Louis Salvail. How to convert the flavor of a quantum bit commitment. In *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 60–77. Springer, 2001. doi:10.1007/3-540-44987-6_5.
- 28 Claude Crépeau. Quantum oblivious transfer. *Journal of Modern Optics*, 41(12):2445–2454, 1994. doi:10.1080/09500349414552291.
- 29 W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, November 1976. doi:10.1109/TIT.1976.1055638.
- 30 Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In *CRYPTO 2012*, volume 7417 of *LNCS*, pages 794–811. Springer, 2012. doi:10.1007/978-3-642-32009-5_46.
- 31 Junbin Fang, Dominique Unruh, Jun Yan, and Dehua Zhou. How to base security on the perfect/statistical binding property of quantum bit commitment?, 2020. arXiv:2020/621.
- 32 Oded Goldreich. A note on computational indistinguishability. *Information Processing Letters*, 34(6):277–281, 1990. doi:10.1016/0020-0190(90)90010-U.
- 33 Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, pages 25–32. ACM, 1989. doi:10.1145/73007.73010.
- 34 Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. In *27th FOCS (sfcs 1986)*, pages 174–187, October 1986. doi:10.1109/SFCS.1986.47.
- 35 Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in minicrypt. In *EUROCRYPT 2021*, volume 12697 of *LNCS*, pages 531–561. Springer, 2021. doi:10.1007/978-3-030-77886-6_18.
- 36 Jonas Haferkamp, Philippe Faist, Naga B. T. Kothakonda, Jens Eisert, and Nicole Yunger Halpern. Linear growth of quantum circuit complexity. *Nature Physics*, 18(5):528–532, May 2022. doi:10.1038/s41567-022-01539-6.
- 37 Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. doi:10.1137/S0097539793244708.
- 38 Carl W. Helstrom. Quantum detection and estimation theory. *Journal of Statistical Physics*, 1(2):231–252, June 1969. doi:10.1007/BF01007479.
- 39 A.S Holevo. Statistical decision theory for quantum systems. *Journal of Multivariate Analysis*, 3(4):337–394, 1973. doi:10.1016/0047-259X(73)90028-6.
- 40 R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *Proceedings of the Twenty-First STOC, STOC '89*, pages 12–24, New York, NY, USA, 1989. ACM. doi:10.1145/73007.73009.

- 41 R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the Twenty-First STOC*, STOC '89, pages 44–61, New York, NY, USA, 1989. ACM. doi:10.1145/73007.73012.
- 42 Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In *CRYPTO 2008*, volume 5157 of *LNCS*, pages 572–591. Springer, 2008. doi:10.1007/978-3-540-85174-5_32.
- 43 Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. Qip = pspace. *J. ACM*, 58(6), December 2011. doi:10.1145/2049697.2049704.
- 44 Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *CRYPTO 2018*, volume 10993 of *LNCS*, pages 126–152. Springer, 2018. doi:10.1007/978-3-319-96878-0_5.
- 45 Akinori Kawachi, Takeshi Koshihara, Harumichi Nishimura, and Tomoyuki Yamakami. Computational indistinguishability between quantum states and its cryptographic application. *Journal of Cryptology*, 25(3):528–555, July 2012. doi:10.1007/s00145-011-9103-4.
- 46 Joe Kilian. A general completeness theorem for two party games. In *Proceedings of the Twenty-Third STOC*, STOC '91, pages 553–560, New York, NY, USA, 1991. ACM. doi:10.1145/103418.103475.
- 47 Joe Kilian, Eyal Kushilevitz, Silvio Micali, and Rafail Ostrovsky. Reducibility and completeness in private computations. *SIAM Journal on Computing*, 29(4):1189–1208, 2000. doi:10.1137/S0097539797321742.
- 48 Hirofumi Kobayashi. General properties of quantum zero-knowledge proofs. In *TCC*, volume 4948 of *LNCS*, pages 107–124. Springer, 2008. doi:10.1007/978-3-540-78524-8_7.
- 49 William Kretschmer. Quantum pseudorandomness and classical complexity. In *16th TQC*, volume 197 of *LIPICs*, pages 2:1–2:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.TQC.2021.2.
- 50 Mohammad Mahmoody and Rafael Pass. The curious case of non-interactive commitments - on the power of black-box vs. non-black-box use of primitives. In *CRYPTO 2012*, volume 7417 of *LNCS*, pages 701–718. Springer, 2012. doi:10.1007/978-3-642-32009-5_41.
- 51 Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *CRYPTO 2022*, volume 13507 of *LNCS*, pages 269–295. Springer, 2022. doi:10.1007/978-3-031-15802-5_10.
- 52 Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, January 1991. doi:10.1007/BF00196774.
- 53 Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. doi:10.1017/CB09780511976667.
- 54 Shien Jin Ong and Salil P. Vadhan. An equivalence between zero knowledge and commitments. In *TCC*, volume 4948 of *LNCS*, pages 482–500. Springer, 2008. doi:10.1007/978-3-540-78524-8_27.
- 55 R. Ostrovsky and A. Wigderson. One-way functions are essential for non-trivial zero-knowledge. In *The 2nd Israel Symposium on Theory and Computing Systems*, pages 3–17, Los Alamitos, CA, USA, June 1993. IEEE Computer Society. doi:10.1109/ISTCS.1993.253489.
- 56 RENATO RENNEN. Security of quantum key distribution. *International Journal of Quantum Information*, 06(01):1–127, 2008. doi:10.1142/S0219749908003256.
- 57 Adi Shamir. On the generation of cryptographically strong pseudorandom sequences. *ACM Trans. Comput. Syst.*, 1(1):38–44, February 1983. doi:10.1145/357353.357357.
- 58 Adi Shamir. $Ip = pspace$. *J. ACM*, 39(4):869–877, October 1992. doi:10.1145/146585.146609.
- 59 Salil P. Vadhan. An unconditional study of computational zero knowledge. *SIAM Journal on Computing*, 36(4):1160–1214, 2006. doi:10.1137/S0097539705447207.
- 60 John Watrous. Limits on the power of quantum statistical zero-knowledge. In *43rd Symposium on Foundations of Computer Science (FOCS 2002)*, 16-19 November 2002, Vancouver, BC, Canada, *Proceedings*, page 459. IEEE Computer Society, 2002. doi:10.1109/SFCS.2002.1181970.

- 61 Stefan Wolf and Jürg Wullschleger. Oblivious transfer is symmetric. In *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 222–232. Springer, 2006. doi:10.1007/11761679_14.
- 62 Jun Yan. General properties of quantum bit commitments, 2022. URL: <https://eprint.iacr.org/2020/1488>.
- 63 Jun Yan, Jian Weng, Dongdai Lin, and Yujuan Quan. Quantum bit commitment with application in quantum zero-knowledge proof (extended abstract). In *ISAAC*, volume 9472 of *LNCS*, pages 555–565. Springer, 2015. doi:10.1007/978-3-662-48971-0_47.
- 64 Andrew C. Yao. Theory and application of trapdoor functions. In *23rd FOCS (1982)*, pages 80–91, November 1982. doi:10.1109/SFCS.1982.45.