

One-Way Functions Imply Secure Computation in a Quantum World

James Bartusek* Andrea Coladangelo[†] Dakshita Khurana[‡] Fermi Ma[§]

Abstract

We prove that quantum-hard one-way functions imply *simulation-secure* quantum oblivious transfer (QOT), which is known to suffice for secure computation of arbitrary quantum functionalities. Furthermore, our construction only makes *black-box* use of the quantum-hard one-way function.

Our primary technical contribution is a construction of *extractable and equivocal* quantum bit commitments based on the black-box use of quantum-hard one-way functions in the standard model. Instantiating the Crépeau-Kilian (FOCS 1988) framework with these commitments yields simulation-secure QOT.

*UC Berkeley. Email: bartusek.james@gmail.com

[†]UC Berkeley. Email: andrea.coladangelo@gmail.com

[‡]UIUC. Email: dakshita@illinois.edu

[§]Princeton University and NTT Research. Email: fermima@alum.mit.edu

Contents

1	Introduction	3
1.1	Our Results	4
1.2	Related Work	6
1.3	Concurrent and Independent Work	6
2	Technical Overview	8
2.1	Recap: Quantum Oblivious Transfer from Commitments	8
2.2	Our Construction: A High-Level Overview	10
2.3	Making Any Quantum (or Classical) Commitment Equivocal	11
2.4	An Extractability Compiler for Equivocal Commitments	13
2.5	Putting it Together: From Commitments to Secure Computation.	14
3	Preliminaries	15
3.1	Commitments	16
3.2	Quantum Rewinding Lemma	18
3.3	Quantum Entropy and Leftover Hashing	18
4	A Quantum Equivocality Compiler	19
4.1	Equivocality	20
4.2	Extractability	22
5	Quantum Extractable Commitments	25
5.1	Extractability	25
5.2	Statistical Hiding	28
6	Quantum Oblivious Transfer from Extractable and Equivocal Commitments	30
6.1	Definitions for Oblivious Transfer with Quantum Communication	30
6.2	Our Construction	30
6.3	Receiver Security	31
6.4	Sender Security	33
7	Acknowledgments	35

1 Introduction

The complexity of cryptographic primitives is central to the study of cryptography. Much of the work in the field focuses on establishing *reductions* between different primitives, typically building more sophisticated primitives from simpler ones. Reductions imply relative measures of complexity among different functionalities, and over the years have resulted in an expansive hierarchy of assumptions and primitives, as well as separations between them.

One-way functions (OWFs) lie at the center of cryptographic complexity: their existence is the minimal assumption necessary for nearly all classical cryptography [LR86, IL89, ILL89]. One-way functions are equivalent to so-called “minicrypt” primitives like pseudorandom generators, pseudorandom functions and symmetric encryption; but provably cannot imply key exchange when used in a black-box way [IR90, BM09].¹ Thus, the existence of key exchange is believed to be a stronger assumption than the existence of one-way functions. Oblivious transfer (OT) is believed to be *even stronger*: it implies key exchange, but cannot be obtained from black-box use of a key exchange protocol [MMP14].

The importance of OT stems from the fact that it can be used to achieve secure computation, which is a central cryptographic primitive with widespread applications. In a nutshell, secure computation allows mutually distrusting participants to compute any public function over their joint private inputs while revealing no private information beyond the output of the computation.

The Quantum Landscape. The landscape of cryptographic possibilities changes significantly when participants have quantum computation and communication capabilities. For one, *unconditionally* secure key distribution — commonly known as *quantum key distribution* (QKD) — becomes possible [BB84]. Moreover, *quantum* oblivious transfer (QOT) is known to be achievable from special types of commitments, as we discuss next.

Crépeau and Kilian [CK88] first proposed a protocol for QOT using quantum bit commitments. The central idea in these QKD and QOT protocols is the use of (what are now known as) “BB84 states”. These are single qubit states encoding either 0 or 1 in either the computational or Hadamard basis. Crucially, measuring (or essentially attempting to copy the encoded bit) in the wrong basis completely destroys information about the encoded bit. Then [BBCS92] presented a transmission-error resistant version of the [CK88] protocol. These protocols did not come with a proof of security, but subsequently Mayers and Salvail [MS94] proved that the [CK88] protocol is secure against a restricted class of attackers that only perform single-qubit measurements. This was later improved by Yao [Yao95], who extended the [MS94] result to handle general quantum adversaries.

By an unfortunate historical accident, the aforementioned security proofs claimed the [CK88] QOT could be *information-theoretically secure*, since at the time it was believed that information-theoretic quantum bit commitment was possible [BCJL93]. Several years later, Mayers [May97] and Lo and Chau [LC97] independently proved the impossibility of information-theoretic quantum bit commitment, and as a consequence, the precise security of [CK88] QOT was once again unclear. This state of affairs remained largely unchanged until 2009, when Damgård, Fehr, Lunemann, Salvail, and Schaffner [DFL⁺09] proved that bit commitment schemes satisfying certain additional properties, namely *extraction and equivocation*, suffice to instantiate [CK88] QOT. [DFL⁺09] called their commitments *dual-mode* commitments, and provided a construction based on the quantum hardness of the learning with errors (QLWE) assumption. We remark that assumptions about the hardness of specific problems like QLWE are qualitatively even worse than general assumptions like QOWFs and QOT. Thus, the following basic question remains:

Do quantum-hard one-way functions suffice for quantum oblivious transfer?

¹In particular, [IR90, BM09] showed that there cannot exist a key exchange protocol that only has oracle access to the input/output behavior of a one-way function, and makes no additional assumptions. Then [Dac16] ruled out the possibility of certain types of key exchange protocols that also make use of the *code* of a one-way function. Constructions of key exchange from one-way functions have eluded researchers for decades. This, combined with the aforementioned negative results, is considered to be evidence that key exchange is a qualitatively stronger primitive than one-way functions in the classical regime. In fact, Impagliazzo [Imp95] stipulates that we live in one of five possible worlds, of which *Minicrypt* is one where classical one-way functions exist but classical public-key cryptography does not.

Quantum OT: The Basis of Secure Quantum Computation. There is a natural extension of secure computation to the quantum world, where Alice and Bob wish to compute a *quantum* circuit on (possibly entangled) *quantum* input states. This setting, usually referred to as secure *quantum* computation, has been previously studied and in fact has a strong tradition in the quantum cryptography literature.

[CGS02, BCG⁺06] constructed unconditional maliciously-secure *multi-party* quantum computation with honest majority. The setting where half (or more) of the players are malicious requires computational assumptions due to the impossibility of unconditionally secure quantum bit commitment [May97, LC97].

In this computational setting, [DNS10, DNS12] showed the feasibility of two-party quantum computation (2PQC) assuming post-quantum OT. More recently, [DGJ⁺20] constructed maliciously-secure general multi-party quantum computation (MPQC) secure against a *dishonest* majority from any maliciously-secure post-quantum multi-party computation (MPC) protocol for classical functionalities, which can itself be obtained from post-quantum OT [ABG⁺20].

Nevertheless, the following natural question has remained unanswered:

Can secure (quantum) computation be obtained from quantum-hard one-way functions?

1.1 Our Results

Our main result is the following:

Quantum oblivious transfer can be based on the assumption that quantum-hard one-way functions exist.

In fact, we prove a stronger result: we show that quantum oblivious transfer can be based on the *black-box use of any statistically binding, quantum computationally hiding commitment*. Such commitments can be based on the black-box use of quantum-hard one-way functions. This in turn implies secure two-party computation of classical functionalities, in the presence of quantum computation and communication capabilities, from (black-box use of) quantum-hard one-way functions [Kil88]. The latter can then be used to obtain secure two-party *quantum* computation, by relying on the work of [DNS12]. Quantum OT can also be used to obtain *multi-party* secure computation of all classical functionalities, in the presence of quantum computation and communication capabilities, and additionally assuming the existence of authenticated channels. This follows from the techniques in [Kil88, CvT95, DI05, IPS08] which obtain classical MPC based on black-box use of any OT protocol. By relying on [DGJ⁺20], this also implies multi-party secure *quantum* computation.

In summary, our main result implies that: (1) 2PQC can be obtained from (black-box use of) quantum-hard OWFs and (2) assuming the existence of authenticated channels, MPQC can be obtained from (black-box use of) quantum-hard OWFs.

This gives a potential separation between the complexity of cryptographic primitives in the classical and quantum worlds. In the former, (two-party) secure computation provably cannot be based on black-box use of quantum-hard one-way functions. It is only known from special types of enhanced public-key encryption schemes or from the hardness of specific problems, both of which are believed to be much stronger assumptions than one-way functions. But in the quantum world, prior to our work, (two-party) secure computation was only known from the special commitments required in the protocol of [DFL⁺09], which can be based on QLWE following [DFL⁺09], or post-quantum OT (implicit in [HSS11, BS20, ABG⁺20]) — but were not known to be achievable from quantum-hard one-way functions.

On the Significance of the Black-Box use of Cryptography in the Quantum Setting. Making black-box use of a cryptographic primitive refers to only having oracle access to its input/output behavior, without having the ability to examine the actual code (i.e., representation as a sequence of symbols) of the primitive. For instance, proving in zero-knowledge that a committed value satisfies some given predicate often requires explicit knowledge of the commitment algorithm; classifying the resulting proof as making “non-black-box” use of the one-way function. In the literature, constructions that make black-box use of cryptographic primitives are often preferred over those that make non-black-box use of cryptography. Besides their conceptual simplicity and elegance, black-box constructions are also of practical interest since

they avoid expensive NP reductions involving circuits of primitives. Perhaps most importantly, in the case of black-box constructions, one can instantiate the underlying primitive with an arbitrary implementation, including physical implementations via secure hardware or those involving quantum communication.

In many quantum protocols, which involve quantum states being transferred between two or more players, black-box constructions are not only significantly preferable but often become *a necessity*. Let us illustrate this with an example. The GMW protocol [GMW87] first showed that secure multi-party computation can be based on any oblivious transfer protocol; however the protocol involved zero-knowledge proofs involving the description of the (classical) oblivious transfer protocol. Due to the GMW [GMW87] protocol being non-black-box in the underlying OT, our OT protocols cannot be used with GMW to obtain multi-party computation of classical functionalities. We instead need to rely on compilers like [Kil88, CvT95, DI05, IPS08] that only make *black-box use* of the underlying OT protocol. As discussed above, the black-box nature of these compilers makes them applicable irrespective of whether they are instantiated with classical or QOT.

In a similar vein, we believe that our *black-box use* of any statistically binding, quantum computationally hiding commitment in our QOT protocol is of particular significance. For instance, one can substitute our statistically binding, quantum computationally hiding commitment with an unconditionally secure one in the quantum random oracle model [Unr14], resulting in *unconditional quantum OT in the quantum random oracle model*. Moreover if in the future, new constructions of statistically binding, quantum computationally hiding commitments involving quantum communication are discovered based on assumptions weaker than quantum-hard one-way functions, it would be possible to plug those into our protocol compilers to obtain QOT. These applications would not have been possible had we required non-black-box use of the underlying commitment.

Primary Tool: Stand-alone Extractable and Equivocal Commitments. As discussed earlier, [DFL+09] show that simulation-secure QOT can be obtained from commitments satisfying certain properties, namely *extraction* and *equivocation*.

- At a high level, extraction requires that there exist an efficient quantum “extractor” that is able to extract a committed message from any quantum committer.
- Equivocality requires that there exist an efficient quantum “equivocator” capable of simulating an interaction with any quantum receiver such that it can later open the commitment to any message of its choice.

These two properties are crucial for proving simulation security of the [CK88] OT protocol: extraction implies receiver security and equivocality implies sender security². Our key technical contribution is the following:

Extractable and equivocal commitments can be based on the black-box use of quantum-hard one-way functions.

We obtain this result via the following transformations, each of which only makes *black-box use* of the underlying primitives.

- *Step 1: Quantum Equivocal Commitments from Quantum-Hard One-Way Functions.* We describe a generic unconditional compiler to turn any commitment into an equivocal commitment in the plain model. By applying our compiler to Naor’s statistically binding commitment [Nao91] — which can be based on quantum-hard one-way functions — we obtain a statistically binding, equivocal commitment.
- *Step 2: Quantum Extractable Commitments from Quantum Equivocal Commitments.* We show that the [CK88, DFL+09, BF10] framework can be used to obtain an extractable commitment that leverages quantum

²It is important to note that extraction and equivocation are only made possible in an ideal world where a simulator has access to the adversary’s state. Participants in the real protocol cannot access each others’ state, which prevents them from extracting or equivocating.

communication, and can be based on the existence of any quantum equivocal commitment. This combined with the previous step implies the existence of quantum extractable commitments based on the existence of quantum-hard one-way functions.

This is in contrast to existing approaches (eg., [HSS11]) that require classical communication but rely on qualitatively stronger assumptions like classical OT with post-quantum security.

- *Step 3: From Extractable Commitments to Extractable and Equivocal Commitments.* We apply the black-box equivocality compiler from the first step to the quantum extractable commitment obtained above, to produce an extractable and equivocal commitment.

We point out that it is generally straightforward to make a classical commitment equivocal using zero-knowledge proofs, but this approach does not apply to quantum commitment protocols. We therefore devise our own equivocality compiler capable of handling quantum commitments and use it in both Step 1 and Step 3.

Plugging our quantum extractable and equivocal commitments into the [CK88] framework yields a final QOT protocol with an interaction pattern that readers familiar with [BB84, CK88] may find interesting: the sender sends the receiver several BB84 states, after which the receiver proves to the sender that it has honestly measured the sender’s BB84 states by *generating more BB84 states of their own and asking the sender to prove that they have measured the receiver’s BB84 states*. An intriguing open question is whether obtaining QOT from one-way functions *requires* this type of two-way quantum communication or, alternatively, quantum memory.³

1.2 Related Work

For some readers, it may appear that the central claim of this work — that quantum-hard one-way functions suffice for oblivious transfer — has already been established [DFL⁺09, BF10]. Indeed, prior work [DFL⁺09, BF10] showed that statistically binding and computational hiding commitments (which are weaker than extractable and equivocal commitments), known to exist from one-way functions, can be plugged into the [CK88] template to achieve an oblivious transfer protocol satisfying *indistinguishability-based* security.

However, the indistinguishability-based security definition for oblivious transfer is *not* standard in the cryptographic literature. When cryptographers refer to “oblivious transfer”, they almost always mean the standard *simulation-based* security notion. Indeed, the fundamental importance of oblivious transfer in modern cryptography is due to the fact that it is necessary and sufficient for secure computation, but this is only true for the simulation-based notion.

Therefore, the goal of this work is to build simulation-secure oblivious transfer, and thereby establish that quantum-hard one-way functions suffice for secure quantum computation. We elaborate further on the details of simulation security in Section 2.

We also note that in the setting of Universal Composability (UC), Unruh [Unr07] obtained quantum oblivious transfer from quantum UC-secure commitments in the common reference string (CRS) model. By contrast, in this work, we focus on the standalone setting in the plain/standard model, and we do not assume the existence of trusted setup or a common reference string. In this setting, we obtain quantum oblivious transfer from quantum-hard one-way functions.

1.3 Concurrent and Independent Work

In a concurrent and independent work, Grilo, Lin, Song, and Vaikuntanathan [GLSV20] also construct simulation-secure quantum oblivious transfer from quantum-hard one way functions via the intermediate primitive of extractable and equivocal commitments. However, the two works take entirely different approaches to constructing these commitments. We briefly summarize these strategies below.

³Naive approaches to removing one direction of quantum communication appear to require the honest parties to be entangled and subsequently perform quantum teleportation.

This work:

1. Construct equivocal commitments from statistically binding commitments via a new “equivocality compiler” based on Watrous [Wat06] rewinding.
2. Construct extractable commitments from equivocal commitments via a new “extractability compiler” based on the [CK88] template.
3. Construct extractable and equivocal commitments from extractable commitments via the same compiler from Step 1.

[GLSV20]:

1. Construct selective opening secure commitments with inefficient simulation against malicious committers from statistically binding commitments and zero-knowledge proofs.
2. Construct QOT with inefficient simulation against malicious receivers from selective opening secure commitments with inefficient simulation against malicious committers, following the [CK88] QOT template.⁴
3. Construct parallel QOT with inefficient simulation against malicious receivers from (stand-alone) QOT with inefficient simulation against malicious receivers via a new lemma for parallel repetition of protocols.
4. Construct verifiable conditional disclosure of secrets, a new primitive introduced in [GLSV20], from parallel QOT with inefficient simulation against malicious receivers, statistically binding commitments, Yao’s garbled circuits, and zero-knowledge proofs.
5. Construct extractable commitments from verifiable conditional disclosure of secrets, statistically binding commitments, and zero-knowledge proofs.
6. Construct extractable and equivocal commitments from extractable commitments and zero-knowledge proofs.

We believe that our result is easier to understand and conceptually simpler, as we do not need to define additional primitives beyond extractable and/or equivocal commitments. Aside from differences in approach, there are several other places where the results differ:

- **This Work: Black-Box Use of One-Way Functions.** A significant advantage of our work over [GLSV20] is that we construct quantum OT from *black-box use* of statistically binding commitments or one-way functions. The OT in [GLSV20] makes non-black-box use of the underlying one-way function. As discussed above, making black-box use of underlying cryptographic primitives is particularly useful in the quantum setting. Due to the extensive use of zero-knowledge proofs and garbled circuits in [GLSV20], it appears difficult to modify their approach to be black-box in the underlying one-way function.
- **This Work: One-Sided Statistical Security.** Additionally, our oblivious transfer protocol offers *one-sided statistical security*. As written, our quantum OT protocol satisfies statistical security against malicious senders (and computational security against malicious receivers). Moreover, this OT can be reversed following the techniques in eg., [WW06] to obtain a quantum OT protocol that satisfies statistical security against malicious receivers (and computational security against malicious senders). On the other hand, the quantum OT protocol in [GLSV20] appears to achieve computational security against both malicious senders and malicious receivers.

⁴[GLSV20] point out that the conclusions of Steps 1 and 2 together had also been established in prior works of [Unr07, DFL⁺09, BF10].

- **[GLSV20]: Verifiable Conditional Disclosure of Secrets.** Towards achieving their main result, [GLSV20] introduce and construct verifiable conditional disclosure of secrets (vCDS). This primitive may be of independent interest.
- **[GLSV20]: Constant Rounds in the CRS Model.** While both works construct $\text{poly}(\lambda)$ -round protocols in the plain model, [GLSV20] additionally construct a constant round OT protocol in the CRS model based on (non-black-box use of) quantum-hard one-way functions.

In an earlier version of this work, we did not consider the CRS model. After both works were posted to the Cryptology ePrint Archive, we realized that our techniques could be adapted straightforwardly to achieve constant round complexity in the CRS model, while still making black-box use of one-way functions. We elaborate further on this below.

Constant Rounds in the CRS Model via Our Techniques. We now briefly sketch how the protocols in this work can be modified to achieve constant rounds in the CRS model. The following two paragraphs are the only component of this work that was not developed concurrently and independently of [GLSV20].

In a nutshell, the reason why our protocols require super-constant rounds in the plain model is because they perform sequential repetitions of specific Σ -protocols to enable rewinding. In the presence of a CRS, it is possible to repeat our protocols in parallel and simulate by means of a trapdoor embedded in the CRS. This avoids the need for rewinding (and therefore sequential repetition), and would bring down the number of rounds from $O(\lambda)$ to $O(1)$ in the CRS model.

In fact, constant round QOT can also be based on the *black-box use* of quantum-hard one-way functions in the presence of a CRS by slightly modifying our template, as follows. Recall that the first step towards getting a constant round protocol is to replace (all of) our sequential repetitions with parallel repetitions. The challenge for the parallel repetitions can be determined as the outcome of a (one-sided simulatable) coin toss. This coin toss can be carefully designed so that a simulator that programs the CRS can program the output of the coin toss, which would allow the simulator to program the challenge message for the parallel repetition. In the CRS model, such a coin toss can be obtained by having one player commit to a random set of coins via Naor’s commitment [Nao91] (where the first message of the commitment is the CRS). Next, the second player sends their random coins, in the clear, to the first player. Finally, the first player decommits their initial commitment string, and the players set the output of the coin toss to be the XOR of the random coins chosen by each player. A simulator can program the output of the coin toss by programming the CRS to equivocate Naor’s commitment. This programming would allow one to simulate the parallel repeated versions of our protocols. We note that this protocol has the drawback of the CRS not being reusable, unlike the CRS-based protocol in [GLSV20]. We stress that our focus is on the plain model, and we leave formal explorations/optimizations of this and related approaches to building black-box OT in the CRS model to future work.

2 Technical Overview

This work establishes that (1) black-box use of post-quantum one-way functions suffices for post-quantum *extractable and equivocal* commitment schemes and moreover, that (2) [CK88] quantum oblivious transfer instantiated with such commitments is a standard *simulation-secure* oblivious transfer. Crucially, the standard notion of simulation-secure (quantum) oblivious transfer that we achieve is sequentially composable and suffices to achieve general-purpose secure quantum computation. Before explaining our technical approach, we provide a complete review of the original [CK88] protocol.

2.1 Recap: Quantum Oblivious Transfer from Commitments

In quantum oblivious transfer (QOT), a quantum sender holding two classical messages m_0, m_1 engages in an interactive protocol over a quantum channel with a quantum receiver holding a classical choice bit

b . Correctness requires the receiver to learn m_b by the end of the protocol. Informally, security demands that a malicious receiver only learn information about one of m_0, m_1 , and that a malicious sender learn nothing about b . Somewhat more formally, as discussed earlier, our focus is on the standard simulation-based notion of security. This stipulates the existence of an efficient quantum simulator that generates the view of an adversary (sender/receiver) when given access to an ideal OT functionality. In particular, when simulating the view of a malicious sender, this simulator must extract the sender's inputs (m_0, m_1) without knowledge of the receiver's input b . And when simulating the view of a malicious receiver, the simulator must extract the receiver's input b , and then simulate the receiver's view given just m_b .

We recall the construction of quantum oblivious transfer due to [CK88] (henceforth CK88), which combines the information theoretic quantum key distribution protocol of [BB84] (henceforth BB84) with cryptographic bit commitments.

CK88 First Message. The first message of the CK88 protocol exactly follows the beginning of the BB84 protocol. For classical bits y, z , let $|y\rangle_z$ denote $|y\rangle$ if $z = 0$, and $(|0\rangle + (-1)^y |1\rangle)/\sqrt{2}$ if $z = 1$, i.e. the choice of z specifies whether to interpret y as a computational or Hadamard basis vector. Let λ denote the security parameter. The sender samples two random 2λ -bit strings x and θ , and constructs "BB84 states" $|x_i\rangle_{\theta_i}$ for $i \in [2\lambda]$. The sender forwards these 2λ BB84 states $(|x_i\rangle_{\theta_i})_{i \in [2\lambda]}$ to the receiver. Next, the receiver samples a 2λ -bit string $\hat{\theta}$, measures each $|x_i\rangle_{\theta_i}$ in the basis specified by $\hat{\theta}_i$, and obtains a 2λ -bit measurement result string \hat{x} .

CK88 Measurement-Check Subprotocol. At this point, the CK88 and BB84 protocols diverge. Since the BB84 protocol is an interaction between two *honest* parties, it assumes the parties comply with the protocol instructions. However, in the CK88 protocol, a malicious receiver who does not measure these BB84 states will be able to compromise sender privacy later in the protocol. Therefore, the next phase of CK88 is a measurement-check subprotocol designed to catch a malicious receiver who skips the specified measurements. This subprotocol requires the use of a quantum-secure classical commitment scheme; for the purposes of this recap, one should imagine a commitment with idealized hiding and binding properties. The subprotocol proceeds as follows:

- For each $i \in [2\lambda]$, the receiver commits to $(\hat{\theta}_i, \hat{x}_i)$.
- Next, the sender picks a random set T of λ indices from $[2\lambda]$, and challenges the receiver to open the corresponding commitments.
- The receiver sends $(\hat{\theta}_i, \hat{x}_i)$ along with the corresponding opening for each $i \in T$.
- The sender verifies each commitment opening, and furthermore checks that $\hat{x}_i = x_i$ for each $i \in T$ where $\hat{\theta}_i = \theta_i$. If any of these checks fail, the sender aborts.

The rough intuition for the subprotocol is simple: from the receiver's point of view, the BB84 states are maximally mixed and therefore completely hide x_i and θ_i . For any index i that the receiver does not measure, it must guess \hat{x}_i . From the receiver's perspective, the sender checks \hat{x}_i against x_i if two $1/2$ -probability events occur: (1) i is included in T , and (2) $\hat{\theta}_i = \theta_i$. This means a malicious receiver who skips a significant number of measurements will be caught with overwhelming probability.

CK88 Privacy Amplification. If all the subprotocol checks pass, the sender continues to the final stage of the CK88 protocol. For convenience, relabel the λ indices in $[2\lambda] \setminus T$ from 1 to λ ; all indices corresponding to opened commitments are discarded for the remainder of the protocol.

For each $i \in [\lambda]$, the sender reveals the correct measurement basis θ_i . The receiver then constructs the index set I_b — where b is its choice bit for the oblivious transfer — as the set of all $i \in [\lambda]$ where $\theta_i = \hat{\theta}_i$. It sets I_{1-b} to be the remaining indices, and sends (I_0, I_1) to the sender. Note that by the hiding property of

the commitments, the sender should not be able to deduce b from (I_0, I_1) ; furthermore, I_0 and I_1 will both be close to size $\lambda/2$, since for each $i \in [\lambda]$, the receiver committed to $\hat{\theta}_i$ before obtaining θ_i .

On receiving I_0, I_1 , the sender sets $x_0 := (x_i)_{i \in I_0}$ and $x_1 := (x_i)_{i \in I_1}$. The intuition is that if a receiver honestly constructs (I_0, I_1) , it will only have information about x_b corresponding to its choice bit b . However, it turns out that even if the receiver maliciously constructs (I_0, I_1) , at least one of x_0 and x_1 will have high min-entropy from its point of view. Thus, by standard privacy amplification techniques, the sender can complete the oblivious transfer as follows. It samples two universal hash functions h_0 and h_1 , both with ℓ -bit outputs, and uses $h_0(x_0)$ to mask the ℓ -bit message m_0 , and uses $h_1(x_1)$ to mask m_1 . That is, the sender sends $(h_0, h_1, h_0(x_0) \oplus m_0, h_1(x_1) \oplus m_1)$ to the receiver, who can then use x_b to recover m_b . Since x_{1-b} will have high entropy, the leftover hash lemma implies that $h_{1-b}(x_{1-b})$ is statistically close to uniform, which hides m_{1-b} from the receiver.

Simulation-Based Security. Turning this intuition into a proof of simulation-based security of the resulting QOT requires some additional insights [DFL⁺09], and requires the commitments used in the measurement-check subprotocol to satisfy two additional properties: extractability and equivocality. In what follows, we briefly summarize why these properties help achieve simulation-based security.

To argue that the resulting QOT protocol satisfies security against a malicious sender, one must demonstrate the existence of a simulator that simulates the sender’s view by generating messages on behalf of an honest receiver, and extracts both QOT inputs of the sender⁵. Now, the measurement-check subprotocol described above is designed to ensure that at least one of the sender’s inputs is hidden from a receiver. To nevertheless enable the simulator to extract both sender inputs, the idea in [DFL⁺09] is to substitute the commitments used in the measurement-check subprotocol with *equivocal commitments* that allow the simulator to later open these commitments to any value of its choice. This enables the simulator to defer any measurements until after it obtains the set T from the sender, and then selectively measure *only* the states that correspond to indices in T . All other states are left untouched until the sender reveals its measurement bases in the final stage of the CK88 protocol. Upon obtaining the sender’s “correct” measurement bases, the simulator measures all the remaining states in the correct bases, allowing it to learn both the inputs of the sender.

To demonstrate that the resulting QOT protocol satisfies security against a malicious receiver, one must demonstrate the existence of a simulator that simulates the receiver’s view by generating messages on behalf of an honest sender, and extracts the receiver’s choice bit. Again by design, the measurement-check subprotocol ensures that the receiver’s choice bit hidden is hidden from the sender. To nevertheless enable the simulator to extract this choice bit, [DFL⁺09] modify the commitments in the measurement-check subprotocol so that the simulator is able to extract all of the $\{(\hat{\theta}_i, \hat{x}_i)\}_{i \in [2\lambda]}$ from the receiver’s commitments. This enables the simulator to compute which one of the sets I_0, I_1 contain more indices i for which $\theta_i = \hat{\theta}_i$; clearly the set with more indices corresponds to the receiver’s choice bit. In summary, the key tool that enables simulation against a malicious receiver is an *extractable* commitment, that forces the receiver to use commitments for which the simulator can extract the committed value, without ever running the opening phase.

To conclude, following [DFL⁺09] the CK88 protocol can be shown to satisfy simulation-based security as long as the commitments used in the measurement-check subprotocol satisfy both the *extractability* and *equivocality* properties that were informally described above.

With this in mind, we now describe our primary technical contribution: a construction of the required extractable and equivocal commitments based on black-box use of quantum-hard one-way functions.

2.2 Our Construction: A High-Level Overview

The rest of this technical overview describes our *black-box* construction of simultaneously *extractable* and *equivocal* quantum bit commitments from any quantum-hard one-way function.

⁵We refer the reader to Section 6.1 for a formal definition of simulation-based QOT.

The ingredients for our construction are the following:

- A general-purpose “equivocality compiler” that turns any bit commitment scheme — classical or quantum — into an *equivocal* quantum commitment scheme. Moreover, if the original commitment scheme is *extractable*, this compiler outputs an *extractable and equivocal* commitment scheme.
- A general-purpose “extractability compiler” that turns any *equivocal* bit commitment scheme — classical or quantum — into an *extractable but not equivocal* commitment scheme.

Both of these compilers require no additional computational assumptions beyond those of the original commitment schemes. Given these compilers, we build extractable and equivocal commitments via the following steps:

- **Instantiation:** Begin with Naor’s statistically-binding, computationally hiding commitments [Nao91]. Naor’s construction makes black-box use of one-way functions and achieves post-quantum computational hiding assuming post-quantum security of the one-way function.⁶
- **Step 1:** Plug Naor’s commitments into our equivocality compiler to obtain an *equivocal* quantum bit commitment scheme.
- **Step 2:** Feed the resulting equivocal quantum bit commitments into our extractability compiler to obtain an *extractable but not equivocal* quantum bit commitment.
- **Step 3:** Run the equivocality compiler *a second time*, but now starting with the extractable commitments produced by the previous step. This gives the desired *extractable and equivocal* quantum bit commitments.

We briefly remark that if one did not care about obtaining a construction that only makes *black-box* use of one-way functions, it would only be necessary to invoke the equivocality compiler once. In slightly more detail, we could design an alternative extractability compiler that works with any (not necessarily equivocal) bit commitment scheme by additionally relying on post-quantum zero-knowledge proofs [Wat06]; this approach would still give a construction of quantum oblivious transfer from post-quantum one-way functions, but the zero-knowledge proofs would be used to prove statements involving the *description* of some one-way function, which is a non-black-box use of the primitive. We believe that constructing QOT from quantum-hard one-way functions is interesting regardless of whether or not the construction is black-box. However, as discussed in Section 1.1, black-box constructions are significantly preferred over non-black-box ones, especially in the quantum setting. Therefore our focus in the rest of this overview will be on presenting a black-box construction.

Organization. We describe our equivocality compiler in Section 2.3 and our extractability compiler in Section 2.4. In Section 2.5, we briefly discuss implications for secure computation in a quantum world.

2.3 Making Any Quantum (or Classical) Commitment Equivocal

Recall that a quantum commitment protocol is *equivocal* if an efficient quantum algorithm called the *equivocator*, with access to the receiver, can generate commitments that can be opened to any value. More precisely, for any receiver (modeled as an efficient malicious quantum algorithm), there must exist an equivocator

⁶In slightly more detail, Naor’s commitment scheme makes black-box use of any pseudo-random generator (PRG). It is straightforward to verify that if the PRG is post-quantum secure, the commitment satisfies computational hiding against quantum attackers. A black-box construction of pseudo-random generators from one-way functions is due to [HILL99]; Aaronson [Aar09] and Zhandry [Zha12] observed that [HILL99] applies to non-uniform quantum attackers with *classical* advice. This can be extended to handle non-uniform quantum advice by giving the one-way function attacker constructed in the [HILL99] reduction many copies of the PRG attacker’s non-uniform quantum advice (which only requires some polynomial upper bound on the number of times the reduction invokes the PRG attacker).

who can generate a computationally indistinguishable commitment that the equivocator can later open arbitrarily.

In this subsection, we describe a *black-box* compiler for a fairly general task (which may be of independent interest): making any *classical or quantum* commitment equivocal. Recall from Section 2.2 that we will need to invoke our equivocality compiler *twice*, once on a classical bit commitment scheme, and once on an extractable quantum bit commitment scheme; in the latter case, our compiler will need to preserve the extractability of the original commitment. Since classical commitments are a subclass of quantum commitments, our exposition will focus on challenges unique to the quantum setting.

Why Existing Classical Solutions Are Insufficient. Let us briefly relax our goal of using one-way functions in an exclusively black-box way. Then there is a simple equivocality compiler that applies to any (statistically-binding and computationally-hiding) classical commitment. Recall that quantum-hard one-way functions imply post-quantum zero-knowledge proofs for NP [Wat06]. Now, in the opening phase of the commitment, to open to a bit b , the committer will not send the randomness used to commit to b in the clear. Instead, the committer sends a zero-knowledge proof for the NP statement that there exist randomness consistent with the commitment, and opening to b . Equivocation is achieved by simulating the zero-knowledge proof. Unfortunately, this technique fails when the commitment involves *quantum* communication, since the statement to be proven is no longer an NP statement (nor a QMA statement). Therefore, we construct a compiler that only makes black-box use of the underlying commitment, and leverages Watrous’s rewinding lemma in the proof of equivocality [Wat06]. Our ideas are reminiscent of those developed in the classical setting, eg. in [PW09] and references therein.

Our Equivocality Compiler. In our construction, to commit to a bit b , the committer and receiver will perform λ sequential repetitions of the following subprotocol:

- The (honest) committer samples 2 uniformly random bits u_0, u_1 , and commits to each one *twice* using the base commitment scheme. Let the resulting commitments be $\mathbf{c}_0^{(0)}, \mathbf{c}_0^{(1)}, \mathbf{c}_1^{(0)}, \mathbf{c}_1^{(1)}$, where the first two are to u_0 and the second two are to u_1 . Note that since the base commitment scheme can be an arbitrary quantum interactive commitment, each commitment $\mathbf{c}_{b_1}^{(b_2)}$ corresponds to the receiver’s quantum state after the commitment phase of the base commitment.
- The receiver sends the committer a random challenge bit β .
- The committer opens the two base commitments $\mathbf{c}_\beta^{(0)}, \mathbf{c}_\beta^{(1)}$. If the openings are invalid or the revealed messages are different, the receiver aborts the entire protocol.

If these λ executions pass, the receiver is convinced that a majority of the committer’s remaining 2λ unopened commitments are honestly generated, i.e. most pairs of commitments are to the same bit.

Rewriting the (honest) committer’s unopened bits as u_1, \dots, u_λ , the final step of the commitment phase is for the committer to send $h_i := u_i \oplus b$ for each $i \in [\lambda]$ (recall that b is the committed bit).

To decommit, the committer reveals each u_i by picking one of the two corresponding base commitments at random, and opening it. The receiver accepts if each one of the base commitment openings is valid, and the opened u_i satisfies $h_i \oplus u_i = b$ for every i .

The (statistical) binding property of the resulting commitment can be seen to follow from the (statistical) binding of the underlying commitment. For any commitment, define the unique committed value as the majority of $(h_i \oplus u_i)$ values in the unopened commitments, setting u_i to \perp if both committed bits in the i^{th} session differ. Due to the randomized checks by the receiver, any committer that tries to open to a value that differs from the unique committed value will already have been caught in the commit phase, and the commitment will have been rejected with overwhelming probability. A similar argument also allows us to establish that this transformation preserves extractability of the underlying commitment. We now discuss why the resulting commitment is *equivocal*.

Quantum Equivocation. The natural equivocation strategy should have the equivocator (somehow) end up with λ pairs of base commitments where for each $i \in [\lambda]$, the pair of commitments is to u_i and $1 - u_i$ for some random bit u_i . This way, it can send an appropriately distributed string h_1, \dots, h_λ , and later open to any b by opening the commitment to $b \oplus h_i$ for each i .

We construct our equivocator using Watrous’s quantum rewinding lemma [Wat06] (readers familiar with Watrous’s technique may have already noticed our construction is tailored to enable its use).

We give a brief, intuition-level recap of the rewinding technique as it pertains to our equivocator. Without loss of generality, the malicious quantum receiver derives its challenge bit β by performing some binary outcome measurement on the four quantum commitments it has just received (and on any auxiliary states). Our equivocator succeeds (in one iteration) if it can prepare four quantum commitments $\mathbf{c}_0^{(0)}, \mathbf{c}_0^{(1)}, \mathbf{c}_1^{(0)}, \mathbf{c}_1^{(1)}$ where:

1. $\mathbf{c}_\alpha^{(0)}, \mathbf{c}_\alpha^{(0)}$ are commitments to the same random bit,
2. $\mathbf{c}_{1-\alpha}^{(0)}, \mathbf{c}_{1-\alpha}^{(0)}$ are commitments to a random bit and its complement,
3. on input $\mathbf{c}_0^{(0)}, \mathbf{c}_0^{(1)}, \mathbf{c}_1^{(0)}, \mathbf{c}_1^{(1)}$, the receiver produces challenge bit $\beta = \alpha$.

That is, the equivocator is successful if the receiver’s challenge bit β corresponds to the bit α that it can open honestly. Watrous’s [Wat06] rewinding lemma applies if the distribution of β is *independent* of the receiver’s choice of α , which is guaranteed here by the hiding of the base commitments. Thus, the rewinding lemma yields a procedure for obtaining an honest-looking interaction where all three properties above are met. Given the output of the rewinding process, the equivocator has successfully “fooled” the committer on this interaction and proceeds to perform this for all λ iterations. As described above, fooling the committer on all λ iterations enables the equivocator to later open the commitment arbitrarily.

2.4 An Extractability Compiler for Equivocal Commitments

In this subsection, we compile any classical or quantum *equivocal* bit commitment into a quantum *extractable* bit commitment. We stress that even though this compiler is applied to equivocal bit commitments, the resulting commitment is *not* guaranteed to be simultaneously *extractable and equivocal*; we refer the reader to Section 2.2 for details on how this compiler fits into our final construction. Recall that a commitment scheme is *extractable* if for any adversarial quantum committer that successfully completes the commitment phase, there exists an efficient quantum algorithm (called the extractor) which outputs the committed bit.

Our construction roughly follows the CK88 template for oblivious transfer, viewing the OT sender as the committer, and having the receiver use the equivocal commitment to commit to its basis choices and measurement results. We note that this is similar in spirit to a construction of statistically-hiding, computationally-binding commitment from a statistically-binding, computationally-hiding commitment (plus quantum communication) given by [CLS01].

Construction. The committer, who intends to commit to a classical bit b , begins by sampling 2λ -bit strings x and θ . It generates the corresponding 2λ BB84 states $|x_i\rangle_{\theta_i}$ and sends this to the receiver. The receiver picks 2λ random measurement bases $\hat{\theta}_i$, and measures each $|x_i\rangle_{\theta_i}$ in the corresponding basis, obtaining outcomes \hat{x}_i .

Next, the receiver and committer engage in a CK88-style measurement-check subprotocol. That is, they temporarily switch roles (for the duration of the subprotocol), and perform the following steps:

1. The receiver (acting as a committer in the subprotocol), commits to each $\hat{\theta}_i$ and \hat{x}_i (for each $i \in [2\lambda]$) with an *equivocal* commitment.
2. The committer (acting as a receiver in the subprotocol), asks the receiver to open the equivocal commitments for all $i \in T$, where $T \subset [2\lambda]$ is a random set of size λ .

3. The receiver (acting as a committer in the subprotocol) opens the λ commitments specified by T .

Provided the receiver passes the measurement-check subprotocol, the committer generates the final message of the commitment phase as follows:

- Discard the indices in T and relabel the remaining λ indices from 1 to λ .
- Partition $\{x_1, \dots, x_\lambda\}$ into $\sqrt{\lambda}$ strings $\vec{x}_1, \dots, \vec{x}_{\sqrt{\lambda}}$ each of length $\sqrt{\lambda}$.
- Sample $\sqrt{\lambda}$ universal hash functions $h_1, \dots, h_{\sqrt{\lambda}}$ each with 1-bit output.
- Finally, send $(\theta_i)_{i \in [\lambda]}, (h_j, h_j(\vec{x}_j) \oplus b)_{j \in [\sqrt{\lambda}]}$.

This concludes the commitment phase.

To decommit, the committer reveals b and $(\vec{x}_1, \dots, \vec{x}_{\sqrt{\lambda}})$. The receiver accepts if (1) for each j , the bit b and the value \vec{x}_j are consistent with the claimed value of $h_j(\vec{x}_j) \oplus b$ from the commit phase, and (2) for each index $i \in [\lambda]$ where $\theta_i = \hat{\theta}_i$, the x_i from the opening is consistent with \hat{x}_i .

Extraction. The use of equivocal commitments in the measurement-check subprotocol makes extraction simple. Given any malicious committer, we construct an extractor as follows.

The extractor plays the role of the receiver and begins an interaction with the malicious committer. But once the committer sends its 2λ BB84 states, the extractor skips the specified measurements, instead leaving these states unmeasured. Next, instead of performing honest commitments to each $\hat{\theta}_i, \hat{x}_i$, the extractor invokes (for each commitment) the equivocator algorithm of the underlying equivocal commitment scheme. Since the equivocator is guaranteed to produce an indistinguishable commitment from the point of view of any malicious receiver for the equivocal commitment, this dishonest behavior by the extractor will go undetected.

When the malicious committer responds with a challenge set $T \subset [2\lambda]$, the extractor samples uniformly random bases $\hat{\theta}_i$ for each $i \in T$, measures the corresponding BB84 states to obtain \hat{x}_i values, and sends $(\hat{\theta}_i, \hat{x}_i)_{i \in T}$. Moreover, the equivocator (for each commitment) will enable the extractor to generate valid-looking openings for all of these claimed values.

Thus, the malicious committer proceeds with the commitment protocol, and sends

$$(\theta_i)_{i \in [\lambda]}, (h_j, h_j(\vec{x}_j) \oplus b)_{j \in [\sqrt{\lambda}]}$$

to the extractor. These correspond to the λ BB84 states that the extractor has not yet measured, so it can simply read off the bases θ_i , perform the specified measurements, and extract the committer's choice of b .

Statistical Hiding. Intuitively, statistical hiding of the above commitment protocol follows because the measurement-check subprotocol forces the receiver to measure states in arbitrary bases, which destroys information about the corresponding x_i values whenever $\hat{\theta}_i \neq \theta_i$. The formal argument is a straightforward application of a quantum sampling lemma of [BF10], devised in part to simplify analysis of [CK88]-style protocols, and we defer further details to the body of the paper.

2.5 Putting it Together: From Commitments to Secure Computation.

Plugging the compilers of Sections 2.3 and 2.4 into the steps described in Section 2.2 yields a black-box construction of simultaneously extractable and equivocal quantum bit commitments from quantum-hard one-way functions. Following [DFL⁺09], these commitments can be plugged into CK88 to obtain maliciously simulation-secure QOT (see Section 6 for further details). Finally, going from QOT to arbitrary secure computation (in a black-box way) follows from a number of prior works, which we briefly survey here for the sake of completeness (see also the discussion in Section 1.1).

- **Secure Computation of Classical Functionalities.** It is well known that maliciously simulation-secure (classical) oblivious transfer can be used in a black-box way to build two-party (classical) computation [Kil88]. This also extends to the multi-party setting tolerating upto all-but-one corruptions [IPS08]. Since the constructions of [Kil88] and [IPS08] only make black-box use of the underlying oblivious transfer and the simulators are straight-line, their correctness/security guarantees continue to hold if parties are quantum and the oblivious transfer uses quantum communication. Therefore, if all parties are quantum and all pairs of parties are connected via authenticated quantum channels, they can securely compute any classical functionality assuming quantum-hard one way functions (and authenticated channels in the multi-party setting).
- **Secure Computation of Quantum Functionalities.** [DNS12] constructed a secure two-party quantum computation protocol assuming black-box access to any secure two-party classical computation protocol. [DGJ⁺20] proved the analogous statement in the multi-party setting: secure multi-party classical computation implies (in a black-box way) secure multi-party quantum computation. Therefore, by invoking the results in the previous bullet, quantum-hard one-way functions also suffice for arbitrary secure quantum computation (additionally assuming authenticated channels in the multi-party setting).

While the results in the second bullet technically subsume the first, stating them separately enables direct comparison with the classical setting. In particular, one takeaway is that while oblivious transfer is necessary and sufficient for secure computation in the classical world, quantum-hard one-way functions suffice for *exactly the same task* when parties (and their communication channels) are quantum.

3 Preliminaries

Notation. We will write density matrices/quantum random variables (henceforth, QRVs) in lowercase bold font, e.g. \mathbf{x} . A quantum register X will be written in uppercase (grey) serif font. A collection of (possibly entangled) QRVs will be written as $(\mathbf{x}, \mathbf{y}, \mathbf{z})$.

Throughout this paper, λ will denote a cryptographic security parameter. We say that a function $\mu(\lambda)$ is *negligible* if $\mu(\lambda) = 1/\lambda^{\omega(1)}$.

The trace distance between two QRVs \mathbf{x} and \mathbf{y} will be written as $\|\mathbf{x} - \mathbf{y}\|_1$. Recall that the trace distance captures the maximum probability that two QRVs can be distinguished by any (potentially inefficient) procedure. We therefore say that two infinite collections of QRVs $\{\mathbf{x}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\mathbf{y}_\lambda\}_{\lambda \in \mathbb{N}}$ are *statistically indistinguishable* if there exists a negligible function $\mu(\lambda)$ such that $\|\mathbf{x}_\lambda - \mathbf{y}_\lambda\|_1 \leq \mu(\lambda)$, and we will frequently denote this with the shorthand $\{\mathbf{x}_\lambda\}_{\lambda \in \mathbb{N}} \approx_s \{\mathbf{y}_\lambda\}_{\lambda \in \mathbb{N}}$.

Non-Uniform Quantum Advice. We will consider non-uniform quantum polynomial-time (QPT) algorithms *with quantum advice*, denoted by $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, where each \mathcal{A}_λ is the classical description of a $\text{poly}(\lambda)$ -size quantum circuit, and each ρ_λ is some (not necessarily efficiently computable) non-uniform $\text{poly}(\lambda)$ -qubit quantum advice. We remark that “non-uniform quantum polynomial-time algorithms” often means non-uniform *classical advice*, but the cryptographic applications in this work will require us to explicitly consider quantum advice.

Therefore, *computational indistinguishability* will be defined with respect to non-uniform QPT distinguishers with quantum advice. That is, two infinite collections of QRVs $\{\mathbf{x}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\{\mathbf{y}_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable if there exists a negligible function $\mu(\cdot)$ such that for all QPT distinguishers $\mathcal{D} = \{\mathcal{D}_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$,

$$|\Pr[\mathcal{D}_\lambda(\sigma_\lambda, \mathbf{x}_\lambda) = 1] - \Pr[\mathcal{D}_\lambda(\sigma_\lambda, \mathbf{y}_\lambda) = 1]| \leq \mu(\lambda).$$

We will frequently denote this with the shorthand $\{\mathbf{x}_\lambda\}_{\lambda \in \mathbb{N}} \approx_c \{\mathbf{y}_\lambda\}_{\lambda \in \mathbb{N}}$.

We will also often omit the subscript “ $\lambda \in \mathbb{N}$ ”, and simply write, for example, $\{\mathbf{x}_\lambda\}$.

3.1 Commitments

3.1.1 Bit Commitments

We define bit commitments with quantum players and quantum communication. First, we fix some notation.

A bit commitment scheme is a two-phase interactive protocol between a quantum interactive committer $\mathcal{C} = (\mathcal{C}_{\text{com}}, \mathcal{C}_{\text{open}})$ and a quantum interactive receiver $\mathcal{R} = (\mathcal{R}_{\text{com}}, \mathcal{R}_{\text{open}})$. In the commit phase, $\mathcal{C}_{\text{com}}(1^\lambda, b)$ for bit $b \in \{0, 1\}$ interacts with $\mathcal{R}_{\text{com}}(1^\lambda)$, after which \mathcal{C}_{com} outputs a state \mathbf{x}_{com} and \mathcal{R}_{com} outputs a state \mathbf{y}_{com} . We denote this interaction by $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}}) \leftarrow \langle \mathcal{C}(1^\lambda, b), \mathcal{R}(1^\lambda) \rangle$.

To be precise, without loss of generality, we model each party as a sequence of unitaries and measurements acting on a local register, and a shared message register. We think of $\langle \mathcal{C}(1^\lambda, b), \mathcal{R}(1^\lambda) \rangle$ as the distribution over pairs of states $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}})$ induced by the measurements of the two parties in the protocol, i.e. each sequence of measurement outcomes by the parties in the protocol defines a leftover joint pure state $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}})$ (after the commit phase), with an associated probability, namely the probability of the particular sequence of measurement outcomes.

In the open phase, $\mathcal{C}_{\text{open}}(\mathbf{x}_{\text{com}})$ interacts with $\mathcal{R}_{\text{open}}(\mathbf{y}_{\text{com}})$, after which $\mathcal{R}_{\text{open}}$ either outputs a bit b' or \perp . We will denote this receiver's output by $\text{OUT}_{\mathcal{R}}(\mathcal{C}_{\text{open}}(\mathbf{x}_{\text{com}}), \mathcal{R}_{\text{open}}(\mathbf{y}_{\text{com}}))$.

First, we give the standard notions of hiding and binding.

Definition 3.1 (Hiding Commitment). A bit commitment scheme is computationally (resp. statistically) *hiding* if the following holds. For any polynomial-size (resp. unbounded-size) receiver $\mathcal{R}_{\text{com}}^* = \{\mathcal{R}_{\text{com}, \lambda}^*, \rho_\lambda\}$ interacting in the commit phase of the protocol, let $\text{OUT}_{\mathcal{R}}(\mathcal{C}_{\text{com}}(1^\lambda, b), \mathcal{R}_{\text{com}, \lambda}^*(\rho_\lambda))$ denote a bit output by $\mathcal{R}_{\text{com}, \lambda}^*$ after interaction with an honest \mathcal{C}_{com} committing to message b . Then for every polynomial-size (resp. unbounded-size) receiver $\mathcal{R}_{\text{com}}^* = \{\mathcal{R}_{\text{com}, \lambda}^*, \rho_\lambda\}$, there exists a negligible function $\nu(\cdot)$ such that

$$|\Pr[\text{OUT}_{\mathcal{R}}(\mathcal{C}_{\text{com}}(1^\lambda, 0), \mathcal{R}_{\text{com}, \lambda}^*(\rho_\lambda)) = 1] - \Pr[\text{OUT}_{\mathcal{R}}(\mathcal{C}_{\text{com}}(1^\lambda, 1), \mathcal{R}_{\text{com}, \lambda}^*(\rho_\lambda)) = 1]| = \nu(\lambda).$$

Definition 3.2 (Statistically Binding Commitment). A bit commitment scheme is statistically binding if for every unbounded-size committer $\mathcal{C}^* = (\mathcal{C}_{\text{com}}^*, \mathcal{C}_{\text{open}}^*)$, there exists a negligible function $\nu(\cdot)$ such that with probability at least $1 - \nu(\lambda)$ over $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}}) \leftarrow \langle \mathcal{C}_{\text{com}}^*, \mathcal{R}(1^\lambda) \rangle$, there exists a bit $b \in \{0, 1\}$ such that

$$\Pr[\text{OUT}_{\mathcal{R}}(\mathcal{C}_{\text{open}}^*(1^\lambda, \mathbf{x}_{\text{com}}), \mathcal{R}_{\text{open}}(1^\lambda, \mathbf{y}_{\text{com}})) = b] \leq \nu(\lambda).$$

As mentioned above, $\langle \mathcal{C}(1^\lambda, b), \mathcal{R}(1^\lambda) \rangle$ is the distribution over pairs of states $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}})$ induced by the measurements performed by the parties in the protocol. Thus, when we write “with probability $1 - \nu(\lambda)$ over $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}}) \leftarrow \langle \mathcal{C}_{\text{com}}^*, \mathcal{R}(1^\lambda) \rangle$ ”, this probability is taken over the measurement outcomes of the two parties in the protocol. Our definition then says that with overwhelming probability over such sampling, there exists a bit b such that the malicious committer, given \mathbf{x}_{com} in the open phase, can only open to b with negligible probability.

We will not consider (plain) computationally binding commitments in this work. We remark that there are subtleties to consider when defining computationally binding commitment secure against quantum attackers, and that the “right” definition appears to be Unruh’s notion of collapse-binding [Unr16]. However, for this work, computational notions of binding will only arise in the context of *extractable commitments*, which can be seen as a strong form of computationally binding commitments.

Post-quantum bit commitments satisfying computational hiding (as in Definition 3.1) and statistical binding (as in Definition 3.2) can be constructed from any post-quantum pseudorandom generator (PRG) [Nao91].⁷ Watrous [Wat06] considers PRGs built from (post-quantum) one-way *permutations* via a “quantum Goldreich-Levin Theorem” of [AC02]. However, Aaronson [Aar09] and Zhandry [Zha12] later pointed out that the original [HILL99] construction of PRGs from one-way *functions* extends to non-uniform quantum adversaries with *classical* advice. This can be extended to handle non-uniform quantum advice by giving the

⁷In Naor’s protocol, the receiver first sends a uniformly random $u \leftarrow \{0, 1\}^{\lambda^3}$, and the committer commits to bit b by sending $(b \cdot u) \oplus G(s)$, where $s \leftarrow \{0, 1\}^\lambda$ and G is a length-tripling PRG.

one-way function attacker constructed in the [HILL99] reduction many copies of the PRG attacker’s non-uniform quantum advice (which only requires some polynomial upper bound on the number of times the reduction invokes the PRG attacker). Thus, post-quantum computationally hiding and statistically binding commitments are known from post-quantum one-way functions. (We note that statistically binding commitments can alternatively be implemented non-interactively with quantum communication from one-way functions [YWLQ15].)

We remark that achieving general definitions for quantum binding is a notoriously subtle task that has been studied in many prior works, and that our definition is more restrictive than other definitions in the literature (e.g. [FUZZ20]). However, our goal is to give a definition that (i) we can satisfy and (ii) suffices for our eventual construction of Quantum Oblivious Transfer. Definition 3.2 directly captures what it means for a quantum bit commitment to be statistically binding to a *classical* bit.

Next, we define notions of extractability and equivocation for commitments.

3.1.2 Extractable Commitments

Informally, a commitment is extractable if for every QPT adversarial committer \mathcal{C}^* , there exists a QPT extractor $\mathcal{E}_{\mathcal{C}^*}$ that in an ideal world extracts the message committed by \mathcal{C}^* while also simulating \mathcal{C}^* ’s view. Looking ahead, extractable commitments will enable a challenger that runs an experiment against an adversary to extract messages committed by the adversary without having to execute the opening phase. We formally define this notion below.

Definition 3.3 (Extractable Commitment). Let $\{\rho_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$ denote auxiliary quantum states that may be entangled. A bit commitment scheme is *extractable* if for every polynomial-size quantum adversarial committer $\mathcal{C}^* = \{\mathcal{C}_{\text{com},\lambda}^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a QPT extractor $\mathcal{E}_{\mathcal{C}^*}$ such that for every polynomial-size adversarial opening strategy $\{\mathcal{C}_{\text{open},\lambda}^*\}_{\lambda \in \mathbb{N}}$ and every polynomial-size quantum distinguisher $\mathcal{D}^* = \{\mathcal{D}_\lambda^*, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\nu(\cdot)$ such that:

$$|\Pr[\mathcal{D}_\lambda^*(\sigma_\lambda, \text{Real}) = 1] - \Pr[\mathcal{D}_\lambda^*(\sigma_\lambda, \text{Ideal}) = 1]| = \nu(\lambda)$$

for Real and Ideal distributions defined below (where we omit indexing by λ for the sake of presentation).

- Denote by Real the distribution consisting of $(\mathbf{x}_{\text{final}}, b)$ where $\mathbf{x}_{\text{final}}$ denotes the final state of $\mathcal{C}_{\text{open}}^*$ and $b \in \{0, 1, \perp\}$ denotes the output by $\mathcal{R}_{\text{open}}$ after the open phase.
- To define the distribution Ideal, first run the extractor $\mathcal{E}_{\mathcal{C}^*}$ on input ρ to obtain outputs $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}}, b^*)$, where \mathbf{x}_{com} and \mathbf{y}_{com} are the final states of committer and receiver respectively after the commit phase. Then run $\mathcal{C}_{\text{open}}^*(\mathbf{x}_{\text{com}})$ and $\mathcal{R}(\mathbf{y}_{\text{com}})$ to produce $(\mathbf{x}_{\text{final}}, b)$, where $\mathbf{x}_{\text{final}}$ is the final state of the committer. If $b \notin \{\perp, b^*\}$, output FAIL and otherwise output $(\mathbf{x}_{\text{final}}, b)$.

3.1.3 Equivocal Commitments

Informally, a commitment is equivocal if for every QPT adversarial receiver \mathcal{R}^* , there exists a QPT *equivocator* $\mathcal{Q}_{\mathcal{R}^*}$ that generates a view for \mathcal{R}^* in the commit phase *without any input*, and then receives an input bit b once the commit phase ends. The equivocator is then able to open the commitment generated in the commit phase to bit b such that: \mathcal{R}^* cannot distinguish between its interaction with $\mathcal{Q}_{\mathcal{R}^*}$ and its interaction with an honest committer that committed to b all along, and opened its commitment honestly in the opening phase. Looking ahead, equivocal commitments will enable a challenger that runs an experiment against an adversary to generate commitments that can later be opened arbitrarily. We will now formally describe this primitive.

Definition 3.4 (Equivocal Commitments). Let $\{\rho_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$ denote auxiliary quantum states that may be entangled. A bit commitment scheme is *equivocal* if for every polynomial-size quantum receiver $\mathcal{R}^* =$

$\{\mathcal{R}_{\text{com},\lambda}^*, \mathcal{R}_{\text{open},\lambda}^*, \rho_\lambda\}$ there exists a QPT equivocator $\mathcal{Q}_{\mathcal{R}^*} = (\mathcal{Q}_{\mathcal{R}^*,\text{com}}, \mathcal{Q}_{\mathcal{R}^*,\text{open}})$ such that for every polynomial-size quantum distinguisher $\mathcal{D}^* = \{\mathcal{D}_\lambda^*, \sigma_\lambda\}$ there exists a negligible function $\nu(\cdot)$ such that for every bit $b \in \{0, 1\}$,

$$|\Pr[\mathcal{D}^*(\sigma_\lambda, \text{Real}_b) = 1] - \Pr[\mathcal{D}^*(\sigma_\lambda, \text{Ideal}_b) = 1]| = \nu(\lambda)$$

for Real_b and Ideal_b distributions as defined below (where we omit indexing by λ for the sake of presentation).

- **Real_b:** Let $\mathcal{R}_{\text{com}}^*(\rho)$ interact with $\mathcal{C}_{\text{com}}(1^\lambda, b)$ in the commit phase, after which \mathcal{C}_{com} obtains state \mathbf{x}_{com} and \mathcal{R}^* obtains state \mathbf{y}_{com} . Then, let $\mathcal{R}_{\text{open}}^*(\mathbf{y}_{\text{com}})$ interact with $\mathcal{C}_{\text{open}}(\mathbf{x}_{\text{com}})$ and output the final state $\mathbf{y}_{\text{final}}$ of $\mathcal{R}_{\text{open}}^*$.
- **Ideal_b:** First run the equivocator $\mathcal{Q}_{\mathcal{R}^*,\text{com}}(\rho)$ to obtain state $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}})$. Then, let $\mathcal{Q}_{\mathcal{R}^*,\text{open}}(b, \mathbf{x}_{\text{com}})$ interact with $\mathcal{R}_{\text{open}}^*(\mathbf{y}_{\text{com}})$ and output the final state $\mathbf{y}_{\text{final}}$ of $\mathcal{R}_{\text{open}}^*$.

3.2 Quantum Rewinding Lemma

We will make use of the following lemma from [Wat06].

Lemma 3.5. *Let \mathcal{Q} be a general quantum circuit with n input qubits that outputs a classical bit b and m qubits. For an n -qubit state $|\psi\rangle$, let $p(|\psi\rangle)$ denote the probability that $b = 0$ when executing \mathcal{Q} on input $|\psi\rangle$. Let $p_0, q \in (0, 1)$ and $\epsilon \in (0, 1/2)$ be such that:*

- For every n -qubit state $|\psi\rangle$, $p_0 \leq p(|\psi\rangle)$,
- For every n -qubit state $|\psi\rangle$, $|p(|\psi\rangle) - q| < \epsilon$,
- $p_0(1 - p_0) \leq q(1 - q)$,

Then, there is a general quantum circuit $\widehat{\mathcal{Q}}$ of size $O\left(\frac{\log(1/\epsilon)}{4 \cdot p_0(1-p_0)} |\mathcal{Q}|\right)$, taking as input n qubits, and returning as output m qubits, with the following guarantee. For an n qubit state $|\psi\rangle$, let $\mathcal{Q}_0(|\psi\rangle)$ denote the output of \mathcal{Q} on input $|\psi\rangle$ conditioned on $b = 0$, and let $\widehat{\mathcal{Q}}(|\psi\rangle)$ denote the output of $\widehat{\mathcal{Q}}$ on input $|\psi\rangle$. Then, for any n -qubit state $|\psi\rangle$,

$$\text{TD}\left(\mathcal{Q}_0(|\psi\rangle), \widehat{\mathcal{Q}}(|\psi\rangle)\right) \leq 4\sqrt{\epsilon} \frac{\log(1/\epsilon)}{p_0(1-p_0)}.$$

3.3 Quantum Entropy and Leftover Hashing

Classical Min-Entropy. For a classical random variable X , its min-entropy $\mathbf{H}_\infty(X)$ is defined as

$$\mathbf{H}_\infty(X) := -\log\left(\max_x \Pr[X = x]\right).$$

In cryptographic settings, we are often interested in the min-entropy of a random variable X sampled from a joint distribution (X, Y) , where Y is side information available to an adversary/distinguisher. Following [RW05], we define the conditional min-entropy of X given Y as

$$\mathbf{H}_\infty(X | Y) := -\log\left(\max_{x,y} \Pr[X = x | Y = y]\right).$$

That is, $\mathbf{H}_\infty(X | Y)$ is (the negative log of) the maximum probability of guessing the outcome of X , maximized over the possible outcomes of Y .

Quantum Conditional Min-Entropy. Let ρ_{XY} denote a bipartite quantum state over registers XY . Following [Ren08, KRS09], the conditional min-entropy of ρ_{XY} given Y is then defined to be

$$\mathbf{H}_\infty(\rho_{XY} | Y) := \sup_{\mathbf{y}} \max\{h \in \mathbb{R} : 2^{-h} \cdot I_X \otimes \mathbf{y}_Y - \rho_{XY} \geq 0\}.$$

In this work, we will exclusively consider the case where the ρ_{XY} is a joint distribution of the form (X, \mathbf{y}) where X is a classical random variable. In other words, ρ_{XY} can be written as

$$\sum_x \Pr[X = x] |x\rangle \langle x| \otimes \mathbf{y}_x.$$

In this case, we will write $\mathbf{H}_\infty(\rho_{XY} | Y)$ as $\mathbf{H}_\infty(X | \mathbf{y})$. We remark that in this particular setting, $\mathbf{H}_\infty(X | \mathbf{y})$ can be interpreted as the (negative log of) the maximum probability of guessing X given quantum state \mathbf{y} [KRS09].

Leftover Hash Lemma with Quantum Side Information. We now state a generalization of the leftover hash lemma to the setting of quantum side information.

Lemma 3.6 ([RK05]). *Let \mathcal{H} be a family of universal hash functions from \mathcal{X} to $\{0, 1\}^\ell$, i.e. for any $x \neq x'$, $\Pr_{h \leftarrow \mathcal{H}}[h(x) = h(x')] = 2^{-\ell}$. Then for joint random variables (X, \mathbf{y}) where X is a classical random variable over \mathcal{X} and \mathbf{y} is a quantum random variable,*

$$\|(h, h(X), \mathbf{y}) - (h, u, \mathbf{y})\|_1 \leq \frac{1}{2^{1+\frac{1}{2}(\mathbf{H}_\infty(X|\mathbf{y})-\ell)}},$$

where $h \leftarrow \mathcal{H}$ and $u \leftarrow \{0, 1\}^\ell$.

4 A Quantum Equivocality Compiler

In this section, we show a generic black-box compiler that takes any quantum-secure bit commitment scheme and produces a quantum-secure *equivocal* bit commitment scheme.

The compiler is described in Protocol 1, where (Commit, Decommit) denotes some statistically binding and computationally hiding bit commitment scheme satisfying Definitions 3.1 and 3.2. We describe how to equivocally commit to a single bit, and note that commitment to an arbitrary length string follows by sequential repetition.

Furthermore, we show that if the underlying commitment (Commit, Decommit) is *extractable* according to Definition 3.3, then the resulting commitment is both extractable and equivocal (Definition 3.4).

These results are captured in the following theorems.

Theorem 4.1. *For $\mathcal{X} \in \{\text{quantum extractability, statistical binding}\}$ and $\mathcal{Y} \in \{\text{computationally, statistically}\}$, if Commit is a \mathcal{Y} -hiding quantum bit commitment satisfying \mathcal{X} , then Protocol 1 is a \mathcal{Y} -equivocal bit commitment satisfying \mathcal{X} .*

Definitions of statistical binding and computational hiding for quantum commitments can be found in Definition 3.2 and Definition 3.1. Definitions of quantum extractability and quantum equivocality are given in Definition 3.3 and Definition 3.4.

These theorems follow from establishing statistical binding, equivocality, and extractability of the commitment in Protocol 1, as we do next. First, we note that if Commit is statistically binding, then Protocol 1 is statistically binding. For any adversarial committer strategy, consider the λ unopened pairs of commitments after the commit phase. Since Commit is statistically binding, we can assume that each of the 2λ commitments is binding to a particular bit, except with negligible probability. Now, if any single pair contains binding commitments to the same bit d_i , then the committer will only be able to open its Protocol 1 commitment to the bit $d_i \oplus e_i$. Thus, to violate binding, the adversarial committer will have to have committed to different bits in each of the λ unopened pairs. However, in this case, the committer will be caught and the receiver will abort except with probability $1/2^\lambda$.

Protocol 1

Committer \mathcal{C} Input: Bit $b \in \{0, 1\}$.

The Protocol: Commit Phase

1. \mathcal{C} samples uniformly random bits $d_{i,j}$ for $i \in [\lambda]$ and $j \in \{0, 1\}$.
2. For every $i \in [\lambda]$, \mathcal{C} and \mathcal{R} sequentially perform the following steps.
 - (a) \mathcal{C} and \mathcal{R} execute four sessions sequentially, namely:
 - $\mathbf{x}_{0,0}, \mathbf{y}_{0,0} \leftarrow \text{Commit}(\mathcal{C}(d_{i,0}), \mathcal{R})$,
 - $\mathbf{x}_{0,1}, \mathbf{y}_{0,1} \leftarrow \text{Commit}(\mathcal{C}(d_{i,0}), \mathcal{R})$,
 - $\mathbf{x}_{1,0}, \mathbf{y}_{1,0} \leftarrow \text{Commit}(\mathcal{C}(d_{i,1}), \mathcal{R})$ and
 - $\mathbf{x}_{1,1}, \mathbf{y}_{1,1} \leftarrow \text{Commit}(\mathcal{C}(d_{i,1}), \mathcal{R})$.
 - (b) \mathcal{R} sends a choice bit $c_i \leftarrow \{0, 1\}$.
 - (c) \mathcal{C} and \mathcal{R} execute two decommitments, obtaining the opened bits:
 - $u \leftarrow \text{Decommit}(\mathcal{C}(\mathbf{x}_{c_i,0}), \mathcal{R}(\mathbf{y}_{c_i,0}))$ and
 - $v \leftarrow \text{Decommit}(\mathcal{C}(\mathbf{x}_{c_i,1}), \mathcal{R}(\mathbf{y}_{c_i,1}))$.
 If $u \neq v$, \mathcal{R} aborts. Otherwise, \mathcal{C} and \mathcal{R} continue.
3. For $i \in [\lambda]$, \mathcal{C} sets $e_i = b \oplus d_{i,1-c_i}$ and sends $\{e_i\}_{i \in [\lambda]}$ to \mathcal{R} .

The Protocol: Decommit Phase

1. \mathcal{C} sends b to \mathcal{R} . In addition,
 - For $i \in [\lambda]$, \mathcal{C} picks $\alpha_i \leftarrow \{0, 1\}$ and sends it to \mathcal{R} .
 - \mathcal{C} and \mathcal{R} execute $\widehat{d}_i \leftarrow \text{Decommit}(\mathcal{C}(\mathbf{x}_{1-c_i, \alpha_i}), \mathcal{R}(\mathbf{y}_{1-c_i, \alpha_i}))$.
2. \mathcal{R} accepts the decommitment and outputs b if for every $i \in [\lambda]$, $\widehat{d}_i = b \oplus e_i$.

Figure 1: Equivocal Bit Commitment.

4.1 Equivocality

The equivocal simulator $(\mathcal{Q}_{\mathcal{R}^*, \text{com}}, \mathcal{Q}_{\mathcal{R}^*, \text{open}})$ is obtained via the use of the quantum rewinding lemma (Lemma 3.5) [Wat06]. For the purposes of defining the simulation strategy, it will be sufficient (w.l.o.g.) to consider a restricted receiver \mathcal{R}^* as follows, for the i^{th} sequential step of the protocol. In our simulation, the state of \mathcal{R}^* will be initialized to the final state at the end of simulating the $(i-1)^{\text{th}}$ step.

1. \mathcal{R}^* takes a quantum register W , representing its auxiliary quantum input. \mathcal{R}^* will use two additional quantum registers that function as work space: V , which is an arbitrary (polynomial-size) register, and A , which is a single qubit register. The registers V and A are initialized to their all-zero states before the protocol begins.
2. Let M denote the polynomial-size register used by \mathcal{C} to send messages to \mathcal{R}^* . After carrying out step 2(a) by running on registers (W, V, A, M) , \mathcal{R}^* measures the register A to obtain a bit c_i , for Step 2(b), which it sends back to \mathcal{C} .
3. Next, \mathcal{R}^* computes the decommitment phases (with messages from \mathcal{C} placed in register M) according to Step 2(c). \mathcal{R}^* outputs registers (W, V, A, M) .

Any polynomial-time quantum receiver can be modeled as a receiver of this restricted form followed by some polynomial-time post-processing of the restricted receiver’s output. The same post-processing can be applied to the output of the simulator that will be constructed for the given restricted receiver.

Following [Wat06], we define a simulator that uses two additional registers, C and Z. C is a one qubit register, while Z is an auxiliary register used to implement the computation that will be described next. Consider a quantum procedure Q_{partial} that implements the strategy described in Protocol 2 using these registers.

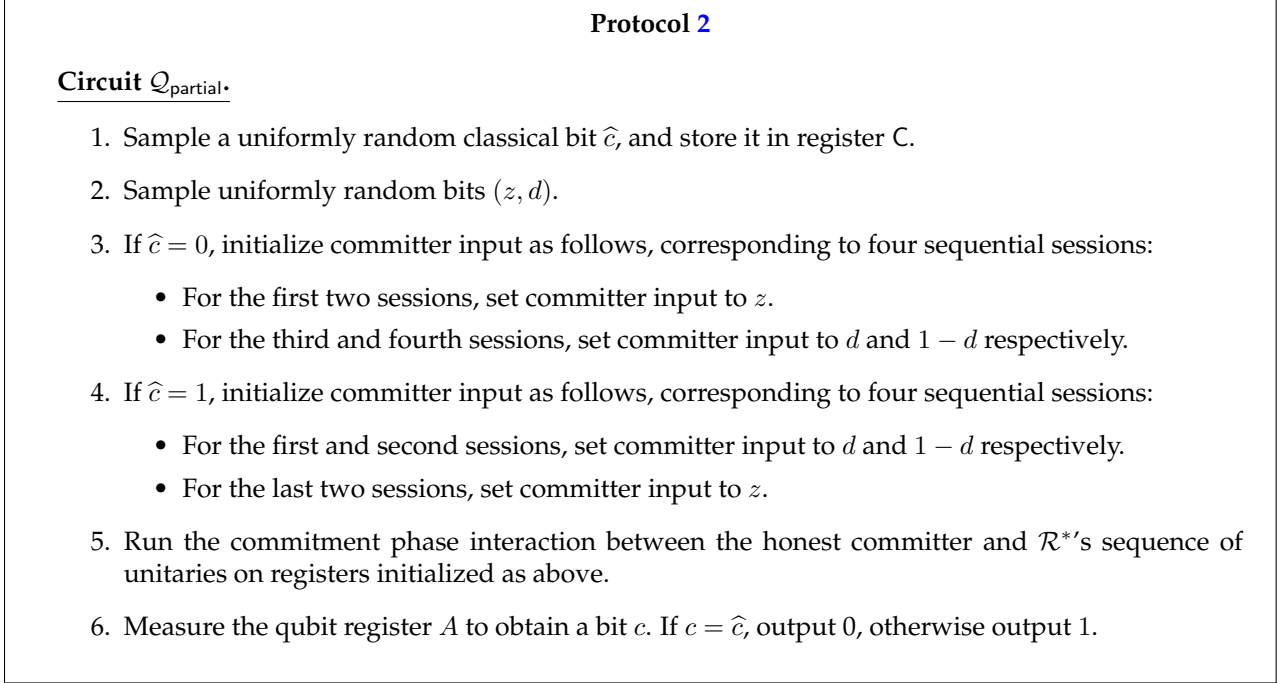


Figure 2: Equivocal Simulator.

Next, we would like to apply Watrous’s quantum rewinding lemma to the Q_{partial} circuit. In order to do this, we will argue that the probability $p(\psi)$ that this circuit outputs 0 is such that $|p(\psi) - \frac{1}{2}| = \text{negl}(\lambda)$, regardless of the auxiliary input $|\psi\rangle$ to \mathcal{R}^* . This follows from the fact that the commitments are (statistically/computationally) hiding. In more detail, by definition, Step 5 produces a distribution on the \mathcal{R}^* ’s side that is identical to the distribution generated by \mathcal{R}^* in its interaction with the committer. If $|p(\psi) - \frac{1}{2}|$ were non-negligible, then the sequence of unitaries applied by \mathcal{R}^* could be used to distinguish commitments generated according to the case $\hat{c} = 0$ from commitments generated according to the case $\hat{c} = 1$, leading to a contradiction.

Now consider the state of the residual qubits of Q_{partial} conditioned on a measurement of its output qubit being 0. The output state of the general quantum circuit \hat{Q} resulting from applying Watrous’s quantum rewinding lemma will have negligible trace distance from this state. This state is over all of the registers discussed above, so the simulator $Q_{\text{com}, \mathcal{R}^*}$ must further process this state as:

- Measure the register C, obtaining challenge c .
- Compute decommitment information corresponding to challenge c , as in Step 2(c) of the protocol (recall that this information is stored in the message register M).
- Output registers (W, V, A, M) . All remaining registers are traced out.

The simulator $\mathcal{Q}_{\mathcal{R}^*, \text{com}}$ executes all λ sequential interactions in this manner, and then samples $e_1, \dots, e_\lambda \leftarrow \{0, 1\}^\lambda$, as the committer messages for Step 3 of Protocol 1. It runs the receiver's unitary on the resulting protocol, and outputs the resulting registers (W, V, A, M) . It additionally outputs private state $\text{st} = (c_1, d_1, \dots, c_\lambda, d_\lambda)$ where c_i, d_i were sampled during the i th execution of Protocol 2.

The simulator $\mathcal{Q}_{\mathcal{R}^*, \text{open}}(b, \text{st}, \mathbf{w}, \mathbf{v}, \mathbf{a}, \mathbf{m})$ parses st as $(c_1, d_1, \dots, c_\lambda, d_\lambda)$. For every $i \in [\lambda]$ it does the following:

- Let $\widehat{d}_i = b \oplus e_i$.
- If $c_i = 0$, it executes the decommitment phase for the $((\widehat{d}_i \oplus d_i) + 2)^{\text{th}}$ session.
- If $c_i = 1$, it executes the decommitment phase for the $(\widehat{d}_i \oplus d_i)^{\text{th}}$ session.

$\mathcal{Q}_{\mathcal{R}^*, \text{open}}$ then executes the receiver's algorithm on these decommitments and outputs the resulting state. Note that each decommitment will be to the bit $\widehat{d}_i = b \oplus e_i$.

To complete the proof of equivocality, we must establish that the distributions Real_b (the view of the receiver interacting with an honest committer) and Ideal_b (the view of the receiver interacting with the equivocator) of Definition 3.4, are indistinguishable. This follows from the (statistical/computational) hiding of the commitment scheme, via an identical argument to the one used above. In particular, if the equivocal simulator produces a distribution that is distinguishable from the real distribution, then there exists a session $i \in [\lambda]$ such that the distribution in the real and ideal experiments upto the $i - 1^{\text{th}}$ session are indistinguishable, but upto the i^{th} session are distinguishable. This contradicts the above guarantee given by the quantum rewinding lemma, since for any i , the post-processed residual qubits of $\mathcal{Q}_{\text{partial}}$ are indistinguishable from the state of \mathcal{R}^* after the i^{th} sequential session in the real protocol (due to the hiding of the commitment scheme).

4.2 Extractability

Next, we prove that Protocol 1 satisfies extractability as long as the underlying commitment (Commit, Decommit) is extractable; in other words, this compiler preserves extractability. Consider the following extractor $\mathcal{E}_{\mathcal{C}^*}$.

- For $i \in [\lambda]$:
 - Execute four sequential commitment sessions with \mathcal{C}^* , where the extractor of Commit is run on all sessions. Obtain outputs $(\rho_{\mathcal{C}^*}, \text{st}_{\mathcal{R}, i, 0}, d'_{i, 0}, \text{st}_{\mathcal{R}, i, 1}, d'_{i, 1})$, where $\rho_{\mathcal{C}^*}$ is the final state of the committer after engaging in all four sequential sessions, and $\text{st}_{\mathcal{R}, i, 0}, \text{st}_{\mathcal{R}, i, 1}$ are receiver states output by the extractor corresponding to the first and third sessions.
 - Corresponding to Step 2(b), compute and send $c_i \leftarrow \{0, 1\}$.
 - Execute Step 2(c) identically to Protocol 1.
- Executes Step 3 of Protocol 1, receiving bits $\{e_i\}_{i \in [\lambda]}$. Fix b^* to be the most frequently occurring bit in $\{e_i \oplus d'_{i, 1 - c_i}\}_{i \in [\lambda]}$, and output the final state of \mathcal{C}^* , the receiver states $\{\text{st}_{\mathcal{R}, i, 0}, \text{st}_{\mathcal{R}, i, 1}\}_{i \in [\lambda]}$, and the extracted bit b^* .

Since the extractor is applied sequentially to parts of the committer, we have that the extractor runs in time polynomial in λ . Next, we will prove indistinguishability between the distributions Real and Ideal defined by the above extractor via a hybrid argument, and based on the definition of extractability of the underlying commitment (Commit, Decommit). In more detail, recall that Real denotes the distribution $(\rho_{\mathcal{C}^*, \text{final}}, b)$ where $\rho_{\mathcal{C}^*, \text{final}}$ denotes the final state of \mathcal{C}^* and b the output of the receiver, and Ideal denotes the final committer state and opened bit after the opening phase of the scheme is run on the output of the extractor.

Note that there are a total of 4λ commitment sessions. Denote the real experiment by $\text{Hybrid}_{0, 1, 1}$. For each $i \in [\lambda], j \in [0, 1], k \in [0, 1]$, define $\text{Hybrid}_{i, j, k}$ to be the distribution obtained as follows.

Commit Phase. Set $a = 1$ and do the following:

1. If $a = \lambda + 1$, obtain $\{e_i\}_{i \in [\lambda]}$ from \mathcal{C}^* and end.
2. If $a < i$,
 - (a) Execute four sessions sequentially, namely:
 - $(\mathbf{x}_{0,0}, \mathbf{y}_{0,0}, d_{0,0}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,0,0}^*}(\boldsymbol{\rho})$
 - $(\mathbf{x}_{0,1}, \mathbf{y}_{0,1}, d_{0,1}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,0,1}^*}(\mathbf{x}_{0,0})$
 - $(\mathbf{x}_{1,0}, \mathbf{y}_{1,0}, d_{1,0}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,1,0}^*}(\mathbf{x}_{0,1})$
 - $(\mathbf{x}_{1,1}, \mathbf{y}_{1,1}, d_{1,1}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,1,1}^*}(\mathbf{x}_{1,0})$
 - (b) Send choice bit $c_a \leftarrow \{0, 1\}$ to \mathcal{C}^* .
 - (c) Execute two decommitments, obtaining the opened bits:
 - $(\mathbf{x}', u_a) \leftarrow \text{Decommit}(\mathcal{C}^*(\mathbf{x}_{1,1}), \mathcal{R}(\mathbf{y}_{c_a,0}))$ and
 - $(\mathbf{x}'', v_a) \leftarrow \text{Decommit}(\mathcal{C}^*(\mathbf{x}'), \mathcal{R}(\mathbf{y}_{c_a,1}))$.
 - Abort if $u_a \neq v_a$.
 - (d) Update state $\boldsymbol{\rho} = \mathbf{x}''$, and set $a = a + 1$.
3. If $a = i$,
 - (a) Execute four sessions sequentially, namely:
 - Set $(\mathbf{x}_{0,0}, \mathbf{y}_{0,0}, d_{0,0}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,0,0}^*}(\boldsymbol{\rho})$.
 - If $(j, k) = (0, 0)$, set $(\mathbf{x}_{0,1}, \mathbf{y}_{0,1}) \leftarrow \text{Commit}(\mathcal{C}^*(\mathbf{x}_{0,0}), \mathcal{R})$ and $d_{0,1}^{(a)} = \tilde{\perp}$.
Else set $(\mathbf{x}_{0,1}, \mathbf{y}_{0,1}, d_{0,1}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,0,1}^*}(\mathbf{x}_{0,0})$.
 - If $j = 0$, set $(\mathbf{x}_{1,0}, \mathbf{y}_{1,0}) \leftarrow \text{Commit}(\mathcal{C}^*(\mathbf{x}_{0,1}), \mathcal{R})$ and $d_{1,0}^{(a)} = \tilde{\perp}$.
Else set $(\mathbf{x}_{1,0}, \mathbf{y}_{1,0}, d_{1,0}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,1,0}^*}(\mathbf{x}_{0,1})$.
 - If $j = 0$ or $k = 0$, set $(\mathbf{x}_{1,1}, \mathbf{y}_{1,1}) \leftarrow \text{Commit}(\mathcal{C}^*(\mathbf{x}_{1,0}), \mathcal{R})$ and $d_{1,1}^{(a)} = \tilde{\perp}$.
Else set $(\mathbf{x}_{1,1}, \mathbf{y}_{1,1}, d_{1,1}^{(a)}) \leftarrow \mathcal{E}_{\mathcal{C}_{i,1,1}^*}(\mathbf{x}_{1,0})$.
 - (b) Send choice bit $c_a \leftarrow \{0, 1\}$ to \mathcal{C}^* .
 - (c) Execute two decommitments, obtaining the opened bits:
 - $(\mathbf{x}', u_a) \leftarrow \text{Decommit}(\mathcal{C}^*(\mathbf{x}_{1,1}), \mathcal{R}(\mathbf{y}_{c_a,0}))$ and
 - $(\mathbf{x}'', v_a) \leftarrow \text{Decommit}(\mathcal{C}^*(\mathbf{x}'), \mathcal{R}(\mathbf{y}_{c_a,1}))$.
 - Abort if $u_a \neq v_a$.
 - (d) Update state $\boldsymbol{\rho} = \mathbf{x}''$, and set $a = a + 1$.
4. If $a > i$,
 - (a) Execute four sessions sequentially, namely:
 - $\mathbf{x}_{0,0}, \mathbf{y}_{0,0} \leftarrow \text{Commit}(\mathcal{C}^*(\boldsymbol{\rho}), \mathcal{R})$,
 - $\mathbf{x}_{0,1}, \mathbf{y}_{0,1} \leftarrow \text{Commit}(\mathcal{C}^*(\mathbf{x}_{0,0}), \mathcal{R})$,
 - $\mathbf{x}_{1,0}, \mathbf{y}_{1,0} \leftarrow \text{Commit}(\mathcal{C}^*(\mathbf{x}_{0,1}), \mathcal{R})$ and
 - $\mathbf{x}_{1,1}, \mathbf{y}_{1,1} \leftarrow \text{Commit}(\mathcal{C}^*(\mathbf{x}_{1,0}), \mathcal{R})$.
 - (b) Send choice bit $c_a \leftarrow \{0, 1\}$ to \mathcal{C}^* .
 - (c) Execute two decommitments, obtaining the opened bits:
 - $(\mathbf{x}', u_a) \leftarrow \text{Decommit}(\mathcal{C}^*(\mathbf{x}_{1,1}), \mathcal{R}(\mathbf{y}_{c_a,0}))$ and
 - $(\mathbf{x}'', v_a) \leftarrow \text{Decommit}(\mathcal{C}^*(\mathbf{x}'), \mathcal{R}(\mathbf{y}_{c_a,1}))$.
 - Abort if $u_a \neq v_a$.
 - (d) Update state $\boldsymbol{\rho} = \mathbf{x}''$, and set $a = a + 1$.

Decommit Phase.

1. Obtain b from \mathcal{C}^* , set $\rho_0 = \rho$.
2. For $a \in [\lambda]$, obtain α_a from \mathcal{C}^* . Then execute $(\rho_a, b_a) = \text{Decommit}(\langle \mathcal{C}_{\text{open}}^*(\rho_{a-1}), \mathcal{R}_{\text{open}}(\mathbf{y}_{1-c_a, \alpha_a}) \rangle)$, aborting if $b_a = \perp$ for any $a \in [\lambda]$.
3. For every $a \leq i$, if $d_{1-c_a, \alpha_a}^{(a)} \neq \tilde{\perp}$ and $b_a \neq d_{1-c_a, \alpha_a}^{(a)}$, output FAIL.
4. Set b^* to be the majority bit in $\{b_a\}_{a \in [\lambda]}$.

The output of $\text{Hybrid}_{i,j,k}$ is (ρ, b^*) .

Claim 4.2. *There exists a negligible function $\mu(\cdot)$ such that for every $i \in [\lambda]$ and every QPT distinguisher \mathcal{D} ,*

$$|\Pr[\mathcal{D}(\text{Hybrid}_{i-1,1,1}) = 1] - \Pr[\mathcal{D}(\text{Hybrid}_{i,0,1}) = 1]| = \mu(\lambda),$$

$$|\Pr[\mathcal{D}(\text{Hybrid}_{i,0,0}) = 1] - \Pr[\mathcal{D}(\text{Hybrid}_{i,0,1}) = 1]| = \mu(\lambda),$$

$$|\Pr[\mathcal{D}(\text{Hybrid}_{i,0,1}) = 1] - \Pr[\mathcal{D}(\text{Hybrid}_{i,1,0}) = 1]| = \mu(\lambda),$$

and

$$|\Pr[\mathcal{D}(\text{Hybrid}_{i,1,0}) = 1] - \Pr[\mathcal{D}(\text{Hybrid}_{i,1,1}) = 1]| = \mu(\lambda).$$

Proof. Suppose this is not the case, then there exists an adversarial committer \mathcal{C}^* , a distinguisher \mathcal{D} , a polynomial $p(\cdot)$, and an initial committer state ψ that corresponds to a state just before the beginning of commitment (ι', j', k') where $(\iota, j, k, \iota', j', k') \in \{(i-1, 1, 1, i, 0, 0), (i, 0, 0, i, 0, 1), (i, 0, 1, i, 1, 0), (i, 1, 0, i, 1, 1)\}$ and where

$$|\Pr[\mathcal{D}(\text{Hybrid}_{\iota,j,k}) = 1] - \Pr[\mathcal{D}(\text{Hybrid}_{\iota',j',k'}) = 1]| \geq \frac{1}{p(\lambda)}.$$

Consider a reduction/adversarial committer $\tilde{\mathcal{C}}$ that obtains initial state ψ , then internally runs \mathcal{C}^* , forwarding all messages between an external receiver and \mathcal{C}^* for the (ι', j', k') th commitment. The commit phase then ends, with committer and receiver states resp. $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}})$. \mathcal{R} then begins executing the opening phase with the external receiver, and on input \mathbf{x}_{com} , continues to run the remaining commit sessions with \mathcal{C}^* internally – generating for it the messages on behalf of \mathcal{R} according to $\text{Hybrid}_{\iota,j,k}$. The only modification is that it forwards \mathcal{C}^* 's opening of the (ι', j', k') th commitment (if and when it is executed) to the external challenger, and records the opened bit b' . After the experiment is complete including the rest of the decommit phase, \mathcal{C}^* outputs the final state $\mathbf{x}_{\text{final}} = \rho$ of the committer.

Then there is a distinguisher \mathcal{D}' that obtains from the challenger $(\mathbf{x}_{\text{final}}, b)$ (or the variable FAIL), runs the distinguisher \mathcal{D} on $(\mathbf{x}_{\text{final}}, b)$, and by construction, obtains advantage $\frac{1}{p(\lambda)}$ in distinguishing the real and ideal experiments for the (ι', j', k') commitment, concluding the proof of the claim. \square

To complete the proof, we note that the only difference between $\text{Hybrid}_{\lambda,1,1}$ and Ideal is the way the bit b^* is computed. In more detail, in $\text{Hybrid}_{\lambda,1,1}$, the bit b^* is computed as the majority of $\{e_{1-c_a, \alpha_a}^{(a)}\}_{a \in [\lambda]}$, where $e_{1-c_a, \alpha_a}^{(a)}$ is the bit extracted from commitment e_{1-c_a, α_a} in the a^{th} iteration. Now for every commitment strategy and every $a \in [\lambda]$, by correctness of extraction (which follows from the indistinguishability between real and ideal distributions for every commitment), the probability that $e_{1-c_a, \alpha_a} \neq e_{c_a, \alpha_a}$ and yet the receiver does not abort in Step 2(c) in the i^{th} sequential repetition, is $\leq \frac{1}{2} + \text{negl}(\lambda)$. Thus, this implies that the probability that $\text{Hybrid}_{i,j,k}$ and Ideal differ is at most $2^{-\lambda/2} + \text{negl}(\lambda) = \text{negl}(\lambda)$. Since Real is identical to $\text{Hybrid}_{0,1,1}$, this completes the proof.

5 Quantum Extractable Commitments

We construct extractable commitments by making use of the following building blocks.

- We let $(\text{EqCommit}, \text{EqDecommit})$ denote any quantum statistically binding and equivocal commitment scheme, satisfying Definition 3.2 and Definition 3.4. Such a commitment can be obtained by applying the compiler from last section to Naor's commitment scheme [Nao91].
- For a suitable polynomial $k(\cdot)$, let $h : \{0, 1\}^{k(\lambda)} \times \{0, 1\}^{\lambda^2} \rightarrow \{0, 1\}$ be a universal hash function that is evaluated on a random seed $s \in \{0, 1\}^{k(\lambda)}$ and input $x \in \{0, 1\}^{\lambda^2}$.

Our extractable commitment scheme is described formally in Figure 3. We show how to commit to a single bit, though commitment to any arbitrary length string follows by sequential repetition. Correctness of the protocol follows by inspection. In the remainder of this section, we prove the following theorem.

Theorem 5.1. *Protocol 3 describes a quantum statistically hiding and extractable bit commitment (satisfying Definition 3.2 and Definition 3.3) whenever $(\text{EqCommit}, \text{EqDecommit})$ is instantiated with any quantum statistically binding and equivocal bit commitment scheme (satisfying Definition 3.2 and Definition 3.4).*

Throughout, we will consider non-uniform adversaries, but for ease of exposition we drop the indexing by λ .

5.1 Extractability

Consider any adversarial committer \mathcal{C}^* with advice ρ . The extractor $\mathcal{E}_{\mathcal{C}^*}(\rho)$ is constructed as follows.

1. Run the first message algorithm of \mathcal{C}^* on input ρ , obtaining message ψ .
2. For $i \in [2\lambda^3]$, sequentially execute equivocal commitment sessions with the equivocal simulator $\mathcal{Q}_{R^*, \text{com}}$, where R^* is the part of \mathcal{C}^* that participates as receiver in the i^{th} session. Session i results in output $(\mathbf{z}_i, \mathbf{y}_{\text{com}, i})$, where \mathbf{z}_i is stored by the extractor, and $\mathbf{y}_{\text{com}, i}$ is the current state of \mathcal{C}^* , which is fed as input into the next session.
3. Obtain T from \mathcal{C}^* , and sample $\hat{\theta} \leftarrow \{+, \times\}^{2\lambda^3}$. Let ψ_i denote the i^{th} qubit of ψ , and measure the qubits ψ_i for $i \in T$, each in basis $\hat{\theta}_i$. Let $\{\hat{x}_i\}_{i \in [T]}$ be the results of the measurements.
4. Let \mathbf{x}_{com} be the current state of \mathcal{C}^* . For each $i \in [T]$, execute $\mathcal{Q}_{R^*, \text{open}}((\hat{\theta}_i, \hat{x}_i), \mathbf{z}_i, \mathbf{x}_{\text{com}})$, where R^* is the part of \mathcal{C}^* that participates in the i^{th} opening, and \mathbf{x}_{com} is updated to be the current state of \mathcal{C}^* after each sequential session.
5. If \mathcal{C}^* aborts at any point, abort and output \perp , otherwise continue.
6. Discard tested positions and restrict $\hat{\theta}$ to the indices in \bar{T} . Obtain $\theta \in \{+, \times\}^{\lambda^3}$ from \mathcal{C}^* . Measure the qubits ψ_i in basis θ_i to obtain \hat{x}_i for $i \in \bar{T}$, and then partition \hat{x} into λ different λ^2 -bit strings $\hat{y}_1, \dots, \hat{y}_\lambda$.
7. Obtain $\{s_\ell, v_\ell\}_{\ell \in [\lambda]}$ from \mathcal{C}^* . Let b^* be the most frequently occurring bit in $\{h(s_\ell, \hat{x}^{(\ell)}) \oplus v_\ell\}_{\ell \in [\lambda]}$. Output $(\mathbf{x}_{\text{com}}, \mathbf{y}_{\text{com}}, b^*)$, where \mathbf{x}_{com} is the resulting state of \mathcal{C}^* and $\mathbf{y}_{\text{com}} = (\theta, \hat{\theta}, \hat{x})$.

We now prove that $\mathcal{E}_{\mathcal{C}^*}$ is a secure extractor in the sense of Definition 3.3. We do so via a sequence of hybrids.

Hyb₁. Define distribution Hyb_1 identically to Real (the honest interaction), except that in Step 2, for $i \in [2\lambda^3]$, sequentially execute equivocal commitment sessions using the equivocal simulator $\mathcal{Q}_{R^*, \text{com}}$, as described in the extractor. In Step 4, for every $i \in T$, open the i^{th} commitment to $(\hat{\theta}_i, \hat{x}_i)$ using $\mathcal{Q}_{R^*, \text{open}}$, as described in the extractor.

Protocol 3

Committer \mathcal{C} Input: Bit $b \in \{0, 1\}$.

The Protocol: Commit Phase.

1. \mathcal{C} chooses $x \leftarrow \{0, 1\}^{2\lambda^3}$, $\theta \leftarrow \{+, \times\}^{2\lambda^3}$ and sends $|x\rangle_\theta$ to \mathcal{R} .
2. \mathcal{R} chooses $\hat{\theta} \leftarrow \{+, \times\}^{2\lambda^3}$ and obtains $\hat{x} \in \{0, 1\}^{2\lambda^3}$ by measuring $|x\rangle_\theta$ in basis $\hat{\theta}$.
 \mathcal{R} commits to $\hat{\theta}$ and \hat{x} position-wise: \mathcal{R} and \mathcal{C} execute sequentially $2\lambda^3$ equivocal commitment sessions with \mathcal{R} as committer and \mathcal{C} as receiver. That is, for each $i \in [2\lambda^3]$, they compute $(\mathbf{x}_{\text{com},i}, \mathbf{y}_{\text{com},i}) \leftarrow \text{EqCommit}(\mathcal{R}(\hat{\theta}_i, \hat{x}_i), \mathcal{C})$.
3. \mathcal{C} sends a random test subset $T \subset [2\lambda^3]$ of size λ^3 to \mathcal{R} .
4. For every $i \in T$, \mathcal{R} and \mathcal{C} engage in $(\hat{\theta}_i, \hat{x}_i) \leftarrow \text{EqDecommit}(\mathcal{R}(\mathbf{x}_{\text{com},i}), \mathcal{C}(\mathbf{y}_{\text{com},i}))$, and \mathcal{C} aborts if any commitment fails to open.
5. \mathcal{C} checks that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$. If all tests pass, \mathcal{C} proceeds with the protocol, otherwise, \mathcal{C} aborts.
6. The tested positions are discarded by both parties: \mathcal{C} and \mathcal{R} restrict x and θ , respectively \hat{x} and $\hat{\theta}$, to the λ^3 indices $i \in \bar{T}$. \mathcal{C} sends θ to \mathcal{R} .
7. \mathcal{C} partitions the remaining λ^3 bits of x into λ different λ^2 -bit strings $x^{(1)}, \dots, x^{(\lambda)}$. For each $\ell \in [\lambda]$, sample a seed $s_\ell \leftarrow \{0, 1\}^{k(\lambda)}$ and compute $d_\ell := h(s_\ell, x^{(\ell)})$. Then output $(s_\ell, b \oplus d_\ell)_{\ell \in [\lambda]}$.

The Protocol: Decommit Phase.

1. \mathcal{C} sends b and $(x^{(1)}, \dots, x^{(\lambda)})$ to \mathcal{R} .
2. If either of the following fails, \mathcal{R} rejects and outputs \perp . Otherwise, \mathcal{R} accepts and outputs b .
 - Let $\{s_\ell, v_\ell\}_{\ell \in [\lambda]}$ be the message received by \mathcal{R} in step 7. Check that for all $\ell \in [\lambda]$, $v_\ell = b \oplus h(s_\ell, x^{(\ell)})$.
 - For each $j \in [\lambda^3]$ such that $\hat{\theta}_j = \theta_j$, check that $\hat{x}_j = x_j$.

Figure 3: Extractable Commitment.

Suppose there exists a polynomial $p(\cdot)$ such that

$$\left| \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_1) = 1] - \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_0) = 1] \right| \geq \frac{1}{p(\lambda)}$$

then consider a sequence of sub-hybrids $\text{Hybrid}_{0,0} = \text{Hybrid}_0$ and $\text{Hybrid}_{0,1}, \text{Hybrid}_{0,2}, \dots, \text{Hybrid}_{0,2\lambda^3}$, where for every $j \in [1, 2\lambda^3]$, $\text{Hybrid}_{0,j}$ is identical to Hybrid_0 except for $i \in [j]$, we sequentially execute equivocal commitment sessions using the equivocal simulator $\mathcal{Q}_{R^*, \text{com}}$, as described in the extractor. For $i \in [j+1, 2\lambda^3]$, the commitments are constructed identically to Hybrid_0 . In Step 4, for every $i \in T$ such that $i \leq j$, open the i^{th} commitment to $(\hat{\theta}_i, \hat{x}_i)$ using $\mathcal{Q}_{R^*, \text{open}}$, as described in the extractor. Open remaining commitments $i \in T$ identically to Hybrid_0 . Then, we have that for large enough λ , there exists an index

$i = i(\lambda) \in [2\lambda^3]$ such that:

$$\left| \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_{0,i-1}) = 1] - \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_{0,i}) = 1] \right| \geq \frac{1}{p(\lambda) \cdot 2\lambda^3}$$

This implies that there is an initial state ψ , which is the state of the receiver just prior to the i^{th} equivocal commitment, and a reduction \mathcal{R} initialized with ψ that internally executes \mathcal{R}^* , externally forwarding only the commit phase of the i^{th} commitment. It then ends the commit phase with the external committer. Next, it begins the open phase, executing all other commitments in Step 2 as part of the open phase internally with \mathcal{R}^* . It also sends a randomly sampled T in Step 3, and executes the decommitments in Step 4 internally with \mathcal{R}^* , except forwarding the opening of the i^{th} commitment externally. It completes the rest of the hybrid according to $\text{Hybrid}_{0,i-1}$ and outputs its final state as $\mathbf{x}_{\text{final}}$. Then the existence of \mathcal{D}^* that distinguishes this state $\mathbf{x}_{\text{final}}$ between the cases when the i^{th} commitment was generated honestly vs. equivocated, as described above, contradicts the equivocality of this commitment. Therefore, we conclude that there is a negligible function $\mu(\cdot)$ such that:

$$\left| \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_1) = 1] - \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_0) = 1] \right| = \mu(\lambda)$$

Hyb_2 . This is identical to Hyb_1 , except that the verifier measures qubits of $|x\rangle_\theta$ *only after* obtaining a description of the set T , and *only measures* the qubits $i \in [T]$. The output of this experiment is identical to Hyb_1 , therefore for any QPT distinguisher (\mathcal{D}^*, σ) ,

$$\Pr[\mathcal{D}^*(\sigma, \text{Hyb}_3) = 1] = \Pr[\mathcal{D}^*(\sigma, \text{Hyb}_2) = 1].$$

Moreover, the only difference between Hyb_2 and Ideal is that Ideal outputs FAIL when the message b opened by \mathcal{C}^* is not \perp and differs from the one extracted by $\mathcal{E}_{\mathcal{C}^*}$. Therefore, to derive a contradiction it will suffice to prove that there exists a negligible function $\nu(\cdot)$ such that

$$\Pr[\text{FAIL}|\text{Ideal}] = \nu(\lambda).$$

Consider any sender \mathcal{C}^* that produces a committer state \mathbf{x}_{com} and then decommits to message b' using strings (y_1, \dots, y_λ) during the decommit phase. Let $T' \subseteq [\lambda]$ denote the set of all indices $\ell \in [\lambda]$ such that the corresponding $x^{(\ell)} \neq v_\ell$, where $\widehat{x}^{(\ell)}$ denotes the values obtained by the extractor in Step 6. Then we have the following claim.

Claim 5.2. *There exists a negligible function $\nu(\cdot)$ such that*

$$\Pr[|T'| > \lambda/2] = \nu(\lambda)$$

where the probability is over the randomness of the extractor.

Proof. For every $\ell \in [\lambda]$, we have that (over the randomness of the extractor):

$$\Pr \left[\mathcal{R}_{\text{open}}(\mathbf{y}_{\text{com}}) \text{ outputs } \perp \text{ in } \langle \mathcal{C}_{\text{open}}^*(\mathbf{x}_{\text{com}}), \mathcal{R}_{\text{open}}(\mathbf{y}_{\text{com}}) \rangle \mid x^{(\ell)} \neq \widehat{x}^{(\ell)} \right] \geq \frac{1}{2}.$$

Indeed, the receiver will reject if for some position i for which $x^{(\ell)} \neq \widehat{x}^{(\ell)}$, it holds that $\theta_i = \widehat{\theta}_i$. Since $\widehat{\theta}$ was sampled uniformly at random, this will occur for a single i with independent probability $1/2$. This implies that $\Pr[|T'| > \lambda/2] \leq \frac{1}{2^{\lambda/2}}$, and the claim follows. \square

By construction of $\mathcal{E}_{\mathcal{C}^*}$, $\Pr[\text{FAIL}|\text{Ideal}] < \Pr[|T'| > \lambda/2]$, and therefore it follows that there exists a negligible function $\nu(\cdot)$ such that

$$\Pr[\text{FAIL}|\text{Ideal}] = \nu(\lambda).$$

5.2 Statistical Hiding

In this subsection, we prove that our extractable commitment scheme described in Fig. 3 is statistically hiding, assuming that EqCommit is a statistically binding (equivocal) commitment scheme.

Our proof closely follows techniques of [BF10], who developed techniques for analyzing [CK88]-style oblivious transfer protocols based on quantum sampling techniques. We import a lemma from [BF10].

[BF10] Quantum Sampling. Consider the following interaction between a player P and a challenger C .

1. P prepares an arbitrary quantum state $|\psi\rangle_{XY}$ where X is a $2n$ -qubit register and Y is an arbitrary-size register of P 's choice.
 P sends the contents of X to C along with two length- $2n$ classical strings $\hat{x} \in \{0,1\}^{2n}$ and $\hat{\theta} \in \{+, \times\}^{2n}$.
2. C samples a uniformly random set of n indices $T \subset [2n]$ as well as a uniformly random $\theta \in \{+, \times\}^{2n}$.
For each $i \in T$ it measures the i th qubit of X in basis $\hat{\theta}_i$. If there exists $i \in T$ where $\theta_i = \hat{\theta}_i$ and the measurement outcome does not match \hat{x}_i , then C aborts the experiment and outputs \perp . Otherwise, C sends $\bar{T} := [2n] \setminus T$ to P .
3. P responds with a subset $S \subseteq \bar{T}$.
4. C measures the indices of X specified by S , obtaining a length- $|S|$ bitstring X_S of measurement outcomes; let s denote the number of indices $i \in S$ where $\theta_i \neq \hat{\theta}_i$.

Lemma 5.3 ([BF10] Quantum Sampling Lemma). *For any computationally unbounded player P and $0 < \varepsilon < 1$, the probability the sampling experiment does not abort and*

$$\mathbf{H}_\infty(X_S \mid X_{\bar{S}}, Y) > s - \varepsilon \cdot |S|.$$

is at most $\text{negl}(|S|^2/n)$.

Proof. (Sketch.) This lemma follows readily from techniques developed in [BF10] to relate quantum sampling to classical sampling. Following [BF10], it suffices to analyze the case where $\hat{x} = 0^{2n}$ and $\hat{\theta} = +^{2n}$, since all other settings of \hat{x} and $\hat{\theta}$ follow from an identical argument under a change of basis.

In Step 2 of the quantum sampling experiment, C verifies that for every index in a random choice of n indices $T \subset [2n]$, measuring the corresponding qubit in the computational basis yields 0. As observed in [BF10], this is a natural quantum analogue of the following classical sampling experiment ([BF10, Example 4]) on a length- $2n$ bitstring X :

1. randomly select a size- n subset $T \subset [2n]$,
2. randomly select a string $\theta \in \{0,1\}^{2n}$,
3. verify that $X_i = 0$ on each index $i \in T$ where $\theta_i = 1$, and abort otherwise.

Let $X_{\bar{T}}$ denote the string X restricted to the n indices $\bar{T} = [2n] \setminus T$. For any initial bitstring X , the ‘‘classical error probability’’ is the probability that the sampling experiment succeeds and $X_{\bar{T}}$ has relative Hamming weight $\geq \delta$ is at most $6e^{-n\delta^2/50}$ (bound derived in [BF10, Appendix B.4]).

[BF10, Theorem 3] demonstrates that the ‘‘quantum error probability’’ is at most the square root of the classical error probability. That is, if the verification in Step 2 passes, then the residual state on X restricted to \bar{T} is within trace distance $\sqrt{6e^{-n\delta^2/50}}$ of a superposition of states with relative Hamming weight at most δ .

Suppose for a moment that the state were *exactly* a superposition of states with relative Hamming weight at most δ . Now consider the restriction of X to a subset of indices $S \subset \bar{T}$ where exactly s positions have

$\hat{\theta}_i \neq \theta_i$. The relative Hamming weight of the state when restricted to S is at most $\delta n/|S|$. As long as $\delta n/|S| < 1/2$, by [BF10, Corollary 1], we have

$$\mathbf{H}_\infty(X_S | X_{\bar{S}}, Y) > s - h(\delta n/|S|) \cdot |S|$$

where $h : [0, 1] \rightarrow [0, 1]$ is the binary entropy function defined as $h(p) = p \log(p) + (1 - p) \log(1 - p)$ for $0 < p < 1$ and as 0 for $p = 0$ or 1. For any desired $\varepsilon \in (0, 1)$, there exists an $\varepsilon' \in (0, 1/2)$ such that $h(\varepsilon') = \varepsilon$, so we can plug in $\delta = \varepsilon'|S|/n$ to achieve the stated min-entropy lower bound of $s - \varepsilon \cdot |S|$. \square

Proof of Statistical Hiding. We now prove statistical hiding of our extractable commitment scheme in Fig. 3. The core proof idea is due to [BF10].

Consider a *purification* of the committer's strategy from Fig. 3, where instead of sampling random $x \leftarrow \{0, 1\}^{2\lambda^3}$, $\theta \leftarrow \{+, \times\}^{2\lambda^3}$ and sending $|x\rangle_\theta$ to \mathcal{R} in Step 1, the committer \mathcal{C} generates $2\lambda^3$ EPR pairs $(|00\rangle + |11\rangle)/\sqrt{2}$, and sends \mathcal{R} one half of each pair. This change has no effect on the receiver's view. In this modification of Step 1, the committer will still sample random $\theta \leftarrow \{+, \times\}^{2\lambda^3}$ but does not use it to perform any measurements (and hence x is not yet determined). We now recast Steps 1 to 5 of Fig. 3, i.e., the "measurement-check subprotocol" to correspond with this undetectable modification of the committer's behavior:

1. \mathcal{C} chooses $2\lambda^3$ EPR pairs $(|00\rangle + |11\rangle)/\sqrt{2}$ and sends \mathcal{R} one half of each pair. The committer samples random $\theta \leftarrow \{+, \times\}^{2\lambda^3}$ but does not perform any measurements (and hence x is not yet determined).
2. \mathcal{R} chooses $\hat{\theta} \leftarrow \{+, \times\}^{2\lambda^3}$ and obtains $\hat{x} \in \{0, 1\}^{2\lambda^3}$ by measuring its $2\lambda^3$ EPR pair halves in bases specified by $\hat{\theta}$.
 \mathcal{R} commits to $\hat{\theta}$ and \hat{x} using *statistically binding commitments*.
3. \mathcal{C} sends a random test subset $T \subset [2\lambda^3]$ of size λ^3 to \mathcal{R} .
4. For every $i \in T$, \mathcal{R} decommits to $(\hat{\theta}_i, \hat{x}_i)$, and \mathcal{C} aborts if any decommitment is invalid.
5. \mathcal{C} checks that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$. If all tests pass, \mathcal{C} proceeds with the protocol, otherwise, \mathcal{C} aborts.

Notice that \mathcal{C} 's messages are completely independent of its bit b up to this point. If \mathcal{R} passes the check above, then \mathcal{C} will obtain a λ^3 -bit string x by measuring its qubits (i.e., EPR pair halves) in positions $\bar{T} := [2\lambda^3] \setminus T$ in the bases specified in θ . \mathcal{C} then partitions x into λ different λ^2 -bit strings $x^{(1)}, \dots, x^{(\lambda)}$ and sends the receiver

$$(s_1, b \oplus h(s_1, x^{(1)})), \dots, (s_\lambda, b \oplus h(s_\lambda, x^{(\lambda)})),$$

where each s_1, \dots, s_λ is a random seed for the universal hash function h . By the leftover hash lemma (Lemma 3.6), b will be statistically hidden if each $x^{(\ell)}$ has sufficient independent min-entropy.

To that end, observe (following [BF10, Section 5]) that the committer decides to accept the receiver's commitment openings $(\hat{\theta}_i, \hat{x}_i)_{i \in [T]}$ *only if* a certain "quantum sampling experiment" succeeds. Since the commitments are statistically binding, the values of $\hat{x}, \hat{\theta}$ that \mathcal{C} would accept are statistically determined after \mathcal{R} commits. Let $|\psi\rangle_{XY}$ be the joint state of the committer and receiver where X corresponds to the committer's $2\lambda^3$ registers, and Y corresponds to the receiver's state. Then the committer's check can only pass if the sampling strategy of Lemma 5.3 succeeds for the committed values of $\hat{x}, \hat{\theta}$. We remark that while no efficient \mathcal{C} can actually *perform* this sampling strategy (as it involves breaking the computational hiding of the underlying equivocal commitments), its behavior is statistically indistinguishable from \mathcal{C} that does perform this sampling.

By applying Lemma 5.3, we conclude that each $x^{(\ell)}$ must therefore have min-entropy $\Omega(\lambda^2)$ conditioned on \mathcal{R} 's view and on all $x^{(j)}$ for $j \neq \ell$, since by Hoeffding's inequality, θ and $\hat{\theta}$ differ in at least $\Omega(\lambda^2)$ places

within the λ^2 indices corresponding to $x^{(\ell)}$ with $1 - \text{negl}(\lambda)$ probability. Applying Lemma 3.6 (leftover hash lemma) one by one to each $x^{(\ell)}$, this implies that each $(s_\ell, h(s_\ell, x^{(\ell)}))$ is statistically indistinguishable from uniform even given the receiver's view, and thus b is statistically hidden.

6 Quantum Oblivious Transfer from Extractable and Equivocal Commitments

6.1 Definitions for Oblivious Transfer with Quantum Communication

An oblivious transfer with quantum communication is a protocol between a quantum interactive sender S and a quantum interactive receiver \mathcal{R} , where the sender S has input $m_0, m_1 \in \{0, 1\}^\lambda$ and the receiver \mathcal{R} has input $b \in \{0, 1\}$. After interaction the sender outputs (m_0, m_1) and the receiver outputs (b, m_b) .

Let $\mathcal{F}(\cdot, \cdot)$ be the following functionality. $\mathcal{F}(b, \cdot)$ takes as input either (m_0, m_1) or abort from the sender, returns end to the sender, and outputs m_b to the receiver in the non-abort case and \perp in the abort case. $\mathcal{F}(\cdot, (m_0, m_1))$ takes as input either b or abort from the receiver, returns m_b to the receiver, and returns end to the sender in the non-abort case, and returns \perp to the sender in the abort case.

Definition 6.1. We let $\langle S(m_0, m_1), R(b) \rangle$ denote an execution of the OT protocol with sender input (m_0, m_1) and receiver input bit b . We denote by $\rho_{\text{out}, S^*} \langle S^*(\rho), R(b) \rangle$ and $\text{OUT}_R \langle S^*(\rho), R(b) \rangle$ the final state of a non-uniform malicious sender $S^*(\rho)$ and the output of the receiver $R(b)$ at the end of an interaction (leaving the indexing by λ implicit). We denote by $\rho_{\text{out}, R^*} \langle S(m_0, m_1), R^*(\rho) \rangle$ and $\text{OUT}_S \langle S(m_0, m_1), R^*(\rho) \rangle$ the final state of a non-uniform malicious receiver $R^*(\rho)$ and the output of the sender $S(m_0, m_1)$ at the end of an interaction. We require OT to satisfy the following security properties:

- **Receiver Security.** For every QPT non-uniform malicious sender S^* , there exists a simulator Sim_{S^*} such that the following holds. For any non-uniform advice ρ, σ where ρ and σ may be entangled, bit $b \in \{0, 1\}$, and QPT non-uniform distinguisher D^* , $\text{Sim}_{S^*}(\rho)$ sends inputs (m_0, m_1) or abort to the ideal functionality $\mathcal{F}_{\text{OT}}(b, \cdot)$, and outputs a final state $\rho_{\text{Sim}, \text{out}, S^*}$. The output of the ideal functionality to the receiver in this experiment is denoted by OUT_R . It must hold that

$$\left| \Pr [D^*(\sigma, (\rho_{\text{Sim}, \text{out}, S^*}, \text{OUT}_R)) = 1] - \Pr [D^*(\sigma, (\rho_{\text{out}, S^*} \langle S^*(\rho), R(b) \rangle, \text{OUT}_R \langle S^*(\rho), R(b) \rangle)) = 1] \right| = \text{negl}(\lambda).$$

- **Sender Security.** For every QPT non-uniform malicious receiver R^* , there exists a simulator Sim_{R^*} such that the following holds. For any non-uniform advice ρ, σ where ρ and σ may be entangled, pair of sender inputs (m_0, m_1) , and QPT non-uniform distinguisher D^* , $\text{Sim}_{R^*}(\rho)$ sends bit b or abort to the ideal functionality $\mathcal{F}_{\text{OT}}(m_0, m_1, \cdot)$, and outputs a final state $\rho_{\text{Sim}, \text{out}, R^*}$. The output of the ideal functionality to the sender in this experiment is denoted by OUT_S . It must hold that

$$\left| \Pr [D^*(\sigma, (\rho_{\text{Sim}, \text{out}, R^*}, \text{OUT}_S)) = 1] - \Pr [D^*(\sigma, (\rho_{\text{out}, R^*} \langle S(m_0, m_1), R^*(\rho) \rangle, \text{OUT}_S \langle S(m_0, m_1), R^*(\rho) \rangle)) = 1] \right| = \text{negl}(\lambda).$$

6.2 Our Construction

We construct simulation-secure quantum oblivious transfer by making use of the following building blocks.

- Let $(\text{EECommit}, \text{EEDecommit})$ denote any quantum bit commitment scheme satisfying extractability (Definition 3.3) and equivocality (Definition 3.4). Such a commitment scheme may be obtained by applying the compiler from Section 4 to the extractable commitment constructed in Section 5.

- Let $h : \{0, 1\}^{k(\lambda)} \times \mathcal{X} \rightarrow \{0, 1\}^\lambda$ be a universal hash with seed length $k(\lambda) = \text{poly}(\lambda)$ and domain \mathcal{X} the set of all binary strings of length at most 8λ .

Our QOT protocol is described in Protocol 4, which is essentially the [CK88] protocol instantiated with our extractable and equivocal commitment scheme.

Protocol 4

Sender S Input: Messages $m_0, m_1 \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$
Receiver R Input: Bit $b \in \{0, 1\}$

The Protocol:

1. S chooses $x \leftarrow \{0, 1\}^{16\lambda}$ and $\theta \leftarrow \{+, \times\}^{16\lambda}$ and sends $|x\rangle_\theta$ to R .
2. R chooses $\hat{\theta} \leftarrow \{+, \times\}^{16\lambda}$ and obtains $\hat{x} \in \{0, 1\}^{16\lambda}$ by measuring $|x\rangle_\theta$ in basis $\hat{\theta}$. Then, S and R execute 16λ sessions of EECCommit sequentially with R acting as committer and S as receiver. In session i , R commits to the bits $\hat{\theta}_i, \hat{x}_i$.
3. S sends a random test subset $T \subset [16\lambda]$ of size 8λ to R .
4. For each $i \in T$, R and S sequentially execute the i 'th EEDecommit, after which S receives the opened bits $\hat{\theta}_i, \hat{x}_i$.
5. S checks that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$. If all tests pass, S accepts, otherwise, S rejects and aborts.
6. The tested positions are discarded by both parties: S and R restrict x and θ , respectively \hat{x} and $\hat{\theta}$, to the 8λ indices $i \in \bar{T}$. S sends θ to R .
7. R partitions the positions of \bar{T} into two parts: the "good" subset $I_b = \{i : \theta_i = \hat{\theta}_i\}$ and the "bad" subset $I_{1-b} = \{i : \theta_i \neq \hat{\theta}_i\}$. R sends (I_0, I_1) to S .
8. S samples seeds $s_0, s_1 \leftarrow \{0, 1\}^{k(\lambda)}$ and sends $(s_0, h(s_0, x_0) \oplus m_0, s_1, h(s_1, x_1) \oplus m_1)$, where x_0 is x restricted to the set of indices I_0 and x_1 is x restricted to the set of indices I_1 .
9. R decrypts s_b using \hat{x}_b , the string \hat{x} restricted to the set of indices I_b .

Figure 4: Quantum Oblivious Transfer.

Theorem 6.2. *The protocol in Figure 4 is a QOT protocol satisfying Definition 6.1 whenever (EECommit, EEDecommit) is instantiated with a quantum bit commitment satisfying extractability (Definition 3.3) and equivocality (Definition 3.4).*

We prove that the resulting QOT protocol satisfies standard simulation-based notions of receiver and sender security.

6.3 Receiver Security

Consider any adversarial sender S^* with advice ρ . The simulator $\text{Sim}_{S^*}(\rho)$ is constructed as follows.

1. Run the first message algorithm of S^* on input ρ to obtain message ψ .

2. Execute 16λ sequential sessions of EECCommit. In each session, run the equivocator $\mathcal{Q}_{\mathcal{R}^*, \text{com}}$, where \mathcal{R}^* denotes the portion of S^* that participates as receiver in the i^{th} sequential EECCommit session.
3. Obtain test subset T of size 8λ from S^* .
4. For each $i \in T$, sample $\hat{\theta}_i \leftarrow \{+, \times\}$. Obtain \hat{x}_i by measuring the i^{th} qubit of ψ in basis $\hat{\theta}_i$. For each $i \in T$, sequentially execute the equivocal simulator $\mathcal{Q}_{\mathcal{R}^*, \text{open}}$ on input $(\hat{\theta}_i, \hat{x}_i)$ and the state obtained from $\mathcal{Q}_{\mathcal{R}^*, \text{com}}$.
5. If S^* continues, discard positions indexed by T . Obtain θ_i for $i \in \bar{T}$ from S^* , and compute x_i for $i \in \bar{T}$ by measuring the i^{th} qubit of ψ in basis θ_i .
6. For every $i \in \bar{T}$, sample bit $d_i \leftarrow \{0, 1\}$. Partition the set \bar{T} into two subsets (I_0, I_1) , where for every $i \in \bar{T}$, place $i \in I_0$ if $d_i = 0$ and otherwise place $i \in I_1$. Send (I_0, I_1) to S .
7. Obtain (y_0, y_1) from S . Set x_0 to be x restricted to the set of indices I_0 and x_1 to be x restricted to the set of indices I_1 . For $b \in \{0, 1\}$, parse $y_b = (s_b, t_b)$ and compute $m_b = t_b \oplus h(s_b, x_b)$.
8. If S^* aborts anywhere in the process, send abort to the ideal functionality. Otherwise, send (m_0, m_1) to the ideal functionality. Output the final state of S^* .

Next, we establish receiver security according to Definition 6.1. Towards a contradiction, suppose there exists a bit $b \in \{0, 1\}$, a non-uniform QPT sender (S^*, ρ) , a non-uniform QPT distinguisher (D^*, σ) , and polynomial $\text{poly}(\cdot)$ s.t.

$$\left| \Pr [D^*(\sigma, (\rho_{\text{Sim, out}, S^*}, \text{OUT}_R)) = 1] - \Pr [D^*(\sigma, (\rho_{\text{out}, S^*}(S^*(\rho), R(b)), \text{OUT}_R(S^*(\rho), R(b)))) = 1] \right| \geq \frac{1}{\text{poly}(\lambda)}.$$

Fix any such b , sender (S^*, ρ) and distinguisher (D^*, σ) . We derive a contradiction via an intermediate hybrid experiment, defined as follows with respect to bit b and sender (S^*, ρ) .

Hyb. In this hybrid, we generate the QOT receiver commitments via the equivocal simulator $\mathcal{Q}_{\mathcal{R}^*}$ (where \mathcal{R}^* is derived from the malicious QOT sender S^*), and otherwise follow the honest QOT receiver's algorithm.

1. Run the first message algorithm of S^* on input ρ to obtain message ψ .
2. Choose $\hat{\theta} \leftarrow \{+, \times\}^{16\lambda}$ and obtain $\hat{x} \in \{0, 1\}^{16\lambda}$ by measuring ψ in basis $\hat{\theta}$. Execute 16λ sequential sessions of EECCommit. In each session, run the equivocator $\mathcal{Q}_{\mathcal{R}^*, \text{com}}$, where \mathcal{R}^* denotes the portion of S^* that participates as receiver in the i^{th} sequential EECCommit session.
3. Obtain test subset T of size 8λ from S^* .
4. For each $i \in T$, sequentially execute the equivocal simulator $\mathcal{Q}_{\mathcal{R}^*, \text{open}}$ on input $(\hat{\theta}_i, \hat{x}_i)$ and the state obtained from $\mathcal{Q}_{\mathcal{R}^*, \text{com}}$.
5. If S^* continues, discard positions indexed by T . Obtain θ_i for $i \in \bar{T}$ from S^* .
6. Partition the set \bar{T} into two subsets: the "good" subset $I_b = \{i : \theta_i = \hat{\theta}_i\}$ and the "bad" subset $I_{1-b} = \{i : \theta_i \neq \hat{\theta}_i\}$. Send (I_0, I_1) to S .
7. Obtain (y_0, y_1) from S . Set x_b to be \hat{x} restricted to the set of indices I_b , and compute and set $m_b = t_b \oplus h(s_b, x_b)$. If S^* aborts anywhere in the process, let \perp be the output of the receiver, otherwise let m_b be the output of the receiver.

The output of Hyb is the joint distribution of the final state of S^* and the output of the receiver. Receiver security then follows from the following two claims.

Claim 6.3. $\Pr [D^*(\sigma, (\rho_{\text{Sim, out}, S^*}, \text{OUT}_R)) = 1] \equiv \Pr [D^*(\sigma, \text{Hyb}) = 1]$.

Proof. The only differences in the simulated distribution are (1) that measurements of S^* 's initial message ψ are delayed (which cannot be noticed by S^*), and (2) a syntactic difference in that the ideal functionality is queried to produce the receiver's output. \square

Claim 6.4. *There exists a negligible function $\nu(\cdot)$ such that*

$$\left| \Pr [D^*(\sigma, \text{Hyb}) = 1] - \Pr [D^*(\sigma, (\rho_{\text{out}, S^*} \langle S^*(\rho), R(b) \rangle, \text{OUT}_R \langle S^*(\rho), R(b) \rangle)) = 1] \right| = \nu(\lambda).$$

Proof. The only difference between the two distributions is that in the first, the receiver generates commitments according to the honest commit algorithms of EECCommit while in the second, commitments in step 2 are generated via the equivocal simulator \mathcal{Q}_{R^*} of EECCommit. Therefore, this claim follows by the equivocality of (EECommit, EEDecommit). \square

6.4 Sender Security

Consider any adversarial receiver R^* with quantum auxiliary state ρ . The simulator $\text{Sim}_{R^*}(\rho)$ is constructed as follows:

1. Choose $x \leftarrow \{0, 1\}^{16\lambda}$ and $\theta \leftarrow \{+, \times\}^{16\lambda}$. Send $|x\rangle_\theta$ to R^* .
2. Execute 16λ sequential sessions of EECCommit. In the i^{th} session for $i \in [16\lambda]$, run the extractor $\mathcal{E}_{C_i^*, \text{com}}$ where C_i^* denotes the portion of R^* that participates as committer in the i^{th} sequential EECCommit session. Obtain from $\mathcal{E}_{C_i^*, \text{com}}$ the extracted values $(\hat{\theta}_i, \hat{x}_i)$.
3. Sample and send a random test subset $T \subset [16\lambda]$ of size 8λ to R^* .
4. For each index in T execute the opening phases EEDecommit with R^* , and abort if any fail. Discard the opened values, and relabel the remaining indices from 1 to 8λ .
5. Check that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$, and if not abort.
6. Send θ to R^* .
7. Obtain R^* 's message, which specifies sets I_0 and $I_1 = [8\lambda] \setminus I_0$. Let S be the set of indices in I_0 such that $\theta_i \neq \hat{\theta}_i$. If $|S| \geq 3\lambda/2$ then set $b = 1$ and otherwise set $b = 0$; intuitively, if $|S| \geq 3\lambda/2$, we treat I_0 as the "bad" subset I_{1-b} where $\theta_i \neq \hat{\theta}_i$ for many indices i .
8. Obtain the output m_b of \mathcal{F}_{OT} on input bit b . Set $m_{1-b} = 0^\lambda$.
9. Let x_0, x_1 denote x restricted to indices I_0, I_1 respectively. Sample seeds $s_0, s_1 \leftarrow \{0, 1\}^{k(\lambda)}$ and send $s_0, h(s_0, x_0) \oplus m_0, s_1, h(s_1, x_1) \oplus m_1$. Output the final state of R^* .

We will now establish sender security according to Definition 6.1. Towards a contradiction, suppose there exists a pair of messages $(m_0, m_1) \in \{0, 1\}^\lambda$, a non-uniform QPT receiver (R^*, ρ) , a non-uniform QPT distinguisher (D^*, σ) , and a polynomial $\text{poly}(\cdot)$ such that

$$\left| \Pr [D^*(\sigma, (\rho_{\text{Sim, out}, R^*}, \text{OUT}_S)) = 1] - \Pr [D^*(\sigma, (\rho_{\text{out}, R^*} \langle S(m_0, m_1), R^*(\rho) \rangle, \text{OUT}_S \langle S(m_0, m_1), R^*(\rho) \rangle)) = 1] \right| \geq \frac{1}{\text{poly}(\lambda)}.$$

Fix any such (m_0, m_1) , receiver (R^*, ρ) and distinguisher (D^*, σ) . We derive a contradiction via an intermediate hybrid experiment, defined as follows. Like the simulator, this hybrid simulates and extracts from the receiver's commitments, but does not yet replace the sender string m_{1-b} with 0^λ .

Hyb. Consider any adversarial receiver R^* with advice ρ . The output of Hyb depends on (m_0, m_1) and is generated as follows.

1. Choose $x \leftarrow \{0, 1\}^{16\lambda}$ and $\theta \leftarrow \{+, \times\}^{16\lambda}$. Send $|x\rangle_\theta$ to R^* .
2. Execute 16λ sequential sessions of EECCommit. In the i^{th} session for $i \in [16\lambda]$, run the extractor $\mathcal{E}_{C_i^*, \text{com}}$ where C_i^* denotes the portion of R^* that participates as committer in the i^{th} sequential EECCommit session. Obtain from $\mathcal{E}_{C_i^*, \text{com}}$ the extracted values $\hat{\theta}_i, \hat{x}_i$.
3. Sample and send a random test subset $T \subset [16\lambda]$ of size 8λ to R^* .
4. For each index in T execute the opening phases EEDecommit with R^* , and abort if any fail. Discard the opened values, and relabel the remaining indices from 1 to 8λ .
5. Check that $x_i = \hat{x}_i$ whenever $\theta_i = \hat{\theta}_i$, and if not abort.
6. Obtain R^* 's message, which specifies sets I_0 and I_1 .
7. Let x_0, x_1 denote x restricted to indices I_0, I_1 respectively. Sample seeds $s_0, s_1 \leftarrow \{0, 1\}^{k(\lambda)}$ and send $s_0, h(s_0, x_0) \oplus m_0, s_1, h(s_1, x_1) \oplus m_1$. Output the final state of R^* .

This hybrid distribution is indistinguishable to the real distribution due to the extractability (Definition 3.3) of the commitment. Indeed, the values extracted from the simulator are used in place of the values opened by R^* . Conditioned on the opening phases succeeding, the joint distribution of the view of R^* and these values will be equal with all but negligible probability. This can be proved via a sequential hybrid argument similar to the one in the proof of Theorem 4.1.

Now, the only difference between this hybrid experiment and the simulation is the value of the string m_{1-b} , where b is the bit extracted in the simulated game. Thus, to complete the proof, it suffices to argue that in the simulation, the string $h(s_{1-b}, x_{1-b})$, which is used to mask m_{1-b} , is statistically close to a uniformly random string.

The remaining argument is essentially identical in structure to the argument of statistical hiding of the extractable commitment in Section 5.2, which in turn is a direct application of quantum sampling techniques developed by [BF10]. We therefore sketch the following statistical steps, highlighting where the argument differs from Section 5.2.

The idea, as in Section 5.2, is to consider a purification of the strategy for the challenger implementing Hyb. In particular, suppose instead of the challenger generating $|x\rangle_\theta$ by sampling x, θ at random, it generates 16λ EPR pairs and sends R^* half of each pair, keeping the other half of each pair for itself. The challenger still samples $\theta \leftarrow \{+, \times\}^{16\lambda}$, but does not yet measure to determine x . When the challenger implementing Hyb needs to check $x_i = \hat{x}_i$, it obtains x_i by measuring its i th EPR pair half in basis $\hat{\theta}_i$. Instead of invoking statistical binding to extract the committed values of $\hat{x}, \hat{\theta}$ as done in Section 5.2, we can simply run the EECCommit extractor to extract $\hat{x}, \hat{\theta}$. Now, if R^* can successfully open its commitments and the challenger implementing Hyb does not abort, this corresponds to a successful execution of the quantum sampling experiment in Lemma 5.3 for $n = 16\lambda$ and the extracted values of $\hat{x}, \hat{\theta}$.

Now, consider the two subsets $I_0, I_1 \subset [8\lambda]$ sent by R^* . For $c \in \{0, 1\}$, let $\text{wt}(I_c)$ be the number of indices in I_c such that $\theta_i \neq \hat{\theta}_i$. If $\text{wt}(I_0) > 3\lambda/2$, then b will be set to 1, and by Lemma 5.3 for $S = I_0$ and $s = \text{wt}(I_0) > 3\lambda/2$ followed by an application of the leftover hash lemma (Lemma 3.6), $(s_0, h(s_0, x_0))$ will be statistically close to a uniformly random string. Otherwise, we have $\text{wt}(I_0) \leq 3\lambda/2$. By Hoeffding's inequality, with probability $1 - \text{negl}(\lambda)$, the strings θ and $\hat{\theta}$ must differ in more than 3λ positions, which means in this case we have $\text{wt}(I_1) > 3\lambda/2$. As a result, we have $b = 0$, and by the same argument as above, $(s_1, h(s_1, x_1))$ is statistically close to a uniformly random string.

Finally, Theorems 4.1, Theorem 5.1, and Theorem 6.2 give the following.

Corollary 6.5. *Quantum oblivious transfer (QOT) satisfying Definition 6.1 can be based on black-box use of statistically binding bit commitments, or on black-box use of quantum-hard one-way functions.*

7 Acknowledgments

The authors are grateful to the Simons Institute programs on *Lattices: Algorithms, Complexity and Cryptography*, and *The Quantum Wave in Computing* for fostering this collaboration. Thanks also to Alex Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan for discussions about similarities and differences with [GLSV20]. A.C. is supported by DoE under grant DOD ONR Federal. This material is based on work supported in part by DARPA under Contract No. HR001120C0024 (for DK). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

References

- [Aar09] Scott Aaronson. Quantum copy-protection and quantum money. In *2009 24th Annual IEEE Conference on Computational Complexity*, pages 229–242. IEEE, 2009.
- [ABG⁺20] Amit Agarwal, James Bartusek, Vipul Goyal, Dakshita Khurana, and Giulio Malavolta. Post-quantum multi-party computation. *Cryptology ePrint Archive*, Report 2020/1395, 2020. <https://eprint.iacr.org/2020/1395>.
- [AC02] Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 323–334. Springer, 2002.
- [BB84] Charles H Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179, 1984.
- [BBCS92] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. Practical quantum oblivious transfer. In Joan Feigenbaum, editor, *CRYPTO’91*, volume 576 of *LNCS*, pages 351–366. Springer, Heidelberg, August 1992.
- [BCG⁺06] Michael Ben-Or, Claude Crépeau, Daniel Gottesman, Avinatan Hassidim, and Adam Smith. Secure multiparty quantum computation with (only) a strict honest majority. In *47th FOCS*, pages 249–260. IEEE Computer Society Press, October 2006.
- [BCJL93] Gilles Brassard, Claude Crépeau, Richard Jozsa, and Denis Langlois. A quantum bit commitment scheme provably unbreakable by both parties. In *34th FOCS*, pages 362–371. IEEE Computer Society Press, November 1993.
- [BF10] Niek J. Bouman and Serge Fehr. Sampling in a quantum population, and applications. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 724–741. Springer, Heidelberg, August 2010.
- [BM09] Boaz Barak and Mohammad Mahmoody-Ghidary. Merkle puzzles are optimal - an $O(n^2)$ -query attack on any key exchange from a random oracle. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 374–390. Springer, Heidelberg, August 2009.

- [BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 269–279. ACM, 2020.
- [CGS02] Claude Crépeau, Daniel Gottesman, and Adam Smith. Secure multi-party quantum computation. In *34th ACM STOC*, pages 643–652. ACM Press, May 2002.
- [CK88] Claude Crépeau and Joe Kilian. Achieving oblivious transfer using weakened security assumptions (extended abstract). In *29th FOCS*, pages 42–52. IEEE Computer Society Press, October 1988.
- [CLS01] Claude Crépeau, Frédéric Légaré, and Louis Salvail. How to convert the flavor of a quantum bit commitment. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 60–77. Springer, Heidelberg, May 2001.
- [CvT95] Claude Crépeau, Jeroen van de Graaf, and Alain Tapp. Committed oblivious transfer and private multi-party computation. In Don Coppersmith, editor, *CRYPTO’95*, volume 963 of *LNCS*, pages 110–123. Springer, Heidelberg, August 1995.
- [Dac16] Dana Dachman-Soled. Towards non-black-box separations of public key encryption and one way function. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 169–191. Springer, Heidelberg, October / November 2016.
- [DFL⁺09] Ivan Damgård, Serge Fehr, Carolin Lunemann, Louis Salvail, and Christian Schaffner. Improving the security of quantum protocols via commit-and-open. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 408–427. Springer, Heidelberg, August 2009.
- [DGJ⁺20] Yfke Dulek, Alex B. Grilo, Stacey Jeffery, Christian Majenz, and Christian Schaffner. Secure multi-party quantum computation with a dishonest majority. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 729–758. Springer, Heidelberg, May 2020.
- [DI05] Ivan Damgård and Yuval Ishai. Constant-round multiparty computation using a black-box pseudorandom generator. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 378–394. Springer, Heidelberg, August 2005.
- [DNS10] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 685–706. Springer, Heidelberg, August 2010.
- [DNS12] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012. Proceedings*, volume 7417 of *LNCS*, pages 794–811. Springer, 2012.
- [FUYZ20] Junbin Fang, Dominique Unruh, Jun Yan, and Dehua Zhou. How to base security on the perfect/statistical binding property of quantum bit commitment? Cryptology ePrint Archive, Report 2020/621, 2020. <https://eprint.iacr.org/2020/621>.
- [GLSV20] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in minicrypt. *CoRR*, abs/2011.14980, 2020.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to prove all NP-statements in zero-knowledge, and a methodology of cryptographic protocol design. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 171–185. Springer, Heidelberg, August 1987.

- [HILL99] 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.
- [HSS11] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 411–428. Springer, Heidelberg, August 2011.
- [IL89] Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *30th FOCS*, pages 230–235. IEEE Computer Society Press, October / November 1989.
- [ILL89] Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24. ACM Press, May 1989.
- [Imp95] Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 134–147. IEEE, 1995.
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer - efficiently. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 572–591. Springer, Heidelberg, August 2008.
- [IR90] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *CRYPTO’88*, volume 403 of *LNCS*, pages 8–26. Springer, Heidelberg, August 1990.
- [Kil88] Joe Kilian. Founding cryptography on oblivious transfer. In *20th ACM STOC*, pages 20–31. ACM Press, May 1988.
- [KRS09] Robert König, Renato Renner, and Christian Schaffner. The operational meaning of min-and max-entropy. *IEEE Transactions on Information theory*, 55(9):4337–4347, 2009.
- [LC97] Hoi-Kwong Lo and Hoi Fung Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78(17):3410, 1997.
- [LR86] Michael Luby and Charles Rackoff. Pseudo-random permutation generators and cryptographic composition. In *18th ACM STOC*, pages 356–363. ACM Press, May 1986.
- [May97] Dominic Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical review letters*, 78(17):3414, 1997.
- [MMP14] Mohammad Mahmoody, Hemanta K. Maji, and Manoj Prabhakaran. On the power of public-key encryption in secure computation. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 240–264. Springer, Heidelberg, February 2014.
- [MS94] Dominic Mayers and Louis Salvail. Quantum oblivious transfer is secure against all individual measurements. In *Proceedings Workshop on Physics and Computation. PhysComp’94*, pages 69–77. IEEE, 1994.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, 4(2):151–158, January 1991.
- [PW09] Rafael Pass and Hoeteck Wee. Black-box constructions of two-party protocols from one-way functions. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 403–418. Springer, Heidelberg, March 2009.
- [Ren08] Renato Renner. Security of quantum key distribution. *International Journal of Quantum Information*, 6(01):1–127, 2008.

- [RK05] Renato Renner and Robert König. Universally composable privacy amplification against quantum adversaries. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 407–425. Springer, Heidelberg, February 2005.
- [RW05] Renato Renner and Stefan Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In Bimal K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 199–216. Springer, Heidelberg, December 2005.
- [Unr07] Dominique Unruh. Random oracles and auxiliary input. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 205–223. Springer, Heidelberg, August 2007.
- [Unr14] Dominique Unruh. Quantum position verification in the random oracle model. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2014.
- [Unr16] Dominique Unruh. Computationally binding quantum commitments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 497–527. Springer, Heidelberg, May 2016.
- [Wat06] John Watrous. Zero-knowledge against quantum attacks. In Jon M. Kleinberg, editor, *38th ACM STOC*, pages 296–305. ACM Press, May 2006.
- [WW06] Stefan Wolf and Jürg Wullschleger. Oblivious transfer is symmetric. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 222–232. Springer, Heidelberg, May / June 2006.
- [Yao95] Andrew Chi-Chih Yao. Security of quantum protocols against coherent measurements. In *27th ACM STOC*, pages 67–75. ACM Press, May / June 1995.
- [YWLQ15] Jun Yan, Jian Weng, Dongdai Lin, and Yujuan Quan. Quantum bit commitment with application in quantum zero-knowledge proof (extended abstract). In Khaled M. Elbassioni and Kazuhisa Makino, editors, *Algorithms and Computation - 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9-11, 2015, Proceedings*, volume 9472 of *Lecture Notes in Computer Science*, pages 555–565. Springer, 2015.
- [Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012.