

Post-Quantum Multi-Party Computation

Amit Agarwal*

James Bartusek†

Vipul Goyal‡

Dakshita Khurana§

Giulio Malavolta¶

Abstract

We initiate the study of multi-party computation for classical functionalities (in the plain model) with security against malicious polynomial-time quantum adversaries. We observe that existing techniques readily give a polynomial-round protocol, but our main result is a construction of *constant-round* post-quantum multi-party computation. We assume mildly super-polynomial quantum hardness of learning with errors (LWE), and polynomial quantum hardness of an LWE-based circular security assumption. Along the way, we develop the following cryptographic primitives that may be of independent interest:

- A spooky encryption scheme for relations computable by quantum circuits, from the quantum hardness of an LWE-based circular security assumption. This yields the first quantum multi-key fully-homomorphic encryption scheme with classical keys.
- Constant-round zero-knowledge secure against multiple parallel quantum verifiers from spooky encryption for relations computable by quantum circuits.

To enable this, we develop a new straight-line non-black-box simulation technique against *parallel* verifiers that does not clone the adversary's state. This forms the heart of our technical contribution and may also be relevant to the classical setting.

- A constant-round post-quantum non-malleable commitment scheme, from the mildly super-polynomial quantum hardness of LWE.

*UIUC. amita2@illinois.edu

†UC Berkeley. bartusek.james@gmail.com

‡CMU. vipul@cmu.edu

§UIUC. dakshita@illinois.com

¶Max Planck Institute for Security and Privacy. giulio.malavolta@hotmail.it

Contents

1	Introduction	1
1.1	Our Results	1
2	Technical Overview	3
2.1	Background	3
2.2	A New Parallel No-Cloning Non-Black-Box Simulation Technique	5
2.3	Quantum AFS-Spooky Encryption	9
2.4	Post-Quantum Non-malleable Commitments	12
2.5	Putting Things Together	14
2.6	Related Work	14
3	Preliminaries	15
3.1	Quantum Computation	16
3.2	Notation for Interactive Protocols	17
3.3	Witness Indistinguishability	18
3.4	Sigma Protocol for NP	18
3.5	Non-Interactive Commitment	19
3.6	Compute and Compare Obfuscation	19
3.7	Function-Hiding Secure Function Evaluation	19
3.8	Quantum Rewinding Lemma	21
4	Quantum Multi-Key Fully-Homomorphic Encryption	21
4.1	Learning with Errors and Lattice Trapdoors	21
4.2	Definition	22
4.3	Background	24
4.4	Construction	25
4.5	Quantum Spooky Encryption	27
5	Quantum-Secure Multi-Committer Extractable Commitment	30
5.1	Definition	31
5.2	Construction	32
5.3	Hiding	32
5.4	Extractability	34
6	Quantum-Secure Multi-Verifier Zero-Knowledge	38
6.1	Definition	39
6.2	Construction	40
6.3	Soundness	41
6.4	Zero-Knowledge	41
7	Quantum-Secure Non-Malleable Commitments	45
7.1	Definition	45
7.2	Non-Malleable Commitments for Small Tags	46
7.2.1	Construction	47
7.2.2	Analysis	48
7.3	Tag Amplification	56
7.3.1	Construction	56
7.3.2	Analysis.	56

8	Quantum-Secure Multi-Party Coin-Flipping	57
8.1	Definition	58
8.2	Construction	58
8.3	Security	59
9	Quantum-Secure Multi-Party Computation	65
9.1	Definition	65
9.2	Construction	66
10	Acknowledgments	66
A	Simple Polynomial-Round Extractable Commitments	73
B	An explicit quantum attack against a classically-secure ZK protocol	74
C	Tag Amplification: Remaining Analysis	74
D	Multi-Committer Extractable Commitments against Arbitrary Distinguishers	80

1 Introduction

Secure multi-party computation (MPC) allows a set of parties to compute a joint function of their inputs, revealing only the output of the function while keeping their inputs private. General secure MPC, initiated in works such as [Yao86, GMW87, BGW88, CCD88], has played a central role in modern theoretical cryptography. The last few years have seen tremendous research optimizing MPC in various ways, enabling a plethora of practical applications that include joint computations on distributed medical data, privacy-preserving machine learning, e-voting, distributed key management, among others. The looming threat of quantum computers naturally motivates the problem of constructing protocols with *provable security against quantum adversaries*.

After Watrous’ breakthrough work on zero-knowledge against quantum adversaries [Wat09], the works of [DL09, LN11, HSS11] considered variants of quantum-secure computation protocols, in the *two-party* setting. Very recently, Bitansky and Shmueli [BS20] obtained the first *constant-round* classical zero-knowledge arguments with security against quantum adversaries. Their techniques (and those of [AP20] in a concurrent work) are based on the recent non-black-box simulation technique of [BKP19], who constructed two-message *classically-secure* weak zero-knowledge in the plain model. Unfortunately, it is unclear whether these protocols compose under parallel repetition. As a result, they become largely inapplicable to the constant-round multi-party setting.

There has also been substantial effort in constructing protocols for securely computing quantum circuits [DNS10, DNS12, DGJ⁺20] (see Section 2.6 for further discussion). But to the best of our knowledge, generic multi-party computation protocols with classical communication and security against quantum adversaries have only been studied in models with *trusted pre-processing or setup*. To make things even worse, [DGJ⁺20] construct a maliciously-secure multi-party protocol for computing quantum circuits, assuming the existence of maliciously-secure post-quantum classical MPC. This means that the only available implementations of such a building block require trusted pre-processing or a common reference string.

Post-Quantum MPC. In this work we initiate the study of MPC protocols that allow classical parties to securely compute general classical functionalities, and where security is guaranteed against *malicious quantum adversaries*. Our focus is on MPC in the *plain model, with a dishonest majority*: Fully classical participants interact with each other with no access to trusted/pre-processed parameters or a common reference string. Multi-party protocols achieving security in this natural setting do not seem to have been previously analyzed in *any* number of rounds. We stress that the challenges of proving post-quantum security of MPC protocols stretch far beyond the appropriate instantiations of the cryptographic building blocks (e.g. avoiding factoring or discrete logarithm-based cryptosystems): in fact, it is possible to devise protocols [AP20, ARU14] based on entirely post-quantum assumptions that are secure against classical adversaries but completely insecure against quantum adversaries.

Because quantum information behaves very differently from classical information, designing post-quantum protocols often requires new techniques to achieve provable security. As an example, a common strategy to prove classical security of MPC protocols is to define a simulator that can extract the inputs of the corrupted parties by “rewinding” them, i.e. taking a snapshot of the state of the adversary and splitting the protocol execution into multiple branches. However, when the adversary is a quantum machine, this technique becomes largely inapplicable since the no-cloning theorem (one of the fundamental principles of quantum mechanics) prevents us from creating two copies of an arbitrary quantum state. One of our key contributions is a new *parallel no-cloning non-black-box simulation technique* that extends the work of [BS20], to achieve security against multiple parallel quantum verifiers.

1.1 Our Results

We begin by summarizing our main result: Classical multi-party computation with security against quantum circuits in the plain model. Here, parties communicate classically via authenticated point-to-point channels as well as broadcast channels, where everyone can send messages in the same round. In each round, all parties simultaneously exchange messages. The network is assumed to be synchronous with

rushing adversaries, i.e. adversaries may generate their messages for any round after observing the messages of all honest parties in that round, but before observing the messages of honest parties in the next round. The (quantum) adversary may corrupt upto all but one of the participants. In this model, we obtain the following main result.

Theorem 1.1 (Informal). *Assuming mildly super-polynomial quantum hardness of LWE and AFS-spooky encryption for relations computable by polynomial-size quantum circuits, there exists a constant-round classical MPC protocol (in the plain model) maliciously secure against quantum polynomial-time adversaries.*

In more detail, our protocol is secure against any adversary $A = \{A_\lambda, \rho_\lambda\}_\lambda$, where each A_λ is the (classical) description of a polynomial-size quantum circuit and ρ_λ is some (possibly inefficiently computable) non-uniform quantum advice. Beyond being interesting in its own right, our plain-model protocol may serve as a useful stepping stone to obtaining interesting protocols for securely computing quantum circuits in the plain model, as evidenced by the work of [DG]⁺20]. This protocol is constructed in Sections 8 and 9.

By “mildly” super-polynomial quantum hardness of LWE, we mean to assume that there exists a constant $c \in \mathbb{N}$, such that for large enough security parameter $\lambda \in \mathbb{N}$, no quantum polynomial time algorithm can distinguish LWE samples from uniform with advantage better than $\text{negl}(\lambda^{\text{ilog}(c,\lambda)})$, where $\text{ilog}(c, \lambda)$ denotes the c -times iterated logarithm $\log \log \dots_c \text{times}(\lambda)$. We note that this is weaker than assuming the quasi-polynomial quantum hardness of LWE, i.e. the assumption that quantum polynomial-time adversaries cannot distinguish LWE samples from uniform with advantage better than $2^{-(\log \lambda)^c}$ for some constant $c > 1$.

A key technical ingredient in our work is an additive function sharing (AFS) spooky encryption scheme [DHRW16] for relations computable by quantum circuits. An AFS-spooky encryption scheme has a publicly-computable algorithm that, on input a set of ciphertexts $\text{Enc}(\text{pk}_1, m_1), \dots, \text{Enc}(\text{pk}_n, m_n)$ encrypted under *independently sampled* public keys and a (possibly quantum) circuit C , computes a new set of ciphertexts

$$\text{Enc}(\text{pk}_1, y_1), \dots, \text{Enc}(\text{pk}_n, y_n) \text{ s.t. } \bigoplus_{i=1}^n y_i = C(m_1, \dots, m_n).$$

In Section 4 we show how to construct AFS-spooky encryption for relations computable by quantum circuits, under an LWE-based circular security assumption. We refer the reader to Section 4.4 for the exact circular security assumption we need, which is similar to the one used in [Mah18]. As a corollary, this immediately yields the first multi-key fully-homomorphic encryption [LTV12] for quantum circuits with classical key generation and classical encryption of classical messages.

Theorem 1.2 (Informal). *Under an appropriate LWE-based circular security assumption, there exists an AFS-spooky encryption scheme for relations computable by polynomial-size quantum circuits with classical key generation and classical encryption of classical messages.*

Our most important technical contribution is a construction of constant-round zero-knowledge arguments against parallel quantum verifiers, and constant-round extractable commitments against parallel quantum committers. Here, we develop a novel *parallel no-cloning non-black-box simulation* technique. This uses as a starting point the recently introduced no-cloning technique of [BS20, AP20], which in turns builds on the classical non-black-box technique of Bitansky, Khurana and Paneth [BKP19].

We point out that we do not obtain protocols that compose under *unbounded* parallel repetition. Instead we build a bounded variant in the multi-party setting (that we also refer to as multi-verifier zero-knowledge and multi-committer extractable commitments) that suffices for our application to constant round MPC. Our technique makes crucial use of AFS-spooky encryption for relations computable by classical circuits. Parallel extractable commitments and zero-knowledge are formally constructed and analyzed in Sections 5 and 6, respectively.

Theorem 1.3 (Informal). *Assuming the quantum polynomial hardness of LWE and the existence of AFS-spooky encryption for relations computable by polynomial-size quantum circuits, there exists:*

- A constant-round classical argument for NP that is computational-zero-knowledge against parallel quantum polynomial-size verifiers.
- A constant-round classical commitment that is extractable against parallel quantum polynomial-size committers.

In addition, we initiate the study of post-quantum non-malleable commitments. Specifically, we construct and rely on constant-round post-quantum non-malleable commitments based on the super-polynomial hardness assumption described above. The formal construction and analysis can be found in Section 7.

Theorem 1.4 (Informal). *Assuming the mildly super-polynomial quantum hardness of LWE and the existence of fully-homomorphic encryption for quantum circuits, there exists a constant-round non-malleable commitment scheme secure against quantum polynomial-size adversaries.*

We also obtain quantum-secure non-malleable commitments in $O(\text{ilog}(c, \lambda))$ rounds for any constant $c \in \mathbb{N}$ based on any (polynomially) quantum-secure extractable commitment. In particular, plugging in these commitments instead of our constant round non-malleable commitments gives an $O(\text{ilog}(c, \lambda))$ round quantum-secure MPC from any quantum AFS-spooky encryption scheme.

2 Technical Overview

2.1 Background

Our starting point is any constant-round post-quantum maliciously-secure MPC protocol in the (programmable) common random string (CRS) model. A long line of work has studied constant-round MPC in the CRS model [CLOS02, IPS08, AJL⁺12, MW16, BL18a, GS18], and these protocols can all be instantiated with primitives that are plausibly quantum-secure. One method of arguing that the resulting protocol is post-quantum secure is to demonstrate that (i) the simulator does not rewind or clone the adversary’s state, and (ii) the reductions to the underlying quantum-secure primitives used to establish indistinguishability of the real and simulated world do not rewind or clone the adversary’s state. As a concrete example, since the simulators and reductions in [GS18] are non-rewinding and non-cloning, implicit in their work is the theorem that any two-message post-quantum maliciously-secure oblivious transfer (OT) in the CRS model with a straight-line simulator implies a two-round post-quantum maliciously-secure MPC protocol in the CRS model. Such an OT is known for example from the quantum hardness of learning with errors [PVW08].

Thus, a natural approach to achieving post-quantum MPC in the plain model is to then securely implement a multi-party functionality that generates the aforementioned CRS. Specifically, we would like a set of n parties to jointly execute a *coin-flipping protocol*. Such a protocol outputs a uniformly random string that may then be used to implement a post-quantum MPC protocol in the CRS model. The programmability requirement on the CRS roughly translates to ensuring that for any quantum adversary, there exists a simulator that on input a random string s , can force the output of the coin-flipping protocol to be equal to s . A protocol satisfying this property is often referred to as a *fully-simulatable* multi-party coin-flipping protocol.

Post-Quantum Multi-Party Coin-Flipping. Existing constant-round protocols [Wee10, Goy11] for multi-party coin-flipping against classical adversaries make use of the following template. Each participant first commits to a uniformly random string using an appropriate perfectly binding commitment.¹ In a later phase, all participants reveal the values they committed to, without actually revealing the randomness used for commitment. Additionally, each participant proves (in zero-knowledge) to every other participant that they opened to the same value that they originally committed to. If all zero-knowledge arguments verify, the protocol output is computed as the sum of the openings of all participants.

But not every classically secure zero-knowledge argument based on post-quantum assumptions is post-quantum secure. Building on prior work [AP20], in Appendix B, we outline a ZK argument that is classically secure, and is based entirely on post-quantum assumptions (LWE), but is not post-quantum secure.

¹We actually require this commitment to also satisfy a property called *non-malleability*, which we discuss later in this section.

To highlight challenges in constructing constant-round protocols, we elaborate on the template discussed above and outline a simple polynomial-round coin tossing protocol. Readers familiar with this template for multi-party coin-tossing may skip a page.

A Simple Protocol in Polynomially Many Rounds. In order to motivate the challenges involved in constructing a post-quantum *constant-round* multiparty coin tossing protocol, we first outline a simple protocol that requires *polynomially many* rounds, and follows from ideas in existing work. Our starting point is the polynomial-round post-quantum zero-knowledge protocol due to Watrous [Wat09]. Ideas developed in [BS20] almost immediately convert this to a post-quantum extractable commitment, assuming polynomial hardness of LWE (or, more generally, any post-quantum oblivious transfer). For completeness, we outline how this is done in Appendix A.

Next, it is possible to use the resulting post-quantum secure extractable commitment to obtain post-quantum multi-party fully-simulatable coin flipping, that admits a straight-line simulator in the dishonest majority setting. The protocol requires rounds that grow linearly with the number of parties and polynomially with the security parameter, and is described in Figure 1. At a very high level, the protocol requires each party to sample uniform randomness. Then each party sequentially commits (via an extractable commitment) to the randomness it sampled. In the next step, all parties broadcast their randomness in the clear, together with (sequential) zero-knowledge proofs by each party that the broadcasted randomness is consistent with the randomness that was previously committed.

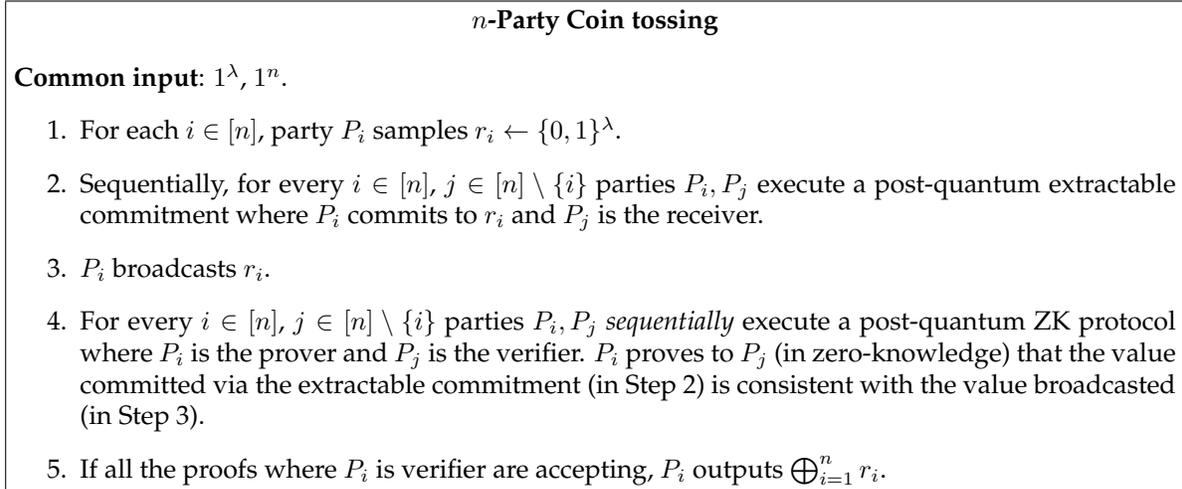


Figure 1: Multiparty Coin Tossing

Recall that the simulator Sim of any coin-flipping protocol obtains a uniformly random string r^* from the ideal functionality, and must force this value as the output. We will briefly describe the construction of Sim for the case where \mathcal{A} controls $n - 1$ parties and Sim plays the role of the only honest party P_1 (the same technique can be easily extended to the case where \mathcal{A} controls any arbitrary subset of parties). To do so, The Sim for the protocol in Figure 1 samples r_i uniformly at random on behalf of each honest party P_i , and commits to r_i in Step 2 following honest sender strategy. At the same time, Sim runs Ext to (sequentially) extract the value committed by every corrupted party in Step 2. This allows the simulator to compute $\bigoplus_{i \in \mathbb{M}} r_i$, where \mathbb{M} denotes the set of corrupted parties. In Step 3, the simulator broadcasts values r'_i on behalf of honest parties such that $\bigoplus_{i \in [n] \setminus \mathbb{M}} r'_i = \bigoplus_{i \in \mathbb{M}} r_i \oplus r^*$. Finally, it invokes the simulator of the ZK protocol to produce proofs on behalf of honest parties. It is easy to see that the output would indeed end up being the intended output r^* .

Notice that replacing Watrous' polynomial-round ZK protocol with the constant-round ZK of [BS20, AP20] only decreases the rounds to linear in the number of parties. To decrease the number of rounds to

constant, it is clear that one would need to find a way to execute the commitment sessions (Step 2) and ZK sessions (Step 4) in parallel. While the recent work of Bitansky and Shmueli [BS20] builds constant-round post-quantum zero-knowledge, their protocol and its guarantees turn out to be insufficient for the parallel setting. In this setting, a single prover would typically need to interact in parallel with $(n - 1)$ different verifiers, a subset or all of which may be adversarial. It should be possible for a simulator to *simultaneously* simulate the view of multiple parallel verifiers. In addition, the argument should continue to satisfy soundness, even if a subset of verifiers colludes with a (cheating) prover.

Post-Quantum Parallel Zero-Knowledge. We overcome this barrier by building the first constant-round zero-knowledge argument secure against *parallel quantum verifiers* from quantum polynomial hardness of an LWE-based circular security assumption. This improves upon the work of [BS20, AP20] who provided arguments with provable security only against a single quantum verifier. Very roughly, the approach in [BS20, AP20] relies on a modification of the [BKP19] homomorphic trapdoors paradigm. We do not assume familiarity with the details of this protocol or paradigm, and will in fact discuss a (variant of) this in the next subsection. For now, we simply point out that in this paradigm, the verifier generates an initial FHE ciphertext and public key, as well as some additional information to enable simulation. The simulator *homomorphically evaluates* the verifier’s (quantum) circuit over the initial FHE ciphertext and then uses the result of this evaluation to recover secrets that will enable simulation.

However, when a prover interacts with several verifiers at once, each verifier will generate its own FHE ciphertexts. In a nutshell, in the parallel setting the simulator can no longer perform individual homomorphic evaluations corresponding to each verifier, due to no-cloning. To address this issue, we develop a novel **parallel no-cloning** simulation strategy. This is our key technical contribution: we develop a novel technique that enables the simulator to *peel away* secret keys of this FHE scheme layer-by-layer. An overview of this technique can be found in Section 2.2.

Our technique also relies on a strong variant of quantum fully-homomorphic encryption that allows for homomorphic operations under multiple keys at once. The encryption scheme that we use is a quantum generalization of the notion of additive function sharing (AFS) *spooky encryption* [DHRW16]. As a contribution of independent interest, we build the first AFS-spooky encryption (that also implies multi-key FHE) for quantum circuits from a circular variant of the LWE assumption. We give an overview of our construction in Section 2.3.

Post-Quantum Non-malleable Commitments. Our construction of zero-knowledge against parallel quantum verifiers gives rise to a coin-flipping protocol that is secure as long as at least one participant is honest, and all committed strings are independent of each other. However, ensuring such independence is not straightforward, even in the classical setting. In fact, upon seeing an honest party’s commitment string c , a malicious, rushing adversary may be able to produce a string c' that commits to a related message. This is known as a malleability attack, and can be prevented by relying on *non-malleable commitments*. In this work, we devise the first post-quantum non-malleable commitments based on slightly superpolynomial hardness of LWE. An overview of our construction can be found in Section 2.4.

Finally, we discuss how to combine all these primitives to build our desired coin-tossing protocol, and a few additional subtleties that come up in the process, in Section 2.5.

2.2 A New Parallel No-Cloning Non-Black-Box Simulation Technique

In the following we give a high-level overview of our constant-round zero-knowledge protocol secure against parallel quantum verifiers. In favor of a simpler exposition, we first describe a *parallel extractable commitment* protocol. A parallel extractable commitment is a commitment where a single receiver interacts in parallel with multiple committers, each committing to its own independent message. The main challenge in this setting is to simulate the view of an adversary corrupting several of these committers, while *simultaneously* recovering all committed messages. Once we build a parallel extractable commitment, obtaining a parallel zero-knowledge protocol becomes a simple exercise (that we discuss towards the end of this overview).

Throughout the following overview we only consider adversaries that are (i) *non-aborting*, i.e. they never interrupt the execution of the protocol, and (ii) *explainable*, i.e. their messages always lie in the support of honestly generated messages, though they can select their random coins and inputs arbitrarily. We further simplify our overview by only considering (iii) *classical* adversaries, while being mindful to avoid any kind of state cloning during extraction. In the end of this overview we discuss how to remove these simplifications.

Cryptographic Building Blocks. Before delving into the description of our protocol, we introduce the technical tools needed for our construction. A fully-homomorphic encryption (FHE) scheme [Gen09] allows one to compute any function (in its circuit representation) over some encrypted message $\text{Enc}(\text{pk}, m)$, without the need to decrypt it first. We say that an FHE is multi-key [LTV12] if it supports the homomorphic evaluation of circuits even over messages encrypted under *independently sampled* public keys:

$$\{\text{Enc}(\text{pk}_i, m_i)\}_{i \in [n]} \xrightarrow{\text{Eval}((\text{pk}_1, \dots, \text{pk}_n), C, \cdot)} \text{Enc}((\text{pk}_1, \dots, \text{pk}_n), C(m_1, \dots, m_n)).$$

Clearly, decrypting the resulting ciphertext should require the knowledge of all of the corresponding secret keys $(\text{sk}_1, \dots, \text{sk}_n)$. Other than semantic security, we require that the scheme is compact, in the sense that the size of the evaluated ciphertext is proportional to $|C(m_1, \dots, m_n)|$ (and possibly the number of parties n) but does not otherwise depend on the size of C .

The second tool that we use is compute and compare obfuscation [WZ17, GKW17]. A compute and compare program $\text{CC}[f, u, z]$ program is defined by a function f , a lock value u , and an output z . On input a string x , the program returns z if and only if $f(x) = u$. The obfuscator Obf is guaranteed to return an obfuscated program $\widetilde{\text{CC}}$ that is indistinguishable from a program that rejects any input, as long as u has sufficient entropy conditioned on f and z . Finally, we use a conditional disclosure of secret (CDS)² scheme. Recall that this is an interactive protocol parametrized by an NP relation \mathcal{R} where both the sender and the receiver share a statement x and in addition, the sender has a secret message m . At the end of the interaction, the receiver obtains m if and only if it knows a valid witness w such that $\mathcal{R}(x, w) = 1$.

A Strawman Solution. We now describe a naive extension of the [BS20, AP20] approach to the parallel setting (where a receiver interacts with multiple committers), and highlight its pitfalls. We do not assume familiarity with [BS20, AP20].

To commit to messages (m_1, \dots, m_n) , the committers and the receiver engage in the following protocol.

- Each committer samples a key pair of a multi-key FHE scheme $(\text{pk}_i, \text{sk}_i)$, a uniform trapdoor td_i , and a uniform lock value lk_i , and sends to the receiver:
 1. A commitment $c_i = \text{Com}(\text{td}_i)$.
 2. An FHE encryption $\text{Enc}(\text{pk}_i, \text{td}_i)$.
 3. An obfuscation $\widetilde{\text{CC}}_i$ of the program $\text{CC}[\text{Dec}(\text{sk}_i, \cdot), \text{lk}_i, (\text{sk}_i, m_i)]$.
- The receiver engages each committer in a (parallel) execution of a CDS protocol where the i 'th committer sends lk_i if the receiver correctly guesses a valid pre-image of c_i .

At a high level, the fact that the protocol hides the message m_i is ensured by the following argument. Since the receiver cannot invert c_i , it cannot guess td_i and therefore the CDS protocol will return 0. This in turn means that the lock lk_i is hidden from the receiver, and consequently that the obfuscated program is indistinguishable from a null program. This is, of course, an informal explanation, and we refer the reader to [BKP19, BS20, AP20] for a formal security analysis.

We now turn to the description of the extractor. The high-level strategy is the following: Upon receiving the first message from all committers, the extractor uses the FHE encryption $\text{Enc}(\text{pk}_i, \text{td}_i)$ and the code of

²In the body of the paper we actually resort to a slightly stronger tool, namely a secure function evaluation protocol with statistical circuit privacy.

the adversary to run the CDS protocol homomorphically (on input td_i) to recover an FHE encryption of lk_i . Then the extractor feeds it as an input to the obfuscated program \widetilde{CC}_i , which returns (sk_i, m_i) .

Unfortunately this approach has a major limitation: It implicitly assumes that each corrupted party is a local algorithm. In other words, we are assuming that the adversary consists of individual subroutines (one per corrupted party), which may not necessarily be the case. As an example, if the adversary were to somehow implement a strategy where corrupted machines do not respond until *all* receiver messages have been delivered, then the above homomorphic evaluation would get stuck and return no output. It is also worth mentioning that what makes the problem challenging is our inability to clone the state of the adversary. If we were allowed to clone its state, then we could extract messages one by one, by running a separate thread under each FHE key.

Multi-Key Evaluation. A natural solution to circumvent the above issue is to rely on multi-key FHE evaluation. Using this additional property, the extractor can turn the ciphertexts $\text{Enc}(pk_1, td_1), \dots, \text{Enc}(pk_n, td_n)$ into a single encryption

$$\text{Enc}((pk_1, \dots, pk_n), (td_1, \dots, td_n))$$

under the hood of all public keys (pk_1, \dots, pk_n) . Given this information, the extractor can homomorphically evaluate all instances of the CDS protocol at once, using the code of the adversary, no matter how intricate. This procedure allows the extractor to obtain the encryption of each lock value $\text{Enc}((pk_1, \dots, pk_n), lk_i)$. In the single committer setting, we could then feed this into the corresponding obfuscated program and call it a day.

However, in the parallel setting, even given multi-key FHE, it is unclear how to proceed. If the compute and compare program \widetilde{CC}_i tried to decrypt such a ciphertext, it would obtain (at best) an encryption under the remaining public keys. Glossing over the fact that the structure of single-key and multi-key ciphertexts might be incompatible, it is unlikely that

$$\text{Dec}(sk_i, \text{Enc}((pk_1, \dots, pk_n), lk_i)) = lk_i$$

which is what we would need to trigger the compute and compare program. The general problem here is that each compute and compare program cannot encode information about other secret keys, thus making it infeasible to decrypt multi-key ciphertexts. One approach to resolve this issue would be to ask all committers to jointly obfuscate a compute and compare program that encodes all secret keys at once. However, this seems to require a general-purpose MPC protocol, which is what we are trying to build in the first place. Therefore, we outline a different approach by imagining a special kind of multi-key fully homomorphic encryption scheme.

A spooky encryption³ scheme [DHRW16] is an FHE scheme that supports a special *spooky evaluation* algorithm, that generates no-signaling correlations among independently encrypted messages. We will restrict attention to a sub-class of no-signaling relations called *additive function sharing* (AFS) relations, and we will call the scheme AFS-spooky. More concretely, on input a circuit C and n independently generated ciphertexts (under independently generated public keys), the algorithm Spooky.Eval produces

$$\{\text{Enc}(pk_i, m_i)\}_{i \in [n]} \xrightarrow{\text{Spooky.Eval}((pk_1, \dots, pk_n), C, \cdot)} \{\text{Enc}(pk_i, y_i)\}_{i \in [n]} \text{ s.t. } \bigoplus_{i=1}^n y_i = C(m_1, \dots, m_n).$$

It is not hard to see that AFS-spooky encryption is a special case of multi-key FHE where multi-key ciphertexts have the following structure

$$\text{Enc}((pk_1, \dots, pk_n), m) = \{\text{Enc}(pk_i, y_i)\}_{i \in [n]} \text{ s.t. } \bigoplus_{i=1}^n y_i = m.$$

This additional structure is going to be our main leverage for constructing an efficient extractor.

³As a historical remark, while the name is inspired by Einstein's quote "spooky action at a distance" referring to entangled quantum states, the concept of spooky encryption (as defined in [DHRW16]) is entirely classical.

The Extractor. Going back to our extractor, our next technical insight is to look for a mechanism to *peel away* encryption layers one by one from an AFS-spooky (multi-key) ciphertext. Our extractor will achieve this via careful *homomorphic* evaluation of the independently generated programs $(\widetilde{\mathbf{CC}}_1, \dots, \widetilde{\mathbf{CC}}_n)$, as described below.

- First, homomorphically execute the code of the adversary using the AFS-spooky scheme to obtain

$$\text{ct}_1 = \text{Enc}((\text{pk}_1, \dots, \text{pk}_n), \text{lk}_1), \dots, \text{ct}_n = \text{Enc}((\text{pk}_1, \dots, \text{pk}_n), \text{lk}_n),$$

as described above.

- Parse ct_n as a collection of individual ciphertexts

$$\text{Enc}((\text{pk}_1, \dots, \text{pk}_n), \text{lk}_n) = \{\text{Enc}(\text{pk}_i, y_i)\}_{i \in [n]} = \{\text{Enc}(\text{pk}_i, y_i)\}_{i \in [n-1]} \cup \underbrace{\{\text{Enc}(\text{pk}_n, y_n)\}}_{\tilde{\text{ct}}_n}.$$

Note that we can interpret the first $n - 1$ elements as an AFS-spooky ciphertext encrypted under $(\text{pk}_1, \dots, \text{pk}_{n-1})$:

$$\tilde{\text{ct}} = \{\text{Enc}(\text{pk}_i, y_i)\}_{i \in [n-1]} = \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \bigoplus_{i=1}^{n-1} y_i \right) = \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \tilde{y} \right)$$

where $\tilde{y} = \bigoplus_{i=1}^{n-1} y_i$.

- Let Γ be the following function

$$\Gamma(\zeta) : \text{Spooky.Eval}(\text{pk}_n, \zeta \oplus \cdot, \tilde{\text{ct}}_n)$$

which homomorphically computes the XOR of ζ with the plaintext of $\tilde{\text{ct}}_n$. Compute the following nested AFS-spooky correlation

$$\begin{aligned} \hat{\text{ct}} &= \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_{n-1}), \Gamma, \tilde{\text{ct}}) \\ &= \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \text{Spooky.Eval}(\text{pk}_n, \tilde{y} \oplus \cdot, \tilde{\text{ct}}_n) \right) \end{aligned} \quad (1)$$

$$= \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \text{Enc} \left(\text{pk}_n, \bigoplus_{i=1}^n y_i \right) \right) \quad (2)$$

$$= \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \text{Enc}(\text{pk}_n, \text{lk}_n) \right) \quad (3)$$

by interpreting $\tilde{\text{ct}}_n$ as a single key ciphertext. Here (1) follows by substituting Γ , and (2) follows by correctness of the AFS-spooky evaluation.

- Run the obfuscated compute and compare program homomorphically to obtain an encryption of sk_n and m_n under $(\text{pk}_1, \dots, \text{pk}_{n-1})$

$$\begin{aligned} \text{Spooky.Eval} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \widetilde{\mathbf{CC}}_n, \hat{\text{ct}} \right) &= \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), \widetilde{\mathbf{CC}}_n(\text{Enc}(\text{pk}_n, \text{lk}_n)) \right) \\ &= \text{Enc} \left((\text{pk}_1, \dots, \text{pk}_{n-1}), (\text{sk}_n, m_n) \right). \end{aligned}$$

- Using the encryption of sk_n under $(\text{pk}_1, \dots, \text{pk}_{n-1})$, update the initial ciphertexts $(\text{ct}_1, \dots, \text{ct}_{n-1})$ by homomorphically decrypting their last component and adding the resulting string. This allows the extractor to obtain

$$\text{Enc}((\text{pk}_1, \dots, \text{pk}_{n-1}), \text{lk}_1), \dots, \text{Enc}((\text{pk}_1, \dots, \text{pk}_{n-1}), \text{lk}_{n-1}).$$

- Recursively apply the procedure described above until $\text{Enc}(\text{pk}_1, \text{lk}_1)$ is recovered, then feed this ciphertext as an input to $\widetilde{\text{CC}}_1$ to obtain (sk_1, m_1) in the clear. Iteratively recover $(\text{sk}_2, \dots, \text{sk}_n)$ by decrypting the corresponding ciphertexts. At this point the extractor knows all secret keys and can decrypt the transcript of the interaction together with the committed messages.

To summarize, this extractor will isolate single-key ciphertexts (albeit in a nested form) by relying on AFS-spooky encryption. These ciphertexts by design will be compatible with compute and compare programs. In turn, evaluating the program *under the encryption* allows us to *escape* from the newly introduced layer. Repeating this procedure recursively eventually leads to a complete recovery of the plaintexts.

We stress that, although the extraction algorithm repeats the nesting operation n times, the additional encryption layer introduced in each iteration is immediately peeled off by executing the obfuscated compute and compare program. Thus the above procedure runs in (strict) polynomial time for *any polynomial* number of parties n .

Parallel Zero Knowledge. The above outline is deliberately simplified and ignores some subtle issues that arise during the analysis of the protocol. As an example, we need to ensure that the adversary is not able to *maul* the commitment of the trapdoor into a CDS encryption to be used in the CDS protocol. This issue also arose in [BS20], and we follow their approach of using non-uniformity in a reduction to the semantic security of the quantum FHE scheme. [BS20] also present the technical tools needed to lift the protocol to the setting of malicious and possibly aborting adversaries (as opposed to explainable), and we roughly follow their approach. However, it is worth pointing out that [BS20] directly construct a zero-knowledge argument, without first constructing and analyzing a stand-alone extractable commitment. Since we use a parallel extractable commitment as a building block in the our coin-flipping protocol, we analyze the above as a stand-alone commitment, which requires a few modifications to the protocol and proof techniques. More discussion about this can be found in Section 5.

Now, we describe how to obtain parallel zero-knowledge (i.e. zero-knowledge against multiple verifiers) from parallel extractable commitments. This is accomplished in a routine manner by enhancing a standard Σ protocol with a stage where each verifier commits to its Σ protocol challenge using a parallel extractable commitment. Using the extractor, the simulator can obtain the challenges ahead of time and can therefore simulate the rest of the transcript, without the need to perform state cloning.

It remains to argue that our extraction strategy does not break down in the presence of quantum adversaries. Observe that the only step that involves the execution of a quantum circuit is the AFS-spooky evaluation of the CDS protocol, under the hood of $(\text{pk}_1, \dots, \text{pk}_n)$. Assuming that we can construct AFS-spooky encryption for relations computable by quantum circuits (which we show in Section 2.3), the remainder of the extraction algorithm only depends on the encryptions of $(\text{lk}_1, \dots, \text{lk}_n)$, which are classical strings. Once the extractor recovers all the secret keys, it can decrypt the (possibly quantum) state of the adversary resulting from the homomorphic evaluation of the CDS, and resume the protocol execution, without the need to clone the adversary's state.

2.3 Quantum AFS-Spooky Encryption

We now turn to the construction of AFS-spooky encryption for relations computable by quantum circuits. The main technical contribution of this section is a construction of multi-key fully-homomorphic encryption for quantum circuits with classical key generation and classical encryption of classical messages. Such schemes were already known in the *single*-key setting, due to [Mah18, Bra18].

Background. At a very high level, these single-key schemes follow a paradigm introduced by Broadbent and Jeffery [BJ15], which makes use of the quantum one-time pad (QOTP). The QOTP is a method of perfectly encrypting arbitrary quantum states with a key that consists of only classical bits. [BJ15] suggest to encrypt a quantum state with a quantum one-time pad (QOTP), and then encrypt the classical bits that comprise the QOTP using a classical fully-homomorphic encryption scheme. One can then apply quantum gates to the encrypted quantum state, and update the classical encryption of the one-time pad appropriately. A key feature of this encryption procedure is that while an encryption of a quantum state necessarily must

be a quantum state, an encryption of classical information does not necessarily have to include a quantum state. Indeed, one can simply give a classical one-time pad encryption of the data, along with a classical fully-homomorphic encryption of the pad.

However, the original schemes presented by Broadbent and Jeffery [BJ15] and subsequent work [DSS16] based on their paradigm left much to be desired. In particular, they required even a classical encryptor to supply quantum “gadgets” encoding their secret key. These gadgets were then used to evaluate a particular non-Clifford gate over encrypted data.⁴ The main innovation in the work of [Mah18] was to remove the need for quantum gadgets, instead showing how to evaluate an appropriate non-Clifford gate using just *classical* information supplied by the encryptor.

Encrypted CNOT Operation. In more detail, evaluating a non-Clifford gate on a ciphertext $(ct, |\phi\rangle)$, where ct is an FHE encryption of a QOTP key and $|\phi\rangle$ is a quantum state encrypted under the QOTP key, involves an operation (referred to as encrypted CNOT) that somehow must “teleport” the bits encrypted in ct into the state $|\phi\rangle$. [Mah18] gave a method for doing this, as long as the ciphertext ct is encrypted under a scheme with some particular properties. Roughly, the scheme must support a “natural” XOR homomorphic operation, it must be circuit private with respect to this homomorphism, and perhaps most stringently, there must exist some trapdoor that can be used to recover the message and the *randomness* used to produce any ciphertext.

[Mah18] observed that the dual-Regev encryption scheme [GPV08] (with large enough modulus-to-noise ratio) does in fact satisfy these properties, as long as one generates the public key matrix \mathbf{A} along with a trapdoor. However, recall that ct was supposed to be encrypted under a fully-homomorphic encryption scheme. [Mah18] resolves this by observing that ciphertexts encrypted under the dual variant of the [GSW13] fully-homomorphic encryption scheme actually already contain a dual-Regev ciphertext. In particular, a dual-GSW ciphertext encrypting a bit μ is a matrix $\mathbf{M} = \mathbf{A}\mathbf{S} + \mathbf{E} + \mu\mathbf{G}$, where \mathbf{G} is the gadget matrix. The final column of \mathbf{M} is $\mathbf{A}\mathbf{s} + \mathbf{e} + \mu[0, \dots, 0, q/2]^T$, which is exactly a dual-Regev ciphertext encrypting μ under public key \mathbf{A} . Note that, crucially, if the dual-GSW public key \mathbf{A} is drawn with a trapdoor, then this trapdoor also functions as a trapdoor for the dual-Regev ciphertext. Thus, an evaluator can indeed perform the encrypted CNOT operation on any ciphertext $(ct, |\phi\rangle)$, by first extracting a dual-Regev ciphertext ct' from ct and then proceeding.

Challenges in the Multi-Key Setting. Now, it is natural to ask whether this approach readily extends to the multi-key setting. Namely, does there exist a multi-key FHE scheme where any (multi-key) ciphertext contains within it a dual-Regev ciphertext *with a corresponding trapdoor*? Unfortunately, this appears to be much less straightforward than in the single-key setting, for the following reason. Observe that (dual) GSW homomorphic operations over ciphertexts $\mathbf{M}_i = \mathbf{A}\mathbf{S}_i + \mathbf{E}_i + \mu_i\mathbf{G}$ always maintain the same \mathbf{A} matrix, while updating \mathbf{S}_i , \mathbf{E}_i , and μ_i . Thus, a trapdoor for \mathbf{A} naturally functions as a trapdoor for the dual-Regev ciphertext that constitutes the last column of \mathbf{M}_i . However, LWE-based multi-key FHE schemes from the literature [CM15, MW16, PS16, BHP17] include a *ciphertext expansion* procedure, which allows an evaluator, given public keys pk_1, \dots, pk_n , and a ciphertext ct encrypted under some pk_i , to convert ct into a ciphertext \hat{ct} encrypted under all keys pk_1, \dots, pk_n . Now, even if these public keys are indeed matrices $\mathbf{A}_1, \dots, \mathbf{A}_n$ drawn with trapdoors τ_1, \dots, τ_n , it is unclear how to combine τ_1, \dots, τ_n to produce a trapdoor $\hat{\tau}$ for the “expanded” ciphertext. Indeed, the expanded ciphertext generally can no longer be written as some $\mathbf{A}\mathbf{S} + \mathbf{E} + \mu\mathbf{G}$, since the expansion procedure constructs a highly structured matrix that includes components from the ciphertexts ct_1, \dots, ct_n , as well as auxiliary encryptions of the randomness used to produce the ciphertexts (see e.g. [MW16]).

A Solution Based on Key-Switching. Thus, we take a different approach. Rather than attempting to tweak known ciphertext expansion procedures to also support “trapdoor expansion”, we rely on the notion of key-switching, which is a method of taking a ciphertext encrypted under one scheme and converting it into a ciphertext encrypted under another scheme. The observation, roughly, is that we do not need to explicitly maintain a trapdoor for the multi-key FHE scheme, as long as it is possible to convert a multi-key FHE ciphertext into a dual-Regev ciphertext that *does* explicitly have a trapdoor. In fact, we will consider a

⁴We also remark here that [Goy18] presented a *multi-key* scheme based on this paradigm, but with the same drawbacks. Note that compactness and classical encryption are crucial in our setting, as per the discussion in the previous section.

natural multi-key generalization of dual-Regev, as described below. Key switching is possible as long as the second scheme has sufficient homomorphic properties, namely, it can support homomorphic evaluation of the *decryption circuit* of the first scheme.

Fortunately, the dual-Regev scheme is already linearly homomorphic, and many known classical multi-key FHE schemes [CM15, MW16, PS16, BHP17] support *nearly linear decryption*, which means that decrypting a ciphertext simply consists of applying a linear function (derived from the secret key) and then rounding. Thus, as long as the evaluator has the secret key of the multi-key FHE ciphertext encrypted under a dual-Regev public key with a trapdoor, they can first key-switch the multi-key FHE ciphertext ct into a dual-Regev ciphertext ct' , and then proceed with the encrypted CNOT operation.

It remains to show how an evaluator may have access to such a dual-Regev encryption. Since we are still in the multi-key setting, we will need a ciphertext and corresponding trapdoor expansion procedure for dual-Regev. However, we show that such a procedure is much easier to come by when the scheme only needs to support *linear* homomorphism (as is the case for the dual-Regev scheme) rather than *full* homomorphism. Each party can draw its own dual-Regev public key A_i along with a trapdoor τ_i , and encrypt its multi-key FHE secret key under A_i to produce a ciphertext ct_i . The evaluator can then treat the block-diagonal matrix $\hat{A} = \text{diag}(A_1, \dots, A_n)$ as an “expanded” public key.⁵ Now, the message and randomness used to generate a ciphertext encrypted under \hat{A} may be recovered by applying τ_1 to the first set of entries of the ciphertext, applying τ_2 to the second set of entries and so on. This observation, combined with an appropriate expansion procedure for the ciphertexts ct_i , allows an evaluator to convert any multi-key FHE ciphertext into a multi-key dual-Regev ciphertext with trapdoor. Given a classical multi-key FHE scheme with nearly linear decryption, this suffices to build multi-key quantum FHE with classical key generation and encryption.

Distributed Setup. We showed above how to convert any classical multi-key FHE scheme into a quantum multi-key FHE scheme, as long as the classical scheme has nearly linear decryption. However, most LWE-based classical multi-key FHE schemes operate in the common random string (CRS) model, which assumes that all parties have access to a common source of randomness, generated by a trusted party. Thinking back to our application to parallel extractable commitments, it is clear that this will not suffice, since we have no CRS a priori, and a receiver that generates a CRS maliciously may be able to break hiding of the scheme. Thus, we rely on the multi-key FHE scheme of [BHP17], where instead of assuming a CRS, the parties participate in a distributed setup procedure. In particular, each party (and in our application, each committer) generates some public parameters pp_i , which are then combined publicly to produce a single set of public parameters pp , which can be used by anyone to generate their own public key / secret key pair.

This form of distributed setup indeed suffices to prove the hiding of our parallel commitment, so it remains to show that our approach, combined with [BHP17], yields a quantum multi-key FHE scheme with distributed setup. First, the [BHP17] scheme does indeed enjoy nearly linear decryption, so plugging it into our compiler described above gives a functional quantum multi-key FHE scheme. Next, we need to confirm that our compiler does not destroy the distributed setup property. This follows since each party draws its own dual-Regev public key with trapdoor without relying on any CRS, or even any public parameters.

Quantum AFS-Spooky Encryption. Finally, we show, via another application of key-switching, how to construct a quantum AFS-spooky encryption scheme (with distributed setup). Recall that we only require “spooky” interactions to hold over classical ciphertexts. That is, for any *quantum* circuit C with classical outputs, given ciphertexts ct_1, \dots, ct_n encrypting $|\phi_1\rangle, \dots, |\phi_n\rangle$ respectively under public keys pk_1, \dots, pk_n , an evaluator can produce ciphertexts ct'_1, \dots, ct'_n where ct'_i encrypts y_i under pk_i , and such that $\bigoplus_{i=1}^n y_i = C(|\phi_1\rangle, \dots, |\phi_n\rangle)$.

Now, using our quantum multi-key FHE scheme, it is possible to compute a single (multi-key) ciphertext \hat{ct} that encrypts $C(|\phi_1\rangle, \dots, |\phi_n\rangle)$ under all public keys pk_1, \dots, pk_n . Then, if each party additionally drew a key pair (pk'_i, sk'_i) for a classical AFS-spooky encryption scheme, and released $\tilde{ct}_1, \dots, \tilde{ct}_n$, where

⁵Actually this expansion should be done slightly more carefully, see Section 4.4 for details.

$\tilde{ct}_i = \text{Enc}(\text{pk}'_i, \text{sk}_i)$ encrypts the i -th party's quantum multi-key FHE secret key under their AFS-spooky encryption public key, then the evaluator can homomorphically evaluate the quantum multi-key FHE decryption circuit (which is classical for classical ciphertexts) with \hat{c} hardcoded, where \hat{c} is the multi-key ciphertext defined at the beginning of this paragraph. This circuit on input $\tilde{ct}_1, \dots, \tilde{ct}_n$ produces the desired output ct'_1, \dots, ct'_n . Finally, note that the classical AFS-spooky encryption scheme must also have distributed setup, and we show (see Section 4.5) that one can derive a distributed-setup AFS-spooky encryption scheme from [BHP17] using standard techniques [DHRW16].

2.4 Post-Quantum Non-malleable Commitments

In this section, we describe how to obtain constant-round post-quantum non-malleable commitments under the assumption that there exists a natural number $c > 0$ such that quantum polynomial-time adversaries cannot distinguish LWE samples from uniform with advantage better than $\lambda^{-\text{ilog}(c, \lambda)}$, where $\text{ilog}(c, \lambda) = \log \log \dots \log(\lambda)$ c times and λ denotes the security parameter.

We will focus on perfectly binding and computationally hiding constant-round interactive commitments. Loosely speaking, a commitment scheme is said to be non-malleable if no adversary (also called a man-in-the-middle), when participating as a receiver in an execution of an honest commitment $\text{Com}(m)$, can at the same time generate a commitment $\text{Com}(m')$, such that the message m' is related to the original message m . This is equivalent (assuming the existence of one-way functions with security against quantum adversaries) to a tag-based notion where the commit algorithm obtains as an additional input a tag in $\{0, 1\}^\lambda$, and the adversary is restricted to using a tag, or identity, that is different from the tag used to generate its input commitment. We will rely on tag-based definitions throughout this paper. We will also only focus on the *synchronous setting*, where the commitments proceed in rounds, and the man-in-the-middle sends its own message for a specific round before obtaining an honest party's message for the next round.

Before describing our ideas, we briefly discuss existing work on *classically-secure* non-malleable commitments. Unfortunately, existing constructions of constant-round non-malleable commitments against classical adversaries from standard polynomial hardness assumptions [Bar02, PR05, PR08, LPV08, PPV08, LP09, Wee10, PW10, LP11, Goy11, GLOV12, GRRV14, GPR16, COSV16, COSV17, Khu17, GR19] either rely on rewinding, or use Barak's non-black-box simulation technique, both of which require the reduction to perform state cloning. As such, known techniques fail to prove quantum security of these constructions.

We now discuss our techniques for constructing post-quantum non-malleable commitments. Just like several classical approaches, we will proceed in two steps.

- We will obtain simple “base” commitment schemes for very small tag/identity spaces from slightly superpolynomial hardness assumptions.
- Then assuming polynomial hardness of LWE against quantum adversaries, and making use of constant-round post-quantum zero-knowledge arguments, we will convert non-malleable commitments for a small tag space into commitments for a larger tag space, while only incurring a constant round overhead.

For the base schemes, there are known classical constructions [PW10] that assume hardness of LWE against 2^{λ^δ} -size adversaries, where λ denotes the security parameter and $0 < \delta < 1$ is a constant. We observe that these constructions can be proven secure in the quantum setting, resulting in schemes that are suitable for tag spaces of $O(\log \log \lambda)$ tags.

Tag Amplification. Since an MPC protocol could be executed among up to $\text{poly}(\lambda)$ parties where $\text{poly}(\cdot)$ is an arbitrary polynomial, we end up requiring non-malleable commitments suitable for tag spaces of $\text{poly}(\lambda)$. This is obtained by combining classical tools for amplifying tag spaces [DDN91] with constant round post-quantum zero-knowledge protocols. Our tag amplification protocol, on input a scheme with tag space $2t$, outputs a scheme with tag space 2^t , for any $t \leq \text{poly}(\lambda)$. This follows mostly along the lines of existing classical protocols, and as such we do not discuss the protocol in detail here. Our protocol can be found in Section 7.3.

Base Schemes from $\lambda^{-\text{ilog}(c,\lambda)}$ Hardness. Returning to the question of constructing appropriate base schemes, we also improve the assumption from 2^{λ^δ} -quantum hardness of LWE (that follows based on [PW10]) to the mildly superpolynomial hardness assumption discussed at the beginning of this subsection. Recall that we will only need to assume that there exists an (explicit) natural number $c > 0$ such that quantum polynomial time adversaries cannot distinguish LWE samples from uniform with advantage better than $\text{negl}(\lambda^{\text{ilog}(c,\lambda)})$ where $\text{ilog}(c, \lambda) = \log \log \dots_c \text{times} \log(\lambda)$. Our base scheme will only be suitable for identities in $\text{ilog}(c+1, \lambda)$, where $c > 0$ is a natural number, independent of λ . We will then repeatedly apply the tag amplification process referred to above to boost the tag space to 2^λ , by adding only a constant number of rounds.

To build our base scheme, we take inspiration from the classically secure non-malleable commitments of Khurana and Sahai [KS17]. However, beyond considering quantum as opposed to classical adversaries, our protocol and analysis will have the following notable differences from [KS17]:

- The work of [KS17] relies on sub-exponential hardness (i.e. 2^{λ^δ} security), which is stronger than the type of superpolynomial hardness we assume. This is primarily because [KS17] were restricted to two rounds, but we can improve parameters while allowing for a larger constant number of rounds.
- [KS17] build a reduction that rewinds an adversary to the beginning of the protocol, and executes the adversary several times, repeatedly sampling the adversary's initial state. This may be undesirable in the quantum setting.⁶ On the other hand, we have a simpler fully straight-line reduction that only needs to run the adversary once.

Specifically, following [KS17], we will establish an *erasure channel* between the committer and receiver that transmits the committed message to the receiver with probability ϵ . To ensure that the commitment satisfies hiding, ϵ is chosen to be a value that is negligible in λ . At the same time, the exact value of ϵ is determined by the identity (tag) of the committer. Recall that $\text{tag} \in [1, \text{ilog}(c+1, \lambda)]$. We will set $\epsilon = \eta^{-\text{tag}}$ where $\eta = \lambda^{\text{ilog}(c+1, \lambda)}$ is a superpolynomial function of λ .

Next, for simplicity, we restrict ourselves to a case where the adversary's tag (which we denote by tag') is smaller than that of the honest party (which we denote by tag). In this case, the adversary's committed message is transmitted with probability $\epsilon' = \eta^{-\text{tag}'}$, whereas the honest committer's message is transmitted with probability only $\epsilon = \eta^{-\text{tag}}$, which is smaller than ϵ' .

We set this up so that the transcript of an execution transmits the adversary's message with probability ϵ' (over the randomness of the honest receiver), and on the other hand, an honestly committed message will remain hidden except with probability $\epsilon < \epsilon'$ (over the randomness of the honest committer). This gap in the probability of extraction will help us argue non-malleability, using a proof strategy that bears resemblance to the proof technique in [BL18b] (who relied on stronger assumptions to achieve such a gap in the non-interactive setting).

We point out one subtlety in our proof that does not appear in [BL18b]. We must rule out a man-in-the-middle adversary that on the one hand, does not commit to a related message if its message was successfully transmitted, but on the other hand, can successfully perform a mauling attack if its message was not transmitted. To rule out such an adversary, just like [KS17], we will design our erasure channel so that the adversary cannot distinguish transcripts where his committed message was transmitted from those where it wasn't.

Finally, our erasure channel can be cryptographically established in a manner similar to prior work [KS17, KKS18, BFJ⁺20] via an indistinguishability-based variant of two-party secure function evaluation, that can be based on quantum hardness of LWE. Specifically, we would like to ensure that the SFE error is (significantly) smaller than the transmission probabilities of our erasure channels: therefore, we will set parameters so that SFE error is $\lambda^{-\text{ilog}(c,\lambda)}$. We refer the reader to Section 7 for additional details about our construction.

On Super-Constant Rounds from Polynomial Hardness. We also observe that for any $t(\lambda) \leq \text{poly}(\lambda)$, non-malleable commitments for tag space of size $t(\lambda)$ can be obtained in $O(t(\lambda))$ rounds based on any extractable commitment using ideas from [DDN91, CR87], where only one party speaks in every round. These admit

⁶In particular this state may not always be efficiently sampleable, in which case it would be difficult to build an efficient reduction.

a straight-line reduction, and can be observed to be quantum-secure. As such, based on quantum polynomial hardness of LWE and quantum FHE, we can obtain a base protocol for $O(\log \log \dots c \text{ times } \log \lambda)$ tags requiring $O(\log \log \dots c \text{ times } \log \lambda)$ rounds, for any constant $c \in \mathbb{N}$. Applying our tag-amplification compiler to this base protocol makes it possible to increase the tag space to 2^λ while only adding a constant number of rounds. Therefore, this technique gives $O(\log \log \dots c \text{ times } \log \lambda)$ round non-malleable commitments for exponentially large tags from quantum polynomial hardness. It also yields constant round non-malleable commitments for a constant number of tags from polynomial hardness.

2.5 Putting Things Together

Finally, we show how to combine the primitives described above to obtain a constant-round coin-flipping protocol that supports straight-line simulation. As we saw above, in the setting of multi-verifier zero-knowledge, simultaneously simulating the view of multiple parties without rewinding can be quite challenging, so a careful protocol and proof is needed.

Recall the outline presented at the beginning of this section, where each party first commits to a uniformly random string, then broadcasts the committed message, and finally proves in ZK that the message broadcasted is equal to the previously committed message. If all proofs verify, then the common output is the XOR of all broadcasted strings. Recall also that the coin-tossing protocol should be *fully-simulatable*. This means that a simulator should be able to force the common output to be a particular uniformly drawn string given to it as input.

It turns out that in order to somehow force a particular output, the simulator should be able to *simultaneously extract in advance* all the messages that adversarial parties committed to. In particular, we require commitments where a simulator can extract from multiple committers committing in parallel. Here, we will rely on our parallel extractable commitment described above. Note that we will also need to simulate the subsequent zero-knowledge arguments given by the malicious parties in parallel, and thus we instantiate these with our parallel zero-knowledge argument described above. However, an issue remains. What if an adversary could somehow *maul* an honest party's commitment to a related message and then broadcast that commitment as their own? This could bias the final outcome away from uniformly random.

Thus, we need to introduce some form of non-malleability into the protocol. Indeed, we will add another step at the beginning where each party commits to its message c_i and some randomness r_i using our post-quantum many-to-one non-malleable commitment.⁷ Each party will then commit to c_i again with our extractable commitment, using randomness r_i . Finally, each party proves in zero-knowledge that the previous commitments were consistent.

This protocol can be proven to be fully simulatable. Intuitively, even though the simulator changes the behavior of honest players in order to extract from the adversary's commitments and then later force the appropriate output, the initial non-malleable commitments given by the adversary must not change in a meaningful way, due to the guarantee of non-malleability. However, additional subtleties arise in the proof of security. In particular, during the hybrids the simulator will first have to simulate the honest party zero-knowledge arguments, before changing the honest party commitments in earlier stages. However, when changing an honest party's commitment, we need to rely on non-malleability to ensure that the malicious party commitments will not also change in a non-trivial way. Here, we use a proof technique that essentially invokes soundness of the adversary's zero-knowledge arguments at an earlier hybrid but allows us to nevertheless rely on non-malleable commitments to enforce that the adversary behaves consistently in all future hybrids. More discussion and a formal analysis can be found in Section 8.

2.6 Related Work

Classical secure multi-party computation was introduced and shown to be achievable in the two-party setting by [Yao82] and in the multi-party setting by [GMW87]. Since these seminal works, there has been

⁷Above we described a construction of one-to-one non-malleable commitment, though a hybrid argument [LPV08] shows that one-to-one implies many-to-one.

considerable interest in reducing the round complexity of classical protocols. In the setting of malicious security against a dishonest majority, [Lin03] gave the first *constant*-round protocol for two-party computation, and [KOS03] gave the first constant-round protocol for multi-party computation. Since then, there has been a long line of work improving on the exact round complexity and assumptions necessary for classical multi-party computation (see e.g. [Pas04, GMPP16]).

Post-quantum classical protocols. The above works generally focus on security against *classical* polynomial-time adversaries. Another line of work, most relevant to the present work, has considered the more general goal of proving the security of classical protocols against arbitrary *quantum* polynomial-time adversaries.

This study was initiated by van de Graaf [VDG98], who observed that the useful rewinding technique often used to prove zero-knowledge in the classical setting may be problematic in the quantum setting. In a breakthrough work, Watrous [Wat09] showed that several well-known classical zero-knowledge protocols are in fact zero-knowledge against quantum verifiers, via a careful rewinding argument. However, these protocols require a polynomial number of rounds to achieve negligible security against quantum attackers. Later, Unruh [Unr12] developed a more powerful rewinding technique that suffices to construct classical zero-knowledge *proofs of knowledge* secure against quantum adversaries, though still in a polynomial number of rounds. In a recent work, [BS20] managed to construct a constant-round post-quantum zero-knowledge protocol, under assumptions similar to those required to obtain classical fully-homomorphic encryption. In another recent work, [AP20] constructed a constant-round protocol that is zero-knowledge against quantum verifiers under the quantum LWE assumption, though soundness holds against only classical provers.

There has also been some work on the more general question of post-quantum secure computation. In particular, [DL09] used the techniques developed in [Wat09] to build a two-party coin-flipping protocol, and [LN11, HSS11] constructed general two-party computation secure against quantum adversaries, in a polynomial number of rounds. More recently, [BS20] gave a *constant*-round two-party coin-flipping protocol, with full simulation of one party. However, prior to this work, nothing was known in the most general setting of post-quantum multi-party computation (in the plain model).

Finally, as mentioned at the beginning of Section 2.1, there exist post-quantum classical protocols in the literature, as long as some form of trusted setup is available.

Quantum protocols. Yet another line of work focuses on protocols for securely computing *quantum* circuits. General multi-party quantum computation was shown to be achievable in the information-theoretic setting (with honest majority) in the works of [CGS02, BCG⁺06]. In the computational setting, [DNS10] gave a two-party protocol secure against a quantum analogue of semi-honest adversaries, and [DNS12] extended security of two-party quantum computation to the malicious setting. In a recent work [DGJ⁺20] constructed a maliciously secure multi-party protocol for computing quantum circuits, assuming the existence of a maliciously secure post-quantum classical MPC protocol. We remark that all of the above protocols operate in a polynomial number of rounds.

3 Preliminaries

Various parts of this section are taken nearly verbatim from [BS20]. All algorithms of cryptographic functionalities in this work are implicitly efficient and classical (i.e. require no quantum computation or a quantum communication channel), unless noted otherwise. We rely on the standard notions of classical Turing machines and Boolean circuits:

- We say that a Turing machine (or algorithm) is PPT if it is probabilistic and runs in polynomial time.
- We sometimes think about PPT Turing machines as polynomial-size uniform families of circuits (as these are equivalent models). A polynomial-size circuit family C is a sequence of circuits $C = \{C_\lambda\}_{\lambda \in \mathbb{N}}$, such that each circuit C_λ is of polynomial size $\lambda^{O(1)}$ and has $\lambda^{O(1)}$ input and output bits. We say that the family is uniform if there exists a polynomial-time deterministic Turing machine M that on input 1^λ outputs C_λ .

- For a PPT Turing machine (algorithm) M , we denote by $M(x; r)$ the output of M on input x and random coins r . For such an algorithm, and any input x , we may write $m \in M(x)$ to denote the fact that m is in the support of $M(x; \cdot)$.

Miscellaneous notation.

- For a distribution \mathcal{D} that may explicitly take its random coins r as input, we denote by $x \leftarrow \mathcal{D}$ the process of sampling from \mathcal{D} , and denote by $x := \mathcal{D}(r)$ the fixed outcome x that results from sampling from \mathcal{D} with random coins r .
- We denote by U_λ the uniform distribution over $\{0, 1\}^\lambda$.
- Given an NP language \mathcal{L} with associated relation $\mathcal{R}_\mathcal{L}$, and an instance x , we let $\mathcal{R}_\mathcal{L}(x)$ denote the set $\{w : \mathcal{R}_\mathcal{L}(x, w) = 1\}$.
- For some natural number c and security parameter λ , we use $\text{ilog}(c, \lambda)$ to denote $\underbrace{\log \log \cdots \log(\lambda)}_{c \text{ times}}$.
- We will use $\Delta(\mathcal{X}, \mathcal{Y})$ to denote the statistical distance between two distributions \mathcal{X} and \mathcal{Y} .

3.1 Quantum Computation

We use standard notions from quantum computation.

- We say that a Turing machine (or algorithm) is QPT if it is quantum and runs in polynomial time.
- We sometimes think about QPT Turing machines as polynomial-size uniform families of quantum circuits (as these are equivalent models). We call a polynomial-size quantum circuit family $C = \{C_\lambda\}_{\lambda \in \mathbb{N}}$ uniform if there exists a polynomial-time deterministic Turing machine M that on input 1^λ outputs C_λ .
- Classical communication channels in the quantum setting are identical to classical communication channels in the classical setting, except that when a set of qubits is sent through a classical communication channel, then the qubits are automatically measured in the standard basis, and the measured (now classical-state) qubits are then sent through the channel.
- A quantum interactive algorithm (in a 2-party setting) has input divided into two registers and output divided into two registers. For the input qubits, one register is for an input message from the other party, and a second register is for a potential inner state the machine holds. For the output, one register is for the message to be sent to the other party, and another register is for a potential inner state for the machine to keep to itself.

Quantum Adversarial Model. We would like to consider security definitions that not only achieve quantum security, but are also composable and can be used modularly inside other protocols. For this we think by default of security against polynomial-size quantum adversaries with non-uniform polynomial-size quantum advice (i.e. an arbitrary quantum mixed state that is not necessarily efficiently generatable).

An adversary will be usually denoted by $A^* = \{A_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, where $\{A_\lambda^*\}_{\lambda \in \mathbb{N}}$ is a polynomial-size non-uniform sequence of quantum circuits, and $\{\rho_\lambda\}_{\lambda \in \mathbb{N}}$ is some polynomial-size sequence of mixed quantum states. All adversaries are implicitly unrestricted in their behaviour (i.e. they are fully malicious and can arbitrarily deviate from protocols). We conclude with notions regarding indistinguishability in the quantum setting.

- A function $f : \mathbb{N} \rightarrow [0, 1]$ is:
 - negligible if for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$.

– noticeable if there exists $c \in \mathbb{N}, N \in \mathbb{N}$ s.t. for every $n \geq N, f(n) \geq n^{-c}$.

- A quantum random variable is simply a random variable that can have values that are quantum states. That is, a quantum random variable induces a probability distribution over a (possibly infinite) set of quantum states. Such quantum random variables can also be thought of as a mixed quantum state, which is simply a distribution over quantum states.
- For two quantum random variables X and Y , quantum distinguisher D with quantum mixed state ρ as auxiliary input, and $\mu \in [0, 1]$, we write $X \approx_{D(\rho), \mu} Y$ if

$$|\Pr[D(X; \rho) = 1] - \Pr[D(Y; \rho) = 1]| \leq \mu.$$

- Two ensembles of quantum random variables $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are said to be computationally indistinguishable, denoted by $X \approx_c Y$, if for every polynomial-size non-uniform quantum distinguisher with quantum advice $D = \{D_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$,

$$X_\lambda \approx_{D_\lambda(\rho_\lambda), \mu(\lambda)} Y_\lambda.$$

- The trace distance between two quantum distributions X, Y , denoted by $\text{TD}(X, Y)$, is a generalization of statistical distance to the quantum setting and represents the maximal distinguishing advantage between two quantum distributions by an unbounded quantum algorithm. We thus say that $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are statistically indistinguishable (and write $X \approx_s Y$), if for every unbounded non-uniform quantum distinguisher $D = \{D_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, $\text{TD}(X_\lambda, Y_\lambda) \leq \mu(\lambda)$.

3.2 Notation for Interactive Protocols

Throughout, we will be considering interactive protocols, generally defined by a set of classical interactive Turing machines $\{M_i\}_{i \in [n]}$. We denote by $\tau \leftarrow \langle \{M_i(y_i)\}_{i \in [n]} \rangle(x)$ the public transcript of their interaction on common input x , where each M_i has private input y_i . More precisely, τ consists of the messages sent between the $\{M_i\}_{i \in [n]}$, and is a random variable over the random coins of each M_i . We let $\text{VIEW}_{M_j}(\langle \{M_i(y_i)\}_{i \in [n]} \rangle(x))$ denote the view of some party M_j that results from this interaction, which consists of the portion of the transcript τ that includes messages sent by or received by M_j , along with M_j 's private state st at the end of the interaction. If M_j is a quantum machine, then st may be a quantum state. If M_j is defined to have some specific output at the end of the interaction, we denote this by $\text{OUT}_{M_j}(\langle \{M_i(y_i)\}_{i \in [n]} \rangle(x))$.

Definition 3.1 (Explainable Transcript). Let $\{M_i\}_{i \in [n]}$ be a (classical) interactive protocol, and consider some subset of participants $\{M_i\}_{i \in I}$. We say that a transcript τ consisting of messages sent by and received by $\{M_i\}_{i \in I}$ is explainable with respect to set I if there exists some $\{M_i^*\}_{i \notin I}$, inputs $\{y_i\}_{i \in I}$ and random coins $\{r_i\}_{i \in I}$ such that τ is consistent with the transcript of an execution $\langle \{M_i^*\}_{i \notin I}, \{M_i(y_i; r_i)\}_{i \in I} \rangle$.

Handling Abort and Misbehaviour. We set a general convention to handle publicly checkable misbehavior by parties in any interactive protocol.

- For security parameter λ , for each message in the protocol, it will be known (publicly) based on λ , what is the length of each message (or upper and lower bounds on that length). If a party sends a message in an incorrect length, the receiving party fixes it locally and trivially; if the message is too long, it cuts the message in a suitable place, and if it's too short then pads with zeros.
- Whenever a party aborts, all other parties ends communication and output \perp .

3.3 Witness Indistinguishability

We use classical constant-round proof systems for NP (where both honest prover and verifier are classical efficient algorithms) that are witness-indistinguishable against quantum verifiers. That is, transcripts generated by the prover for two witnesses to the same instance are indistinguishable to quantum attackers.

Definition 3.2 (WI Proof System for NP). *A witness-indistinguishable proof system for a language $\mathcal{L} \in \text{NP}$ is a pair (P, V) of classical PPT interactive Turing machines. P and V interact on common input 1^λ and x , and P additionally takes a private input w . At the end of the interaction, V outputs a bit indicating whether it accepts or rejects. The proof system should satisfy the following properties.*

1. **Perfect Completeness:** For any $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, $w \in \mathcal{R}_{\mathcal{L}}(x)$,

$$\Pr[\text{OUT}_V(P(w), V)(1^\lambda, x) = 1] = 1.$$

2. **Statistical Soundness:** For any non-uniform unbounded prover $P^* = \{P_\lambda^*\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for any security parameter $\lambda \in \mathbb{N}$ and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,

$$\Pr[\text{OUT}_V(P_\lambda^*, V)(1^\lambda, x) = 1] = \mu(\lambda).$$

3. **Witness Indistinguishability:** For every non-uniform quantum polynomial-size verifier $V^* = \{V_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, for any two sequences of witnesses $\{w_\lambda\}_{\lambda \in \mathbb{N}}$, $\{v_\lambda\}_{\lambda \in \mathbb{N}}$ s.t. for every $\lambda \in \mathbb{N}$, w_λ and v_λ are both witnesses for the same $x_\lambda \in \mathcal{L} \cap \{0, 1\}^\lambda$, we have,

$$\{\text{VIEW}_{V_\lambda^*}(P(w_\lambda), V_\lambda^*(\rho_\lambda))(1^\lambda, x)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{VIEW}_{V_\lambda^*}(P(v_\lambda), V_\lambda^*(\rho_\lambda))(1^\lambda, x)\}_{\lambda \in \mathbb{N}}.$$

3.4 Sigma Protocol for NP

Definition 3.3 (Sigma Protocol for NP). *A sigma protocol for an NP relation \mathcal{R} is a pair $(P = (P_1, P_2), V = (V_1, V_2))$ of classical PPT Turing machines with the following syntax. Given an instance x and witness w , $P_1(x, w)$ outputs a string α and a prover state st . $V_1(1^{|x|})$ is public-coin, and outputs a uniformly random string β . Next, $P_2(x, w, \text{st}, \alpha, \beta)$ outputs a string γ and finally, $V_2(x, \alpha, \beta, \gamma)$ either accepts or rejects. The proof system should satisfy the following properties.*

1. **Completeness:** For any $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, $w \in \mathcal{R}_{\mathcal{L}}(x)$,

$$\Pr[\text{OUT}_V(P(w), V)(x) = 1] = 1.$$

2. **Statistical Soundness:** For any non-uniform unbounded prover $P^* = \{P_\lambda^*\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for any $\lambda \in \mathbb{N}$ and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,

$$\Pr[\text{OUT}_V(P_\lambda^*, V)(x) = 1] \leq \mu(\lambda).$$

3. **Special Zero-Knowledge:** There exists a PPT simulator Sim such that for any $\{x_\lambda, w_\lambda\}_{\lambda \in \mathbb{N}}$ where $|x_\lambda| = \lambda$ and $(x_\lambda, w_\lambda) \in \mathcal{R}$, and $\{\beta_\lambda\}_{\lambda \in \mathbb{N}}$ where $|\beta_\lambda| = \lambda$,

$$\{(\alpha, \gamma) \mid (\alpha, \text{st}) \leftarrow P_1(x_\lambda, w_\lambda), \gamma \leftarrow P_2(x_\lambda, w_\lambda, \text{st}, \alpha, \beta_\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{(\alpha, \gamma) \leftarrow \text{Sim}(x_\lambda, \beta_\lambda)\}_{\lambda \in \mathbb{N}}.$$

Observe that due to the prover's first message being generated independently of the verifier's message, this implies that for any $\{x_\lambda, w_\lambda\}_{\lambda \in \mathbb{N}}$ where $|x_\lambda| = \lambda$ and $(x_\lambda, w_\lambda) \in \mathcal{R}$,

$$\{\alpha \leftarrow P_1(x_\lambda, w_\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\alpha \mid (\alpha, \gamma) \leftarrow \text{Sim}(x_\lambda, 0^\lambda)\}_{\lambda \in \mathbb{N}}.$$

We refer to this as **First-Message Indistinguishability**.

Sigma protocols are known to follow from classical zero-knowledge proof systems such as the (parallel repetition) of the 3-coloring protocol [GMW91], which is in turn based on non-interactive perfectly-binding and computationally hiding commitments.

3.5 Non-Interactive Commitment

Definition 3.4 (Quantum-secure Non-interactive Commitment). *A non-interactive commitment is defined by a PPT algorithm Com that takes as input security parameter 1^λ and $x \in \{0, 1\}^*$, and outputs a commitment c . The commitment algorithm satisfies:*

1. **Perfect Binding:** *For any $x, x' \in \{0, 1\}^*$ of the same length, if $c \in \text{Com}(1^\lambda, x), c \in \text{Com}(1^\lambda, x')$, then $x = x'$.*
2. **Quantum Computational Hiding:** *For any pair of $\text{poly}(\lambda)$ -length strings $x_0 = \{x_{0,\lambda}\}_{\lambda \in \mathbb{N}}, x_1 = \{x_{1,\lambda}\}_{\lambda \in \mathbb{N}}$, we have,*

$$\{\text{Com}(1^\lambda, x_{0,\lambda})\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Com}(1^\lambda, x_{1,\lambda})\}_{\lambda \in \mathbb{N}}.$$

Instantiations. Non-interactive commitments with quantum hiding are known based on various standard assumptions, including LWE [GHKW17].

3.6 Compute and Compare Obfuscation

We start by defining the class of *compute and compare circuits*.

Definition 3.5 (Compute and compare). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^\lambda$ be a circuit, and let $u \in \{0, 1\}^\lambda$ and $z \in \{0, 1\}^*$ be two strings. Then $\text{CC}[f, u, z](x)$ is a circuit that returns z if $f(x) = u$, and 0 otherwise.*

We now define compute and compare (CC) obfuscators with perfect correctness. In what follows Obf is a PPT algorithm that takes as input a CC circuit $\text{CC}[f, u, z]$ and outputs a new circuit $\widetilde{\text{CC}}$. (We assume that the CC circuit $\text{CC}[f, u, z]$ is given in some canonical description from which f, u , and z can be read.)

Definition 3.6 (CC obfuscator). *An algorithm Obf is a compute and compare obfuscator if it satisfies:*

1. **Perfect correctness:** *For any circuit $f : \{0, 1\}^n \rightarrow \{0, 1\}^\lambda, u \in \{0, 1\}^\lambda, z \in \{0, 1\}^*$,*

$$\Pr \left[\forall x \in \{0, 1\}^n : \widetilde{\text{CC}}(x) = \text{CC}[f, u, z](x) \mid \widetilde{\text{CC}} \leftarrow \text{Obf}(\text{CC}[f, u, z]) \right] = 1 .$$

2. **Simulation:** *There exists a PPT simulator Sim such that for any polynomial-size quantum circuit family $f = \{f_\lambda\}_{\lambda \in \mathbb{N}}$ and polynomial-length output string $z = \{z_\lambda\}_{\lambda \in \mathbb{N}}$,*

$$\{\widetilde{\text{CC}}|u \leftarrow U_\lambda, \widetilde{\text{CC}} \leftarrow \text{Obf}(\text{CC}[f_\lambda, u, z_\lambda])\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Sim}(1^{|f_\lambda|}, 1^{|z_\lambda|}, 1^\lambda)\}_{\lambda \in \mathbb{N}}.$$

Instantiations. Compute and compare obfuscators with almost perfect correctness are constructed in [GKW17, WZ17] based on quantum LWE, and recently with perfect correctness in [GKVV19] based on quantum LWE.

3.7 Function-Hiding Secure Function Evaluation

We define secure function evaluation protocols with statistical circuit privacy and quantum input privacy.

Definition 3.7 (2-Message Function Hiding SFE). *A two-message secure function evaluation protocol ($\text{SFE.Gen}, \text{SFE.Enc}, \text{SFE.Eval}$) has the following syntax:*

- $\text{dk} \leftarrow \text{SFE.Gen}(1^\lambda)$: *a probabilistic algorithm that takes a security parameter 1^λ and outputs a secret key dk .*
- $\text{ct} \leftarrow \text{SFE.Enc}(\text{dk}, x)$: *a probabilistic algorithm that takes a string $x \in \{0, 1\}^*$ and outputs a ciphertext ct .*
- $\widehat{\text{ct}} \leftarrow \text{SFE.Eval}(C, \text{ct})$: *a probabilistic algorithm that takes a classical circuit C and ciphertext ct and outputs an evaluated ciphertext $\widehat{\text{ct}}$.*

- $\hat{x} = \text{SFE.Dec}(\text{dk}, \hat{\text{ct}})$: a deterministic algorithm that takes a ciphertext $\hat{\text{ct}}$ and outputs a string \hat{x} .

For any polynomial-size family of classical circuits $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ the scheme satisfies:

- **Perfect Correctness:** For any $\lambda \in \mathbb{N}$, $x \in \{0, 1\}^*$ and circuit $C \in \mathcal{C}_\lambda$,

$$\Pr[\text{SFE.Dec}_{\text{dk}}(\hat{\text{ct}}) = C(x) | \text{dk} \leftarrow \text{SFE.Gen}(1^\lambda), \text{ct} \leftarrow \text{SFE.Enc}_{\text{dk}}(x), \hat{\text{ct}} \leftarrow \text{SFE.Eval}(C, \text{ct})] = 1$$

- **Quantum Input Privacy:** For polynomial $\ell(\lambda)$ and polynomial-size quantum adversary $\mathbf{A}^* = \{\mathbf{A}_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for every two length $\ell(\lambda)$ messages $\{x_{0,\lambda}\}_{\lambda \in \mathbb{N}}$, $\{x_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ for every $\lambda \in \mathbb{N}$:

$$\Pr[\mathbf{A}_\lambda^*(\text{ct}) = b | \text{dk} \leftarrow \text{SFE.Gen}(1^\lambda), \text{ct} \leftarrow \text{SFE.Enc}_{\text{dk}}(x)] \leq \frac{1}{2} + \mu(\lambda)$$

- **Statistical Circuit Privacy:** There exist unbounded algorithms, probabilistic Sim and deterministic Ext such that for every $x \in \{0, 1\}^*$, $\text{ct} \in \text{SFE.Enc}(x)$, the extractor outputs $\text{Ext}(\text{ct}) = x$ and:

$$\left\{ \text{SFE.Eval}(C, \text{ct}^*) \right\}_{\substack{\lambda \in \mathbb{N}, C \in \mathcal{C}_\lambda, \\ \text{ct}^* \in \{0,1\}^{\text{poly}(\lambda)}}} \approx_s \left\{ \text{Sim}(C(\text{Ext}(\text{ct}^*); 1^\lambda)); 1^\lambda \right\}_{\substack{\lambda \in \mathbb{N}, C \in \mathcal{C}_\lambda, \\ \text{ct}^* \in \{0,1\}^{\text{poly}(\lambda)}}}$$

Specifically, there exists a constant $c > 0$ such that for large enough λ , the statistical distance between the two distributions is at most $2^{-\lambda^c}$.

We will use the following claim in our analysis. This follows directly from the statistical circuit privacy property.

Claim 3.1 (Evaluations of Agreeing Circuits are Statistically Close). *Let $\text{ct}^* = \{\text{ct}_\lambda^*\}_{\lambda \in \mathbb{N}}$ be any (possibly non-ciphertext) $\text{poly}(\lambda)$ length string and let $\mathcal{C}_0 = \{\mathcal{C}_{0,\lambda}\}_{\lambda \in \mathbb{N}}$, $\mathcal{C}_1 = \{\mathcal{C}_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ be two families of circuits such that for all $\lambda \in \mathbb{N}$, $\mathcal{C}_{0,\lambda}$ and $\mathcal{C}_{1,\lambda}$ have identical truth tables. Then*

$$\left\{ \text{SFE.Eval}(C_0, \text{ct}^*) \right\}_{\lambda \in \mathbb{N}, C_0 \in \mathcal{C}_{0,\lambda}} \approx_s \left\{ \text{SFE.Eval}(C_1, \text{ct}^*) \right\}_{\lambda \in \mathbb{N}, C_1 \in \mathcal{C}_{1,\lambda}}$$

Specifically, there exists a constant $c > 0$ such that for large enough λ , the statistical distance between the two distributions is at most $2^{-\lambda^c}$.

Secure function evaluation schemes satisfying Definition 3.7 for functions in NC1 are known based on quantum hardness of LWE [BD18b].

We also define a superpolynomially secure variant of 2-message function hiding SFE where the quantum input privacy property restricts adversaries to having smaller than inverse superpolynomial advantage, for a small superpolynomial function.

Definition 3.8 (2-Message Function Hiding SFE). *A two-message SFE protocol with superpolynomial security is identical to the definition in Definition 3.7, except that it modifies the quantum input privacy requirement as follows: There exists a constant $c > 0$ such that for polynomial $\ell(\lambda)$ and polynomial-size quantum adversary $\mathbf{A}^* = \{\mathbf{A}_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for every two length $\ell(\lambda)$ messages $\{x_{0,\lambda}\}_{\lambda \in \mathbb{N}}$, $\{x_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ for every $\lambda \in \mathbb{N}$:*

$$\Pr[\mathbf{A}_\lambda^*(\text{ct}) = b | \text{dk} \leftarrow \text{SFE.Gen}(1^\lambda), \text{ct} \leftarrow \text{SFE.Enc}_{\text{dk}}(x)] \leq \frac{1}{2} + \mu(\lambda^{\text{ilog}(c,\lambda)})$$

Secure function evaluation schemes satisfying Definition 3.8 for functions in NC1 can be based on quantum slightly superpolynomial hardness of LWE [BD18b]. Specifically, we assume that QPT distinguishers have advantage at most $\text{negl}(\lambda^{\text{ilog}(c,\lambda)})$ in distinguishing LWE samples from uniformly random matrices.

3.8 Quantum Rewinding Lemma

We will make use of the following lemma from [Wat09] and re-worded in [BS20].

Lemma 3.1. *There is a quantum algorithm R that gets as input:*

- A general quantum circuit Q with n input qubits that outputs a classical bit b and an additional m qubits.
- An n -qubit state $|\psi\rangle$.
- A number $t \in \mathbb{N}$.

R executes in time $t \cdot \text{poly}(|Q|)$ and outputs a distribution over m -qubit states $D_\psi := R(Q, |\psi\rangle, t)$ with the following guarantees.

For an n -qubit state $|\psi\rangle$, denote by Q_ψ the conditional distribution of the output distribution $Q(|\psi\rangle)$, conditioned on $b = 0$, and denote by $p(\psi)$ the probability that $b = 0$. If there exist $p_0, q \in (0, 1), \epsilon \in (0, \frac{1}{2})$ such that:

- Amplification executes for enough time: $t \geq \frac{\log(1/\epsilon)}{4 \cdot p_0(1-p_0)}$,
- There is some minimal probability that $b = 0$: For every n -qubit state $|\psi\rangle, p_0 \leq p(\psi)$,
- $p(\psi)$ is input-independent, up to ϵ distance: For every n -qubit state $|\psi\rangle, |p(\psi) - q| < \epsilon$, and
- q is closer to $\frac{1}{2}$: $p_0(1-p_0) \leq q(1-q)$,

then for every n -qubit state $|\psi\rangle$,

$$\text{TD}(Q_\psi, D_\psi) \leq 4\sqrt{\epsilon} \frac{\log(1/\epsilon)}{p_0(1-p_0)}.$$

4 Quantum Multi-Key Fully-Homomorphic Encryption

4.1 Learning with Errors and Lattice Trapdoors

The (decisional) learning with errors problem (LWE), introduced by [Reg05], is parameterized by a modulus q , positive integers n, m , and an error distribution χ . It asks to distinguish between the distributions $(\mathbf{A}, \mathbf{As} + \mathbf{e} \bmod q)$ and (\mathbf{A}, \mathbf{u}) , where \mathbf{A} is uniformly random in $\mathbb{Z}_q^{m \times n}$, \mathbf{s} is uniformly random in \mathbb{Z}_q^n , \mathbf{u} is uniformly random in \mathbb{Z}_q^m , and \mathbf{e} is chosen from χ^m . As shown in [Reg05, PRS17], for *any* sufficiently large modulus q , the LWE problem where χ is a discrete Gaussian distribution with parameter $\sigma = \alpha q \geq 2\sqrt{n}$ (i.e. the distribution over \mathbb{Z} where the probability of x is proportional to $e^{-\pi(|x|/\sigma)^2}$), is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\alpha)$ in *worst case* dimension n lattices. One can truncate the discrete Gaussian distribution to have support only over integers bounded in absolute value by $\sigma \cdot \omega(\sqrt{\log(\lambda)})$ while only introducing a negligible difference. Thus, we will use the fact that χ may be a B -bounded distribution, for some value B .

We will make use of the notion of a lattice trapdoor, defined in the following theorem [MP12].

Theorem 4.1 ([Ajt99, MP12]). *There is an efficient randomized algorithm $\text{GenTrap}(1^n, 1^m, q)$ that, given any integers $n \geq 1, q \geq 2$, and sufficiently large $m = O(n \log q)$, outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a trapdoor $\tau_{\mathbf{A}}$ such that the distribution of \mathbf{A} is negligibly (in n) far from the uniform distribution. Moreover, there is an efficient deterministic algorithm Invert that on input $\mathbf{A}, \tau_{\mathbf{A}}$, and $\mathbf{As} + \mathbf{e}$, where \mathbf{s} is arbitrary in \mathbb{Z}_q^n and $\|\mathbf{e}\| \leq q/(O(n \log q))$, returns \mathbf{s} and \mathbf{e} with overwhelming probability over $(\mathbf{A}, \tau_{\mathbf{A}}) \leftarrow \text{GenTrap}(1^n, 1^m, q)$.*

In fact, we'll need a slightly stronger version of the above statement. In particular, we will actually need the correctness of Invert to hold *perfectly* rather than statistically over the randomness of GenTrap . This can be arranged by slightly tweaking the GenTrap procedure.

Lemma 4.1. *There exist algorithms GenTrap and Invert as described in Theorem 4.1 where Invert returns s, e with probability 1.*

Proof. (Sketch) Call a matrix-trapdoor pair $(\mathbf{A}, \tau_{\mathbf{A}})$ “functional” if \mathbf{A} is full rank (rank n) mod q , $\tau_{\mathbf{A}}$ is an $m \times m$ matrix such that $\tau_{\mathbf{A}} \cdot \mathbf{A} = 0 \pmod{q}$,⁸ each entry of $\tau_{\mathbf{A}}$ is “small enough”, and $\tau_{\mathbf{A}}$ is full rank over the rationals. Such a functional matrix-trapdoor pair may be used to invert *any* vector $\mathbf{v} = \mathbf{A}\mathbf{s} + \mathbf{e}$, for small enough \mathbf{e} , as follows. Left multiply \mathbf{v} by $\tau_{\mathbf{A}}$ over \mathbb{Z}_q , and then left multiply the result by $\tau_{\mathbf{A}}^{-1}$ over the rationals, which recovers \mathbf{e} . Then subtract \mathbf{e} from \mathbf{v} and recover \mathbf{s} by linear algebra. Now observe that the four conditions for $(\mathbf{A}, \tau_{\mathbf{A}})$ to be functional are all efficiently checkable. Thus, the modified GenTrap algorithm can operate as follows. Sample $(\mathbf{A}, \tau_{\mathbf{A}})$ as before, then check if it is functional, and if not replace $(\mathbf{A}, \tau_{\mathbf{A}})$ with some fixed functional pair. Since GenTrap only outputs a non-functional pair with negligible probability, this modification maintains the requirement that the distribution of \mathbf{A} is negligibly close to uniform. \square

4.2 Definition

Definition 4.1 (Quantum Multi-Key Fully-Homomorphic Encryption (QMFHE)). *A quantum multi-key fully-homomorphic encryption scheme is given by six algorithms (QMFHE.KeyGen, QMFHE.Enc, QMFHE.QEnc, QMFHE.Eval, QMFHE.Dec, QMFHE.QDec) with the following syntax.*

- $(\text{pk}, \text{sk}) \leftarrow \text{QMFHE}(1^\lambda)$: A PPT algorithm that given a security parameter, samples a classical public key and a classical secret key.
- $c \leftarrow \text{QMFHE.Enc}(\text{pk}, b)$: A PPT algorithm that takes as input a bit b and outputs a classical ciphertext.
- $|\phi\rangle \leftarrow \text{QMFHE.QEnc}(\text{pk}, |\psi\rangle)$: A QPT algorithm that takes as input a qubit $|\psi\rangle$ and outputs a ciphertext represented in qubits.
- $\widehat{c}, |\widehat{\phi}\rangle \leftarrow \text{QMFHE.Eval}((\text{pk}_1, \dots, \text{pk}_n), C, (|\phi_1\rangle, \dots, |\phi_n\rangle))$: A QPT algorithm that takes as input
 1. A set of n public keys.
 2. A general quantum circuit with $\ell_1 + \dots + \ell_n$ input qubits and ℓ' output qubits, out of which m are measured.
 3. A set of n ciphertexts where $|\phi_i\rangle$ encrypts an ℓ_i -qubit state under public key pk_i . Some of the ℓ_i ciphertexts are possibly classical ciphertexts (generated by the classical encryption algorithm) encrypting classical bits.

The evaluation algorithm outputs a classical ciphertext \widehat{c} encrypting m bits (under keys $\text{pk}_1, \dots, \text{pk}_n$), plus a quantum ciphertext $|\widehat{\phi}\rangle$ encrypting an $(\ell' - m)$ -qubit quantum state (under keys $\text{pk}_1, \dots, \text{pk}_n$).

- $b \leftarrow \text{QMFHE.Dec}((\text{sk}_1, \dots, \text{sk}_n), c)$: A PPT algorithm that takes as input a set of n secret keys and a classical ciphertext c and outputs a bit.
- $|\psi\rangle \leftarrow \text{QMFHE.QDec}((\text{sk}_1, \dots, \text{sk}_n), |\phi\rangle)$: A QPT algorithm that takes as input a set of n secret keys and a quantum ciphertext $|\phi\rangle$ and outputs a qubit.

The scheme satisfies the following.

1. **Quantum Semantic Security:** *The encryption algorithm maintains quantum semantic security.*
2. **Compactness:** *There exists a polynomial $\text{poly}(\cdot)$ s.t. for every quantum circuit C with ℓ' output qubits and an encryption of an input for C , the output size of the evaluation algorithm is $\text{poly}(\lambda, \ell')$, where λ is the security parameter of the scheme.*

⁸The trapdoor generation procedure presented in [MP12] actually results in an “inhomogeneous” trapdoor, where it holds that $\tau_{\mathbf{A}} \cdot \mathbf{A} = \mathbf{G}$, for the gadget matrix \mathbf{G} . However, one can derive a trapdoor satisfying $\tau_{\mathbf{A}} \cdot \mathbf{A} = 0$ from an inhomogeneous trapdoor.

3. **Classicality-Preserving Quantum Homomorphism:** Let $C = \{C_\lambda\}_{\lambda \in \mathbb{N}}$ be a polynomial-size quantum circuit, where C_λ has $\ell_1(\lambda) + \dots + \ell_n(\lambda)$ input qubits and $\ell'(\lambda)$ output qubits, of which $m(\lambda)$ are measured. Let $|\phi_1\rangle, \dots, |\phi_n\rangle = \{|\phi_i\rangle_\lambda, \dots, |\phi_n\rangle_\lambda\}_{\lambda \in \mathbb{N}}$ be an input state for C , let $(\text{pk}_1, \text{sk}_1), \dots, (\text{pk}_n, \text{sk}_n) = \{(\text{pk}_1, \text{sk}_1)_\lambda, \dots, (\text{pk}_n, \text{sk}_n)_\lambda\}_{\lambda \in \mathbb{N}}$ be pairs of public and secret keys ($\forall i \in [n], \lambda \in \mathbb{N}, (\text{pk}_i, \text{sk}_i)_\lambda \in \text{QMFHE.KeyGen}(1^\lambda)$) and let $r_1, \dots, r_n = \{(r_1)_\lambda, \dots, (r_n)_\lambda\}_{\lambda \in \mathbb{N}}$ be n random strings for the encryption algorithm. Then there exists a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\text{TD}(\rho_{0,\lambda}, \rho_{1,\lambda}) \leq \mu(\lambda),$$

where ρ_0, ρ_1 are quantum distributions defined as follows:

- $\rho_{0,\lambda}$: For each $i \in [n]$, encrypt each classical bit of $|\phi_i\rangle$ with $\text{QMFHE.Enc}(\text{pk}_i, \cdot)$ and the rest with $\text{QMFHE.QEnc}(\text{pk}_i, \cdot)$ (using randomness r_i). Execute $\text{QMFHE.Eval}((\text{pk}_1, \dots, \text{pk}_n), C, \cdot)$ on the n encryptions to get $\widehat{c}, |\widehat{\phi}\rangle$, where \widehat{c} is a classical ciphertext encrypting $m(\lambda)$ bits. Then output $\text{QMFHE.Dec}((\text{sk}_1, \dots, \text{sk}_n), \widehat{c}), \text{QMFHE.QDec}((\text{sk}_1, \dots, \text{sk}_n), |\widehat{\phi}\rangle)$.
- $\rho_{1,\lambda}$: Output $C(|\phi_1\rangle, \dots, |\phi_n\rangle)$.

Known classical LWE-based constructions of multi-key fully-homomorphic encryption [CM15, MW16, PS16, BHP17] do not quite satisfy the above syntax.⁹ Instead, they relax the syntax to allow for some notion of setup. In this work, we will be interested in the notion of *distributed setup* which was achieved in the classical setting by [BHP17].

Definition 4.2 (QMFHE with Distributed Setup). A QMFHE scheme QMFHE has distributed setup if it includes the following algorithm.

- $\text{QMFHE.Setup}(1^\lambda, 1^n, i)$: A PPT algorithm that takes as input the security parameter, a number of parties, and an index $i \in [n]$, and outputs a string pp_i .

We then define the public parameters of the scheme $\text{pp} = (\text{pp}_1, \dots, \text{pp}_n)$ and assume that all other algorithms take pp as input.

Remark 4.1. This notion of distributed setup gives rise to a stronger notion of semantic security, which considers *rushing* adversaries that may generate $\{\text{pp}_j\}_{j \neq i}$ maliciously, possibly depending on pp_i . More formally, in the security game the adversary first picks n and an $i \in [n]$ and sends these to its challenger. The challenger then runs $\text{pp}_i \leftarrow \text{QMFHE.Setup}(1^\lambda, 1^n, i)$, and returns pp_i to the adversary. Then, the adversary generates $\{\text{pp}_j\}_{j \in [n] \setminus \{i\}}$ arbitrarily and sends these to its challenger. Finally, the challenger draws a public key secret key pair based on $\{\text{pp}_i\}_{i \in [n]}$, and the semantic security game continues as usual. This notion of semantic security was achieved in the classical setting by [BHP17].

In this work, we also consider a more stringent requirement on the operation of the QMFHE.Dec algorithm, which we call *nearly linear decryption of classical ciphertexts*. Essentially, this states that decrypting a classical ciphertext c encrypted under keys $\text{sk}_1, \dots, \text{sk}_n$ amounts to computing a linear function \mathcal{L}_c (defined by c) on the concatenated secret keys $[\text{sk}_1 \mid \dots \mid \text{sk}_n]$ modulo some integer q , and then rounding.

Definition 4.3 (QMFHE with Nearly Linear Decryption of Classical Ciphertexts). A QMFHE scheme QMFHE has nearly linear decryption of classical ciphertexts if the QMFHE.Dec algorithm operates as follows.

- $\text{QMFHE.Dec}((\text{sk}_1, \dots, \text{sk}_n), c)$: There is an efficiently computable linear function \mathcal{L}_c (determined by c) and an (even) integer q such that the decryption procedure computes

$$\mathcal{L}_c(\text{sk}_1, \dots, \text{sk}_n) = b \cdot q/2 + e \pmod q$$

(where $e < q/4$) and returns $b \in \{0, 1\}$. Equivalently, we can define linear functions $\mathcal{L}_c^{(1)}, \dots, \mathcal{L}_c^{(n)}$ such that the decryption procedure computes

$$\sum_{i \in [n]} \mathcal{L}_c^{(i)}(\text{sk}_i) = b \cdot q/2 + e \pmod q.$$

⁹Though there are NTRU-based constructions that do [LTV12, AJM20].

Finally, we remark that we do not consider an additional security property often found in classical constructions of multi-key FHE, which roughly stipulates that partial decryptions of other parties may be simulated. This property is most relevant when considering the direct application of multi-key FHE to MPC, but we will not need it in this work.

4.3 Background

We follow the template given by [Mah18] for constructing a quantum fully-homomorphic encryption scheme. [Mah18] essentially showed that a classical fully-homomorphic scheme FHE can be converted into a quantum fully-homomorphic scheme, as long as it has a few additional properties, necessary for computing the so-called *encrypted CNOT* operation (properties 2-4 in Definition 4.4).

Unfortunately, no known fully-homomorphic schemes immediately satisfy these properties. However, [Mah18] observed that the dual Regev (non-fully-homomorphic) encryption scheme of [GPV08] does satisfy these properties, and moreover, and that there exists a fully-homomorphic encryption scheme FHE (the dual version of [GSW13]) with an efficient procedure for converting an FHE ciphertext into a dual Regev ciphertext encrypting the same message. In fact, the conversion procedure presented in [Mah18] simply consists of taking the last column of the FHE ciphertext. This suffices to give a quantum fully-homomorphic encryption scheme, since before every encrypted CNOT operation, the evaluator can convert any FHE ciphertext needed during the operation into a dual Regev ciphertext, and then proceed. The fourth property below is needed in order to convert an evaluated dual Regev ciphertext back into an FHE ciphertext upon completion of the encrypted CNOT operation.¹⁰

This motivated the definition of a quantum-capable fully-homomorphic encryption scheme given in [Mah18]. Such a scheme admits an efficient procedure that converts ciphertexts into ciphertexts of an alternate encryption scheme AHE that satisfies the properties necessary to carry out the encrypted CNOT operation. [Mah18] showed that any such FHE scheme gives rise to an FHE scheme that can additionally encrypt quantum states and evaluate quantum circuits. Below we give the analogous definition for *multi-key* fully-homomorphic encryption, and follow with a sketch of the analogous conversion from quantum-capable multi-key fully-homomorphic encryption to quantum multi-key fully-homomorphic encryption.

Definition 4.4 (Quantum-Capable Multi-Key Fully-Homomorphic Encryption Scheme). *Let MFHE be a classical multi-key fully-homomorphic encryption scheme. MFHE is quantum-capable if i) its KeyGen procedure outputs a public key pk , secret key sk , and “trapdoor” τ , and ii) there exists an alternate encryption scheme AHE such that the following properties holds.*

1. *There exists an algorithm MFHE.Convert that takes as input a set of public keys $\text{pk}_1, \dots, \text{pk}_t$ and a ciphertext c encrypted under $\text{pk}_1, \dots, \text{pk}_t$, and outputs an encryption \hat{c} under AHE with public keys $\text{pk}_1, \dots, \text{pk}_t$, where c and \hat{c} encrypt the same value.*
2. *There exists an invertible operation \oplus_H (which may depend on $\text{pk}_1, \dots, \text{pk}_t$) on AHE ciphertexts such that, for all $x_0, x_1, a \in \{0, 1\}$, $\text{AHE.Enc}((\text{pk}_1, \dots, \text{pk}_t), x_0) \oplus_H a \cdot \text{AHE.Enc}((\text{pk}_1, \dots, \text{pk}_t), x_1)$ is an AHE encryption of $x_0 \oplus a \cdot x_1$ under $\text{pk}_1, \dots, \text{pk}_t$.*
3. *There exists a distribution \mathcal{D} (which may depend on $\text{pk}_1, \dots, \text{pk}_t$) such that for all ciphertexts c that can arise during homomorphic evaluation,¹¹*

$$\begin{aligned} & \{ \text{AHE.Enc}((\text{pk}_1, \dots, \text{pk}_t), x; r) \mid (x, r) \leftarrow \mathcal{D} \} \\ & \approx_s \{ \text{AHE.Enc}((\text{pk}_1, \dots, \text{pk}_t), x; r) \oplus_H c \mid (x, r) \leftarrow \mathcal{D} \}, \end{aligned}$$

and there is an efficient procedure for generating the superposition $\sum_{x,r} \sqrt{\mathcal{D}(x,r)} |x, r\rangle$.

¹⁰For technical reasons, the *randomness* in the dual Regev ciphertext must also be recovered, which motivates the need for a trapdoor rather than merely a secret key.

¹¹This set will consist of all ciphertexts with noise below some fixed bound.

4. There exists an efficient function f such that for any $c = \text{AHE.Enc}((pk_1, \dots, pk_t), x; r)$, $f((\tau_1, \dots, \tau_t), c) = (x, r)$.

From quantum-capability to multi-key quantum FHE. We now sketch, following [Mah18]’s approach in the single-key setting, how a quantum-capable multi-key FHE scheme gives rise to a full-fledged quantum multi-key FHE scheme. The following description makes use of the quantum one-time pad (QOTP), which is a method of *perfectly* encrypting arbitrary quantum states $|\phi\rangle$ using classical bits k . We refer the reader to [Mah18] for details about how the QOTP is constructed and proven secure. We do not present details about how individual quantum gates are evaluated, or the inner workings of the encrypted CNOT operation, electing instead to present a high-level picture. Again we refer the reader to [Mah18] for all of these details. The following assumes a quantum-capable multi-key FHE scheme QCMFHE, and describes a quantum multi-key FHE scheme QMFHE.

- **Key generation.** This procedure generates $(pk', sk', \tau) \leftarrow \text{QCMFHE.KeyGen}(1^\lambda)$, computes a ciphertext $ct^{(\tau)} \leftarrow \text{QCMFHE.Enc}(pk', \tau)$, and sets the public key of QMFHE to $pk := (pk', ct^{(\tau)})$ and the secret key to $sk := sk'$.
- **Encryption.** To encrypt a quantum state $|\phi\rangle$, sample a random QOTP key k and release the ciphertext $ct = (\text{QCMFHE.Enc}(pk', k), \text{QOTP}(k, |\phi\rangle))$. To encrypt a classical string m , sample a classical one-time pad key k and release $ct = (\text{QCMFHE.Enc}(pk', k), k \oplus m)$.
- **Homomorphic evaluation.** This operation takes as input t public keys $\{pk_i = (pk'_i, ct_i^{(\tau)})\}_{i \in [t]}$ and t ciphertexts $\{ct_i = (c_i, |c\rangle_i)\}_{i \in [t]}$. It first expands each c_i into a multi-key ciphertext c'_i encrypted under all public keys pk'_1, \dots, pk'_t , and gathers all components into a quantum multi-key ciphertext $(\hat{c}, |\hat{c}\rangle) = ((c'_1, \dots, c'_t), (|c\rangle_1, \dots, |c\rangle_t))$. It also expands and concatenates the ciphertexts $\{ct_i^{(\tau)}\}_{i \in [t]}$ to produce an evaluation key $\hat{ct}^{(\tau)}$ that encrypts the trapdoors (τ_1, \dots, τ_t) under all public keys pk'_1, \dots, pk'_t . Next, it applies a quantum circuit gate by gate on the ciphertext, as follows.
 - If the gate is of a particular type, namely, it is a Clifford operator, then homomorphically evaluating the gate can be done via a *parallel* procedure, where a classical circuit is applied homomorphically over \hat{c} to produce \hat{c}' , and a quantum circuit is applied *directly* to $|c\rangle$ to produce $|\hat{c}'\rangle$.
 - Any universal gate set for quantum computation must contain at least one non-Clifford operator, and [Mah18] includes the Toffoli gate. [Mah18] gives a procedure for homomorphically applying the Toffoli gate that involves parallel operations as above along with an encrypted CNOT operation, which requires the following manipulation. First, QCMFHE.Convert is run on \hat{c} to produce an AHE ciphertext \hat{d} . This ciphertext is used to define a quantum circuit that is applied to $|c\rangle$ to produce $|c'\rangle$ along with a (measured) AHE ciphertext \hat{d}' . Finally, the function f (defined in property 4 of Definition 4.4), with ciphertext \hat{d}' hard-coded, is applied homomorphically over $\hat{ct}^{(\tau)}$ to produce a classical ciphertext c' encrypting the message and randomness from \hat{d}' .
- **Decryption.** Given a ciphertext $(\hat{c}, |\hat{c}\rangle)$ encrypting a qubit under pk_1, \dots, pk_t , and corresponding QCMFHE secret keys sk'_1, \dots, sk'_t , this operation runs $\text{QCMFHE.Dec}((sk'_1, \dots, sk'_t), \hat{c})$ to produce a key k , and then uses k to decrypt the one-time padded state $|\hat{c}\rangle$. The same procedure works if $|\hat{c}\rangle$ was instead a classical string $k \oplus m$.

4.4 Construction

Existing classical multi-key fully-homomorphic encryption schemes [LTV12, CM15, MW16, PS16, BHP17, AJJM20] in the literature do not appear to admit a simple conversion procedure necessary for quantum-capability, such as the one enjoyed by dual-GSW in the single-key setting. However, we show that indeed there exists a general conversion procedure that works for *any* multi-key fully-homomorphic encryption scheme

with *nearly linear decryption* (see Definition 4.3). This method is essentially key-switching (see for example [BV11, BDGM19]), and relies on the existence of a multi-key *linearly*-homomorphic encryption scheme. This multi-key linearly homomorphic scheme is an extension of dual Regev encryption, and is implicit in the construction that follows.

Let MFHE be a classical multi-key fully-homomorphic encryption scheme with nearly linear decryption. Consider the following scheme QCMFHE, which is identical to MFHE except that it has a different KeyGen algorithm and it additionally supports a Convert algorithm. Let q be an even k -bit modulus, let $\mathbf{g} = (1, 2, \dots, 2^k)$, and for $y \in \mathbb{Z}_q$, let $\mathbf{g}^{-1}(y) \in \{0, 1\}^k$ be the binary expansion of y , i.e. it holds that $\mathbf{g}^\top \cdot \mathbf{g}^{-1}(y) = y$.

- QCMFHE.KeyGen(1^λ):
 1. Compute $(\text{MFHE.pk}, \text{MFHE.sk}) \leftarrow \text{MFHE.KeyGen}(1^\lambda)$, where $\text{MFHE.sk} \in \mathbb{Z}_q^\ell$ and $\ell = \text{poly}(\lambda)$.
 2. Let n, m be positive integers and χ be a B -bounded error distribution, where $n, m, B = \text{poly}(\lambda)$.
 3. Draw $(\mathbf{B}, \tau) \leftarrow \text{GenTrap}(1^n, 1^m, q)$, $\mathbf{b} \leftarrow \mathbb{Z}_q^n$, and set $\mathbf{A} = \begin{pmatrix} \mathbf{B} \\ \mathbf{b}^\top \end{pmatrix}$.
 4. Parse $\text{MFHE.sk} \in \mathbb{Z}_q^\ell$ as $\mu_1, \dots, \mu_\ell \in \mathbb{Z}_q$, and for each $i \in [\ell]$, compute the following.
 - (a) Draw $\mathbf{S} \leftarrow \mathbb{Z}_q^{n \times k}$ and $\mathbf{E} \leftarrow \chi^{(m+1) \times k}$.
 - (b) Set $\mathbf{C}_i := \mathbf{A} \cdot \mathbf{S} + \mathbf{E} + \mu_i \cdot \mathbf{g}^\top \cdot \mathbf{u}_{m+1}$, where \mathbf{u}_{m+1} is the $(m+1)$ -dimensional vector with all 0s except the final coordinate is 1.
 5. Output $\text{pk} := (\text{MFHE.pk}, \mathbf{C}_1, \dots, \mathbf{C}_\ell)$, $\text{sk} := \text{MFHE.sk}$, and τ .
- QCMFHE.Convert($(\text{pk}_1, \dots, \text{pk}_t), c$): Let the linear function \mathcal{L}_c determined by c consist of coefficients $a_{1,1}, \dots, a_{1,\ell}, \dots, a_{t,1}, \dots, a_{t,\ell}$. Parse each pk_i to obtain $\mathbf{C}_{i,1}, \dots, \mathbf{C}_{i,\ell}$ and define $\hat{\mathbf{C}}_{i,j}$ as follows. Let $\bar{\mathbf{C}}_{i,j}$ be the first m rows of $\mathbf{C}_{i,j}$ and $\mathbf{c}_{i,j}$ be the last row. Then $\hat{\mathbf{C}}_{i,j} \in \mathbb{Z}_q^{(m\ell+1) \times k}$ is the matrix with 0s everywhere except that the $(i-1)m+1, \dots, im$ rows are set to $\bar{\mathbf{C}}_{i,j}$ and the last row is set to $\mathbf{c}_{i,j}$.
Output

$$\sum_{i \in [t], j \in [\ell]} \hat{\mathbf{C}}_{i,j} \cdot \mathbf{g}^{-1}(a_{i,j}).$$

We assume that the parameters of MFHE are instantiated in a particular way, namely, the modulus q is set such that for any well-formed ciphertext c (encrypted under a set of t public keys) that may arise during homomorphic evaluation, $\mathcal{L}_c(\text{sk}_1, \dots, \text{sk}_t) = q/2 + e \pmod q$, where $q \geq \omega(\text{poly}(\lambda)) \cdot |e|$. Recall that the linear function \mathcal{L}_c is guaranteed to exist by the nearly linear decryption property.

Theorem 4.2. *Assuming the existence of a multi-key fully-homomorphic encryption scheme MFHE with nearly linear decryption and a particular circular security property, there exists a quantum multi-key fully-homomorphic encryption scheme QMFHE. Moreover, QMFHE satisfies the following properties.*

1. The QMFHE.Setup algorithm is equivalent to MFHE.Setup.
2. If MFHE is perfectly correct, then QMFHE satisfies Classicality-Preserving Quantum Homomorphism.

Proof. QMFHE is obtained by first applying the construction described in Section 4.4 to MFHE to obtain QCMFHE, followed by the construction sketched in Section 4.3.

First, we argue that QCMFHE is indeed quantum-capable. Consider the output of the QCMFHE.Convert algorithm. It is straightforward to verify that if c is a well-formed encryption under public keys $\text{pk}_1, \dots, \text{pk}_t$ of the bit μ , where each pk_i may be parsed as $\begin{pmatrix} \mathbf{B}_i \\ \mathbf{b}_i^\top \end{pmatrix}$, then the resulting vector may be written as

$$\begin{pmatrix} \mathbf{B}_1 & & & \\ & \ddots & & \\ & & \mathbf{B}_t & \\ \mathbf{b}_1^\top & \dots & \mathbf{b}_t^\top & \end{pmatrix} \cdot \mathbf{s}^* + \mathbf{e}^* + \frac{q}{2} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \mu \end{pmatrix},$$

for some $\mathbf{s}^* \in \mathbb{Z}_q^{tn}$, $\mathbf{e}^* \in \mathbb{Z}^{(m+1)\ell}$. This is exactly an encryption of μ under the dual Regev scheme with public key

$$\begin{pmatrix} \mathbf{B}_1 & & & \\ & \ddots & & \\ & & \mathbf{B}_t & \\ \mathbf{b}_1^\top & \dots & \mathbf{b}_t^\top & \end{pmatrix}.$$

Thus the AHE scheme we use in Definition 4.4 is identical to the scheme used in [Mah18]. This shows that QCMFHE satisfies the first requirement in Definition 4.4, and the fact that it satisfies also the second requirement is immediate.

To confirm that QCMFHE satisfies the third requirement, we take a closer look at \mathbf{e}^* . The distribution \mathcal{D} used by [Mah18] samples μ and \mathbf{s} uniformly at random, and \mathbf{e} from a discrete Gaussian distribution with “large enough” parameter B' . This requirement will hold if B' is super-polynomially larger than the entries of \mathbf{e}^* (see Lemma 3.3 and Section 5.3 of [Mah18] for more details). Note that the modulus q is super-polynomially larger than each entry of \mathbf{e}^* , by the assumption on parameters of MFHE. Indeed, all but the last entry are bounded by $\ell \cdot k \cdot B = \text{poly}(\lambda)$, and the last entry is bounded by $t \cdot \ell \cdot k \cdot B$ plus the error that results from the nearly linear decryption, which is super-polynomially smaller than q . This allows us to define B' to be large enough such that the third requirement will hold.

To confirm that QCMFHE satisfies the fourth requirement, note that \mathbf{s}^* may be written as a concatenation of t n -dimensional vectors, and that the i 'th such vector may be recovered by using τ_i , by Lemma 4.1. This process also recovers all but the last entry of \mathbf{e}^* . The last entry of \mathbf{e}^* may then be recovered by subtracting the public key times \mathbf{s}^* and rounding the last element of the resulting vector.

The above shows that QCMFHE is quantum-capable according to Definition 4.4. Next, we discuss security of the scheme QMFHE obtained by applying the construction sketched in Section 4.3. Observe that QMFHE.KeyGen outputs a public key that contains a MFHE public key, a dual Regev public key, an encryption of the MFHE secret key under the dual Regev public key, and an encryption of the dual Regev *trapdoor* under the MFHE secret key. Since QMFHE encryption involves encrypting a QOTP key under MFHE and using that key to *perfectly* hide the message, it follows that security of QMFHE reduces to the security of MFHE in the presence of the particular two-cycle of keys described above (which at the very least relies on LWE to ensure security of dual Regev). Thus, as stated in the theorem, security follows from a particular circular security property of MFHE.¹²

It remains to argue that the two extra properties promised by the theorem statement hold. First, note that the constructions given in Section 4.4 and Section 4.3 do not alter any MFHE.Setup algorithm that may exist. Next, the second property boils down to showing that for *every* choice of random coins used in QMFHE.KeyGen, homomorphic evaluation of quantum (or classical) circuits will be statistically correct. Perfect correctness of any Clifford operation follows directly from perfect correctness of MFHE. Statistical correctness of the encrypted CNOT operation follows from properties 2 and 3 of Definition 4.4 (this analysis can be found in [Mah18]), plus the *perfect* correctness of property 4, which is ensured by using the variant of the GenTrap algorithm promised by Lemma 4.1. \square

4.5 Quantum Spooky Encryption

We define the notion of spooky encryption for (classical) relations computable by quantum circuits, generalizing the purely classical notion from [DHRW16]. In favor of a simpler exposition we present the additive function sharing (AFS) variant of the notion, but we note that considering more general relations (in the same spirit as [DHRW16]) is also possible.

Definition 4.5 (Quantum AFS-Spooky Encryption). *A quantum AFS-spooky encryption scheme is given by six algorithms (Spooky.KeyGen, Spooky.Enc, Spooky.QEnc, Spooky.Eval, Spooky.Dec, Spooky.QDec) with the same syntax as the corresponding QMFHE algorithms defined in Definition 4.1, except for the following differences.*

¹²This property is similar to the one needed by [Mah18] in the single-key setting, in the sense that encryption of a dual Regev trapdoor is part of the circular security requirement.

- $b \leftarrow \text{QMFHE.Dec}(\text{sk}, c)$: A PPT algorithm that takes as input a secret key and a classical ciphertext c and outputs a bit. (This algorithm takes only one secret key, as opposed to n secret keys in QMFHE.)
- $\widehat{c}_1, \dots, \widehat{c}_n, |\widehat{\phi}\rangle \leftarrow \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_n), C, (|\phi_1\rangle, \dots, |\phi_n\rangle))$: A QPT algorithm that takes as input
 1. A set of n public keys.
 2. A general quantum circuit with $\ell_1 + \dots + \ell_n$ input qubits and ℓ' output qubits, out of which m are measured.
 3. A set of n ciphertexts where $|\phi_i\rangle$ encrypts an ℓ_i -qubit state under pk_i . Some of the ℓ_i ciphertexts are possibly classical ciphertexts (generated by the classical encryption algorithm) encrypting classical bits.

The evaluation algorithm outputs n classical ciphertexts $(\widehat{c}_1, \dots, \widehat{c}_n)$ each encrypting m bits under the corresponding pk_i , plus a quantum ciphertext $|\widehat{\phi}\rangle$ encrypting an $(\ell' - m)$ -qubit quantum state (under keys $\text{pk}_1, \dots, \text{pk}_n$). (This algorithm outputs n classical ciphertexts, as opposed to one in QMFHE.)

The scheme satisfies the same properties of quantum semantic security and compactness as defined in Definition 4.1. In the following we present the notion of correctness for quantum AFS-spooky encryption.

- **Correctness of Spooky Evaluation:** Let $C = \{C_\lambda\}_{\lambda \in \mathbb{N}}$ be a polynomial-size quantum circuit, where C_λ has $\ell_1(\lambda) + \dots + \ell_n(\lambda)$ input qubits and $\ell'(\lambda)$ output qubits, of which $m(\lambda)$ are measured. Let $|\phi_1\rangle, \dots, |\phi_n\rangle = \{|\phi_1\rangle_\lambda, \dots, |\phi_n\rangle_\lambda\}_{\lambda \in \mathbb{N}}$ be an input state for C , let $(\text{pk}_1, \text{sk}_1), \dots, (\text{pk}_n, \text{sk}_n) = \{(\text{pk}_1, \text{sk}_1)_\lambda, \dots, (\text{pk}_n, \text{sk}_n)_\lambda\}_{\lambda \in \mathbb{N}}$ be pairs of public and secret keys ($\forall i \in [n], \lambda \in \mathbb{N}, (\text{pk}_i, \text{sk}_i)_\lambda \in \text{Spooky.KeyGen}(1^\lambda)$) and let $r_1, \dots, r_n = \{(r_1)_\lambda, \dots, (r_n)_\lambda\}_{\lambda \in \mathbb{N}}$ be n random strings for the encryption algorithm. Then there exists a negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

$$\text{TD}(\rho_{0,\lambda}, \rho_{1,\lambda}) \leq \mu(\lambda),$$

where ρ_0, ρ_1 are quantum distributions defined as follows:

- $\rho_{0,\lambda}$: For each $i \in [n]$, encrypt each classical bit of $|\phi_i\rangle$ with $\text{Spooky.Enc}(\text{pk}_i, \cdot)$ and the rest with $\text{Spooky.QEnc}(\text{pk}_i, \cdot)$ (using randomness r_i). Execute $\text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_n), C, \cdot)$ on the n encryptions to get $(\widehat{c}_1, \dots, \widehat{c}_n), |\widehat{\phi}\rangle$, where $(\widehat{c}_1, \dots, \widehat{c}_n)$ are classical ciphertexts each encrypting $m(\lambda)$ bits. Then output

$$\bigoplus_{i=1}^n \text{Spooky.Dec}(\text{sk}_i, \widehat{c}_i), \text{Spooky.QDec}((\text{sk}_1, \dots, \text{sk}_n), |\widehat{\phi}\rangle).$$

- $\rho_{1,\lambda}$: Output $C(|\phi_1, \dots, \phi_n\rangle)$.

(Classical) AFS-Spooky Encryption with Distributed Setup. As a stepping stone towards the main result of this section, we show how to construct spooky encryption for *classical* relations with a distributed setup. More precisely, assuming the hardness of the LWE problem, we show an instantiation of spooky encryption for any polynomial-size (classical) circuit where the parties jointly compute the public parameters of the system via a local algorithm Spooky.Setup (with the same syntax as Definition 4.2). This stands in contrast with the scheme of [DHRW16], where the common reference string is assumed to be sampled by a trusted party. Before describing the construction, we recall a useful lemma (rephrased) from [DHRW16].

Lemma 4.2 ([DHRW16]). *Let Spooky be an AFS-spooky encryption scheme that supports (i) single key additive homomorphism and (ii) two-key spooky multiplication. Then the same scheme supports the AFS-spooky evaluation of all polynomial-size (classical) circuits.*

It follows that it suffices to construct a spooky encryption that supports a single multiplication over an arbitrary pair of keys. We do this by showing that the scheme from [BHP17] supports two-key spooky multiplication. This follows from the fact that the decryption circuit is identical to that of [CM15, MW16], which was shown to support two-key spooky multiplication in [DHRW16]. For completeness, we recall the modified algorithms in the following.

- Spooky.Enc(pk, m): Same as MFHE.Enc but append an extra $\delta = 0$ to the resulting ciphertext.
- Spooky.Dec(sk, c): Let \mathcal{L}_c be the linear function defined by c , compute

$$v = \mathcal{L}_c(\text{sk}) + \delta \pmod{q}$$

and return 0 if $|v| < q/4$ and 1 otherwise.

- Spooky.Eval((pk₁, pk₂), C, (c₁, c₂)): Compute

$$\hat{c} \leftarrow \text{MFHE.Eval}\left(\left(\text{pk}_1, \text{pk}_2\right), \prod, (c_1, c_2)\right)$$

and let $\mathcal{L}_{\hat{c}}^{(1)}$ and $\mathcal{L}_{\hat{c}}^{(2)}$ be the linear functions defined by the resulting \hat{c} . Sample a uniform δ from \mathbb{Z}_q and return $(\mathcal{L}_{\hat{c}}^{(1)}, \delta)$ and $(\mathcal{L}_{\hat{c}}^{(2)}, -\delta)$.

As discussed above, the scheme is quantum semantically secure assuming the hardness of the LWE problem. Correctness follows, for the same choice of parameters of [BHP17], by an invocation of the following lemma.

Lemma 4.3 ([DHRW16]). *Fix a modulus $q \in \mathbb{Z}$, a bit $b \in \{0, 1\}$ and a value $v \in \mathbb{Z}_q$ such that $v = q/2 \cdot b + e \pmod{q}$, for some $|e| < q/4$. Sample v_1 and v_2 uniformly at random from \mathbb{Z}_q constrained on the fact that $v_1 + v_2 = v \pmod{q}$, and let $b_i = 0$ if $|v_i| < q/4$ and $b_i = 1$ otherwise. Then*

$$\Pr[b_1 \oplus b_2 = b] > 1 - 2(|e| + 1)/q$$

over the random choice of v_1 and v_2 .

Quantum AFS-Spooky Encryption with Distributed Setup. Finally, we show how to combine a classical AFS-spooky encryption scheme Spooky (with distributed setup) with a quantum multi-key fully-homomorphic encryption QMFHE (with distributed setup) to obtain a quantum AFS-spooky encryption scheme QSpooky. Since both of the building blocks have a distributed setup, then so does the resulting encryption scheme. The scheme is described below.

- QSpooky.Setup(1^λ): Compute $\text{pp} \leftarrow \text{QMFHE.Setup}(1^\lambda)$ and $\tilde{\text{pp}} \leftarrow \text{Spooky.Setup}(1^\lambda)$ and return $(\text{pp}, \tilde{\text{pp}})$.
- QSpooky.KeyGen($1^\lambda, \text{pp}$): Sample

$$(\text{pk}', \text{sk}') \leftarrow \text{QMFHE.KeyGen}(1^\lambda, \text{pp}) \text{ and } (\tilde{\text{pk}}, \tilde{\text{sk}}) \leftarrow \text{Spooky.KeyGen}(1^\lambda, \tilde{\text{pp}})$$

and compute $\tilde{c} \leftarrow \text{Spooky.Enc}(\tilde{\text{pk}}, \text{sk}')$. Return $\text{pk} := (\text{pk}', \tilde{\text{pk}}, \tilde{c})$ as the public key and $\text{sk} := (\text{sk}', \tilde{\text{sk}})$ as the secret key.

- QSpooky.Enc(pk, m): Return $\text{QMFHE.Enc}(\text{pk}', m)$.
- QSpooky.QEnc(pk, $|\psi\rangle$): Return $\text{QMFHE.QEnc}(\text{pk}', |\psi\rangle)$.
- QSpooky.Eval((pk₁, ..., pk_n), C, (| ϕ_1 ⟩, ..., | ϕ_n ⟩)): Compute

$$(\hat{c}, |\hat{\phi}\rangle) \leftarrow \text{QMFHE.Eval}((\text{pk}'_1, \dots, \text{pk}'_n), C, (|\phi_1\rangle, \dots, |\phi_n\rangle))$$

and let $(\tilde{c}_1, \dots, \tilde{c}_n)$ be the corresponding element of each public key. Compute

$$(\hat{c}_1, \dots, \hat{c}_n) \leftarrow \text{Spooky.Eval}((\tilde{\text{pk}}_1, \dots, \tilde{\text{pk}}_n), \text{QMFHE.Dec}(\cdot, \hat{c}), (\tilde{c}_1, \dots, \tilde{c}_n))$$

and return $(\hat{c}_1, \dots, \hat{c}_n, |\hat{\phi}\rangle)$.

- $\text{QSpooky.Dec}(\text{sk}, c)$: Return $\text{Spooky.Dec}(\tilde{\text{sk}}, c)$.
- $\text{QSpooky.QDec}((\text{sk}_1, \dots, \text{sk}_n), |\phi\rangle)$: Return $\text{QMFHE.QDec}((\text{sk}'_1, \dots, \text{sk}'_n), |\phi\rangle)$.

The following theorem establishes our claim.

Theorem 4.3. *Assuming that QMFHE is a quantum multi-key fully-homomorphic encryption scheme and that Spooky is a classical AFS-spooky encryption, QSpooky is a quantum AFS-spooky encryption scheme.*¹³

Proof. Assuming quantum semantic security of Spooky, the changes in the key generation algorithm do not affect the security of the scheme. Then quantum semantic security follows from an invocation of the quantum semantic security of QMFHE. Correctness of spooky evaluation follows from the classicality-preserving homomorphism of QMFHE and from the correctness of classical spooky evaluation of Spooky. \square

5 Quantum-Secure Multi-Committer Extractable Commitment

In this section, we follow the outline presented in Section 2.2 to construct a commitment scheme that allows for simultaneous extraction from multiple parallel committers. The protocol is somewhat more involved than the high-level description given earlier, so we briefly highlight the differences.

First, the committer is instructed to (non-interactively) commit to its message and trapdoor at the very beginning of the protocol. We use these commitments to take advantage of non-uniformity in the reductions between hybrids in the extractability proof. In particular, hybrids that come before the step where the simulator goes “under the hood” of the FHE may still need access to the trapdoor and commitment, and this can be given to any reduction via non-uniform advice consisting of each committer’s first message and corresponding openings.

Next, the CDS described earlier is replaced with a function-hiding secure function evaluation (SFE) protocol. In order to rule out the malleability attack mentioned in Section 2.2, where a malicious receiver mauls the AFS-spooky encryption of the committer’s trapdoor into an SFE encryption of the trapdoor, we do the following. The first message sent by the receiver to each committer C_i will actually be a commitment to some key k_i of a generic secret-key encryption scheme. After C_i sends its AFS-spooky encryption ciphertext and compute and compare obfuscation, the receiver prepares and sends a secret-key encryption of an arbitrary message. Then, the receiver’s input to the SFE consists of the opening to its earlier commitment k_i , and the SFE checks if the secret-key encryption sent by the receiver is actually an encryption of the committer’s trapdoor under secret key k_i . If so, it returns the lock and otherwise it returns \perp . This setup ensures that a malicious receiver cannot maul the AFS-spooky encryption of the committer’s trapdoor, for the following reason. If it could, then a non-uniform reduction to the semantic security of AFS-spooky encryption may obtain the receiver’s committed k_i as advice and decrypt the receiver’s secret-key encryption to obtain the trapdoor. Of course, this assumes the receiver actually acted explainably in sending a valid commitment at the beginning of the protocol, and this is ensured by the opening check performed under the SFE. We note that this mechanism is somewhat different than what was presented in [BS20], as they directly build a zero-knowledge argument (i.e. without first constructing a stand-alone extractable commitment) and are able to take advantage of witness indistinguishability to enforce explainable behavior.

Compliant Distinguishers. Finally, we discuss the issue of *committer* explainability. Recall from the high-level overview that a simulator is able to extract from a committer by homomorphically evaluating its code on an AFS-spooky encryption ciphertext *generated by the committer*. Thus, if the committer acts arbitrarily maliciously and does not return a well-formed ciphertext, the extraction may completely fail. Again, [BS20] address this issue by only analyzing their commitment within the context of a larger zero-knowledge argument protocol, and having the verifier prove to the prover using a witness indistinguishable proof that it performed the commitment explainably.

¹³In fact, we also need the quantum spooky encryption scheme to be *multi-hop*, which follows if the classical AFS-spooky scheme is multi-hop (which is satisfied by [DHRW16]).

Thus, without adding zero-knowledge and performing [GK96]-style analysis to handle non-explainable and aborting committers, we will only obtain extractability against explainable committers. However, since we will be using this protocol inside larger protocols where participants are not assumed to be acting explainably, restricting the class of committers we consider in our definition is problematic. We instead consider arbitrary committers but restrict the class of *distinguishers* (who are supposed to decide whether they received the view of a committer interacting in the real protocol or the view of a committer interacting with the extractor) to those that always output 0 on input a non-explainable transcript. In other words, any advantage these distinguishers may have must be coming from their behavior on input explainable views. Even though checking whether a particular view is explainable or not is not efficient, it turns out that this definition lends itself quite nicely to composition, since one can use witness indistinguishability/zero-knowledge to construct provably compliant distinguishers between hybrids for the larger protocols.

For completeness, and because post-quantum multi-committer extractable commitments may be of independent interest, we also show in Appendix D how to add zero-knowledge within the extractable commitment protocol itself to obtain security against arbitrary committers.

5.1 Definition

Definition 5.1 (Quantum-Secure Multi-Committer Extractable Commitment). *A quantum-secure multi-committer extractable commitment scheme is a pair (C, R) of classical PPT interactive Turing machines. In the commit phase, R interacts with n copies $\{C_i\}_{i \in [n]}$ of C (who do not interact with each other) on common input 1^λ and 1^n , with each C_i additionally taking a private input $m_i \in \{0, 1\}^*$. This produces a transcript τ , which may be parsed as a set of n transcripts $\{\tau_i\}_{i \in [n]}$, one for each set of messages exchanged between R and C_i . In the decommitment phase, each C_i outputs m_i along with its random coins r_i , and R on input $(1^\lambda, \tau_i, m_i, r_i)$ either accepts or rejects. The scheme should satisfy the following properties.*

- **Perfect Correctness:** For any $\lambda, n \in \mathbb{N}, i \in [n]$,

$$\Pr[R(1^\lambda, \tau_i, m_i, r_i) = 1 \mid \{\tau_i\}_{i \in [n]} \leftarrow \langle R, C_1(m_1; r_1), \dots, C_n(m_n; r_n) \rangle(1^\lambda, 1^n)] = 1.$$

- **Perfect Binding:** For any $\lambda \in \mathbb{N}$ and string $\tau \in \{0, 1\}^*$, there does not exist (m, r) and (m', r') with $m \neq m'$ such that $R(1^\lambda, \tau, m, r) = R(1^\lambda, \tau, m', r') = 1$.
- **Quantum Computational Hiding:** For any non-uniform quantum polynomial-size receiver $R^* = \{R_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, any polynomial $\ell(\cdot)$, and any sequence of sets of strings $\{m_{\lambda,1}^{(0)}, \dots, m_{\lambda,n}^{(0)}\}_{\lambda, n \in \mathbb{N}}, \{m_{\lambda,1}^{(1)}, \dots, m_{\lambda,n}^{(1)}\}_{\lambda, n \in \mathbb{N}}$ where each $|m_{\lambda,i}^{(b)}| = \ell(\lambda)$,

$$\begin{aligned} & \{\text{VIEW}_{R_\lambda^*}(\langle R_\lambda^*(\rho_\lambda), C_1(m_{\lambda,1}^{(0)}), \dots, C_n(m_{\lambda,n}^{(0)}) \rangle(1^\lambda, 1^n))\}_{\lambda, n \in \mathbb{N}} \\ & \approx_c \{\text{VIEW}_{R_\lambda^*}(\langle R_\lambda^*(\rho_\lambda), C_1(m_{\lambda,1}^{(1)}), \dots, C_n(m_{\lambda,n}^{(1)}) \rangle(1^\lambda, 1^n))\}_{\lambda, n \in \mathbb{N}}. \end{aligned}$$

The extractability property will require the following two definitions. First, for any adversary $C^* = \{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset $I \subseteq [n]$ of n committers, any honest party messages $\{m_i\}_{i \notin I}$, and any security parameter $\lambda \in \mathbb{N}$, define $\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_i)\}_{i \notin I} \rangle(1^\lambda, 1^n))$ to consist of the following.

1. The view of C_λ^* on interaction with the honest receiver R and set $\{C_i(m_i)\}_{i \notin I}$ of honest parties; this view includes a set of transcripts $\{\tau_i\}_{i \in I}$ and a state st .
2. A set of strings $\{m_i\}_{i \in I}$, where each m_i is defined relative to τ_i as follows. If there exists m'_i, r_i such that $R(1^\lambda, \tau_i, m'_i, r_i) = 1$, then $m_i = m'_i$, otherwise, $m_i = \perp$.

Next, we consider distinguishers $D = \{D_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$ that take as input a sample $(\{\tau_i\}_{i \in I}, \text{st}, \{m_i\}_{i \in I})$ from the distribution just described. We say that D is compliant if whenever $\{\tau_i\}_{i \in I}$ is not an explainable transcript with respect to the set I , D outputs 0 with overwhelming probability (over the randomness of D).

- **Multi-Committer Extractability:** *There exists a quantum expected-polynomial-time extractor Ext such that for any compliant non-uniform polynomial-size quantum distinguisher $D = \{D_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$, such that for all adversaries $C^* = \{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset of n committers, namely, $\{C_i\}_{i \in I}$ for some set $I \subseteq [n]$, the following holds for all polynomial-size sequences of inputs $\{\{m_{i,\lambda}\}_{i \notin I}\}_{\lambda \in \mathbb{N}}$ and $\lambda \in \mathbb{N}$.*

$$\begin{aligned} & \left| \Pr[D_\lambda(\text{VIEW}_{C_\lambda^*}^{\text{msg}}((R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i,\lambda})\}_{i \notin I})(1^\lambda, 1^n)), \sigma_\lambda) = 1] \right. \\ & \left. - \Pr[D_\lambda(\text{Ext}(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda), \sigma_\lambda) = 1] \right| \leq \mu(\lambda). \end{aligned}$$

Remark 5.1. Observe that the above definition of quantum computational hiding does not consider potentially malicious committers that interact in the protocol to try to gain information about commitments made by other committers. This is without loss of generality, since all communication occurs between R and some C_i . In particular, no messages are sent between any C_i and C_j .

5.2 Construction

Ingredients: All of the following are assumed to be quantum-secure.

- A non-interactive perfectly-binding commitment Com .
- A secret-key encryption scheme (Enc, Dec) .¹⁴
- A compute-and-compare obfuscator Obf .
- A quantum AFS-spooky encryption scheme with distributed setup $(\text{Spooky.Setup}, \text{Spooky.KeyGen}, \text{Spooky.Enc}, \text{Spooky.QEnc}, \text{Spooky.Eval}, \text{Spooky.Dec}, \text{Spooky.QDec})$.
- A two-message function-hiding secure function evaluation scheme $(\text{SFE.Gen}, \text{SFE.Enc}, \text{SFE.Eval}, \text{SFE.Dec})$.

5.3 Hiding

Perfect correctness and perfect binding are immediate, so we move to quantum computational hiding.

Lemma 5.1. *Protocol 2 is quantum computational hiding.*

Proof. Fix any non-uniform quantum polynomial-size receiver $R^* = \{R_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, a polynomial $\ell(\cdot)$, and two sequences of sets $\{m_{\lambda,1}^{(0)}, \dots, m_{\lambda,n}^{(0)}\}_{\lambda,n \in \mathbb{N}}$, $\{m_{\lambda,1}^{(1)}, \dots, m_{\lambda,n}^{(1)}\}_{\lambda,n \in \mathbb{N}}$ where each $|m_{\lambda,i}^{(b)}| = \ell(\lambda)$. Consider the following sequence of hybrids for each $i \in [n]$, where each alters the view of R^* in its interaction with C_i . The lemma follows immediately once we show that for all $i \in [n]$, $\mathcal{H}_{i,0} \approx_c \mathcal{H}_{i,6}$.

- $\mathcal{H}_{i,0}$: $\{\text{VIEW}_{R_\lambda^*}(\langle R_\lambda^*, C_i(m_{\lambda,i}^{(0)}) \rangle)(1^\lambda, 1^n)\}_{\lambda \in \mathbb{N}}$.
- $\mathcal{H}_{i,1}$: Same as $\mathcal{H}_{i,0}$ except that in Step 1, $c_i^{(\text{msg})}$ and $c_i^{(\text{td})}$ are commitments to 0.
- $\mathcal{H}_{i,2}$: Same as $\mathcal{H}_{i,1}$ except that in Step 7, C_i computes SFE.Eval on the circuit C_\perp that always outputs \perp .
- $\mathcal{H}_{i,3}$: Same as $\mathcal{H}_{i,2}$ except that in Step 5, the compute-and-compare obfuscation is simulated: $\widetilde{\text{CC}} \leftarrow \text{Sim}^{\text{CC}}(1^{|\text{Spooky.Dec}(\text{sk}_i, \cdot)|}, 1^{|\text{sk}_i| + |\ell(\lambda)|}, 1^\lambda)$.
- $\mathcal{H}_{i,4}$: Same as $\mathcal{H}_{i,3}$ except that in Step 5, the compute-and-compare obfuscation is performed honestly with respect to message $m_{\lambda,i}^{(1)}$.

¹⁴We use the syntax that for key k , a ciphertext of message m is computed as $\text{ct} \leftarrow \text{Enc}(k, m)$ and decrypted as $m := \text{Dec}(k, \text{ct})$.

Protocol 2

Common input: $1^\lambda, 1^n$.

C_i 's additional input: A string m_i .

1. Each C_i computes $td_i \leftarrow U_\lambda$ and sends $c_i^{(\text{msg})} \leftarrow \text{Com}(1^\lambda, m_i), c_i^{(\text{td})} \leftarrow \text{Com}(1^\lambda, td_i)$ to R.
2. For each $i \in [n]$, R computes $k_i, r_i \leftarrow U_\lambda$ and sends $c_i^{(\text{key})} := \text{Com}(1^\lambda, k_i; r_i)$ to C_i .
3. Each C_i computes and sends $pp_i \leftarrow \text{Spooky.Setup}(1^\lambda)$ to R.
4. R defines $pp := \{pp_i\}_{i \in [n]}$, and sends pp to each C_i . Each C_i checks that the pp_i it received matches the pp_i it sent in Step 3, and if not, it aborts.
5. Each C_i computes
 - $lk_i \leftarrow U_\lambda$,
 - $(pk_i, sk_i) \leftarrow \text{Spooky.KeyGen}(1^\lambda, pp)$,
 - $ct_i \leftarrow \text{Spooky.Enc}(pk_i, td_i)$,
 - and $\widetilde{CC}_i \leftarrow \text{Obf}(\text{CC}[\text{Spooky.Dec}(sk_i, \cdot), lk_i, (sk_i, m_i)])$,
 and sends $(pk_i, ct_i, \widetilde{CC}_i)$ to R.
6. For each $i \in [n]$, R computes $ct_i^{(\text{td})} \leftarrow \text{Enc}(k_i, 0^\lambda)$, $dk_i \leftarrow \text{SFE.Gen}(1^\lambda)$, and $ct_i^{(\text{SFE})} \leftarrow \text{SFE.Enc}(dk_i, (k_i, r_i))$ and sends $(ct_i^{(\text{td})}, ct_i^{(\text{SFE})})$ to C_i .
7. Define the circuit $C[c_i^{(\text{key})}, ct_i^{(\text{td})}, td_i, lk_i](\cdot)$ to take as input (k_i, r_i) , check if $c_i^{(\text{key})}$ opens to k_i with opening r_i and if $td_i = \text{Dec}(k_i, c_i^{(\text{td})})$, and if so output lk_i , and otherwise output \perp . Each C_i computes and sends $\widehat{ct}_i^{(\text{SFE})} \leftarrow \text{SFE.Eval}(C[c_i^{(\text{key})}, ct_i^{(\text{td})}, td_i, lk_i], ct_i^{(\text{SFE})})$.

Figure 2: A constant-round quantum-secure multi-committer extractable commitment.

- $\mathcal{H}_{i,5}$: Same as $\mathcal{H}_{i,4}$ except that in Step 7, the SFE.Eval is performed honestly.
- $\mathcal{H}_{i,6}$: Same as $\mathcal{H}_{i,5}$ except that in Step 1, $c_i^{(\text{msg})}$ is a commitment to $m_{\lambda,i}^{(1)}$ and $c_i^{(\text{td})}$ is a commitment to td_i . Note that is this exactly $\{\text{VIEW}_{\mathcal{R}_\lambda^*}(\langle \mathcal{R}_\lambda^*, C_i(m_{\lambda,i}^{(1)}) \rangle)(1^\lambda, 1^n)\}_{\lambda \in \mathbb{N}}$.

Now we argue indistinguishability between each hybrid.

- $\mathcal{H}_{i,0} \approx_c \mathcal{H}_{i,1}$: This follows directly from the quantum computational hiding of Com .
- $\mathcal{H}_{i,1} \approx_s \mathcal{H}_{i,2}$: We consider two cases. First, conditioned on C_i aborting in Step 4, the hybrids are trivially indistinguishable. Next, conditioned on C_i not aborting in Step 4, we show below that with overwhelming probability (over the randomness of C_i and R), the circuit $C[c_i^{(\text{key})}, ct_i^{(\text{td})}, td_i, lk_i]$ is functionally equivalent to C_\perp . Given this, the indistinguishability of hybrids $\mathcal{H}_{i,1}$ and $\mathcal{H}_{i,2}$ follows directly from the circuit privacy of SFE.

Assuming that the circuits are not functionally equivalent with noticeable probability, we construct a non-uniform $A = \{A_\lambda, \rho_\lambda^A\}_{\lambda \in \mathbb{N}}$ that breaks the distributed-setup quantum semantic security of Spooky (see Definition 4.2). In the security game, A interacts with a challenger to generate $pp := \{pp_i\}_{i \in [n]}$ for

n parties. Then, the challenger draws a public key pk_i based on pp , a random $td_i \leftarrow U_\lambda$, and outputs an encryption ct_i of td_i under pk_i . A wins if it returns td_i , which would clearly break semantic security. Now, we describe the distribution ρ_λ^A that A receives as non-uniform advice (this distribution will ultimately be fixed to the advice state that gives A the best advantage). It will be generated as follows.

1. Run R_λ^* on ρ_λ , and feed to R_λ^* the first messages $\{c_i^{(msg)}, c_i^{(td)}\}_{i \in [n]}$ it expects from $\{C_i\}_{i \in [n]}$ (which are commitments to 0).
2. Continue running R_λ^* until it outputs its set of messages $\{c_i^{(key)}\}_{i \in [n]}$.
3. Output the inner state of R_λ^* , the messages exchanged so far, and the following. For $c_i^{(key)}$, check (inefficiently) if it is a commitment to some k_i and if so, output k_i .

Finally, we describe A. A_λ receives from its challenger the i 'th public parameters pp_i . It then runs R_λ^* on the state it received as advice and pp_i . R_λ^* returns $pp = \{pp_j\}_{j \in [n]}$, where by assumption pp includes the same pp_i that it took as input. A_λ then forwards $\{pp_j\}_{j \in [n] \setminus \{i\}}$ to its challenger, who returns with a public key pk_i and a ciphertext ct_i . At this point, A_λ generates $\{(pk_j, ct_j, \widetilde{CC}_j)\}_{j \in [n] \setminus \{i\}}$ honestly and for party i , fixes (pk_i, ct_i) along with $\widetilde{CC}_i \leftarrow \text{Sim}^{\text{CC}}(1^{|\text{Spooky.Dec}(sk_i, \cdot)|}, 1^{|\text{sk}_i| + |\ell(\lambda)|}, 1^\lambda)$. It then continues to run R_λ^* on input all of these tuples.

When R_λ^* returns $ct_i^{(td)}$, A_λ checks if it received some k_i as part of its non-uniform advice, and if so, it decrypts $ct_i^{(td)}$ using key k_i to recover a message td_i . It returns td_i to the challenger, who then determines if A_λ succeeded.

Note that, by the simulation security of compute-and-compare obfuscation, the probability that A_λ succeeds in this game is negligibly close to the probability it succeeds if it gave R_λ^* an honest compute-and-compare obfuscation \widetilde{CC}_i . This follows because the lock value lk_i is completely independent of R_λ^* 's view through Step 6. Finally, the probability that A_λ succeeds in returning td is at least the probability that $c_i^{(key)}$ is a well-formed commitment to k_i and $ct_i^{(td)}$ is an encryption of td under key k_i , which is exactly the probability that the circuits described above are *not* functionally equivalent. Thus A_λ has non-negligible advantage in this game, a contradiction.

- $\mathcal{H}_{i,2} \approx_c \mathcal{H}_{i,3}$: This follows directly from the simulation security of compute-and-compare obfuscation, since at this point, the lock value lk_i is independent of the rest of the distribution.
- $\mathcal{H}_{i,3} \approx_c \mathcal{H}_{i,4}$: Same argument as $\mathcal{H}_{i,2} \approx_c \mathcal{H}_{i,3}$.
- $\mathcal{H}_{i,4} \approx_c \mathcal{H}_{i,5}$: Same argument as $\mathcal{H}_{i,1} \approx_c \mathcal{H}_{i,2}$.
- $\mathcal{H}_{i,5} \approx_c \mathcal{H}_{i,6}$: Same argument as $\mathcal{H}_{i,0} \approx_c \mathcal{H}_{i,1}$.

□

5.4 Extractability

Lemma 5.2. *Protocol 2 is multi-committer extractable.*

Proof. In the following we describe the extractor. For notational convenience we assume that the set of corrupted parties I is of size $|I| = \ell$ and we assume without loss of generality that $I = [\ell]$.

$\text{Ext}(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda)$:

1. Set ρ_λ to be the inner state of C_λ^* . Begin running C_λ^* until it outputs $\{c_i^{(msg)}, c_i^{(td)}\}_{i \in I}$.
2. Compute the commitments $\{c_i^{(key)}\}_{i \in I}$ as specified in the protocol.

3. Receive $\{\text{pp}_i\}_{i \in I}$ from C_λ^* , draw $\text{pp}_i \leftarrow \text{Spooky.Setup}(1^\lambda)$ for each $i \notin I$, and send $|I|$ copies of $\text{pp} := \{\text{pp}_i\}_{i \in [n]}$ to C_λ^* .
4. Receive $\{\text{pk}_i, \text{ct}_i, \widetilde{\text{CC}}_i\}_{i \in I}$.
5. Let $|\phi\rangle$ be the inner state of C_λ^* at this point, compute $\text{ct}_{|\phi}\leftarrow \text{Spooky.QEnc}(\text{pk}_1, |\phi\rangle)$.
6. For each $i \in I$, compute $\tilde{\text{ct}}_i^{(\text{td})} \leftarrow \text{Spooky.Eval}(\text{pk}_i, \text{Enc}(k_i, \cdot), \text{ct}_i)$ and $\tilde{\text{ct}}_i^{(\text{SFE})} \leftarrow \text{Spooky.Enc}(\text{pk}_i, \text{ct}_i^{(\text{SFE})})$ where $\text{ct}_i^{(\text{SFE})}$ is computed as specified in the protocol.
7. Let C_{Final}^* be the quantum circuit (derived from the adversary) that, on input the ciphertexts $\{\text{ct}_i^{(\text{td})}, \text{ct}_i^{(\text{SFE})}\}_{i \in I}$ and the quantum state of the adversary, computes the messages of the corrupted parties (corresponding to Step 7 of the protocol) and the updated quantum state of the adversary. Compute the following spooky evaluation procedure:

$$(\bar{\text{ct}}_1, \dots, \bar{\text{ct}}_\ell, \widehat{\text{ct}}_{|\phi}) \leftarrow \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_\ell), C_{\text{Final}}^*, (\text{ct}_{|\phi}), \{\tilde{\text{ct}}_i^{(\text{td})}, \tilde{\text{ct}}_i^{(\text{SFE})}\}_{i \in I}))$$

where each $\bar{\text{ct}}_i = (\bar{\text{ct}}_i^{(1)}, \dots, \bar{\text{ct}}_i^{(\ell)})$ are classical ciphertexts encrypted under $\text{pk}_1, \dots, \text{pk}_\ell$.

8. For all $i \in I$ compute

$$(\widehat{\text{ct}}_i^{(1)}, \dots, \widehat{\text{ct}}_i^{(\ell)}) \leftarrow \text{Spooky.Eval}\left((\text{pk}_1, \dots, \text{pk}_\ell), \text{SFE.Dec}\left(\text{dk}_i, \bigoplus_{j=1}^{\ell} \cdot\right), \bar{\text{ct}}_i\right).$$

9. For all $i \in [\ell]$ define $C_{\text{Nest}, i}$ to be the classical circuit that, on input a set of strings $\{y_j\}_{j \in [i-1]}$, a public key pk , and a ciphertext ct , computes

$$\text{ct}' \leftarrow \text{Spooky.Eval}(\text{pk}, y_1 \oplus \dots \oplus y_{i-1} \oplus \cdot, \text{ct}).$$

The circuit returns $\widetilde{\text{CC}}_i(\text{ct}')$.

10. For all $i \in [\ell, \dots, 2]$ compute iteratively

$$(\tilde{\text{ct}}_{\text{sk}, i}, \tilde{\text{ct}}_{m, i}) \leftarrow \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_{i-1}), C_{\text{Nest}, i}(\cdot, \text{pk}_i, \widehat{\text{ct}}_i^{(i)}), (\widehat{\text{ct}}_i^{(1)}, \dots, \widehat{\text{ct}}_i^{(i-1)}))$$

and for all $j \in [i-1]$ update the variables

$$(\widehat{\text{ct}}_j^{(1)}, \dots, \widehat{\text{ct}}_j^{(i-1)}) \leftarrow \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_{i-1}), C_{\text{Rec}, i}[\widehat{\text{ct}}_j^{(i)}], (\widehat{\text{ct}}_j^{(1)}, \dots, \widehat{\text{ct}}_j^{(i-1)}), \tilde{\text{ct}}_{\text{sk}, i})$$

where $C_{\text{Rec}, i}[\widehat{\text{ct}}_j^{(i)}]$ is the circuit that takes as input $2i-2$ strings (z_1, \dots, z_{i-1}) and (s_1, \dots, s_{i-1}) and computes

$$\bigoplus_{k=1}^{i-1} z_k \oplus \text{Spooky.Dec}\left(\bigoplus_{k=1}^{i-1} s_k, \widehat{\text{ct}}_j^{(i)}\right).$$

11. At the end of the iteration compute $(\text{sk}_1, m_1) \leftarrow \widetilde{\text{CC}}_1(\widehat{\text{ct}}_1^{(1)})$, then for all $i \in [2, \dots, \ell]$ compute

$$\text{sk}_i \leftarrow \bigoplus_{k=1}^{i-1} \text{Spooky.Dec}(\text{sk}_k, \tilde{\text{ct}}_{\text{sk}, i}^{(k)})$$

and

$$m_i \leftarrow \bigoplus_{k=1}^{i-1} \text{Spooky.Dec}(\text{sk}_k, \tilde{\text{ct}}_{m, i}^{(k)}).$$

12. Use the extracted keys (sk_1, \dots, sk_ℓ) to decrypt the state of the adversary from $\widehat{ct}_{|\phi}$ and the ciphertexts $(\widehat{ct}_1^{(\text{SFE})}, \dots, \widehat{ct}_\ell^{(\text{SFE})})$ from the ciphertexts $(\bar{ct}_1, \dots, \bar{ct}_\ell)$ as defined in Step 7 of the extractor. Return the transcript together with the state and the extracted messages (m_1, \dots, m_ℓ) .

We are now going to show that the transcript output by the extractor is computationally indistinguishable (with respect to compliant distinguishers) from that resulting from the real execution of the protocol. We do this by defining a sequence of hybrid distributions (for all $i \in I$) where we modify the interaction with the i -th corrupted party. In some of the following hybrids, the simulator inefficiently extracts the messages m_i and the trapdoors td_i from the messages $(c_i^{(\text{msg})}, c_i^{(\text{td})})$ of the corrupted parties. Note that this implies that each distribution is not necessarily computable in polynomial time. However, these hybrids should be thought of as mental experiments, which are going to be helpful in arguing about the indistinguishability of the simulator (which instead runs in strict quantum polynomial time).

- \mathcal{H}_0 : $\{\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i,\lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n))\}_{\lambda \in \mathbb{N}}$. Recall that this distribution includes the messages $\{m_i\}_{i \in I}$ committed by the transcript (if they exist).
- $\mathcal{H}_{i,1}$: Same as \mathcal{H}_0 except that $ct_i^{(\text{SFE})}$ is computed as $\text{SFE.Enc}(dk_i, (0^\lambda, 0^\lambda))$.
- $\mathcal{H}_{i,2}$: Same as $\mathcal{H}_{i,1}$ except that $c_i^{(\text{key})}$ is computed as a commitment to 0^λ .
- $\mathcal{H}_{i,3}$: Same as $\mathcal{H}_{i,2}$ except that $ct_i^{(\text{td})}$ is computed as $\text{Enc}(k_i, td_i)$, where td_i is extracted (inefficiently) from $c_i^{(\text{td})}$.
- $\mathcal{H}_{i,4}$: Same as $\mathcal{H}_{i,3}$ except that $c_i^{(\text{key})}$ is computed as a commitment to k_i , as specified in the protocol.
- $\mathcal{H}_{i,5}$: Same as $\mathcal{H}_{i,4}$ except that $ct_i^{(\text{SFE})}$ is computed as $\text{SFE.Enc}(dk_i, (k_i, r_i))$, as specified in the protocol.

We then define the last hybrid below. Note that the distribution induced by this hybrid is computable in (quantum) polynomial time.

- \mathcal{H}_6 : This is the output of the extractor as described above.

Now we argue indistinguishability between each hybrid. All reductions for distinguishers between hybrids below may receive the state of C^* after Step 1 along with the corresponding committed values $\{m_i, td_i\}_{i \in I}$ (if they exist) as non-uniform advice. Also, we only consider distinguishers that succeed with non-negligible probability, even conditioned on the event that the transcript received is *explainable*, i.e., each message lies in the support of the corresponding algorithm. This is because we consider only distinguishers that are *compliant*, i.e., they output 0 with overwhelming probability if the transcript is not explainable.

- $\mathcal{H}_0 \approx_c \mathcal{H}_{i,1}$: This follows from the quantum input privacy of the SFE protocol. The reduction takes the transcript of the protocol after executing Step 1 (including the state of the adversary and the messages $\{m_i\}_{i \in I}$) as non-uniform advice and continues to run the protocol honestly through Step 5. It then sets $(0, 0)$ and (k_i, r_i) as the challenge messages for SFE. In Step 6, the reduction sets $ct_i^{(\text{SFE})} = ct^*$, where ct^* is the challenge ciphertext. The rest of the protocol proceeds without changes. The reduction returns whatever the distinguisher returns.

Clearly if $ct^* = \text{SFE.Enc}(dk_i, (0, 0))$, then the distribution is identical to $\mathcal{H}_{i,2}$. On the other hand if $ct^* = \text{SFE.Enc}(dk_i, (k_i, r_i))$, then the distribution induced by the reduction is identical to $\mathcal{H}_{i,1}$. This implies that the two hybrids are computationally indistinguishable.

- $\mathcal{H}_{i,1} \approx_c \mathcal{H}_{i,2}$: This follows from an invocation of the (non-uniform) quantum computational hiding of the commitment scheme.
- $\mathcal{H}_{i,2} \approx_c \mathcal{H}_{i,3}$: This follows from a (non-uniform) reduction to the quantum semantic security of the secret-key encryption scheme, where td_i (together with the transcript so far and the messages $\{m_i\}_{i \in I}$) is given as non-uniform advice to the reduction.

- $\mathcal{H}_{i,3} \approx_c \mathcal{H}_{i,4}$: Same argument as $\mathcal{H}_{i,1} \approx_c \mathcal{H}_{i,2}$.
- $\mathcal{H}_{i,4} \approx_c \mathcal{H}_{i,5}$: Same argument as $\mathcal{H}_0 \approx_c \mathcal{H}_{i,1}$.
- $\mathcal{H}_{i,5} \approx_s \mathcal{H}_6$: We are going to argue that, conditioned on the event that the messages of the corrupted parties are explainable, the two hybrids are identical, except if an error in the evaluation (and consequently in the decryption) of the AFS-spooky encryption scheme occurs. Once that is established, statistical indistinguishability follows from the correctness of the AFS-spooky encryption scheme (which holds for *all choices* of the random coins used in the setup, key generation, and encryption algorithms). To substantiate this claim, recall that

$$(\widehat{\text{ct}}_{|\phi}, \bar{\text{ct}}_1, \dots, \bar{\text{ct}}_\ell) \leftarrow \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_\ell), \text{C}_{\text{Final}}^*, (\text{ct}_{|\phi}, \{\tilde{\text{ct}}_i^{(\text{td})}, \tilde{\text{ct}}_i^{(\text{SFE})}\}_{i \in I}))$$

where

$$\begin{aligned} \tilde{\text{ct}}_i^{(\text{SFE})} &= \text{Spooky.Enc}(\text{pk}_i, \text{ct}_i^{(\text{SFE})}) \\ &= \text{Spooky.Enc}(\text{pk}_i, \text{SFE.Enc}(\text{dk}_i, (k_i, r_i))) \end{aligned}$$

and

$$\begin{aligned} \tilde{\text{ct}}_i^{(\text{td})} &= \text{Spooky.Eval}(\text{pk}_i, \text{Enc}(k_i, \cdot), \text{ct}_i) \\ &= \text{Spooky.Eval}(\text{pk}_i, \text{Enc}(k_i, \cdot), \text{Spooky.Enc}(\text{pk}_i, \text{td}_i)) \\ &= \text{Spooky.Enc}(\text{pk}_i, \text{Enc}(k_i, \text{td}_i)). \end{aligned}$$

Therefore, by definition of $\text{C}_{\text{Final}}^*$ we have that for all $i \in [\ell]$ and $j \in [\ell]$

$$\bar{\text{ct}}_i^{(j)} = \text{Spooky.Enc}(\text{pk}_j, x_i^{(j)})$$

such that

$$\bigoplus_{j=1}^{\ell} x_i^{(j)} = \widehat{\text{ct}}_i^{(\text{SFE})} = \text{SFE.Eval}(\text{C}[c_i^{(\text{key})}, \text{ct}_i^{(\text{td})}, \text{td}_i, \text{lk}_i], \text{ct}_i^{(\text{SFE})}).$$

Recall that

$$(\widehat{\text{ct}}_i^{(1)}, \dots, \widehat{\text{ct}}_i^{(\ell)}) = \text{Spooky.Eval}\left((\text{pk}_1, \dots, \text{pk}_\ell), \text{SFE.Dec}\left(\text{dk}_i, \bigoplus_{j=1}^{\ell} \cdot\right), \bar{\text{ct}}_i\right)$$

and therefore for all $i \in [\ell]$ and $j \in [\ell]$ we have that $\widehat{\text{ct}}_i^{(j)} = \text{Spooky.Enc}(\text{pk}_j, y_i^{(j)})$ such that

$$\begin{aligned} \bigoplus_{j=1}^{\ell} y_i^{(j)} &= \text{SFE.Dec}\left(\text{dk}_i, \bigoplus_{j=1}^{\ell} x_i^{(j)}\right) \\ &= \text{SFE.Dec}\left(\text{dk}_i, \text{SFE.Eval}(\text{C}[c_i^{(\text{key})}, \text{ct}_i^{(\text{td})}, \text{td}_i, \text{lk}_i], \text{ct}_i^{(\text{SFE})})\right) \\ &= \text{lk}_i \end{aligned}$$

by the perfect correctness of the SFE protocol. Now recall that

$$(\tilde{\text{ct}}_{\text{sk}, \ell}, \tilde{\text{ct}}_{m, \ell}) = \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_{\ell-1}), \text{C}_{\text{Nest}, \ell}(\cdot, \text{pk}_\ell, \widehat{\text{ct}}_\ell^{(\ell)}), (\widehat{\text{ct}}_\ell^{(1)}, \dots, \widehat{\text{ct}}_\ell^{(\ell-1)}))$$

which implies that the two ciphertexts encode the output of the obfuscated program $\widetilde{\text{CC}}_\ell(\text{ct}'_\ell)$ where

$$\begin{aligned}\text{ct}'_\ell &= \text{Spooky.Eval}(\text{pk}_\ell, y_\ell^{(1)} \oplus \dots \oplus y_\ell^{(\ell-1)} \oplus \cdot, \widehat{\text{ct}}_\ell^{(\ell)}) \\ &= \text{Spooky.Enc}(\text{pk}_\ell, y_\ell^{(1)} \oplus \dots \oplus y_\ell^{(\ell)}) \\ &= \text{Spooky.Enc}(\text{pk}_\ell, \text{lk}_\ell).\end{aligned}$$

By the perfect correctness of the compute-and-compare obfuscation, the two variables $(\tilde{\text{ct}}_{\text{sk},\ell}, \tilde{\text{ct}}_{m,\ell})$ are AFS-spooky encryptions of (sk_ℓ, m_ℓ) , under $(\text{pk}_1, \dots, \text{pk}_\ell)$. This implies that the variables $(\widehat{\text{ct}}_1, \dots, \widehat{\text{ct}}_{\ell-1})$ are correctly updated to

$$\begin{aligned}\widehat{\text{ct}}_i &= (\widehat{\text{ct}}_i^{(1)}, \dots, \widehat{\text{ct}}_i^{(\ell-1)}) \\ &= \text{Spooky.Eval}((\text{pk}_1, \dots, \text{pk}_{\ell-1}), \text{C}_{\text{Rec},\ell}[\widehat{\text{ct}}_i^{(\ell)}], (\widehat{\text{ct}}_i^{(1)}, \dots, \widehat{\text{ct}}_i^{(\ell-1)}, \tilde{\text{ct}}_{\text{sk},\ell}))\end{aligned}$$

where $\widehat{\text{ct}}_i^{(j)} = \text{Spooky.Enc}(\text{pk}_j, \tilde{y}_i^{(j)})$ such that

$$\begin{aligned}\bigoplus_{j=1}^{\ell-1} \tilde{y}_i^{(j)} &= \bigoplus_{k=1}^{\ell-1} z_k \oplus \text{Spooky.Dec}\left(\bigoplus_{k=1}^{\ell-1} s_k, \widehat{\text{ct}}_i^{(\ell)}\right) \\ &= \bigoplus_{k=1}^{\ell-1} y_i^{(k)} \oplus \text{Spooky.Dec}\left(\text{sk}_\ell, \widehat{\text{ct}}_i^{(\ell)}\right) \\ &= y_i^{(1)} \oplus \dots \oplus y_i^{(\ell)} \\ &= \text{lk}_i.\end{aligned}$$

by the definition of $\text{C}_{\text{Rec},\ell}[\widehat{\text{ct}}_i^{(\ell)}]$. Recursively applying the above procedure, we obtain that

$$\begin{aligned}\widetilde{\text{CC}}_1(\widehat{\text{ct}}_1) &= \widetilde{\text{CC}}_1(\text{Spooky.Enc}(\text{pk}_1, \text{lk}_1)) \\ &= (\text{sk}_1, m_1)\end{aligned}$$

by the perfect correctness of the compute-and-compare obfuscation. It follows that the extractor successfully recomputes sk_1 , which allows it to iteratively recover $(\text{sk}_2, \dots, \text{sk}_\ell)$ from $(\tilde{\text{ct}}_{\text{sk},2}, \dots, \tilde{\text{ct}}_{\text{sk},\ell})$. Consequently, the decrypted transcript, the (possibly quantum) state of the adversary, and the messages (m_1, \dots, m_ℓ) are distributed identically as in the previous hybrid, conditioned on the fact that no error occurs during the evaluation algorithm. □

6 Quantum-Secure Multi-Verifier Zero-Knowledge

In this section, we use standard techniques to derive a multi-verifier zero-knowledge protocol from our multi-committer extractable commitment. We follow the approach given in [GK96] to upgrade a commitment-challenge-response Σ protocol to a full-fledged zero-knowledge protocol in constant rounds. In particular, the (multiple) verifiers will each commit to their challenge before the Σ protocol is executed, using our multi-committer extractable commitment scheme. A simulator will then be able to extract the challenge from all verifiers *simultaneously* and proceed to simulate each Σ protocol.

As in [BS20], a couple of subtleties arise in the proof. First, the extractable commitment guarantee does not hold against arbitrary malicious verifiers, as captured by our notion of simulation indistinguishability against *compliant* distinguishers. Thus, we have the verifier attach a witness indistinguishable proof (WI) that it acted explainably during the commitment phase, and indeed committed to the challenge that is

sent during the Σ protocol. However, in the proof of soundness, the verifier's initial commitment must be switched to a commitment to 0, since the reduction will receive the Σ protocol challenge from its challenger. This requires the verifier to prove a different statement under the WI, which must only be possible when interacting with a cheating prover. Details can be found in the description of Protocol 3.

Simulation Strategy. Following [BS20], we construct a zero-knowledge simulator that makes use of two non-rewinding sub-routines. Given an arbitrary malicious (multi-)verifier V^* , we consider the following two distributions. First, consider the real distribution over the final state of V^* on interaction with the honest prover, except that any time V^* aborts, the distribution outputs only a \perp symbol. We refer to this as $\text{RealNoAbort}_\perp(V^*)$. Next, consider the real distribution except that any time V^* *does not* abort, the distribution outputs only a \perp symbol. We refer to this as $\text{RealAbort}_\perp(V^*)$.

As a stepping stone towards proving zero-knowledge, we construct an entirely straight-line simulator SimNoAbort_\perp such that $\text{SimNoAbort}_\perp(V^*)$ is indistinguishable from $\text{RealNoAbort}_\perp(V^*)$. By entirely straight-line, we mean that not only does SimNoAbort_\perp not rewind V^* , it never even re-starts V^* from the beginning. Analogously, we also construct a simulator SimAbort_\perp such that $\text{SimAbort}_\perp(V^*)$ is indistinguishable from $\text{RealAbort}_\perp(V^*)$, and SimAbort_\perp is entirely straight-line.

Now, we combine the above simulators into a straight-line simulator SimComb_\perp that succeeds with probability negligibly close to $1/2$. SimComb_\perp simply chooses uniformly at random whether to run SimNoAbort_\perp or SimAbort_\perp and outputs the resulting view if the sub-routine is successful and \perp otherwise. Finally, we invoke the Watrous rewinding lemma to amplify the success probability of SimComb_\perp , resulting in the final simulator Sim .

We will actually make use of the sub-routines SimNoAbort_\perp and SimAbort_\perp explicitly in later sections, where the entirely straight-line nature of these procedures will be useful. In particular, we use both simulators in constructing non-malleable commitments (Section 7) and just SimNoAbort_\perp in the coin-flipping protocol in Section 8 (since we define an alternate/simpler abort generation procedure in that protocol).

6.1 Definition

Definition 6.1 (Quantum-Secure Multi-Verifier Zero-Knowledge Argument for NP). *A quantum-secure multi-verifier zero-knowledge argument for a language $\mathcal{L} \in \text{NP}$ is a pair (P, V) of classical PPT interactive Turing machines. P interacts with n copies $\{V_i\}_{i \in [n]}$ of V (who do not interact with each other) on common input 1^λ and 1^n , with each V_i additionally taking an input $x_i \in \mathcal{L}$, and P additionally taking inputs $\{x_i, w_i \in \mathcal{R}_\mathcal{L}(x_i)\}_{i \in [n]}$. At the end of the interaction, each V_i outputs a bit, indicating whether it accepts or rejects.*

1. **Perfect Completeness:** For any $\lambda, n \in \mathbb{N}, i \in [n], x \in \mathcal{L} \cap \{0, 1\}^\lambda$, and $w \in \mathcal{R}_\mathcal{L}(x)$,

$$\Pr[\text{OUT}_{V_i}(P(x, w), V_i(x))(1^\lambda, 1^n) = 1] = 1.$$

2. **Quantum Computational Soundness:** For any non-uniform quantum polynomial-size prover $P^* = \{P_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for all $\lambda, n \in \mathbb{N}, i \in [n]$, and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,

$$\Pr[\text{OUT}_{V_i}(P_\lambda^*(\rho_\lambda), V_i(x))(1^\lambda, 1^n) = 1] \leq \mu(\lambda).$$

3. **Quantum Computational Zero-Knowledge:** There exists a quantum expected polynomial-time simulator Sim such that for any non-uniform quantum polynomial-size adversary $V^* = \{V_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset of n verifiers, namely, $\{V_i\}_{i \in I}$ for some set $I \subseteq [n]$,

$$\left\{ \text{VIEW}_{V_\lambda^*} \left\langle \begin{array}{l} P(\{x_i, w_i\}_{i \in [n]}), \\ V_\lambda^*(\{x_i\}_{i \in I}, \rho_\lambda), \\ \{V_i(x_i)\}_{i \notin I} \end{array} \right\rangle (1^\lambda, 1^n) \right\}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}} \\ \approx_c \{ \text{Sim}(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda) \}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}}$$

where $\lambda \in \mathbb{N}, x_i \in \mathcal{L} \cap \{0, 1\}^\lambda, w_i \in \mathcal{R}_\mathcal{L}(x_i)$.

6.2 Construction

Ingredients: All of the following are assumed to be quantum-secure.

- A non-interactive perfectly-binding commitment Com .
- A multi-committer extractable commitment $\text{eCom} = (\text{eCom.C}, \text{eCom.R})$.
- A WI proof system $\text{WI} = (\text{WI.P}, \text{WI.V})$.
- A sigma protocol for NP $\Sigma = (\Sigma.P, \Sigma.V)$.

Remark 6.1. Observe that since the sigma protocol Σ is public-coin, Protocol 3 is publicly-verifiable. That is, any third party, upon observing the transcript of interaction between P and V, can deduce whether V accepted or not. This fact will be used in Section 8.

Protocol 3

Common input: 1^λ and 1^n .

V_i 's additional input: $x_i \in \mathcal{L}$.

P's additional input: $\{x_i, w_i \in \mathcal{R}_{\mathcal{L}}(x_i)\}_{i \in [n]}$.

1. For each $i \in [n]$, P computes and sends $c_i \leftarrow \text{Com}(1^\lambda, w_i)$ to V_i .
2. Each V_i computes a challenge $\beta_i \leftarrow \Sigma.V_1(1^{|x_i|})$. Then, P and $\{V_i\}_{i \in [n]}$ interact, with P taking the role of eCom.R and V_i taking the role of $\text{eCom.C}_i(\beta_i)$, to produce $\{\tau_i\}_{i \in [n]} \leftarrow \langle \text{eCom.R}, \{\text{eCom.C}_i(\beta_i)\}_{i \in [n]} \rangle(1^\lambda, 1^n)$.
3. For each $i \in [n]$, P computes $(\alpha_i, \text{st}_i) \leftarrow \Sigma.P_1(x_i, w_i)$ and sends α_i to V_i .
4. Each V_i sends β_i .
5. For each $i \in [n]$, P and V_i interact (in parallel) with P taking the role of WI.V and V_i taking the role of WI.P to give P a WI proof that
 - τ_i is explainable, and opens to β_i ,
 - **OR**, c_i opens to a non-witness $z_i \notin \mathcal{R}_{\mathcal{L}}(x_i)$.
6. For each $i \in [n]$, P and V_i interact (in parallel) with P taking the role of WI.P and V_i taking the role of WI.V to give V_i a WI proof that
 - c_i opens to some string z_i ,
 - **OR**, $x_i \in \mathcal{L}$.
7. For each $i \in [n]$, P computes and sends $\gamma_i = \Sigma.P_2(x_i, w_i, \text{st}_i, \alpha_i, \beta_i)$ to V_i .
8. Each V_i accepts if $\Sigma.V_2(x_i, \alpha_i, \beta_i, \gamma_i) = 1$.

Figure 3: A constant-round quantum-secure multi-verifier zero-knowledge argument for $\mathcal{L} \in \text{NP}$.

6.3 Soundness

Lemma 6.1. *Protocol 3 has quantum computational soundness.*

Proof. Assume towards contradiction that there exists a non-uniform quantum polynomial-size prover $P^* = \{P_\lambda^*, \rho_\lambda^*\}_{\lambda \in \mathbb{N}}$ that with noticeable probability, convinces V_i to accept on input instances $\{x_\lambda\}_{\lambda \in \mathbb{N}}$ where $|x_\lambda| = \lambda$ and $x_\lambda \notin \mathcal{L}$. Let $V = V_i$, and by averaging, we can assume that P_λ^* sends a fixed first message c_λ to V . Furthermore, since P_λ^* succeeds in convincing V to accept instances $x_\lambda \notin \mathcal{L}$ with noticeable probability, the statistical soundness of the WI in Step 6 of the protocol implies that c_λ must be a well-formed commitment, that is, $c_\lambda = \text{Com}(1^\lambda, z_\lambda; s_\lambda)$ for some (z_λ, s_λ) . Now, consider the following sequence of computationally indistinguishable hybrid distributions.

- \mathcal{H}_0 : $\{\text{VIEW}_{P^*}(P^*(\rho_\lambda), V)(1^\lambda, x_\lambda)\}_{\lambda \in \mathbb{N}}$.
- \mathcal{H}_1 : Same as \mathcal{H}_0 except that V uses (z_λ, s_λ) as the witness for the second part of the WI statement given in Step 5 of the protocol. $\mathcal{H}_0 \approx_c \mathcal{H}_1$ follows from the witness indistinguishability of WI, where the reduction is given (z_λ, s_λ) as non-uniform advice.
- \mathcal{H}_2 : Same as \mathcal{H}_1 except that in Step 2, V takes the role of $\text{eCom.C}(0^\lambda)$ rather than $\text{eCom.C}(\beta)$. $\mathcal{H}_1 \approx_c \mathcal{H}_2$ follows from the computational hiding of eCom .

Using P^* , we construct a cheating prover $\Sigma.P^* = \{\Sigma.P_\lambda^*, \Sigma.\rho_\lambda^*\}_{\lambda \in \mathbb{N}}$ for the sigma protocol. The non-uniform advice $\Sigma.\rho_\lambda^*$ is generated as follows. Run $P^*(\rho_\lambda^*)$ until it outputs its first message c_λ . Extract from c_λ the message z_λ committed and the corresponding opening s_λ and define the resulting advice to consist of the state of P^* at this point, along with (z_λ, s_λ) .

$\Sigma.P_\lambda^*(\Sigma.\rho_\lambda^*)$:

1. Interact with P_λ^* , taking the role of $\text{eCom.C}(0^\lambda)$.
2. Continue running P_λ^* , obtaining the message α .
3. Send α to $\Sigma.V$, receive β , and send β to P_λ^* .
4. Interact with P_λ^* to give P_λ^* a WI proof as in Step 5 of the protocol, using witness (z_λ, s_λ) .
5. Interact with P_λ^* to receive a WI proof from P_λ^* as in Step 6 of the protocol.
6. Continue running P_λ^* , obtaining the message γ , and send γ to $\Sigma.V$.

Now note that P_λ^* 's view in this interaction is exactly \mathcal{H}_2 . Thus, since P^* succeeds in convincing V to accept with noticeable probability, and $\mathcal{H}_0 \approx_c \mathcal{H}_2$, it must be the case that $\Sigma.V$ accepts with noticeable probability, a contradiction. □

6.4 Zero-Knowledge

Theorem 6.1. *Protocol 3 is quantum computational zero-knowledge.*

Proof. We begin by describing the two sub-routines SimNoAbort_\perp and SimAbort_\perp mentioned above. Then we combine them into SimComb_λ , which we use to derive the final simulator Sim .

$\text{SimNoAbort}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)$:

1. Set ρ_λ to be the inner state of V_λ^* . For each $i \in I$, compute and send $c_i \leftarrow \text{Com}(1^\lambda, 0)$ to V_λ^* .

- Let $V.\text{eCom}_\lambda^*$ be the portion of V_λ^* that interacts with P in Step 2 above. Note that its state at the beginning of this interaction is $\rho_\lambda, \{c_i\}_{i \in I}$. Compute

$$(\{\tau_i\}_{i \in I}, \text{st}, \{\beta'_i\}_{i \in I}) \leftarrow \text{eCom.Ext}(1^\lambda, 1^n, I, V.\text{eCom}_\lambda^*, (\rho_\lambda, \{c_i\}_{i \in I})).$$

If eCom.Ext produced an abort transcript, then halt and return \perp . Otherwise continue, setting st to be the inner state of V_λ^* .

- For all $i \in I$, compute $(\alpha_i, \gamma_i) \leftarrow \Sigma.\text{Sim}(x_i, \beta'_i)$ and send α_i to V_λ^* .
- V_λ^* returns $\{\beta_i\}_{i \in I}$.
- Take the role of the honest prover WI.P in the $|I|$ WI proofs that V_λ^* gives. If V_λ^* fails to prove any of the statements, then halt and output \perp .
- Give V_λ^* a total of $|I|$ WI proofs using the $|I|$ witnesses that show $\{c_i\}_{i \in I}$ are valid commitments. Then, send $\{\gamma_i\}_{i \in I}$ to V_λ^* .
- Output the inner state of V_λ^* .

$\text{SimAbort}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)$:

- Set ρ_λ to be the inner state of V_λ^* .
- Interact with V_λ^* as the honest prover until the end of Step 6 of the protocol, with exactly 3 differences:
 - The commitments c_i in Step 1 are to 0 rather than w_i .
 - The messages α_i sent in Step 3 are generated by the simulator of the sigma protocol, $(\alpha_i, \gamma_i) \leftarrow \Sigma.\text{Sim}(x_i, 0^\lambda)$.
 - In Step 6, the witnesses used for the WI proofs are for the first statement (that c_i is a valid commitment).
- If at some point during the interaction V_λ^* either aborts or fails in one of its WI proofs, halt and output V_λ^* 's inner state. Otherwise, output \perp .

$\text{SimComb}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)$: With probability $1/2$, execute $\text{SimNoAbort}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)$ and otherwise execute $\text{SimAbort}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)$.

$\text{Sim}(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)$: Let $\overline{\text{SimComb}}_\perp(\cdot) := \text{SimComb}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \cdot)$ be the circuit SimComb_\perp with all inputs hard-coded except for ρ_λ , and output $R(\overline{\text{SimComb}}_\perp, \rho_\lambda, \lambda)$, where R is the algorithm from Lemma 3.1.

Next, we introduce some notation. For $V^* = \{V_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, let $\text{Real}(V^*)$ denote

$$\left\{ \text{VIEW}_{V_\lambda^*} \left\langle \begin{array}{l} P(\{x_i, w_i\}_{i \in [n]}), \\ V_\lambda^*(\{x_i\}_{i \in I}, \rho_\lambda), \\ \{V_i(x_i)\}_{i \notin I} \end{array} \right\rangle (1^\lambda, 1^n) \right\}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}}$$

for $\lambda \in \mathbb{N}$, $x_i \in \mathcal{L} \cap \{0, 1\}^\lambda$, $w_i \in \mathcal{R}_\mathcal{L}(x_i)$, and let

- $\text{RealNoAbort}_\perp(V^*)$ be the distribution $\text{Real}(V^*)$, except that whenever an abort occurs, the distribution outputs \perp , and
- $\text{RealAbort}_\perp(V^*)$ be the distribution $\text{Real}(V^*)$, except that if an abort *does not* occur, the distribution outputs \perp .

We continue by proving two lemmas that will be useful on their own in later sections, and will also be useful in proving the quantum computational zero-knowledge of ZK.

Lemma 6.2. *For any V^* , $\text{RealNoAbort}_\perp(V^*) \approx_c \text{SimNoAbort}_\perp(V^*)$.*

Proof. We prove this via a sequence of hybrids.

- \mathcal{H}_0 : $\text{RealNoAbort}_\perp(V^*)$.
- \mathcal{H}_1 : Same as \mathcal{H}_0 except that after Step 2, the values $\{\beta'_i\}_{i \in I}$ committed by the transcripts $\{\tau_i\}_{i \in I}$ are inefficiently extracted, and after Step 4, if any $\beta'_i \neq \beta_i$, the hybrid aborts (outputs \perp).
- \mathcal{H}_2 : Same as \mathcal{H}_1 except that the transcript of the extractable commitment is simulated. In particular, the hybrid computes

$$(\{\tau_i\}_{i \in I}, \text{st}, \{\beta'_i\}_{i \in I}) \leftarrow \text{eCom.Ext}(1^\lambda, 1^n, I, V.\text{eCom}^*_\lambda, (\rho_\lambda, \{c_i\}_{i \in I})),$$

and proceeds to run the verifier with inner state st .

- \mathcal{H}_3 : Same as \mathcal{H}_2 except that the equality checks introduced in \mathcal{H}_1 are removed, and the sigma protocol is simulated. In particular, for each $i \in I$, the hybrid computes $(\alpha_i, \gamma_i) \leftarrow \Sigma.\text{Sim}(x_i, \beta'_i)$.
- \mathcal{H}_4 : Same as \mathcal{H}_3 except that each commitment for $i \in I$ in the first message sent by the prover is $c_i \leftarrow \text{Com}(1^\lambda, 0)$.
- \mathcal{H}_5 : Same as \mathcal{H}_4 except that the WI proofs given by the prover for $i \in I$ are generated with witnesses showing that c_i is a valid commitment.

Observe that \mathcal{H}_5 is exactly $\text{SimNoAbort}_\perp(V^*)$. Now we show that each consecutive pair of hybrids is indistinguishable.

- $\mathcal{H}_0 \approx_s \mathcal{H}_1$: This follows from the statistical soundness of the WI proved in Step 5.
- $\mathcal{H}_1 \approx_c \mathcal{H}_2$: Assume there exists a distinguisher D for \mathcal{H}_1 and \mathcal{H}_2 that succeeds with non-negligible probability. We build a compliant¹⁵ distinguisher D' that breaks the extractability property of eCom .

First, we fix a sequence of instance-witness pairs $\{\{x_{i,\lambda}, w_{i,\lambda}\}_{i \in I}\}_{\lambda \in \mathbb{N}}$, and first messages $\{\{c_{i,\lambda}\}_{i \in I}\}_{\lambda \in \mathbb{N}}$ for which D succeeds with non-negligible probability. The witnesses $\{\{w_{i,\lambda}\}_{i \in I}\}_{\lambda \in \mathbb{N}}$ will be given as non-uniform advice to D' . Now, D' will take as input either the real or the simulated view with respect to committer $\{V.\text{eCom}^*_\lambda, (\rho_\lambda, \{c_{i,\lambda}\}_{i \in I})\}_{\lambda \in \mathbb{N}}$. This view includes the messages committed and the final state of the committer, which is the state of V^* after Step 2 of the protocol. D' proceeds to simulate the rest of the interaction between P and V^* , making use of the witnesses it received as non-uniform advice during Steps 3 and 7, as well as the committed messages it received from its challenger to implement the check introduced in \mathcal{H}_1 . If V^* aborts or fails to prove any of the WI statements in Step 5, D' outputs 0. Otherwise, it queries D with V^* 's final view and outputs what D outputs.

Observe that D' 's advantage is equivalent to D 's advantage. This follows because i) whenever D' queries D with a transcript, it is a faithful execution of either \mathcal{H}_1 or \mathcal{H}_2 , depending on whether eCom was simulated or not, and ii) whenever D' does *not* query D , it means that V^* failed to prove one of its WI statements, so D 's input would have been \perp . Finally, D' is compliant by the statistical soundness of the WI.

- $\mathcal{H}_2 \approx_c \mathcal{H}_3$: This follows from the special zero-knowledge property of Σ .
- $\mathcal{H}_3 \approx_c \mathcal{H}_4$: This follows from the computational hiding of Com .
- $\mathcal{H}_4 \approx_c \mathcal{H}_5$: This follows from the witness indistinguishability of WI.

¹⁵Recall that such a distinguisher is guaranteed to output 0 with overwhelming probability on input any non-explainable view.

□

Lemma 6.3. For any V^* , $\text{RealAbort}_\perp(V^*) \approx_c \text{SimAbort}_\perp(V^*)$.

Proof. We prove this via a sequence of hybrids.

- \mathcal{H}_0 : $\text{RealAbort}_\perp(V^*)$. Note that if V^* has not aborted at some point during Steps 1-6 of the protocol, this distribution outputs \perp .
- \mathcal{H}_1 : Same as \mathcal{H}_0 except that in Step 3, for each $i \in I$, P sends α_i where $(\alpha_i, \gamma_i) \leftarrow \text{Sim}(x_i, 0^\lambda)$.
- \mathcal{H}_2 : Same as \mathcal{H}_1 except that each commitment for $i \in I$ in the first message sent by the prover is $c_i \leftarrow \text{Com}(1^\lambda, 0)$.
- \mathcal{H}_3 : Same as \mathcal{H}_2 except that the WI proofs for $i \in I$ given by the prover are generated with witnesses showing that c_i is a valid commitment.

Observe that \mathcal{H}_3 is exactly $\text{SimAbort}_\perp(V^*)$. Now we show that each consecutive pair of hybrids is indistinguishable.

- $\mathcal{H}_0 \approx_c \mathcal{H}_1$: This follows from the first-message indistinguishability of Σ .
- $\mathcal{H}_1 \approx_c \mathcal{H}_2$: This follows from the computational hiding of Com .
- $\mathcal{H}_2 \approx_c \mathcal{H}_3$: This follows from the witness indistinguishability of WI.

□

To finish the proof of zero-knowledge, we introduce some more notation.

- Let $\text{Pr}_{\text{Real}}^{\text{Abort}}(V^*)$ be the probability that V^* aborts in the real interaction with the honest prover.
- Let $\text{Pr}_{\text{SimNoAbort}}^{\text{Abort}}(V^*)$ be the probability that V^* aborts in SimNoAbort_\perp (i.e. the outcome is \perp).
- Let $\text{Pr}_{\text{SimAbort}}^{\text{Abort}}(V^*)$ be the probability that V^* aborts in SimAbort_\perp (i.e. the outcome is not \perp).
- Let $\text{SimNoAbort}_\perp(V^*) := \{\text{SimNoAbort}_\perp(1^\lambda, n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)\}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}}$.
- Let $\text{SimAbort}_\perp(V^*) := \{\text{SimAbort}_\perp(1^\lambda, n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)\}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}}$.
- Let $\text{SimComb}_\perp(V^*) := \{\text{SimComb}_\perp(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)\}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}}$.
- Let $\text{Sim}(V^*) := \{\text{Sim}(1^\lambda, 1^n, I, \{x_i\}_{i \in [n]}, V_\lambda^*, \rho_\lambda)\}_{\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]}}$.
- Let $\text{RealNoAbort}(V^*)$ be the distribution $\text{Real}(V^*)$ conditioned on there not being an abort.
- Let $\text{RealAbort}(V^*)$ be the distribution $\text{Real}(V^*)$ conditioned on there being an abort.
- Let $\text{SimNoAbort}(V^*)$ be the distribution SimNoAbort_\perp conditioned on there not being an abort (i.e. conditioned on the output not being \perp).
- Let $\text{SimAbort}(V^*)$ be the distribution SimAbort_\perp conditioned on there being an abort (i.e. conditioned on the output not being \perp).
- Let $\text{SimComb}(V^*)$ be the distribution SimComb_\perp conditioned on the output not being \perp .

Following [BS20], we show that $\text{Real}(V^*) \approx_c \text{SimComb}(V^*)$ via a sequence on hybrids. In particular, we show that

$$\begin{aligned}
\text{Real}(V^*) &\stackrel{(1)}{=} (1 - \Pr_{\text{Real}}^{\text{Abort}}(V^*))\text{RealNoAbort}(V^*) + (\Pr_{\text{Real}}^{\text{Abort}}(V^*))\text{RealAbort}(V^*) \\
&\stackrel{(2)}{\approx_s} (1 - \Pr_{\text{SimNoAbort}}^{\text{Abort}}(V^*))\text{RealNoAbort}(V^*) + (\Pr_{\text{SimAbort}}^{\text{Abort}}(V^*))\text{RealAbort}(V^*) \\
&\stackrel{(3)}{\approx_c} (1 - \Pr_{\text{SimNoAbort}}^{\text{Abort}}(V^*))\text{SimNoAbort}(V^*) + (\Pr_{\text{SimAbort}}^{\text{Abort}}(V^*))\text{RealAbort}(V^*) \\
&\stackrel{(4)}{\approx_c} (1 - \Pr_{\text{SimNoAbort}}^{\text{Abort}}(V^*))\text{SimNoAbort}(V^*) + (\Pr_{\text{SimAbort}}^{\text{Abort}}(V^*))\text{SimAbort}(V^*) \\
&\stackrel{(5)}{\approx_s} \text{SimComb}(V^*),
\end{aligned}$$

where

1. The equality (1) follows by definition.
2. The indistinguishability (2) follows as a corollary of Lemma 6.2 and Lemma 6.3. Indeed, $\text{RealNoAbort}_{\perp}(V^*) \approx_c \text{SimNoAbort}_{\perp}(V^*)$ in particular implies that the difference in the probability that the verifier aborts in the real interaction versus the simulated interaction is negligible, and likewise for $\text{RealAbort}_{\perp}(V^*) \approx_c \text{SimAbort}_{\perp}(V^*)$.
3. The indistinguishability (3) follows as a corollary of Lemma 6.2. This can be seen by considering two cases. First, if the probability that the verifier aborts in the real interaction is negligible, then $\text{RealNoAbort}(V^*) \approx_c \text{SimNoAbort}(V^*)$ directly follows from Lemma 6.2, and the indistinguishability follows. Otherwise, this probability is non-negligible, meaning that $\text{RealAbort}(V^*)$ is efficiently samplable. Thus, a reduction to Lemma 6.2 can sample from the distribution $\text{RealAbort}(V^*)$ whenever it receives \perp from its challenger.¹⁶
4. The indistinguishability (4) follows as a corollary of Lemma 6.3 via a similar analysis as the last step.
5. The indistinguishability (5) follows from the definition of $\text{Sim}(V^*)$ and the claim that the difference between $\Pr_{\text{SimNoAbort}}^{\text{Abort}}$ and $\Pr_{\text{SimAbort}}^{\text{Abort}}$ is negligible (which follows as a corollary of Lemma 6.2 and Lemma 6.3).

Finally, this implies that $\text{Real}(V^*) \approx_c \text{Sim}(V^*)$ by applying Lemma 3.1 for each $(\lambda, \{x_i\}_{i \in [n]}, \{w_i\}_{i \in [n]})$ with the following parameters. Set $Q := \text{SimComb}_{\perp}$ as defined in the description of Sim , and set $\epsilon := \text{negl}(\lambda) + 2^{-\lambda \cdot \frac{3}{4}}$, $p_0 := 1/4$, and $q := 1/2$, as described in [BS20, Proposition 3.5]. This completes the proof of quantum computational zero-knowledge. □

7 Quantum-Secure Non-Malleable Commitments

7.1 Definition

In this section, we define quantum-secure non-malleable commitments w.r.t. commitment. We consider the synchronous setting where there is a quantum man-in-the-middle adversary $\text{MIM} = \{\text{MIM}_{\lambda}, \rho_{\lambda}\}_{\lambda \in \mathbb{N}}$ interacting with a classical honest committer \mathcal{C} with tag $\text{tag}_{\mathcal{C}}$ (where \mathcal{C} commits to value v) in the left session, and interacting with classical honest receiver \mathcal{R} in the right session. The MIM uses tag tag_{MIM} in its interaction with \mathcal{R} . Prior to the interaction, the value v is given to \mathcal{C} as local input.

¹⁶A more formal analysis of this can be found in [BS20, Proposition 3.4].

Then the commit phase is executed. After obtaining an honest left message in any round, the MIM sends its own right message. And after obtaining an honest right message in any round, the MIM sends its own left message. Let $\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v), \text{MIM}(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}})$ denote a random variable that describes the value v' committed by the MIM in the right session, jointly with the view of the MIM in the full (both left and right sessions) experiment. If the $\text{tag}_\mathcal{C}$ used by \mathcal{C} in the left interaction is identical to the tag_{MIM} used by the MIM in the right interaction, then the value v' committed to in the right interaction is defined to be \perp . If the MIM sends a message that causes an honest party to abort in either the left or the right execution, then the value v' committed to in the right interaction is also defined to be \perp .

We will concern ourselves with computationally hiding and statistically binding commitments that additionally satisfy the non-malleability property defined below.

Definition 7.1 (Quantum-Secure Non-Malleable Commitments with respect to Commitment). *For any $\ell = \ell(\lambda)$ and $p = p(\lambda)$, a commitment scheme $\langle \mathcal{C}, \mathcal{R} \rangle$ is said to be quantum secure non-malleable with respect to commitment for tags in $[\ell]$ if for every $v_1, v_2 \in \{0, 1\}^{2p(\lambda)}$, for every quantum polynomial-size $\text{MIM} = \{\text{MIM}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ and every quantum polynomial-size distinguisher $\text{D} = \{\text{D}_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\eta(\cdot)$ such that for all large enough $\lambda \in \mathbb{N}$, and for all $\text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}} \in [\ell]$ where $\text{tag}_\mathcal{C} \neq \text{tag}_{\text{MIM}}$, the following holds:*

$$\left| \Pr[\text{D}_\lambda(\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v_1), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}}), \sigma_\lambda) = 1] \right. \\ \left. - \Pr[\text{D}_\lambda(\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v_2), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}}), \sigma_\lambda) = 1] \right| = \eta(\lambda) \quad (4)$$

We will also consider a more general setting where the MIM interacts with polynomially many committers in the left session, and a single honest receiver in the right session. For any polynomial $n = n(\lambda)$ number of left sessions, we will let $\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(\{v_i\}_{i \in [n]}), \text{MIM}(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}})$ denote a random variable that describes the value v' committed by the MIM in the right session, jointly with the view of the MIM in the full (both left and right sessions) experiment.

Definition 7.2 (Many-one Quantum-Secure Non-Malleable Commitments with respect to Commitment). *For any $\ell = \ell(\lambda)$, $p = p(\lambda)$ and $n = n(\lambda)$, a commitment scheme $\langle \mathcal{C}, \mathcal{R} \rangle$ is said to be quantum secure many-one non-malleable with respect to commitment for tags in $[\ell]$ if for every pair of tuples $(\{v_i^1\}_{i \in [n]}, \{v_i^2\}_{i \in [n]}) \in \{0, 1\}^{2np(\lambda)}$, for every quantum polynomial-size $\text{MIM} = \{\text{MIM}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ and every quantum polynomial-size distinguisher $\text{D} = \{\text{D}_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\eta(\cdot)$ such that for all large enough $\lambda \in \mathbb{N}$, and for all $(\{\text{tag}_i^c\}_{i \in [n]}, \text{tag}^{\text{MIM}})$ where each tag is in $[\ell]$ such that $\text{tag}^{\text{MIM}} \notin \{\text{tag}_i^c\}_{i \in [n]}$, the following holds:*

$$\left| \Pr[\text{D}_\lambda(\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(\{v_i^1\}_{i \in [n]}), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}}), \sigma_\lambda) = 1] \right. \\ \left. - \Pr[\text{D}_\lambda(\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(\{v_i^2\}_{i \in [n]}), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_\mathcal{C}, \text{tag}_{\text{MIM}}), \sigma_\lambda) = 1] \right| = \eta(\lambda) \quad (5)$$

7.2 Non-Malleable Commitments for Small Tags

First, we provide an overview of our scheme for tags in $[N]$ where $N = \lambda^{\text{ilog}(c+1, \lambda)}$. We will assume non-interactive perfectly binding commitments and two-message SFE which can be broken with advantage at most $\text{negl}(\lambda^{\text{ilog}(c, \lambda)})$ by polynomial size quantum circuits. Recall that as discussed in the technical overview, we will have the committer and receiver establish an erasure channel via a two-party input-hiding SFE. This channel will transmit the committer's value with probability ϵ , depending on their tag. Here, we discuss our construction in more detail.

The committer on input $m \in \{0, 1\}^{p(\lambda)}$ sends a perfectly binding, computationally hiding commitment to m , denoted by $\text{Com}(m)$. Next, the committer and receiver run an SFE execution, where the receiver input is a uniformly random r_1 and committer input is m along with uniformly random s_1 that is of the same length as r_1 , and m is transmitted to the receiver if and only if $s_1 = r_1$. The length of r_1 and s_1 is carefully chosen so that the probability that they are equal is $\eta^{-\text{tag}}$, where $\eta = \lambda^{\text{ilog}(c+1, \lambda)}$ is a small superpolynomial value. Additionally, the SFE scheme is such that evaluations of agreeing circuits are (subexponentially)

statistically close. This essentially means that no matter how a malicious committer or receiver may behave, the message m is revealed with probability close to $\eta^{-\text{tag}}$ (with an error of $\text{negl}(\eta^{\text{tag}})$).

Now, let us consider a setting where the MIM uses tag_{MIM} and honest committer uses tag_c such that $\text{tag}_{\text{MIM}} < \text{tag}_c$. In this case, the SFE statistically hides the committed message except with probability roughly $\eta^{-\text{tag}_c}$, and on the other hand, the MIM's message is revealed with probability roughly $\eta^{-\text{tag}_{\text{MIM}}}$, which is greater than $\eta^{-\text{tag}_c}$. Intuitively, this means that any MIM that tries to maul or copy the committed message cannot succeed, at least in executions where the MIM's value was revealed but the honest committer's was not. We generalize this to all transcripts by relying on the fact that no MIM can actually tell whether the MIM's value was revealed, and therefore cannot behave any differently in transcripts where extraction occurred vs where it didn't.

Formally, we will prove that the joint distribution \mathcal{V}_1 of the MIM's view and committed value when the honest commitment is to v_1 , is indistinguishable from the joint distribution \mathcal{V}_2 when the honest commitment is to v_2 . This is done as follows.

- First, we use the input-hiding property of SFE to argue that any distinguisher D that distinguishes \mathcal{V}_1 from \mathcal{V}_2 , *must also distinguish these distributions when restricted to executions where $s_1 = r_1$ in the right execution* where the MIM is the committer. We prove that if this is not the case, then D can be used to guess the input r_1 of the honest receiver in the right execution, contradicting the input-hiding property of SFE.
- Once this is established, we restrict ourselves to transcripts where $s_1 = r_1$ in the right execution.
- We rely on our setting of parameters to ensure that the transcripts where $s_1 = r_1$ in the left execution can only form a negligible fraction of all transcripts where $s_1 = r_1$ in the right execution.
- Roughly, this means that for an overwhelming fraction of transcripts where $s_1 = r_1$ in the right SFE execution, the left SFE execution perfectly erases the honest committer's message.
- As a result, any D that distinguishes between the distributions \mathcal{V}_1 and \mathcal{V}_2 , also distinguishes between these distributions when restricted to $s_1 = r_1$ in the right execution. We note that conditioned on $s_1 = r_1$, the MIM's message can be efficiently extracted, and therefore D can be used to carefully break the (super-polynomial) hiding of the commitment Com .

This completes a sketch of our argument when $\text{tag}_{\text{MIM}} < \text{tag}_c$. In case $\text{tag}_{\text{MIM}} > \text{tag}_c$, this argument does not go through, since the honest committer's message is revealed with probability that is larger than the MIM's message. To deal with this situation, we append another sequential instance of SFE to our commitment, where the probability of extraction varies as a function of $2N - \text{tag}$, instead of as a function of tag . This means that a committer with tag_c will run two instances of SFE, one which transmits the committed message with probability $\eta^{-\text{tag}_c}$, and another that transmits it with probability $\eta^{2N - \text{tag}_c}$. Now, for $\text{tag}_{\text{MIM}} \neq \text{tag}_c$, in at least one of these sessions, the probability that the MIM's message is revealed will be larger than the probability that the committer's message is revealed. Moreover, since all these probabilities of revealing messages are negligible, the other session will not reveal the committer's message except with negligible probability, and therefore, we can switch to a hybrid where the other session never outputs the committer's message. Finally, since proving security against synchronous MIM adversaries suffices for our applications, we only focus on formally proving synchronous security here, but we suspect that similar arguments would suffice to prove security of our construction against non-synchronous adversaries.

7.2.1 Construction

Ingredients and notation: We will assume the existence of

- A non-interactive perfectly-binding quantum computationally hiding commitment scheme Com where there exists a constant $c_1 > 0$ s.t. no QPT adversary has advantage better than $\text{negl}(\lambda^{\log(c_1, \lambda)})$ in the hiding game.

- A two-message SFE satisfying Definition 3.8. This means that there exists a constant $c_2 > 0$ s.t. no QPT adversary has advantage better than $\text{negl}(\lambda^{\text{ilog}(c_2, \lambda)})$ in the quantum input privacy game.
- A quantum-secure zero-knowledge argument for NP (ZK.P, ZK.V). (We do not require multi-verifier zero-knowledge for this section.)

Let $c = \max(c_1, c_2)$, $\eta = \lambda^{\text{ilog}(c+1, \lambda)}$, and $N = \text{ilog}(c+1, \lambda)$. We describe the protocol for tags or identities in $[N]$ in Protocol 4. Also, define the language

$$\mathcal{L} = \left\{ \left(\begin{array}{c} c, \text{ct}_1, \text{ct}_2, \text{ct}'_1, \\ \text{ct}'_2, \text{tag}, N \end{array} \right) : \exists \left(\begin{array}{c} m, r, s_1, s_2, \\ u_1, u_2 \end{array} \right) \text{ s.t. } \left. \begin{array}{l} |s_1| = \text{tag} \cdot (\log \eta), \\ |s_2| = (2N - \text{tag}) \cdot (\log \eta), \\ c = \text{Com}(1^\lambda, m; r), \\ \text{ct}_1 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_1, (m||r)], \text{ct}'_1; u_1), \\ \text{ct}_2 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_2, (m||r)], \text{ct}'_2; u_2) \end{array} \right\}.$$

where $\text{Id}(\cdot)$ denotes the identity function.

Protocol 4

Common Input: 1^λ and a tag $\in [N]$. Set $t_1 = \text{tag} \cdot (\log \eta)$, and $t_2 = (2N - \text{tag}) \cdot (\log \eta)$.

C's Input: A message $m \in \{0, 1\}^{p(\lambda)}$.

Commit Stage:

1. C samples $r \leftarrow U_\lambda$, and sends $c_1 = \text{Com}(1^\lambda, m; r)$.
2. R samples $\text{dk}_1 \leftarrow \text{SFE.Gen}(1^\lambda)$, $r_1 \leftarrow U_{t_1}$, and sends $\text{ct}_{1,R} \leftarrow \text{SFE.Enc}_{\text{dk}_1}(r_1)$.
3. C samples $s_1 \leftarrow U_{t_1}$, $u_1 \leftarrow U_\lambda$ and sends $\text{ct}_1 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_1, (m||r)], \text{ct}_{1,R}; u_1)$ and s_1 , where $\text{Id}(\cdot)$ is the identity function.
4. R samples $\text{dk}_2 \leftarrow \text{SFE.Gen}(1^\lambda)$, $r_2 \leftarrow U_{t_2}$, and sends $\text{ct}_{2,R} \leftarrow \text{SFE.Enc}_{\text{dk}_2}(r_2)$.
5. C samples $s_2 \leftarrow U_{t_2}$, $u_2 \leftarrow U_\lambda$ and sends $\text{ct}_2 \leftarrow \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_2, (m||r)], \text{ct}_{2,R}; u_2)$ and s_2 , where $\text{Id}(\cdot)$ is the identity function.
6. C runs $\text{ZK.P}(x, w)$ and R runs $\text{ZK.V}(x)$ in an execution of ZK (with common input (1^λ)) for language \mathcal{L} (defined above), where $x = (c_1, \text{ct}_1, \text{ct}_2, \text{ct}_{1,R}, \text{ct}_{2,R}, \text{tag}, N)$ and $w = (m, r, s_1, s_2, u_1, u_2)$.

Figure 4: A constant round non-malleable commitment for tags in $[N]$, where $N = \text{ilog}(c+1, \lambda)$.

7.2.2 Analysis

In the reveal stage, the committer outputs (m, r) and the receiver accepts the decommitment if $c_1 = \text{Com}(1^\lambda, m; r)$. Perfect binding follows due to the perfect binding property of Com , and hiding follows by non-malleability, which we formally prove below.

Lemma 7.1. *Protocol 4 is a non-malleable commitment according to Definition 7.1 for tags in $[N]$.*

Proof. It suffices to show that for every $v_1, v_2 \in \{0, 1\}^{2p(\lambda)}$ and every QPT MIM = $\{\text{MIM}_\lambda, \rho_\lambda\}$ and $D = \{D_\lambda, \rho_\lambda\}$, there exists a negligible function $\eta(\cdot)$ such that for large enough $\lambda \in \mathbb{N}$, for all $\text{tag}_C, \text{tag}_{\text{MIM}} \in [\ell]$

where $\text{tag}_C \neq \text{tag}_{\text{MIM}}$, the following holds.

$$\left| \Pr[D_\lambda(\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v_1), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_C, \text{tag}_{\text{MIM}}), \sigma_\lambda) = 1] \right. \\ \left. - \Pr[D_\lambda(\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v_2), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_C, \text{tag}_{\text{MIM}}), \sigma_\lambda) = 1] \right| = \eta(\lambda)$$

To that end, we define the distributions

$$\{\mathcal{H}_{v_1, \sigma_\lambda} := (\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v_1), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_C, \text{tag}_{\text{MIM}}), \sigma_\lambda)\}_{\lambda \in \mathbb{N}}, \\ \{\mathcal{H}_{v_2, \sigma_\lambda} := (\text{View-Val}_{\text{MIM}_\lambda}(\mathcal{C}(v_2), \text{MIM}_\lambda(\rho_\lambda), \mathcal{R})(1^\lambda, \text{tag}_C, \text{tag}_{\text{MIM}}), \sigma_\lambda)\}_{\lambda \in \mathbb{N}}$$

We also define the following collections of random variables (each indexed by λ). Each is defined w.r.t. a (fixed) adversary $\text{MIM} = \{\text{MIM}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, but we sometimes drop this adversary from notation for convenience. We will also sometimes condition on the MIM aborting in the left execution (where it acts as receiver). By this, we will refer to an execution where the MIM sends a message that causes an honest party to abort.

- Let \Pr_x^{Abort} be the probability that MIM aborts in $\mathcal{H}_{x, \sigma_\lambda}$, where $x \in \{v_1, v_2\}$.
- Let $\mathcal{H}_{x, \sigma_\lambda}^{\text{No Abort}}$ be the distribution $\mathcal{H}_{x, \sigma_\lambda}$ conditioned on there not being an abort.
- Let $\mathcal{H}_{x, \sigma_\lambda}^{\text{Abort}}$ be the distribution $\mathcal{H}_{x, \sigma_\lambda}$ conditioned on there being an abort. Note that by definition, in this distribution, the value committed by the MIM is always set to \perp .

The following distributions will not be used explicitly in the hybrids, but will be convenient to define for the proof.

- Let $\mathcal{H}_{x, \sigma_\lambda, \perp}$ be the distribution $\mathcal{H}_{x, \sigma_\lambda}$ except whenever an abort occurs, the distribution outputs \perp .
- Let $\mathcal{H}_{x, \sigma_\lambda, \perp}^{\text{Abort}}$ be the distribution $\mathcal{H}_{x, \sigma_\lambda}$ except whenever an abort does not occur, the distribution outputs \perp .

We show that $\{\mathcal{H}_{v_1, \sigma_\lambda}\}_{\lambda \in \mathbb{N}} \approx_c \{\mathcal{H}_{v_2, \sigma_\lambda}\}_{\lambda \in \mathbb{N}}$ via a sequence of hybrids. In particular, we show that

$$\begin{aligned} \mathcal{H}_{v_1, \sigma_\lambda} &\stackrel{(1)}{=} (1 - \Pr_{v_1}^{\text{Abort}}) \mathcal{H}_{v_1, \sigma_\lambda}^{\text{No Abort}} + (\Pr_{v_1}^{\text{Abort}}) \mathcal{H}_{v_1, \sigma_\lambda}^{\text{Abort}} \\ &\stackrel{(2)}{\approx_s} (1 - \Pr_{v_2}^{\text{Abort}}) \mathcal{H}_{v_1, \sigma_\lambda}^{\text{No Abort}} + (\Pr_{v_2}^{\text{Abort}}) \mathcal{H}_{v_1, \sigma_\lambda}^{\text{Abort}} \\ &\stackrel{(3)}{\approx_c} (1 - \Pr_{v_2}^{\text{Abort}}) \mathcal{H}_{v_2, \sigma_\lambda}^{\text{No Abort}} + (\Pr_{v_2}^{\text{Abort}}) \mathcal{H}_{v_1, \sigma_\lambda}^{\text{Abort}} \\ &\stackrel{(4)}{\approx_c} (1 - \Pr_{v_2}^{\text{Abort}}) \mathcal{H}_{v_2, \sigma_\lambda}^{\text{No Abort}} + (\Pr_{v_2}^{\text{Abort}}) \mathcal{H}_{v_2, \sigma_\lambda}^{\text{Abort}} \\ &\stackrel{(5)}{=} \mathcal{H}_{v_2, \sigma_\lambda}, \end{aligned}$$

where

1. The equalities (1) and (5) follow by definition.
2. The indistinguishability (2) follows as a corollary of Claim 7.1. Indeed, $\mathcal{H}_{v_1, \sigma_\lambda, \perp} \approx_c \mathcal{H}_{v_2, \sigma_\lambda, \perp}$ in particular implies that the difference in the probability that the MIM aborts in both executions is negligible.
3. The indistinguishability (3) follows as a corollary of Claim 7.1. This can be seen by considering two cases. First, if the probability that the MIM aborts in $\mathcal{H}_{v_1, \sigma_\lambda}$ is negligible, then $\mathcal{H}_{v_1, \sigma_\lambda}^{\text{No Abort}} \approx_c \mathcal{H}_{v_2, \sigma_\lambda}^{\text{No Abort}}$ directly follows from Claim 7.1, and the indistinguishability follows. Otherwise, this probability is non-negligible, meaning that $\mathcal{H}_{v_1, \sigma_\lambda}^{\text{Abort}}$ is efficiently sampleable. Thus, a reduction to Claim 7.1 can sample from the distribution $\mathcal{H}_{v_1, \sigma_\lambda}^{\text{Abort}}$ whenever it receives \perp from its challenger.¹⁷

¹⁷A more formal analysis of this can be found in [BS20, Lemma 3.2].

4. The indistinguishability (4) follows as a corollary of Claim 7.5 via a similar analysis as the last step.

Claim 7.1.

$$\{\mathcal{H}_{v_1, \sigma_\lambda, \perp}\}_{\lambda \in \mathbb{N}} \approx_c \{\mathcal{H}_{v_2, \sigma_\lambda, \perp}\}_{\lambda \in \mathbb{N}}$$

Proof. We will prove this claim via the following sequence of hybrids. We set some notation before defining these hybrids. We will set $t_1 = \text{tag}_C \cdot (\log \eta)$ and $t_2 = (2N - \text{tag}_C) \cdot (\log \eta)$. We also set $t'_1 = \text{tag}_{\text{MIM}} \cdot (\log \eta)$ and $t'_2 = (2N - \text{tag}_{\text{MIM}}) \cdot (\log \eta)$. As a general rule, when referring to some protocol variable y in the left execution, we will use the variable as is (and denote it by y), and in the right execution, we will denote this variable by y' .

We let $\mathcal{H}_{v_1, \sigma_\lambda, \perp} = \mathcal{H}_0$.

\mathcal{H}_1 : In this hybrid, the challenger executes the simulator $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ ¹⁸ for Π_{zk} on V_λ^* , which denotes a wrapper around the portion of the MIM that participates in Step 6 of the protocol, and an instance-advice distribution $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ defined as follows:

- Set the state of MIM_λ to ρ_λ .
- Execute Steps 1-5 of the protocol the same way as in the experiment $\mathcal{H}_{v_1, \sigma_\lambda, \perp}$, and set (x, w, \mathcal{L}) according to Protocol 4 on behalf of \mathcal{C} .
- Let $\sigma_\lambda^{(x_\lambda)}$ denote the joint distribution of the protocol transcript, the state of the MIM at the end of Step 5, and the value v' committed by the MIM in Step 1.

If there is an abort during sampling, then output \perp . Otherwise, the output of this hybrid is the output of $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp$. By Lemma 6.2,

$$\mathcal{H}_0 \approx_c \mathcal{H}_1.$$

\mathcal{H}_2 : This is identical to \mathcal{H}_1 except the following change.

In Step 3, \mathcal{C} sends $\text{ct}_1 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_1, (0^{p(\lambda)+\lambda})], \text{ct}_{1,R}; u_1)$ and s_1 . Here $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ and V_λ^* are defined identically to \mathcal{H}_1 except with the updated ct_1 from Step 3, and the simulator $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ is executed. If there is an abort during sampling, then output \perp . Otherwise, the output of this hybrid is the output of $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp$. We prove in Claim 7.2, that

$$\mathcal{H}_1 \approx_s \mathcal{H}_2.$$

\mathcal{H}_3 : This is identical to \mathcal{H}_2 except the following change.

In Step 5, \mathcal{C} sends $\text{ct}_2 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_2, (0^{p(\lambda)+\lambda})], \text{ct}_{2,R}; u_2)$ and s_2 . Here $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ and V_λ^* are defined identically to \mathcal{H}_2 except with the updated ct_2 from Step 5, and the simulator $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ is executed. If there is an abort during sampling, then output \perp . Otherwise, the output of this hybrid is the output of $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp$. We prove in Claim 7.3, that

$$\mathcal{H}_2 \approx_s \mathcal{H}_3.$$

\mathcal{H}_4 : This is identical to \mathcal{H}_3 except the following change.

In Step 1, \mathcal{C} sets $c_1 = \text{Com}(1^\lambda, 0; r)$. Here $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ and V_λ^* are defined identically to \mathcal{H}_3 except with the updated c_1 from Step 1, and the simulator $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ is executed. If there is an abort during sampling, then output \perp . Otherwise, the output of this hybrid is the output of $\Pi_{\text{zk}}.\text{SimNoAbort}_\perp$. We prove in Claim 7.4, that

$$\mathcal{H}_3 \approx_c \mathcal{H}_4.$$

¹⁸Note that we drop the input I since there is only one verifier in this setting.

Claim 7.2.

$$\Delta(\mathcal{H}_1, \mathcal{H}_2) \leq 2^{-t_1} + \text{negl}(2^{t_1})$$

Proof. Note that the output of $\text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_1, (m||r)], \text{ct}_{1,R}; u_1)$ is identical in both hybrids, unless $s_1 = r_1$. Denote by \mathcal{H}'_1 the distribution that is identical to \mathcal{H}_1 except it outputs \perp when $s_1 = r_1$. Denote by \mathcal{H}'_2 the distribution that is identical to \mathcal{H}_2 except it outputs \perp when $s_1 = r_1$. Now by statistical circuit privacy, we have that there exists a constant $c > 0$ such that $\Delta(\mathcal{H}'_1, \mathcal{H}'_2) \leq 2^{-\lambda^c}$.

Finally, note that in each one of $\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}'_1, \mathcal{H}'_2$,

$$\Pr[s_1 = r_1] \leq 2^{-t_1}.$$

Thus we have,

$$\Delta(\mathcal{H}_1, \mathcal{H}_2) \leq \Delta(\mathcal{H}'_1, \mathcal{H}'_2) + 2 \cdot \Pr[s_1 = r_1] \leq 2^{-t_1} + 2 \cdot 2^{-\lambda^c} \leq 2^{-t_1} + \text{negl}(2^{t_1})$$

where the last equation follows by our setting of t_1 . □

Claim 7.3.

$$\Delta(\mathcal{H}_2, \mathcal{H}_3) \leq 2^{-t_2} + \text{negl}(2^{t_2})$$

Proof. The proof follows nearly identically to that of Claim 7.2. □

Claim 7.4.

$$\mathcal{H}_3 \approx_c \mathcal{H}_4$$

Proof. Throughout this proof, we will use the notation $\Pr[E|\mathcal{H}]$ to refer to the probability that event E occurs in the output of distribution \mathcal{H} .

Suppose

$$\Pr[\text{MIM aborts}|\mathcal{H}_3] = 1 - \text{negl}(\lambda).$$

Then by hiding of the commitment Com

$$\Pr[\text{MIM aborts}|\mathcal{H}_4] = 1 - \text{negl}(\lambda),$$

so both hybrids output \perp except with negligible probability, and are therefore computationally indistinguishable.

Thus for the rest of this proof, we will assume that there exists a polynomial $p(\cdot)$ such that:

$$\Pr[\text{MIM does not abort}|\mathcal{H}_3] \geq \frac{1}{p(\lambda)}.$$

By the hiding of the commitment Com,

$$\Pr[\text{MIM does not abort}|\mathcal{H}_4] \geq \frac{1}{p(\lambda)} - \text{negl}(\lambda).$$

For $x \in [0, 4]$, we will denote by $\mathcal{H}_x^{\text{No Abort}}$ the distribution \mathcal{H}_x conditioned on the MIM not aborting. Recall that whenever the MIM aborts, the two hybrids output \perp . Therefore, it suffices to prove that

$$\mathcal{H}_3^{\text{No Abort}} \approx_c \mathcal{H}_4^{\text{No Abort}}.$$

Now, recall that the variables r'_1, s'_1, t'_1 and so on refer to the *right* execution in each experiment. For $r'_1 \leftarrow U_{t'_1}$ and s'_1 sampled independently of r'_1 , we have that $\Pr[s'_1 = r'_1] = 2^{-t'_1}$. Then, by quantum input

privacy of SFE according to Definition 3.8, there exists a negligible function $\mu(\cdot)$ such that for any $x \in [0, 4]$, $r'_1 \leftarrow U_{t'_1}$ and s'_1 chosen by the MIM in \mathcal{H}_x ,

$$2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \leq \Pr[s'_1 = r'_1 | \mathcal{H}_x^{\text{No Abort}}] \leq 2^{-t'_1} + \mu(\lambda^{\text{ilog}(c,\lambda)}). \quad (6)$$

Similarly, by quantum input privacy of SFE according to Definition 3.8, there exists a negligible function $\mu'(\cdot)$ such that for any $x \in [0, 4]$, $r'_2 \leftarrow U_{t'_2}$ and s'_2 output by the MIM in \mathcal{H}_x ,

$$2^{-t'_2} - \mu'(\lambda^{\text{ilog}(c,\lambda)}) \leq \Pr[s'_2 = r'_2 | \mathcal{H}_x^{\text{No Abort}}] \leq 2^{-t'_2} + \mu'(\lambda^{\text{ilog}(c,\lambda)}) \quad (7)$$

Soundness of the ZK argument, together with setting $x = 0$ in equations (6) and (7) implies that for $i \in [2]$,

$$\begin{aligned} & \Pr \left[\text{SFE.Dec}_{\text{dk}'_i}(\text{ct}'_i) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (s'_i = r'_i) \middle| \mathcal{H}_0^{\text{No Abort}} \right] \\ & \geq (1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_i} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right). \end{aligned} \quad (8)$$

Note that for $i \in [2]$, $\text{SFE.Dec}_{\text{dk}'_i}(\text{ct}'_i) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (s'_i = r'_i)$ can be efficiently checked by a challenger that samples dk'_i and r'_i .

Therefore, by combining Lemma 6.2 with equation (8), we have that for $i \in [2]$,

$$\begin{aligned} & \Pr \left[\text{SFE.Dec}_{\text{dk}'_i}(\text{ct}'_i) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (s'_i = r'_i) \middle| \mathcal{H}_1^{\text{No Abort}} \right] \\ & \geq (1 - \text{negl}(\lambda) - \text{negl}(\lambda)) \cdot \left(2^{-t'_i} - \lambda^{\text{ilog}(c,\lambda)} \right) \\ & \geq (1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_i} - \lambda^{\text{ilog}(c,\lambda)} \right) \end{aligned} \quad (9)$$

where the previous equation, for $i \in [2]$, follows by considering a non-uniform reduction to Claim 6.2 that fixes any transcript where $\text{SFE.Dec}_{\text{dk}'_i}(\text{ct}'_i) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (r'_i = s'_i)$.

Next, we split our analysis into two cases. Depending on whether tag_C is greater or smaller than tag_{MIM} , one of the two cases will always be true.

- **Case 1:** $\text{tag}_C > \text{tag}_{\text{MIM}}$. In this case, $t_1 = \text{tag}_C \cdot (\log \eta)$, $t'_1 = \text{tag}_{\text{MIM}} \cdot (\log \eta)$.

Now combining Claim 7.2 and equation (9) with i set to 1, implies:

$$\begin{aligned} & \Pr \left[\text{SFE.Dec}_{\text{dk}'_1}(\text{ct}'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (s'_1 = r'_1) \middle| \mathcal{H}_2^{\text{No Abort}} \right] \\ & \geq (1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right) - 2^{-t_1} - \text{negl}(2^{t_1}) \end{aligned} \quad (10)$$

Next, we will carefully combine equation (10) with Claim 7.3. First, we note that the check $\text{SFE.Dec}_{\text{dk}'_1}(\text{ct}'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (r'_1 = s'_1)$ is performed before Step 5. Additionally, the only difference between $\mathcal{H}_2^{\text{No Abort}}$ and $\mathcal{H}_3^{\text{No Abort}}$ is in Step 5. As a result, for *every* (fixed) prefix of the transcript until Step 4, the distribution of Steps 5 and 6 generated according to $\mathcal{H}_2^{\text{No Abort}}$ is at most $2^{-t_2} + \text{negl}(2^{-t_2})$ -far from their distribution generated according to $\mathcal{H}_3^{\text{No Abort}}$. This implies:

$$\begin{aligned} & \Pr \left[\text{SFE.Dec}_{\text{dk}'_1}(\text{ct}'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (s'_1 = r'_1) \middle| \mathcal{H}_3^{\text{No Abort}} \right] \\ & \geq (1 - 2^{-t_2} - \text{negl}(2^{t_2})) \cdot \left((1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right) - 2^{-t_1} - \text{negl}(2^{t_1}) \right) \\ & \geq (1 - \text{negl}(\lambda) - 2^{-t_2} - \text{negl}(2^{t_2})) \cdot \left(2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right) - 2^{-t_1} + 2^{-t_1-t_2} - \text{negl}(2^{t_1}) \\ & \geq (1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right) - 2^{-t_1} - \text{negl}(2^{t_1}) \end{aligned} \quad (11)$$

Combining equation (11) with the $\text{negl}(\lambda^{\text{ilog}(c,\lambda)})$ -hiding of the commitment Com,

$$\begin{aligned} & \Pr \left[\text{SFE.Dec}_{dk'_1}(ct'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \wedge (s'_1 = r'_1) \middle| \mathcal{H}_4^{\text{No Abort}} \right] \\ & \geq (1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right) - 2^{-t_1} - \text{negl}(2^{t_1}) - \text{negl}(\lambda^{\text{ilog}(c,\lambda)}). \end{aligned} \quad (12)$$

Combining equations (11), (12) with equation (6), for $x \in [3, 4]$,

$$\begin{aligned} & \Pr \left[\text{SFE.Dec}_{dk'_1}(ct'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') \middle| (s'_1 = r'_1), \mathcal{H}_x^{\text{No Abort}} \right] \\ & \geq \frac{(1 - \text{negl}(\lambda)) \cdot \left(2^{-t'_1} - \mu(\lambda^{\text{ilog}(c,\lambda)}) \right) - 2^{-t_1} - \text{negl}(2^{t_1}) - \text{negl}(\lambda^{\text{ilog}(c,\lambda)})}{2^{-t'_1} + \mu(\lambda^{\text{ilog}(c,\lambda)})} \\ & \geq 1 - \text{negl}(\lambda) \end{aligned} \quad (13)$$

where the last equation follows by recalling that $2^{-t'_1} = \eta^{-\text{tag}_{\text{MIM}}}$, $2^{-t_1} = \eta^{-\text{tag}_c}$ and $\text{tag}_c \geq (\text{tag}_{\text{MIM}} + 1)$, which implies

$$2^{-t_1} = \eta^{-\text{tag}_c} \leq \eta^{-\text{tag}_{\text{MIM}} - 1} = \frac{2^{-t'_1}}{\eta}$$

for $\eta = \lambda^{\text{ilog}(c+1,\lambda)}$.

Let us assume towards a contradiction that there exists a quantum polynomial size distinguisher D and a polynomial poly such that for large enough $\lambda \in \mathbb{N}$:

$$\Pr[D = 1 | \mathcal{H}_3^{\text{No Abort}}] - \Pr[D = 1 | \mathcal{H}_4^{\text{No Abort}}] \geq \frac{1}{\text{poly}(\lambda)}. \quad (14)$$

By quantum input privacy of the SFE scheme, for each $x \in [3, 4]$,

$$|\Pr[D = 1 | (s'_1 \neq r'_1), \mathcal{H}_x^{\text{No Abort}}] - \Pr[D = 1 | (s'_1 = r'_1), \mathcal{H}_x^{\text{No Abort}}]| = \text{negl}(\lambda)$$

which combined with equation (14) implies that there exists a polynomial $q(\cdot)$ such that

$$\Pr[D = 1 | (s'_1 = r'_1), \mathcal{H}_3^{\text{No Abort}}] - \Pr[D = 1 | (s'_1 = r'_1), \mathcal{H}_4^{\text{No Abort}}] \geq \frac{1}{q(\lambda)} \quad (15)$$

Equations (13) and (15) together imply that:

$$\begin{aligned} & \Pr[D = 1 \wedge \text{SFE.Dec}_{dk'_1}(ct'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') | (s'_1 = r'_1), \mathcal{H}_3^{\text{No Abort}}] \\ & - \Pr[D = 1 \wedge \text{SFE.Dec}_{dk'_1}(ct'_1) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') | (s'_1 = r'_1), \mathcal{H}_4^{\text{No Abort}}] \\ & \geq \frac{1}{q(\lambda)} \cdot (1 - \text{negl}(\lambda)) \geq \frac{1}{2q(\lambda)} \end{aligned}$$

This, combined with equation (6) gives a distinguisher that distinguishes $\text{Com}(m)$ and $\text{Com}(0)$ with advantage at least $2^{-2t'_1}$, contradicting the hiding of Com as desired.

- **Case 2:** $\text{tag}_{\text{MIM}} > \text{tag}_c$. A similar analysis as in Case 1 implies that for any distinguisher D distinguishing $\mathcal{H}_3^{\text{No Abort}}$ from $\mathcal{H}_4^{\text{No Abort}}$, there exists a polynomial $q'(\cdot)$ such that

$$\begin{aligned} & \Pr[D = 1 \wedge \text{SFE.Dec}_{dk'_2}(ct'_2) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') | (s'_2 = r'_2), \mathcal{H}_3^{\text{No Abort}}] \\ & - \Pr[D = 1 \wedge \text{SFE.Dec}_{dk'_2}(ct'_2) \rightarrow (m', r') \text{ s.t. } c'_1 = \text{Com}(1^\lambda, m'; r') | (s'_2 = r'_2), \mathcal{H}_4^{\text{No Abort}}] \\ & \geq \frac{1}{q'(\lambda)} \end{aligned}$$

This, combined with equation (7) gives a distinguisher that distinguishes $\text{Com}(m)$ and $\text{Com}(0)$ with advantage at least $2^{-2t'_2}$, contradicting the hiding of Com as desired.

□

Claim 7.1 follows by observing that in \mathcal{H}_4 , we erased all information about either one of v_1 or v_2 . Therefore one can perform the above hybrids in reverse order, while arguing indistinguishability, until one ends up with $\mathcal{H}_{v_2, \sigma_\lambda, \perp}$. □

Claim 7.5.

$$\{\mathcal{H}_{v_1, \sigma_\lambda, \perp}^{\text{Abort}}\}_{\lambda \in \mathbb{N}} \approx_c \{\mathcal{H}_{v_2, \sigma_\lambda, \perp}^{\text{Abort}}\}_{\lambda \in \mathbb{N}}$$

Proof. Note that in this case, the value committed by the MIM is always \perp . Therefore, proving indistinguishability of these distributions is significantly more straightforward than in Claim 7.1. The proof again relies on a sequence of hybrid experiments, that we define below. Recall that when referring to some protocol variable y in the left execution, we will use the variable as is (and denote it by y), and in the right execution, we will denote this variable by y' .

We let $\mathcal{H}_{v_1, \sigma_\lambda, \perp}^{\text{Abort}} = \mathcal{H}_0$.

\mathcal{H}_1 : In this hybrid, the challenger executes one run of the simulator $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ ¹⁹ for Π_{zk} on \mathbf{V}_λ^* , which denotes a wrapper around the portion of the MIM that participates in Step 6 of the protocol, and an instance-advice distribution $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ defined as follows:

- Set the state of MIM_λ to ρ_λ .
- Execute Steps 1-5 of the protocol the same way as in the experiment $\mathcal{H}_{v_1, \sigma_\lambda, \perp}$, and set (x, w, \mathcal{L}) according to Protocol 4 on behalf of \mathcal{C} .
- Let $\sigma_\lambda^{(x_\lambda)}$ denote the joint distribution of the protocol transcript, the state of the MIM at the end of Step 5, and the value v' committed by the MIM in Step 1.

If there is an abort during sampling then output the transcript generated until the abort happens and the state of the MIM. Otherwise, the output of the hybrid is the output of $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$. By Lemma 6.3,

$$\mathcal{H}_0 \approx_c \mathcal{H}_1.$$

\mathcal{H}_2 : This is identical to \mathcal{H}_1 except the following change.

In Step 3, \mathcal{C} sends $\text{ct}_1 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_1, (0^{p(\lambda)+\lambda})], \text{ct}_{1,R}; u_1)$ and s_1 . Here $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ and \mathbf{V}_λ^* are defined identically to \mathcal{H}_1 except with the updated ct_1 from Step 3, and the simulator $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ is executed. If there is an abort during sampling then output the transcript generated until the abort happens and the state of the MIM. Otherwise, the output of the hybrid is the output of $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$. We prove in Claim 7.6, that

$$\mathcal{H}_1 \approx_s \mathcal{H}_2.$$

\mathcal{H}_3 : This is identical to \mathcal{H}_2 except the following change.

In Step 5, \mathcal{C} sends $\text{ct}_2 = \text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_2, (0^{p(\lambda)+\lambda})], \text{ct}_{2,R}; u_2)$ and s_2 . Here $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ and \mathbf{V}_λ^* are defined identically to \mathcal{H}_2 except with the updated ct_2 from Step 5, and the simulator $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ is executed. If there is an abort during sampling then output the transcript generated until the abort happens and the state of the MIM. Otherwise, the output of the hybrid is the output of $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$. We prove in Claim 7.7, that

$$\mathcal{H}_2 \approx_s \mathcal{H}_3.$$

\mathcal{H}_4 : This is identical to \mathcal{H}_3 except the following change.

¹⁹Note that we drop the input I since there is only one verifier in this setting.

In Step 1, C sets $c_1 = \text{Com}(1^\lambda, 0; r)$. Here $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ and V_λ^* are defined identically to \mathcal{H}_3 except with the updated c_1 from Step 1, and the simulator $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ is executed. If there is an abort during sampling then output the transcript generated until the abort happens and the state of the MIM. Otherwise, the output of the hybrid is the output of $\Pi_{\text{zk}}.\text{SimAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$. We prove in Claim 7.8, that

$$\mathcal{H}_3 \approx_c \mathcal{H}_4.$$

Claim 7.6.

$$\Delta(\mathcal{H}_1, \mathcal{H}_2) \leq 2^{-t_1} + \text{negl}(2^{t_1})$$

Proof. Note that the output of $\text{SFE.Eval}(\text{CC}[\text{Id}(\cdot), s_1, (m||r)], \text{ct}_{1,R}; u_1)$ is identical in both hybrids, unless $s_1 = r_1$. Denote by \mathcal{H}'_1 the distribution that is identical to \mathcal{H}_1 except it outputs \perp when $s_1 = r_1$. Denote by \mathcal{H}'_2 the distribution that is identical to \mathcal{H}_2 except it outputs \perp when $s_1 = r_1$. Now by statistical circuit privacy, we have that there exists a constant $c > 0$ such that $\Delta(\mathcal{H}'_1, \mathcal{H}'_2) \leq 2^{-\lambda^c}$.

Finally, note that in each one of $\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}'_1, \mathcal{H}'_2$,

$$\Pr[s_1 = r_1] \leq 2^{-t_1}.$$

Thus we have,

$$\Delta(\mathcal{H}_1, \mathcal{H}_2) \leq \Delta(\mathcal{H}'_1, \mathcal{H}'_2) + 2 \cdot \Pr[s_1 = r_1] \leq 2^{-\lambda^c} + 2 \cdot 2^{-t_1} \leq 2^{-t_1} + \text{negl}(2^{t_1})$$

where the last equation follows by our setting of t_1 . □

Claim 7.7.

$$\Delta(\mathcal{H}_2, \mathcal{H}_3) \leq 2^{-t_2} + \text{negl}(2^{t_2})$$

Proof. The proof follows nearly identically to that of Claim 7.2. □

Claim 7.8.

$$\mathcal{H}_3 \approx_c \mathcal{H}_4$$

Proof. Recall that both hybrids only output transcripts, the value committed by the MIM, and the MIM's state when the MIM aborts at some point. Otherwise both hybrids output \perp . As such, the value committed by the MIM in every transcript in the distributions \mathcal{H}_3 , and \mathcal{H}_4 , is \perp .

It remains to prove that the joint distribution of the transcript and state of the MIM in \mathcal{H}_3 is indistinguishable from \mathcal{H}_4 , which follows immediately by the hiding of Com . □

Claim 7.5 follows by observing that in \mathcal{H}_4 , we erased all information about either one of v_1 or v_2 . Therefore one can perform the above hybrids in reverse order, while arguing indistinguishability, until one ends up with $\mathcal{H}_{v_2, \sigma_\lambda, \perp}^{\text{Abort}}$. □

This concludes the proof of the lemma. □

7.3 Tag Amplification

Beginning with quantum-secure non-malleable commitments that support tags in $[t]$, we describe how to obtain quantum-secure non-malleable commitments that support tags in $[2^{t/2}]$, for any $t \leq \text{poly}(\lambda)$. We will apply this compiler recursively a constant number of times to our base construction to obtain a scheme for tags in $[2^\lambda]$.

The protocol itself follows nearly identically along the lines of existing tag amplification compilers in the literature [Wee10, PW10]. Each larger tag (in $2^{t/2}$) is encoded into a set S_{tag} of t small tags, and this set satisfies the following property. For each pair $\text{tag}, \text{tag}' \in 2^{t/2}$ where $\text{tag} \neq \text{tag}'$, there exists an element in the set $S_{\text{tag}'}$ that does not lie in the set S_{tag} . Now the committer on input a message m and tag tag , generates t commitments, where the i^{th} commitment commits via the underlying non-malleable commitment to the message m using as tag the i^{th} member of the set S_{tag} . The committer then proves to the receiver (in zero-knowledge) that all the commitments were correctly generated according to protocol specifications.

Now, the property of the tag encoding scheme guarantees that the MIM will always end up using *at least one small tag* that is different from all small tags used by the honest committer. As such, we can argue that the value committed by the MIM using this differing tag is independent of the honest committer's input.

We formally prove this via a hybrid argument, where we first rely on soundness of the ZK argument to argue that we can “focus” solely on the values committed by the MIM using the differing tag. Next, we simulate the ZK argument on behalf of the honest committer. Then we modify the values committed via all the honest small tags, while arguing that the value committed by the MIM under the differing tag does not change.

Finally, we point out an interesting feature of our amplification proof: we split the use of the ZK simulator into two cases: one simulator that only outputs non-aborting views (and otherwise outputs \perp), and a separate one that only outputs aborting views (and otherwise outputs \perp).

7.3.1 Construction

We describe our compiler that converts a tag-based non-malleable commitment scheme for tags in $t(\lambda)$, where $t(\lambda) \leq \text{poly}(\lambda)$, into one that supports tags in $2^{t/2}$, while adding only a constant number of rounds. This is formalized in Protocol 5. We will let nmCom denote a non-malleable commitment for tags in $t(\lambda)$, and we denote the message length by $p(\lambda)$.

Ingredients and notation: We will assume the existence of a quantum-secure zero-knowledge argument for NP (ZK.P, ZK.V). (We do not require multi-verifier zero-knowledge for this section.)

7.3.2 Analysis.

In the reveal stage, the committer outputs $(m, r_{1,C})$ and the receiver accepts the decommitment if this produces a valid decommitment of c_1 according to nmCom . The perfect binding property of the scheme in Protocol 5 follows directly from the perfect binding property of the underlying protocol nmCom . Hiding follows from non-malleability, the proof of which is in Appendix C.

Lemma 7.2. *Protocol 5 is a one-one non-malleable commitment according to Definition 7.1 for tags in $[2^{t/2}]$.*

We also have the following lemma, that follows by a standard hybrid argument, due to [LPV08].

Lemma 7.3. *Every quantum-secure non-malleable commitment satisfying Definition 7.1 also satisfies Definition 7.2.*

We conclude this section with the following theorem, that can be obtained by applying the compiler in Protocol 5 ($c + 1$) times to the base non-malleable commitment from Protocol 4, and then applying Lemma 7.3.

Theorem 7.1. *Assuming there exists a constant $c \in \mathbb{N}$ such that at all quantum polynomial size circuits have advantage $\text{negl}(\lambda^{\log(c,\lambda)})$ in distinguishing LWE samples from uniform, and spooky encryption for relations computable by quantum circuits, there exist quantum-secure constant round non-malleable commitments satisfying Definition 7.2.*

Protocol 5

Common Input: Security parameter 1^λ , $\text{tag} \in [2^{t/2}]$ represented as $\{\text{tag}_i\}_{i \in [t/2]}$. Here $\text{tag}_i = i \parallel \text{tag}[i]$ where $\text{tag}[i]$ denotes the i^{th} bit of tag .

C's Input: A string $m \in \{0, 1\}^{p(\lambda)}$.

Commit Stage:

- **Stage 1:** In parallel, for all $i \in [k]$, C runs $\text{nmCom.C}(m; r_{i,C})$ and R runs $R(r_{i,R})$ with common input tag_i . Let c_i and ϕ_i denote the set of all messages generated by C and R respectively in the i^{th} parallel execution.
- **Stage 2:** C executes $\text{ZK.P}(x, w, \mathcal{L})$ and R executes $\text{ZK.V}(x, \mathcal{L})$ where:
 - $x = \{c_i\}_{i \in [t/2]}, w = (m, \{r_{i,C}\}_{i \in [t/2]})$,
 - $\mathcal{L} = \{ \{b_i\}_{i \in [t/2]} \mid \exists (a, \{s_i\}_{i \in [t/2]}) \text{ s.t. } \forall i \in [t/2], \text{nmCom.C}(a, \text{tag}_i, \phi_i; s_i) = b_i \}$
 - where $\text{nmCom.C}(a, \text{tag}_i, \phi_i; s_i)$ denotes the transcript output by C_i on input receiver messages ϕ_i .

Figure 5: A constant round non-malleable commitment for tags in $[2^{t/2}]$.

Finally, we remark that a folklore technique [DDN91] where the committer and receiver participate in rounds, sending \perp in every round, except for the committer sending $\text{ECom}(m; r)$ in round i (where $i = \text{tag}$) using a single-committer extractable commitment yields a one-to-one non-malleable commitment for $\text{tag} \in [N]$ in the synchronous setting, in $O(N)$ rounds, for any $N \leq \text{poly}(\lambda)$. Setting $N = \text{ilog}(c, \lambda)$ for any constant $c \in \mathbb{N}$ yields a protocol with $O(\text{ilog}(c, \lambda))$ rounds for $\text{ilog}(c, \lambda)$ tags. Applying our tag amplification compiler to this scheme $(c+1)$ times, yields a non-malleable commitment for tags in 2^λ against synchronous adversaries, in $O(\text{ilog}(c, \lambda))$ rounds. The underlying extractable commitment can be instantiated using the technique of [BS20] based on polynomial quantum hardness of LWE and polynomial hardness of QFHE. This yields $O(\text{ilog}(c, \lambda))$ round post-quantum non-malleable commitments from polynomial hardness assumptions.

8 Quantum-Secure Multi-Party Coin-Flipping

We now combine the primitives constructed in earlier sections to build a constant-round coin-flipping protocol secure against quantum polynomial-time adversaries. The protocol was described at a high level in Section 2.5, and is given in full detail in Protocol 6. As explained in the overview, each party will first commit to random strings c_i, r_i using a non-malleable commitment, then commit to c_i using our parallel extractable commitment with randomness r_i , and finally broadcast c_i . The parties will output $\bigoplus_{i \in [n]} c_i$ as the common output if all parties manage to prove in zero-knowledge that they behaved honestly throughout the protocol.

Proof Strategy. Our simulator will be structurally similar to the zero-knowledge simulator described in Section 6, in the sense that we build a simulator SimNoAbort_\perp specifically for non-aborting transcripts and a simulator SimAbort_\perp specifically for aborting transcripts. The bulk of the work in SimNoAbort_\perp involves sampling instances and an advice state (consisting of the adversary's view through Step 5 of the protocol) for the final part of the adversary, which in particular interacts with honest parties in order to verify their zero-knowledge arguments in Step 6. These arguments are then simulated by (part of) the zero-knowledge simulator $\text{ZK.SimNoAbort}_\perp$, which takes as input the adversary and the sampled advice state.

However, as alluded to in Section 2.5, once this simulation is performed in the hybrids, it is no longer possible to directly invoke the soundness of the adversary’s zero-knowledge arguments in Step 6, when changing how SimNoAbort_\perp samples the advice state. Thus, we invoke soundness in the very first hybrid to claim that the following “check” never fails, except with negligible probability.

The check fails if the Step 1-5 messages sent by at least one of the malicious parties are not explainable, yet the honest parties do not abort. If this check fails, we simply append check-fail to the transcript. Now, this check will continue to be computed in later hybrids, but we can claim that since (as we show) all later hybrids are indistinguishable from the first hybrid, check-fail must also only appear with negligible probability in these later hybrids. When it comes time to invoke the non-malleability of the honest party commitments, we can use the fact that check-fail appears with negligible probability to show that a malicious party cannot even change its *extractable* commitment based on the changing simulated view. If it could, then since non-malleability implies that its previously sent non-malleable commitment couldn’t change, then it must be the case that its messages are no longer explainable (since its two commitments are no longer consistent). Thus, the check will fail and appear in the hybrid’s output, a contradiction. Of course, turning this intuition into a formal proof requires much care, especially since this check is inefficient. Thus, we will make liberal use of non-uniform fixing arguments.

We also remark that it would be most natural to rely on *many-to-many* non-malleable commitments in this multi-party setting. However, we only have a post-quantum construction of one-one commitments in Section 7. Thus, when invoking non-malleability, a reduction must isolate the commitment of a single malicious party that would constitute a mauling attack. We again use both non-uniformity and the check-fail condition here, showing that any mauling attack would cause check-fail to appear, and thus that there must exist *some* malicious party for which the check fails over specifically its messages. The identity of this party can then be given as non-uniform advice to a reduction.

8.1 Definition

Definition 8.1 (Quantum-Secure Fully-Simulatable Multi-Party Coin-Flipping). *Let $k = k(\lambda)$ be any fixed polynomial. An fully-simulatable n -party k -coin-flipping protocol with quantum security is given by n classical interactive Turing machines (P_1, \dots, P_n) with joint input $(1^\lambda, 1^n)$ and outputs $r_i \in \{0, 1\}^{k(\lambda)} \cup \{\perp\}$.*

Given a coin-flipping protocol and an adversary $A^ = \{A_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ that corrupts a set of parties $\mathbb{S} \subset [n]$, let \mathbb{H} denote $[n] \setminus \mathbb{S}$ and define the random variable $\text{Real}(A_\lambda^*, \rho_\lambda)$ to consist of the outputs of honest parties $\{P_i\}_{i \in \mathbb{H}}$ as well as the view $\text{VIEW}_{A_\lambda^*}(A_\lambda^*(\rho_\lambda), \{P_i\}_{i \in \mathbb{H}})(1^\lambda, 1^n)$ of A_λ^* after executing the protocol in the presence of A_λ^* .*

We require the following security property. Fix any $\mathbb{S} \subset [n]$. There exists a quantum expected polynomial-time simulator Sim , such that for any quantum polynomial-size adversary $A^ = \{A_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ that participates in the protocol, generating joint messages on behalf of all algorithms in \mathbb{S} ,*

$$\{\text{Real}(A_\lambda^*, \rho_\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Sim}(1^\lambda, 1^n, r, A_\lambda^*, \rho_\lambda) \mid r \leftarrow U_{k(\lambda)}\}_{\lambda \in \mathbb{N}}, \quad (16)$$

where for the “protocol output” part of its simulation output, $\text{Sim}(1^\lambda, 1^n, r, A_\lambda^, \rho_\lambda)$ is restricted to output either $r^{|\mathbb{H}|}$ or $\perp^{|\mathbb{H}|}$.*

8.2 Construction

Ingredients: All of the following are assumed to be quantum-secure.

- A many-to-one non-malleable commitment $\text{nmCom} = (\text{nmCom.C}, \text{nmCom.R})$.
- A multi-committer extractable commitment $\text{eCom} = (\text{eCom.C}, \text{eCom.R})$.
- A multi-verifier publicly-verifiable zero-knowledge argument for NP $\text{ZK} = (\text{ZK.P}, \text{ZK.V})$.

Languages. We define two NP languages $\mathcal{L}^{(\text{eCom})}$ and $\mathcal{L}^{(\text{both})}$. Let $(x, y) := \langle \text{nmCom.C}(c; r), \text{nmCom.R}_y \rangle$ denote the transcript of an execution of nmCom where the receiver messages are fixed to y and x is the set

of resulting committer messages. Similarly, let $(x, y) := \langle \text{eCom.C}(c; r), \text{eCom.R}_y \rangle$ denote the transcript of an execution between some eCom.C and eCom.R (which may be part of a larger eCom transcript involving other committers) where the receiver messages are fixed to y and x is the set of resulting committer messages. Then $\mathcal{L}^{(\text{eCom})}$ and $\mathcal{L}^{(\text{both})}$ are defined as follows.

$$\begin{aligned} \mathcal{L}^{(\text{eCom})} &:= \left\{ (x, y) \mid \exists (c, r) \text{ s.t. } (x, y) := \langle \text{eCom.C}(c; r), \text{eCom.R}_y \rangle \right\} \\ \mathcal{L}^{(\text{both})} &:= \left\{ (x, y, x', y', c) \mid \exists (r, s) \text{ s.t. } \begin{aligned} (x, y) &:= \langle \text{nmCom.C}((c, r); s), \text{nmCom.R}_y \rangle, \\ (x', y') &:= \langle \text{eCom.C}(c; r), \text{eCom.R}_{y'} \rangle \end{aligned} \right\} \end{aligned}$$

Protocol 6

Common input: Security parameter 1^λ and number of parties 1^n .

1. For all $i \in [n]$, P_i samples $c_i \leftarrow \{0, 1\}^{k(\lambda)}$, $\{r_{i,j}, s_{i,j}\}_{j \in [n] \setminus \{i\}} \leftarrow \{0, 1\}^{2(n-1)\lambda}$.
2. For all $i \in [n]$, $j \in [n] \setminus \{i\}$, P_i runs $\text{nmCom.C}((c_i, r_{i,j}); s_{i,j})$ with tag i and P_j runs nmCom.R to produce

$$\left(\alpha_{i,j}^{(c)}, \alpha_{i,j}^{(r)} \right) \leftarrow \langle \text{nmCom.C}((c_i, r_{i,j}); s_{i,j}), \text{nmCom.R} \rangle(1^\lambda),$$

where $\alpha_{i,j}^{(c)}$ denotes the committer messages sent by P_i and $\alpha_{i,j}^{(r)}$ denotes the receiver messages sent by P_j .

3. For all $j \in [n]$, P_j runs eCom.R and each P_i for $i \neq j$ runs $\text{eCom.C}_i(c_i; r_{i,j})$ to produce

$$\left(\left\{ \beta_{i,j}^{(c)} \right\}_{i \in [n] \setminus \{j\}}, \left\{ \beta_{i,j}^{(r)} \right\}_{i \in [n] \setminus \{j\}} \right) \leftarrow \langle \{ \text{eCom.C}_i(c_i; r_{i,j}) \}_{i \in [n] \setminus \{j\}}, \text{eCom.R} \rangle(1^\lambda, 1^{n-1})$$

where $\beta_{i,j}^{(c)}$ denotes the committer messages sent by P_i and $\beta_{i,j}^{(r)}$ denotes the receiver messages sent by P_j .

4. For all $i \in [n]$, P_i runs $\text{ZK.P}(\{x_{i,j}, w_{i,j}\}_{j \in [n] \setminus \{i\}})$ and each $\{P_j\}_{j \in [n] \setminus \{i\}}$ runs $\text{ZK.V}_j(x_{i,j})$ in an execution of ZK with common input $(1^\lambda, 1^{n-1})$ for language $\mathcal{L}^{(\text{eCom})}$ (defined above), where $x_{i,j} = (\beta_{i,j}^{(c)}, \beta_{i,j}^{(r)})$ and $w_{i,j} = (c_i, r_{i,j})$.
5. For all $i \in [n]$, P_i broadcasts c_i .
6. For all $i \in [n]$, P_i runs $\text{ZK.P}(\{x_{i,j}, w_{i,j}\}_{j \in [n] \setminus \{i\}})$ and each $\{P_j\}_{j \in [n] \setminus \{i\}}$ runs $\text{ZK.V}_j(x_{i,j})$ in an execution of ZK with common input $(1^\lambda, 1^{n-1})$ for language $\mathcal{L}^{(\text{both})}$ (defined above), where $x_{i,j} = (\alpha_{i,j}^{(c)}, \alpha_{i,j}^{(r)}, \beta_{i,j}^{(c)}, \beta_{i,j}^{(r)}, c_i)$ and $w_{i,j} = (r_{i,j}, s_{i,j})$.
7. For all $i \in [n]$, P_i runs the ZK verification algorithm on all $2n(n-1)$ proofs provided in Steps 4 and 6. If all proofs are accepting, P_i outputs $\bigoplus_{i \in [n]} c_i$, and otherwise outputs \perp .

Figure 6: A quantum-secure constant-round coin-flipping protocol.

8.3 Security

Theorem 8.1. *For any n and polynomial $k(\lambda)$, Protocol 6 is a quantum-secure fully-simulatable n -party k -coin-flipping protocol.*

Proof. Fix a number of parties n , a polynomial $k = k(\lambda)$, and a set $\mathbb{S} \subset [n]$ of malicious parties. We construct a simulator Sim that for every quantum polynomial-size adversary $A^* = \{A_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ corrupting parties in \mathbb{S} , outputs a distribution

$$\{\text{Sim}(1^\lambda, 1^n, r, A_\lambda^*, \rho_\lambda) \mid r \leftarrow U_{k(\lambda)}\}_{\lambda \in \mathbb{N}}$$

that satisfies the conditions of Definition 8.1. Similar to the proof of zero-knowledge in Section 6.4, the simulator Sim will make use of two sub-routines, SimNoAbort_\perp and SimAbort_\perp .

$\text{SimNoAbort}_\perp(1^\lambda, 1^n, r, A_\lambda^*, \rho_\lambda)$:

1. Let i^* denote the smallest index of a party in \mathbb{H} . Define the machine V_λ^* as follows. V_λ^* will act on behalf of verifiers $\{\text{ZK.V}_j\}_{j \in \mathbb{S}}$ in the ZK session in Step 6 of Protocol 6 where party P_{i^*} is the prover. Thus, it consists of the portion of A_λ^* that interacts during this step *as well as* the portion of the honest parties \mathbb{H} that interact in the $n - 1$ sessions where P_{i^*} is *not* the prover.

We will next describe how a particular instance-advice distribution $(\{x_j\}_{j \in \mathbb{S}}, \sigma^{\{x_j\}_{j \in \mathbb{S}}})$ is generated for V_λ^* . Generating this distribution will involve simulating Steps 1-5 of Protocol 6 for adversary A_λ^* . In particular, the advice state $\sigma^{\{x_j\}_{j \in \mathbb{S}}}$ will include the transcript $\tau^{(5)}$ of the entire simulated execution through Step 5, the inner state $\rho^{(5)}$ of A_λ^* at this point, as well as the witnesses $\{r_{i,j}, s_{i,j}\}_{i \in \mathbb{H} \setminus \{i^*\}, j \in [n] \setminus \{i\}}$ to be used by parties $\{P_i\}_{i \in \mathbb{H} \setminus \{i^*\}}$ in Step 6 of Protocol 6. The instances $\{x_j\}_{j \in \mathbb{S}}$ will be a subset of $\tau^{(5)}$.

In particular, for each $j \in \mathbb{S}$, x_j will be set to $(\alpha_{i^*,j}^{(c)}, \alpha_{i^*,j}^{(r)}, \beta_{i^*,j}^{(c)}, \beta_{i^*,j}^{(r)}, c'_{i^*})$, which are the messages exchanged by i^* and j during Steps 2 and 3 when P_{i^*} was acting as a committer, as well the value c'_{i^*} broadcast by P_{i^*} in Step 5.

This instance-advice distribution is generated as follows.

- (a) For each party $\{P_i\}_{i \in \mathbb{H} \setminus \{i^*\}}$, sample $c_i, \{r_{i,j}, s_{i,j}\}_{j \in [n] \setminus \{i\}}$ as in Step 1 of Protocol 6.
- (b) Set $\rho^{(1)} := \rho_\lambda$ to be the inner state of A_λ^* , and interact with A_λ^* to run Step 2 of Protocol 6 honestly, with the only difference being that party P_{i^*} commits to $(0^{k(\lambda)}, 0^\lambda)$. Let $\rho^{(2)}$ be the resulting inner state of A_λ^* , and let $\{\alpha_{i^*,j}^{(c)}, \alpha_{i^*,j}^{(r)}\}_{j \in \mathbb{S}}$ be the messages sent between P_{i^*} and \mathbb{S} in commitments where P_{i^*} was the committer.
- (c) Define the machine C_λ^* as follows. C_λ^* will act on behalf of committers $\{\text{eCom.C}_j\}_{j \in \mathbb{S}}$ in the eCom session in Step 3 of Protocol 6 where party P_{i^*} is the receiver. Thus, it consists of the portion of A_λ^* that interacts during this step *as well as* the portion of the honest parties \mathbb{H} that interact in the $n - 1$ sessions where P_{i^*} is *not* the receiver. The advice state σ_λ given to C_λ^* will include the transcript $\tau^{(2)}$ of the execution so far, the inner state $\rho^{(2)}$ of A_λ^* at this point, and the messages and randomness $\{c_i, r_{i,j}\}_{i \in \mathbb{H}, j \in [n] \setminus \{i\}}$ to be used in the commitments by honest players in Step 3, where $c_{i^*} = 0^{k(\lambda)}$ and $r_{i^*,j}$ is uniformly and independently sampled from the rest of the transcript. The view of C_λ^* at the end of this interaction includes the updated execution transcript $\tau^{(3)}$ as well as the updated inner state $\rho^{(3)}$ of A_λ^* .

Now, compute

$$(\{\tau_j\}_{j \in \mathbb{S}}, \text{st}, \{c'_j\}_{j \in \mathbb{S}}) \leftarrow \text{eCom.Ext}(1^\lambda, 1^{n-1}, \mathbb{S}, C_\lambda^*, \sigma_\lambda),$$

and parse $(\{\tau_j\}_{j \in \mathbb{S}}, \text{st})$ to obtain $\tau^{(3)}$ and $\rho^{(3)}$, where $\tau^{(3)}$ in particular includes the messages $\{\beta_{i^*,j}^{(c)}, \beta_{i^*,j}^{(r)}\}_{j \in \mathbb{S}}$ exchanged by P_{i^*} and \mathbb{S} in commitments where some party $j \in \mathbb{S}$ was the receiver. If eCom.Ext produced an abort transcript, then return \perp , and otherwise continue, setting $\rho^{(3)}$ to be the state of A_λ^* .

- (d) Interact with A_λ^* to run Steps 4 and 5 of the protocol honestly, with the only difference being that party P_{i^*} broadcasts $c'_{i^*} := \bigoplus_{j \in \mathbb{S}} c'_j \bigoplus_{j \in \mathbb{H} \setminus \{i^*\}} c_j \oplus r$ in Step 5.
- (e) If there exists $j \in \mathbb{S}$ such that the output c_j of P_j in Step 5 is not equal to c'_j , then return \perp . Otherwise, let $\tau^{(5)}$ denote the transcript so far and let $\rho^{(5)}$ denote the state of A_λ^* at the end of Step 5. For each $j \in \mathbb{S}$ set $x_j = (\alpha_{i^*,j}^{(c)}, \alpha_{i^*,j}^{(r)}, \beta_{i^*,j}^{(c)}, \beta_{i^*,j}^{(r)}, c'_{i^*})$, and set $\sigma^{\{x_j\}_{j \in \mathbb{S}}} = (\tau^{(5)}, \rho^{(5)}, \{r_{i,j}, s_{i,j}\}_{i \in \mathbb{H} \setminus \{i^*\}, j \in [n] \setminus \{i\}})$.

2. Now, compute

$$\text{VIEW} \leftarrow \text{ZK.SimNoAbort}_{\perp}(1^{\lambda}, 1^{n-1}, \mathbb{S}, \{x_j\}_{j \in \mathbb{S}}, V_{\lambda}^*, \sigma^{\{x_j\}_{j \in \mathbb{S}}}).$$

If the proofs (which are included in VIEW) output by A_{λ}^* are accepting, then output VIEW, and otherwise output \perp .

$\text{SimAbort}_{\perp}(1^{\lambda}, 1^n, r, A_{\lambda}^*, \rho_{\lambda})$:

1. Set ρ_{λ} as the initial state of A_{λ}^* . Execute Steps 1-6 with A_{λ}^* using honest party strategy according to Protocol 6 on behalf of parties $\{P_j\}_{j \in \mathbb{H}}$. If sampling this distribution leads to an abort at any point, halt and return the view of A_{λ}^* . Otherwise, return \perp .

$\text{SimComb}_{\perp}(1^{\lambda}, 1^n, I, A_{\lambda}^*, \rho_{\lambda})$: With probability 1/2, execute $\text{SimNoAbort}_{\perp}(1^{\lambda}, 1^n, I, \{x_i\}_{i \in [n]}, V_{\lambda}^*, \rho_{\lambda})$ and otherwise execute $\text{SimAbort}_{\perp}(1^{\lambda}, 1^n, I, A_{\lambda}^*, \rho_{\lambda})$.

$\text{Sim}(1^{\lambda}, 1^n, I, A_{\lambda}^*, \rho_{\lambda})$: Let $\overline{\text{SimComb}}_{\perp}(\cdot) := \text{SimComb}_{\perp}(1^{\lambda}, 1^n, I, A_{\lambda}^*, \cdot)$ be the circuit SimComb_{\perp} with all inputs hard-coded except for ρ_{λ} , and output $R(\overline{\text{SimComb}}_{\perp}, \rho_{\lambda}, \lambda)$, where R is the algorithm from Lemma 3.1.

This concludes the description of the simulator. Before proceeding to the proof of indistinguishability, we define the following collections of random variables (each indexed by λ). Each is defined with respect to the adversary $A^* = \{A_{\lambda}^*, \rho_{\lambda}\}_{\lambda \in \mathbb{N}}$ that we are considering. Throughout, whenever we say abort, we mean that either one of the parties controlled by the adversary aborts, or it fails to prove one of its statements.

- Let $\text{SimNoAbort}_{\perp}(A^*) := \{\text{SimNoAbort}_{\perp}(1^{\lambda}, 1^n, r, A_{\lambda}^*, \rho_{\lambda})\}_{\lambda \in \mathbb{N}}$.
- Let $\text{SimAbort}_{\perp}(A^*) := \{\text{SimAbort}_{\perp}(1^{\lambda}, 1^n, r, A_{\lambda}^*, \rho_{\lambda})\}_{\lambda \in \mathbb{N}}$.
- Let $\text{SimComb}_{\perp}(A^*) := \{\text{SimComb}_{\perp}(1^{\lambda}, 1^n, r, A_{\lambda}^*, \rho_{\lambda})\}_{\lambda \in \mathbb{N}}$.
- Let $\text{Sim}(V^*) := \{\text{Sim}(1^{\lambda}, 1^n, r, A_{\lambda}^*, \rho_{\lambda})\}_{\lambda \in \mathbb{N}}$.
- Let $\text{RealNoAbort}_{\perp}(A^*)$ be the distribution $\text{Real}(A^*)$, except that whenever an abort occurs, the distribution outputs \perp .
- Let $\text{RealAbort}_{\perp}(A^*)$ be the distribution $\text{Real}(A^*)$, except that if an abort *does not* occur, the distribution outputs \perp .
- Let $\text{SimComb}(A^*)$ be the distribution SimComb_{\perp} conditioned on the output not being \perp .

Next, we prove the following claim.

Claim 8.1.

$$\text{RealNoAbort}_{\perp}(A^*) \approx_c \text{SimNoAbort}_{\perp}(A^*)$$

Proof. This can be proved via the following sequence of hybrids.

- \mathcal{H}_0 : $\text{RealNoAbort}_{\perp}(A^*)$.
- \mathcal{H}_1 : This hybrid is the same as \mathcal{H}_0 , except that it attaches check-fail to the output if the following (inefficient check) on Steps 1-4 of the transcript, fails.

Let i^* denote the smallest index in \mathbb{H} . For $j \in \mathbb{S}$, let $y_j := (\alpha_{j,i^*}^{(c)}, \alpha_{j,i^*}^{(r)}, \beta_{j,i^*}^{(c)}, \beta_{j,i^*}^{(r)}, c_j)$ be the messages exchanged between P_{i^*} and P_j in Steps 2 and 3 when P_j was acting as the committer, along with the message broadcast by P_j in Step 4. The check fails if there exists a $j \in \mathbb{S}$ such that $y_j \notin \mathcal{L}^{(\text{both})}$.

- \mathcal{H}_2 : Let V_{λ}^* be the machine defined in the description of $\text{SimNoAbort}_{\perp}$. Sample instance-advice distribution $(\{x_j\}_{j \in \mathbb{S}}, \sigma^{\{x_j\}_{j \in \mathbb{S}}})$ as described below.

1. Execute Steps 1-4 of the protocol identically to \mathcal{H}_2 . Let $\tau^{(4)}$ denote the transcript generated so far, let $\rho^{(4)}$ denote the state of A_λ^* at the end of Step 4, and let $\{r_{i,j}, s_{i,j}\}_{i \in \mathbb{H} \setminus \{i^*\}, j \in [n] \setminus \{i\}}$ be strings drawn in Step 1 of the protocol.
2. If the check described in \mathcal{H}_2 fails, then attach check-fail to the transcript.
3. For $j \in \mathbb{S}$, let $x_j := (\alpha_{i^*,j}^{(c)}, \alpha_{i^*,j}^{(r)}, \beta_{i^*,j}^{(c)}, \beta_{i^*,j}^{(r)}, c'_{i^*})$ be the messages exchanged between P_{i^*} and P_j in Steps 2 and 3 when P_{i^*} was acting as the committer and P_j was acting as the receiver, along with the message broadcast by P_{i^*} in Step 4. Set $\sigma^{\{x_j\}_{j \in \mathbb{S}}} = (\tau^{(4)}, \rho^{(4)}, \{r_{i,j}, s_{i,j}\}_{i \in \mathbb{H} \setminus \{i^*\}, j \in [n] \setminus \{i\}})$.

Now, compute

$$\text{VIEW} \leftarrow \text{ZK.SimNoAbort}_\perp(1^\lambda, 1^{n-1}, \mathbb{S}, V_\lambda^*, (\{x_j\}_{j \in \mathbb{S}}, \sigma^{\{x_j\}_{j \in \mathbb{S}}}))$$

If the proofs (which are included in VIEW) output by A_λ^* are accepting, then output VIEW, and otherwise output \perp .

- \mathcal{H}_3 : Sample instance-advice distribution identically to \mathcal{H}_2 , except that in Step 2, party P_{i^*} commits to $(0^{k(\lambda)}, 0^\lambda)$.
- \mathcal{H}_4 : Sample instance-advice distribution identically to \mathcal{H}_3 , except that in Step 4, the ZK session where P_{i^*} is the prover is simulated.
- \mathcal{H}_5 : Sample instance-advice distribution identically to \mathcal{H}_4 , except that in Step 3, for all $j \in [n] \setminus \{i^*\}$, party P_{i^*} commits to $0^{k(\lambda)}$.
- \mathcal{H}_6 : Sample instance-advice distribution identically to \mathcal{H}_5 , except that in Step 4, the ZK session where P_{i^*} is the prover is performed honestly, with witnesses $\{(0^{k(\lambda)}, r_{i^*,j})\}_{j \in \mathbb{S}}$.
- \mathcal{H}_7 : Let $C^* = (C_\lambda^*, \sigma_\lambda)_{\lambda \in \mathbb{N}}$ be the machine and corresponding non-uniform advice as defined in the description of SimNoAbort_\perp . Sample instance-advice distribution identically to \mathcal{H}_6 , except that the values committed by A^* in interaction with P_{i^*} are extracted as in the description of Sim. Additionally, this hybrid outputs \perp if the values $\{c_j\}_{j \in \mathbb{S}}$ output by parties $\{P_j\}_{j \in \mathbb{S}}$ in Step 5 do not match the extracted $\{c'_j\}_{j \in \mathbb{S}}$.
- \mathcal{H}_8 : This hybrid is the distribution $\text{SimNoAbort}_\perp(A^*)$. The only differences between \mathcal{H}_7 and \mathcal{H}_8 are the check introduced in \mathcal{H}_1 is removed, and:
 - In \mathcal{H}_7 , the challenger generates P_{i^*} 's message in Step 4 by sampling uniformly random c_{i^*} .
 - In \mathcal{H}_8 , the challenger generates P_{i^*} 's message in Step 4 as $c_{i^*} = \bigoplus_{j \in \mathbb{S}} c'_j \bigoplus_{j \in \mathbb{H} \setminus \{i^*\}} c_j \oplus r$.

Now we show that each consecutive pair of hybrids is indistinguishable. We let BAD_i be the event that in hybrid \mathcal{H}_i , check-fail appears in the output distribution (meaning that the check introduced in \mathcal{H}_1 failed AND there was no abort).

- $\mathcal{H}_0 \approx_s \mathcal{H}_1$: It suffices to show that $\Pr[\text{BAD}_1] = \text{negl}(\lambda)$, which follows directly from the quantum computational soundness of ZK.
- $\mathcal{H}_1 \approx_c \mathcal{H}_2$: This follows from Lemma 6.2.
- $\mathcal{H}_2 \approx_c \mathcal{H}_3$: We define distributions \mathcal{H}'_2 and \mathcal{H}'_3 that are identical to \mathcal{H}_2 and \mathcal{H}_3 respectively, except that \mathcal{H}'_2 and \mathcal{H}'_3 do not perform the additional check described in \mathcal{H}_1 , and as such, never attach check-fail to the output.

To prove that \mathcal{H}_2 and \mathcal{H}_3 are computationally indistinguishable, it suffices to prove that $\mathcal{H}_2 \approx_c \mathcal{H}'_2$, $\mathcal{H}'_2 \approx_c \mathcal{H}'_3$, and $\mathcal{H}'_3 \approx_c \mathcal{H}_3$. The first indistinguishability follows from the fact that $\mathcal{H}_2 \approx_c \mathcal{H}_1$ and $\Pr[\text{BAD}_1] = \text{negl}(\lambda)$, which means that $\Pr[\text{BAD}_2] = \text{negl}(\lambda)$. The second indistinguishability

follows directly from the hiding of nmCom (implied by Definition 7.1). In what follows, we show that $\Pr[\text{BAD}_3] = \text{negl}(\lambda)$, which implies that $\mathcal{H}'_3 \approx_c \mathcal{H}_3$.

Let $\text{BAD}_{3,j}$ be the event that, in hybrid \mathcal{H}_3 , $y_j \notin \mathcal{L}^{(\text{both})}$ (where y_j was defined in \mathcal{H}_1) and yet the hybrid did not abort. Now suppose that there exists a polynomial $p(\cdot)$ such that for large enough $\lambda \in \mathbb{N}$, $\Pr[\text{BAD}_3] \geq 1/p(\lambda)$. Assuming that $n = \text{poly}(\lambda)$, this implies that there exists a polynomial $p'(\cdot)$ such that for large enough $\lambda \in \mathbb{N}$, there exists some j_λ^* such that $\Pr[\text{BAD}_{3,j_\lambda^*}] \geq 1/p'(\lambda)$.

We will use this to contradict many-to-one non-malleability of nmCom, by building a QPT man-in-the-middle adversary $\text{MIM} = \{\text{MIM}_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$ that uses A^* to contradict Definition 7.1.

MIM_λ obtains as non-uniform advice i) the index j_λ^* that maximizes $\Pr[\text{BAD}_{3,j_\lambda^*}]$ and ii) A_λ^* 's advice state ρ_λ . It simulates the first two steps of the coin-flipping protocol in the presence of A_λ^* . During Step 2, it interacts with a challenger on the left committing to either $(c_{i^*}, r_{i^*,j})$ for each $j \in \mathbb{S}$, or to $(0^{k(\lambda)}, 0^\lambda)$ for each $j \in \mathbb{S}$, on behalf of P_{i^*} . It forwards these to A_λ^* on behalf of P_{i^*} and uses the strategy in \mathcal{H}_2 to generate messages on behalf of all other honest parties. When A_λ^* outputs committer messages computed on behalf of $P_{j_\lambda^*}$ in its interaction with P_{i^*} , MIM_λ forwards these to a challenger on the right, and in return obtains receiver messages on behalf of P_{i^*} .

In other words, MIM_λ 's interaction with its challengers generates either the random variable (defined in Definition 7.1)

$$\text{View-Val}_{\text{MIM}_\lambda} \langle \{\mathcal{C}(c_{i^*}, r_{i^*,j})\}_{j \in \mathbb{S}}, \text{MIM}_\lambda(\rho_\lambda), \mathcal{R} \rangle (1^\lambda, \text{tag}_{i^*}, \text{tag}_{j_\lambda^*})$$

or the random variable

$$\text{View-Val}_{\text{MIM}_\lambda} \langle \{\mathcal{C}(0^{k(\lambda)}, 0^\lambda)\}_{j \in \mathbb{S}}, \text{MIM}_\lambda(\rho_\lambda), \mathcal{R} \rangle (1^\lambda, \text{tag}_{i^*}, \text{tag}_{j_\lambda^*}),$$

depending on which strings the challenger on the left is committing to.

We now show the existence of a quantum polynomial-time $D = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ that succeeds in distinguishing these distributions with non-negligible advantage, which contradicts the non-malleability of nmCom as defined in Definition 7.1. The distribution received by D_λ includes the message (c^*, r^*) committed by MIM_λ in its interaction with P_{i^*} on the right, along with the final view VIEW of MIM_λ , which includes A^* 's view after Step 2 of Protocol 6. It then simulates the remainder of the coin-flipping protocol in the presence of A_λ^* , with one difference. Instead of implementing the check introduced in \mathcal{H}_1 , it checks only that $y_{j_\lambda^*} \in \mathcal{L}^{(\text{both})}$, using the message (c^*, r^*) . If this check failed and there was no abort, it outputs 1 and otherwise outputs 0.

Now observe that

$$\begin{aligned} & \Pr \left[D_\lambda(\text{View-Val}_{\text{MIM}_\lambda} \langle \{\mathcal{C}(c_{i^*}, r_{i^*,j})\}_{j \in \mathbb{S}}, \text{MIM}_\lambda(\rho_\lambda), \mathcal{R} \rangle (1^\lambda, \text{tag}_{i^*}, \text{tag}_{j_\lambda^*})) = 1 \right] \\ &= \Pr[\text{BAD}_{2,j_\lambda^*}] \leq \Pr[\text{BAD}_2] = \text{negl}(\lambda), \text{ and} \\ & \Pr \left[D_\lambda(\text{View-Val}_{\text{MIM}_\lambda} \langle \{\mathcal{C}(0^{k(\lambda)}, 0^\lambda)\}_{j \in \mathbb{S}}, \text{MIM}_\lambda(\rho_\lambda), \mathcal{R} \rangle (1^\lambda, \text{tag}_{i^*}, \text{tag}_{j_\lambda^*})) = 1 \right] \\ &= \Pr[\text{BAD}_{3,j_\lambda^*}] \geq 1/p'(\lambda), \end{aligned}$$

which establishes that D has a non-negligible advantage, a contradiction.

- $\mathcal{H}_3 \approx_c \mathcal{H}_4$: This follows from the quantum zero-knowledge of ZK. The non-uniform advice given to the malicious verifier derived from A^* will include the transcript of the first two rounds of Protocol 6 executed with adversary A^* , along with the openings (if they exist) of the commitments made by A^* in Step 2. The final view of this verifier will also include these openings, allowing the reduction to efficiently simulate the remainder of the protocol, in particular using these openings to efficiently implement the check introduced in \mathcal{H}_1 .

- $\mathcal{H}_4 \approx_c \mathcal{H}_5$: We consider a sequence of sub-hybrids $\mathcal{H}_{4,0}, \dots, \mathcal{H}_{4,|\mathbb{S}|}$. Associate the set \mathbb{S} with the set $[1, |\mathbb{S}|]$, and define $\mathcal{H}_{4,j}$ so that in Step 3, P_{i^*} commits to $0^{k(\lambda)}$ when interacting with adversarial parties P_k for $k \leq j$ and commits to c_{i^*} when interacting with adversarial parties P_k for $k > j$. Observe that $\mathcal{H}_4 = \mathcal{H}_{4,0}$ and $\mathcal{H}_5 = \mathcal{H}_{4,|\mathbb{S}|}$.

We now show that for any $j \in [1, |\mathbb{S}|]$, the indistinguishability $\mathcal{H}_{4,j-1} \approx_c \mathcal{H}_{4,j}$ follows from the quantum computational hiding of eCom. Indeed, define a receiver $R^* = (R_\lambda^*, \rho_\lambda)$ that interacts with a single committer committing to either c_{i^*} or $0^{k(\lambda)}$ as follows. It takes as non-uniform advice ρ_λ the transcript of the first two rounds of Protocol 6 executed with adversary A_λ^* , along with the openings (if they exist) of the commitments made by A_λ^* in Step 2. It then simulates the remainder of the protocol, interacting with the challenger to implement P_{i^*} 's messages in Step 3 during the eCom session when P_j is the receiver. Observe that R^* can indeed efficiently simulate the entire protocol, in particular it can implement the check introduced in \mathcal{H}_1 since it has the openings to the commitments given by A_λ^* in Step 2. Any efficient distinguisher that distinguishes between $\mathcal{H}_{4,j-1}$ and $\mathcal{H}_{4,j}$ with non-negligible advantage immediately implies that R^* distinguishes with non-negligible advantage, breaking quantum computational hiding of eCom.

- $\mathcal{H}_5 \approx_c \mathcal{H}_6$: Same argument as $\mathcal{H}_3 \approx_c \mathcal{H}_4$.
- $\mathcal{H}_6 \approx_c \mathcal{H}_7$: Assume that there exists a distinguisher D that can distinguish between the outputs of these hybrids with non-negligible advantage. We build a compliant²⁰ distinguisher D' that breaks the extractability property of eCom.

D' will receive as non-uniform advice i) the transcript of the first two rounds of Protocol 6 executed with adversary A^* , ii) the state of A^* at this point, and iii) the openings (if they exist) of the commitments made by *each* party in Step 2. Note that the non-uniform advice given to the committer C^* defined in the description of SimNoAbort_\perp is a strict subset of this advice. Now, D' will forward this subset (which consists of the state of A^* and the commitment openings of parties $\{P_i\}_{i \in \mathbb{H}}$) to its challenger, and receive either the real or simulated view with respect to committer C^* . It can then efficiently generate the rest of the distribution using its non-uniform advice and the view it received from the challenger, additionally returning an abort if the messages $\{c'_i\}_{i \in \mathbb{S}}$ it received as part of the challenge distribution do not match the messages $\{c_i\}_{i \in \mathbb{S}}$ broadcast in Step 5. If during Step 4, any of the parties $\{P_i\}_{i \in \mathbb{S}}$ fails to prove its ZK statement, D' outputs 0. Otherwise, it queries D with the final distribution and outputs what D outputs.

It remains to show that i) D' 's advantage is negligibly close to D 's advantage, and ii) D' is compliant. The first point requires two observations. First, whenever D' does *not* query D , it means that A^* failed to prove one of its ZK statements, so D 's input would have been \perp anyway. Next, we need to show that when D' *does* query D with a transcript, it is a faithful execution of either \mathcal{H}_6 or \mathcal{H}_7 , depending on whether eCom was simulated or not. If eCom was simulated, the distribution is equivalent to \mathcal{H}_7 . If not, the distribution is equivalent to \mathcal{H}_6 , except for the extra abort condition carried out by the reduction. However, observe that the probability that the reduction produces an abort but \mathcal{H}_6 does not is at most $\Pr[\text{BAD}_6]$. Since $\mathcal{H}_6 \approx_c \mathcal{H}_1$, and $\Pr[\text{BAD}_1] = \text{negl}(\lambda)$, this quantity is negligible. Now it remains to argue that D' is compliant, but this follows directly from the quantum computational soundness of ZK.

- $\mathcal{H}_7 \approx_s \mathcal{H}_8$: First, $\Pr[\text{BAD}_7] = \text{negl}(\lambda)$, since $\mathcal{H}_7 \approx_c \mathcal{H}_1$. Next, switching P_{i^*} 's message in Step 4 is perfectly indistinguishable since r is uniformly random. □

Now, note that $\text{RealAbort}_\perp(A^*) \approx_c \text{SimAbort}_\perp(A^*)$ follows by definition. Then, it follows identically to the proof of Theorem 6.1 that $\text{Real}(A^*) \approx_c \text{SimComb}(A^*)$, and then by applying Lemma 3.1, that $\text{Real}(A^*) \approx_c \text{Sim}(A^*)$. □

²⁰Recall that such a distinguisher is guaranteed to output 0 with overwhelming probability on input any non-explainable view.

9 Quantum-Secure Multi-Party Computation

9.1 Definition

We follow the standard real/ideal world paradigm for defining secure multi-party computation (MPC) as in [Gol04], replacing classical adversaries with quantum adversaries.

Consider n parties P_1, \dots, P_n with inputs x_1, \dots, x_n that wish to interact in a protocol Π to evaluate any functionality f on their joint inputs. The security of protocol Π (with respect to a functionality f) is defined by comparing the real-world execution of the protocol with an ideal-world evaluation of f by a trusted party. Informally, it is required that for every quantum adversary $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ that corrupts some subset of the parties $I \subset [n]$ and participates in the real execution of the protocol, there exists an adversary Sim , also referred to as a simulator, that can *achieve the same effect* in the ideal world. In fact, we provide a strictly stronger definition that allows the adversary A some arbitrary non-uniform quantum advice $\{\rho_\lambda\}_{\lambda \in \mathbb{N}}$ that may even depend on the inputs x_1, \dots, x_n .

We now formally describe the security definition, which only considers the case of fully *malicious* adversaries. Let $\vec{x} = (x_1, \dots, x_n)$ be the set of inputs.

The Real Execution. In the real execution, the n -party protocol Π for computing f is executed in the presence of a quantum polynomial-time adversary $A = \{A_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, where A corrupts some set $I \subset [n]$ of the parties. The honest parties follow the instructions of Π , and A sends all messages of the protocol on behalf of the corrupted parties following any arbitrary quantum polynomial-time strategy.

The interaction of $(A_\lambda, \rho_\lambda)$ in the protocol Π defines a random variable $\text{REAL}_{\Pi, A}(\lambda, \vec{x}, \rho_\lambda)$ whose value is determined by the randomness of the adversary and the honest parties. This random variable contains the output of the adversary (which may be an arbitrary function of its view and in particular may be a quantum state) as well as the outputs of the honest parties.

The Ideal Execution. In the ideal execution, an ideal world adversary Sim interacts with a trusted party, as follows.

- **Send inputs to the trusted party:** Each honest party sends its input to the trusted party. Each corrupt party P_i , (controlled by Sim) may either send its input x_i , or send some other input of the same length to the trusted party. Let x'_i denote the value sent by party P_i .
- **Trusted party sends output to the adversary:** The trusted party computes $f(x'_1, \dots, x'_n) = (y_1, \dots, y_n)$ and sends $\{y_i\}_{i \in I}$ to the adversary.
- **Adversary instructs trusted party to abort or continue:** This is formalized by having the adversary send either an abort or continue message to the trusted party. In the latter case, the trusted party sends to each honest party P_i its output value y_i . In the former case, the trusted party sends the special symbol \perp to each honest party.
- **Outputs:** Sim outputs an arbitrary function of its view, and the honest parties output the values obtained from the trusted party.

The interaction of Sim with the trusted party defines a random variable $\text{IDEAL}_{f, \text{Sim}}(\lambda, \vec{x}, \rho_\lambda)$. Having defined the real and the ideal worlds, we now proceed to define our notion of security.

Definition 9.1. Let f be an n -party functionality, and Π be an n -party protocol. Protocol Π securely computes f if for every quantum polynomial-time real-world adversary $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ corrupting a set of at most $n - 1$ players, there exists a quantum polynomial-time ideal-world adversary Sim such that for any set of inputs $\vec{x} \in (\{0, 1\}^*)^n$ and any non-uniform quantum advice $\rho = \{\rho_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\{\text{REAL}_{\Pi, A}(\lambda, \vec{x}, \rho_\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{IDEAL}_{f, \text{Sim}}(\lambda, \vec{x}, \rho_\lambda)\}_{\lambda \in \mathbb{N}}.$$

9.2 Construction

Given the construction of quantum-secure multi-party coin-flipping from Section 8, it is straightforward to achieve quantum-secure multi-party computation, due to the following lemma adapted from [KOS03]. For completeness, we give a sketch of the proof.

Lemma 9.1. *Given a quantum-secure multi-party coin-flipping protocol and a quantum-secure protocol Π for computing f in the common random string (CRS) model with straight-line black-box simulation, the natural composition of the two is a quantum-secure protocol for computing f with no setup assumptions.*

Proof. (sketch) Consider any adversary $(A_\lambda, \rho_\lambda)$ for the composed protocol. A_λ may be split into two parts: A_1 interacts in the coin-flipping protocol and produces a state st , which is passed to A_2 , who interacts in Π . We now construct a simulator Sim for the composed protocol as follows. It begins by running the straight-line black-box simulator Sim_Π for Π until it outputs a CRS r (note that since Sim_Π is straight-line, this CRS-generation step is independent of the adversary and advice, and does not require a call to the ideal functionality). At this point, Sim runs the simulator for the multi-party coin-flipping protocol on input r , adversary A_1 , and non-uniform advice ρ_λ . This simulation produces a final state st . Finally, Sim completes the execution of Sim_Π on input $A_2(st)$ and outputs what Sim_Π outputs. □

10 Acknowledgments

Part of this work was done during a visit to the Simons Institute Berkeley for the “Lattices: Algorithms, Complexity, and Cryptography” program.

This material is based on work supported in part by DARPA under Contract Nos. HR001120C0024 (for AA and DK) and HR001120C0025 (for VG). 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.

The authors thank Zvika Brakerski and Rishab Goyal for insightful discussions. The authors are also grateful to Daniel Wichs for pointing out the counterexample in Appendix B, which we included with his permission.

References

- [AJJM20] Prabhanjan Ananth, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta. Multikey fhe in the plain model. *Cryptology ePrint Archive, Report 2020/180*, 2020. <https://eprint.iacr.org/2020/180>.
- [AJL⁺12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.
- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In Jiri Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *ICALP 99*, volume 1644 of *LNCS*, pages 1–9, Prague, Czech Republic, July 11–15, 1999. Springer, Heidelberg, Germany.
- [AP20] Prabhanjan Ananth and Rolando L. La Placa. Secure quantum extraction protocols. *Theory of Cryptography Conference, TCC, 2020*. <https://eprint.iacr.org/2019/1323>.

- [ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *55th FOCS*, pages 474–483, Philadelphia, PA, USA, October 18–21, 2014. IEEE Computer Society Press.
- [Bar02] Boaz Barak. Constant-Round Coin-Tossing with a Man in the Middle or Realizing the Shared Random String Model. In *FOCS 2002*, pages 345–355, 2002.
- [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, Berkeley, CA, USA, October 21–24, 2006. IEEE Computer Society Press.
- [BCM⁺18] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh V. Vazirani, and Thomas Vidick. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In Mikkel Thorup, editor, *59th FOCS*, pages 320–331, Paris, France, October 7–9, 2018. IEEE Computer Society Press.
- [BD18a] Zvika Brakerski and Nico Döttling. Two-message statistical sender-private OT from LWE. *IACR Cryptology ePrint Archive*, 2018:530, 2018.
- [BD18b] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 370–390, Panaji, India, November 11–14, 2018. Springer, Heidelberg, Germany.
- [BDGM19] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In *TCC 2019, Part II*, *LNCS*, pages 407–437. Springer, Heidelberg, Germany, March 2019.
- [BFJ⁺20] Saikrishna Badrinarayanan, Rex Fernando, Aayush Jain, Dakshita Khurana, and Amit Sahai. Statistical ZAP arguments. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part III*, volume 12107 of *Lecture Notes in Computer Science*, pages 642–667. Springer, 2020.
- [BGI⁺17] Saikrishna Badrinarayanan, Sanjam Garg, Yuval Ishai, Amit Sahai, and Akshay Wadia. Two-message witness indistinguishability and secure computation in the plain model from new assumptions. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part III*, volume 10626 of *Lecture Notes in Computer Science*, pages 275–303. Springer, 2017.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10, Chicago, IL, USA, May 2–4, 1988. ACM Press.
- [BHP17] Zvika Brakerski, Shai Halevi, and Antigoni Polychroniadou. Four round secure computation without setup. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 645–677, Baltimore, MD, USA, November 12–15, 2017. Springer, Heidelberg, Germany.
- [BJ15] Anne Broadbent and Stacey Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 609–629, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- [BKP19] Nir Bitansky, Dakshita Khurana, and Omer Paneth. Weak zero-knowledge beyond the black-box barrier. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 1091–1102, 2019.

- [BL18a] Fabrice Benhamouda and Huijia Lin. k-round multiparty computation from k-round oblivious transfer via garbled interactive circuits. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 500–532, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.
- [BL18b] Nir Bitansky and Huijia Lin. One-message zero knowledge and non-malleable commitments. In *Theory of Cryptography Conference, TCC 2018, Goa, India, November 11-14, 2018, Proceedings*, 2018.
- [Bra18] Zvika Brakerski. Quantum FHE (almost) as secure as classical. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 67–95, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- [BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. *STOC*, 2020.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106, Palm Springs, CA, USA, October 22–25, 2011. IEEE Computer Society Press.
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (abstract) (informal contribution). In Carl Pomerance, editor, *CRYPTO’87*, volume 293 of *LNCS*, page 462, Santa Barbara, CA, USA, August 16–20, 1988. Springer, Heidelberg, Germany.
- [CGS02] Claude Crépeau, Daniel Gottesman, and Adam Smith. Secure multi-party quantum computation. In *34th ACM STOC*, pages 643–652, Montréal, Québec, Canada, May 19–21, 2002. ACM Press.
- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *34th ACM STOC*, pages 494–503, Montréal, Québec, Canada, May 19–21, 2002. ACM Press.
- [CM15] Michael Clear and Ciaran McGoldrick. Multi-identity and multi-key leveled FHE from learning with errors. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 630–656, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- [COSV16] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Concurrent non-malleable commitments (and more) in 3 rounds. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 270–299. Springer, 2016.
- [COSV17] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. In *Annual International Cryptology Conference*, pages 127–157. Springer, 2017.
- [CR87] Benny Chor and Michael Rabin. Achieving independence in logarithmic number of rounds. pages 260–268, 01 1987.
- [DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-Malleable Cryptography (Extended Abstract). In *STOC 1991*, 1991.
- [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, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference*

- on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part III*, volume 12107 of *Lecture Notes in Computer Science*, pages 729–758. Springer, 2020.
- [DHRW16] Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 93–122, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- [DL09] Ivan Damgård and Carolin Lunemann. Quantum-secure coin-flipping and applications. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 52–69, Tokyo, Japan, December 6–10, 2009. Springer, Heidelberg, Germany.
- [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, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Heidelberg, Germany.
- [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, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 794–811, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
- [DSS16] Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum homomorphic encryption for polynomial-sized circuits. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 3–32, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press.
- [GHKW17] Rishab Goyal, Susan Hohenberger, Venkata Koppula, and Brent Waters. A generic approach to constructing and proving verifiable random functions. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, pages 537–566, 2017.
- [GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, June 1996.
- [GKVV19] Rishab Goyal, Venkata Koppula, Satyanarayana Vusirikala, and Brent Waters. On perfect correctness in (lockable) obfuscation. *Cryptology ePrint Archive*, Report 2019/1010, 2019. <https://eprint.iacr.org/2019/1010>.
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In Chris Umans, editor, *58th FOCS*, pages 612–621, Berkeley, CA, USA, October 15–17, 2017. IEEE Computer Society Press.
- [GLOV12] Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *FOCS*, 2012.
- [GMPP16] Sanjam Garg, Pratyay Mukherjee, Omkant Pandey, and Antigoni Polychroniadou. The exact round complexity of secure computation. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 448–476, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.

- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229, New York City, NY, USA, May 25–27, 1987. ACM Press.
- [GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(3):691–729, 1991.
- [Gol04] Oded Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [Goy11] Vipul Goyal. Constant round non-malleable protocols using one way functions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 695–704, San Jose, CA, USA, June 6–8, 2011. ACM Press.
- [Goy18] Rishab Goyal. Quantum multi-key homomorphic encryption for polynomial-sized circuits. Cryptology ePrint Archive, Report 2018/443, 2018. <https://eprint.iacr.org/2018/443>.
- [GPR16] Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In *STOC*, pages 1128–1141, New York, NY, USA, 2016. ACM.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206, Victoria, BC, Canada, May 17–20, 2008. ACM Press.
- [GR19] Vipul Goyal and Silas Richelson. Non-malleable commitments using goldreich-levin list decoding. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 686–699. IEEE Computer Society, 2019.
- [GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *FOCS 2014*, pages 41–50, 2014.
- [GS18] Sanjam Garg and Akshayaram Srinivasan. Two-round multiparty secure computation from minimal assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 468–499, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [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, Santa Barbara, CA, USA, August 14–18, 2011. Springer, Heidelberg, Germany.
- [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, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Heidelberg, Germany.
- [Khu17] Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, volume 10678 of *Lecture Notes in Computer Science*, pages 139–171. Springer, 2017.

- [KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Statistical witness indistinguishability (and more) in two messages. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 34–65. Springer, 2018.
- [KOS03] Jonathan Katz, Rafail Ostrovsky, and Adam Smith. Round efficiency of multi-party computation with a dishonest majority. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 578–595, Warsaw, Poland, May 4–8, 2003. Springer, Heidelberg, Germany.
- [KS17] Dakshita Khurana and Amit Sahai. How to achieve non-malleability in one or two rounds. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 564–575. IEEE Computer Society, 2017.
- [Lin03] Yehuda Lindell. Parallel coin-tossing and constant-round secure two-party computation. *Journal of Cryptology*, 16(3):143–184, June 2003.
- [LN11] Carolin Lunemann and Jesper Buus Nielsen. Fully simulatable quantum-secure coin-flipping and applications. In Abderrahmane Nitaj and David Pointcheval, editors, *AFRICACRYPT 11*, volume 6737 of *LNCS*, pages 21–40, Dakar, Senegal, July 5–7, 2011. Springer, Heidelberg, Germany.
- [LP09] Huijia Lin and Rafael Pass. Non-malleability Amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC '09*, pages 189–198, 2009.
- [LP11] Huijia Lin and Rafael Pass. Constant-round non-malleable commitments from any one-way function. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 705–714, San Jose, CA, USA, June 6–8, 2011. ACM Press.
- [LPV08] Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkatasubramanian. Concurrent Non-malleable Commitments from Any One-Way Function. In *TCC 2008*, pages 571–588, 2008.
- [LTV12] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 1219–1234, New York, NY, USA, May 19–22, 2012. ACM Press.
- [Mah18] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In Mikkel Thorup, editor, *59th FOCS*, pages 332–338, Paris, France, October 7–9, 2018. IEEE Computer Society Press.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.
- [MW16] Pratyay Mukherjee and Daniel Wichs. Two round multiparty computation via multi-key FHE. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 735–763, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 536–553, 2014.
- [Pas04] Rafael Pass. Bounded-concurrent secure multi-party computation with a dishonest majority. In László Babai, editor, *36th ACM STOC*, pages 232–241, Chicago, IL, USA, June 13–16, 2004. ACM Press.

- [PPV08] Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive One-Way Functions and Applications. In *Advances in Cryptology — CRYPTO '08*, pages 57–74, 2008.
- [PR05] Rafael Pass and Alon Rosen. Concurrent Non-Malleable Commitments. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS '05*, pages 563–572, 2005.
- [PR08] Rafael Pass and Alon Rosen. New and Improved Constructions of Nonmalleable Cryptographic Protocols. *SIAM J. Comput.*, 38(2):702–752, 2008.
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th ACM STOC*, pages 461–473, Montreal, QC, Canada, June 19–23, 2017. ACM Press.
- [PS16] Chris Peikert and Sina Shiehian. Multi-key FHE from LWE, revisited. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 217–238, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Heidelberg, Germany.
- [PW10] Rafael Pass and Hoeteck Wee. Constant-round non-malleable commitments from sub-exponential one-way functions. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 638–655, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press.
- [Unr12] Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.
- [VDG98] Jeroen Van De Graaf. *Towards a Formal Definition of Security for Quantum Protocols*. PhD thesis, CAN, 1998. AAINQ35648.
- [Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, May 2009.
- [Wee10] Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *51st FOCS*, pages 531–540, Las Vegas, NV, USA, October 23–26, 2010. IEEE Computer Society Press.
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In Chris Umans, editor, *58th FOCS*, pages 600–611, Berkeley, CA, USA, October 15–17, 2017. IEEE Computer Society Press.
- [Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd FOCS*, pages 160–164, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *FOCS*, 1986.

A Simple Polynomial-Round Extractable Commitments

In what follows, we describe ideas in [BS20] that can be used to convert a post-quantum zero-knowledge protocol to an extractable commitment scheme (assuming quantum hardness of LWE). Specifically, we start with the (polynomial-round) zero-knowledge protocol in [Wat09] that can be based on any quantum one-way function, and convert it into a (polynomial-round) extractable commitment scheme.

Let $\text{Com}(\alpha; \beta)$ denote a non-interactive perfectly binding, quantum-hiding commitment to classical string α with randomness β . We also let CDS denote a two-party two-message conditional disclosure of secrets protocol, where in the first message, the receiver outputs a statement x , an NP language L , and purportedly commits to an NP witness for this statement. Next, the sender encodes a secret m in such a way that the receiver can recover m if and only if it previously committed to an NP witness for x . The (informal) security property is that the NP witness is hidden from a semi-honest sender, and the sender's secret m is hidden from a malicious receiver whenever $x \notin L$. It is well-known (eg., [OPP14, BGI⁺17]) that this can be achieved by combining a specific type of two-message OT (called statistically sender-private OT, that can itself be based on quantum hardness of LWE [BD18a]) with garbled circuits.

Given these components, a polynomial-round extractable commitment is described in Figure 7.

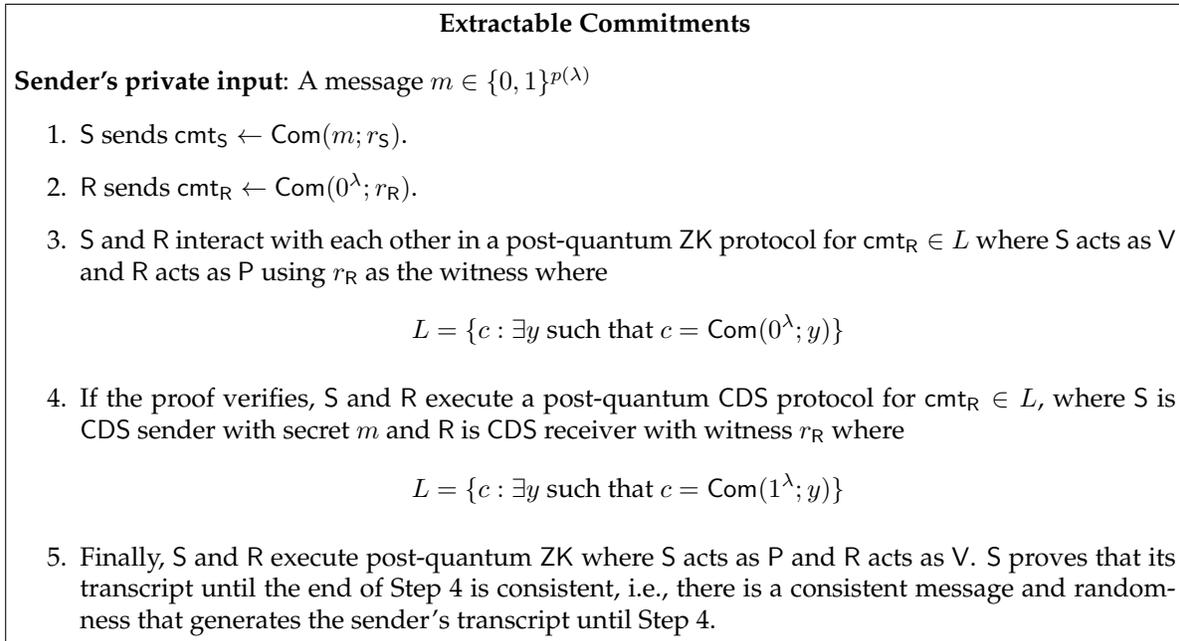


Figure 7: Extractable Commitments

At a high level, the commitment is hiding because any cheating receiver that completes Step 3 will, by the soundness of the ZK protocol, have committed to 0^λ in Step 2. By the perfect binding property of the commitment, this means that the statement of the CDS protocol, in Step 4, is false. Therefore, because the indistinguishability-based security of CDS (as discussed above), the committer's message m remains hidden from a QPT receiver.

The commitment is extractable against quantum committers (which also implies binding), because of the following argument: Consider extractor Ext that in Step 2 generates cmt_R as $\text{Com}(1^\lambda; r_R)$ instead of committing to 0^λ . Next, Ext runs ZK.Sim to simulate the proof in Step 3. After this point, Ext uses r_R as witness in Step 4, which enables it to successfully retrieve m from CDS.

B An explicit quantum attack against a classically-secure ZK protocol

In this section, we will present the construction of ZK protocol for which the zero knowledge property holds against classical verifiers, however, there exists an explicit attack w.r.t a malicious quantum verifier. This example is inspired by the recent construction of quantum extraction schemes for NP relations[AP20] where the setting is as follows: The sender \mathcal{S} and the receiver \mathcal{R} hold an NP instance x . Additionally, \mathcal{S} holds a witness w for the instance x . The desired property is the following: i) Extractability: For any QPT malicious \mathcal{S} , there exists a QPT extractor that can extract a valid witness w' for x , ii) Zero-Knowledge: For any PPT malicious \mathcal{R} , there exists a PPT simulator which can simulate the view of \mathcal{R} without having access to w .

The extraction scheme presented in [AP20] makes use of the “test of quantumness” protocol [BCM⁺18] as a key ingredient which, as the name suggests, is used to attest whether the prover is a quantum machine or a classical one. The authors leverage this test to construct a quantum extraction protocol (i.e it admits a quantum extractor) secure against classical receivers which is referred to as to as cQEXT. We note that the quantum extractor construction presented in [AP20] is straight-line i.e. it does not perform any kind of quantum rewinding [Wat09, Unr12] on \mathcal{S} . Also, the extractor makes only black-box use of the malicious \mathcal{S} i.e. it does not make any use of the circuit representation of a malicious \mathcal{S} . As we will see shortly, these two properties will be crucial in the construction of our counterexample. Our counterexample involves a classical prover $\mathcal{P}(x, w)$ interacting with a classical verifier $\mathcal{V}(x)$ in the following manner:

1. \mathcal{P} and \mathcal{V} engage in a cQEXT protocol where \mathcal{P} acts a cQEXT sender using (x, w) and \mathcal{V} acts as a cQEXT receiver using x .
2. $\mathcal{P}(x, w)$ and $\mathcal{V}(x)$ engage in a standard classical zero knowledge proof protocol which is post-quantum secure.
3. \mathcal{V} outputs 1 if the proof in Step 2 is accepting. Otherwise, it outputs 0.

The soundness of the above protocol follows from the soundness of the ZK protocol in Step 2. Also, zero-knowledge property of the above protocol w.r.t classical verifiers follows from the zero-knowledge property of the cQEXT protocol in Step 1 and ZK protocol in Step 2. However, the protocol is not zero-knowledge against a malicious QPT verifier \mathcal{V}^* for the following simple reason: \mathcal{V}^* can simply execute the extractor algorithm for the cQEXT protocol in Step 1 and therefore retrieve the witness w completely. The reason it will be able to do so without any issue is because the cQEXT extractor is black-box and straight-line. Hence, a malicious QPT verifier, which does not have any rewinding ability or access to the code of prover, can still execute the extractor algorithm seamlessly.

C Tag Amplification: Remaining Analysis

Here, we prove Lemma 7.2.

Let $\text{tag} \in [2^{t/2}]$ denote the tag used by the committer in the left session and $\text{tag}' \in [2^{t/2}]$ be the tag used by the MIM in the right session. Observe that, two sets of decomposed tags $\{\text{tag}_i\}_{i \in [t/2]}$ and $\{\text{tag}'_i\}_{i \in [t/2]}$, derived from two distinct tags, tag and tag' , are such that $\exists \alpha \in [t/2]$ s.t. $\forall i \in [t/2] : \text{tag}'_\alpha \neq \text{tag}_i$

For any values u (respectively v) committed to by C in the left session, denote by u' (respectively v') the value committed to by the MIM in the right session. Additionally, let u_i (resp, v_i) denote the value committed to by C in the i^{th} parallel execution of nmCom as part of the left commitment and let u'_i (resp, v'_i) denote the value committed to by MIM in the i^{th} parallel execution of nmCom as part of the right commitment. The soundness of ZK ensures that when the proof verifies:

$$\Pr[u' \neq u'_\alpha] = \text{negl}(\lambda) \text{ and } \Pr[v' \neq v'_\alpha] = \text{negl}(\lambda) \quad (17)$$

where the probability is over the randomness of honest verifier, and α denotes the first index in the real (resp., simulated) experiments such that for every $i \in [t/2]$, $\text{tag}'_\alpha \neq \text{tag}_i$. Whenever the proof does not

verify, the commitment is not ‘valid’ and u' (resp. v') = \perp .

Next, recall that in the real world, $\text{VIEW}_{\text{MIM}_\lambda} \langle C(u), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$ denotes the joint distribution of the view of MIM along with the value u' committed to in the right session when the left committer obtains input u . Similarly, $\text{VIEW}_{\text{MIM}_\lambda} \langle C(v), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$ denotes the joint distribution of the view of MIM along with the value v' committed to in the right session when the right committer obtains input v .

By Equation (17), whenever the MIM’s proof verifies, v' can be replaced by v'_α in the distribution $\text{VIEW}_{\text{MIM}_\lambda} \langle C(v), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$ to yield a statistically indistinguishable distribution $\text{VIEW}'_{\text{MIM}_\lambda} \langle C(v), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$. Similarly, u' can be replaced by u'_α in the distribution $\text{VIEW}_{\text{MIM}_\lambda} \langle C(u), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$ to yield a statistically indistinguishable distribution $\text{VIEW}'_{\text{MIM}_\lambda} \langle C(u), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$.

It suffices to prove that:

$$\{\text{VIEW}'_{\text{MIM}_\lambda} \langle C(u), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')\}_{\lambda \in \mathbb{N}} \approx_c \text{VIEW}'_{\text{MIM}_\lambda} \langle C(v), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}') \quad (18)$$

To that end, we define the following collections of random variables (each indexed by λ). Each is defined with respect to the adversary $\text{MIM} = \{\text{MIM}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ that we consider. Throughout, when we say abort, we mean that MIM^* aborts before Step 2, or that the MIM fails to provide an accepting proof.

- Let $\text{Pr}_u^{\text{Abort}}(\text{MIM})$ be the probability that MIM aborts in $\text{VIEW}'_{\text{MIM}_\lambda} \langle C(u), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$.
- Let $\text{Pr}_v^{\text{Abort}}(\text{MIM})$ be the probability that MIM aborts in $\text{VIEW}'_{\text{MIM}_\lambda} \langle C(v), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')$.
- Let $\text{VIEW}_u(\text{MIM}) := \{\text{VIEW}'_{\text{MIM}_\lambda} \langle C(u), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')\}_{\lambda \in \mathbb{N}}$.
- Let $\text{VIEW}_v(\text{MIM}) := \{\text{VIEW}'_{\text{MIM}_\lambda} \langle C(v), \text{MIM}(\rho_\lambda), R \rangle(1^\lambda, \text{tag}, \text{tag}')\}_{\lambda \in \mathbb{N}}$.
- Let $\text{VIEW}_u^{\text{No Abort}}(\text{MIM})$ be the distribution $\text{VIEW}_u(\text{MIM})$ conditioned on there not being an abort.
- Let $\text{VIEW}_u^{\text{Abort}}(\text{MIM})$ be the distribution $\text{VIEW}_u(\text{MIM})$ conditioned on there being an abort.
- Let $\text{VIEW}_v^{\text{No Abort}}(\text{MIM})$ be the distribution $\text{VIEW}_v(\text{MIM})$ conditioned on there not being an abort.
- Let $\text{VIEW}_v^{\text{Abort}}(\text{MIM})$ be the distribution $\text{VIEW}_v(\text{MIM})$ conditioned on there being an abort.

The following distributions will not be used explicitly in the hybrids, but will be convenient to define for the proof.

- Let $\text{VIEW}_{v,\perp}(\text{MIM})$ be the distribution $\text{VIEW}_v(\text{MIM})$, except that whenever an abort occurs, the distribution outputs \perp .
- Let $\text{VIEW}_{v,\perp}^{\text{Abort}}(\text{MIM})$ be the distribution $\text{VIEW}_v(\text{MIM})$, except that if an abort *does not* occur, the distribution outputs \perp .
- Let $\text{VIEW}_{u,\perp}(\text{MIM})$ be the distribution $\text{VIEW}_u(\text{MIM})$, except that whenever an abort occurs, the distribution outputs \perp .
- Let $\text{VIEW}_{u,\perp}^{\text{Abort}}(\text{MIM})$ be the distribution $\text{VIEW}_u(\text{MIM})$, except that if an abort *does not* occur, the distribution outputs \perp .

We show that $\text{VIEW}_u(\text{MIM}) \approx_c \text{VIEW}_v(\text{MIM})$ via a sequence on hybrids. In particular, we prove:

$$\begin{aligned}
\text{VIEW}_v(\text{MIM}) &\stackrel{(1)}{=} (1 - \Pr_v^{\text{Abort}}(\text{MIM}))\text{VIEW}_v^{\text{No Abort}}(\text{MIM}) + (\Pr_v^{\text{Abort}}(\text{MIM}))\text{VIEW}_v^{\text{Abort}}(\text{MIM}) \\
&\stackrel{(2)}{\approx_s} (1 - \Pr_u^{\text{Abort}}(\text{MIM}))\text{VIEW}_v^{\text{No Abort}}(\text{MIM}) + (\Pr_u^{\text{Abort}}(\text{MIM}))\text{VIEW}_v^{\text{Abort}}(\text{MIM}) \\
&\stackrel{(3)}{\approx_c} (1 - \Pr_u^{\text{Abort}}(\text{MIM}))\text{VIEW}_u^{\text{No Abort}}(\text{MIM}) + (\Pr_u^{\text{Abort}}(\text{MIM}))\text{VIEW}_v^{\text{Abort}}(\text{MIM}) \\
&\stackrel{(4)}{\approx_c} (1 - \Pr_u^{\text{Abort}}(\text{MIM}))\text{VIEW}_u^{\text{No Abort}}(\text{MIM}) + (\Pr_u^{\text{Abort}}(\text{MIM}))\text{VIEW}_u^{\text{Abort}}(\text{MIM}) \\
&\stackrel{(5)}{=} \text{VIEW}_u(\text{MIM}),
\end{aligned}$$

where

1. The equalities (1) and (5) follow by definition.
2. The indistinguishability (2) follows as a corollary of Claim C.1. Indeed, $\text{VIEW}_{u,\perp}(\text{MIM}) \approx_c \text{VIEW}_{v,\perp}(\text{MIM})$ in particular implies that the difference in the probability that the MIM aborts in the real interaction versus the simulated interaction is negligible.
3. The indistinguishability (3) follows as a corollary of Claim C.1. This can be seen by considering two cases. First, if the probability that the MIM aborts in the real interaction is negligible, then $\text{VIEW}_v^{\text{No Abort}}(\text{MIM}) \approx_c \text{VIEW}_u^{\text{No Abort}}(\text{MIM})$ directly follows from Claim C.1, and the indistinguishability follows. Otherwise, this probability is non-negligible, meaning that $\text{VIEW}_v^{\text{Abort}}(\text{MIM})$ is efficiently sampleable. Thus, a reduction to Claim C.1 can sample from the distribution $\text{VIEW}_v^{\text{Abort}}(\text{MIM})$ whenever it receives \perp from its challenger.²¹
4. The indistinguishability (4) follows as a corollary of Claim C.3 via a similar analysis as the last step.

Claim C.1.

$$\text{VIEW}_{u,\perp}(\text{MIM}) \approx_c \text{VIEW}_{v,\perp}(\text{MIM})$$

Proof. We prove this via a sequence of hybrids. We use \mathcal{H}_k to denote the joint distribution of MIM's view (consisting of commitment and proof transcripts along with MIM's state) and the value that MIM commits to in the right session of Hybrid k , using tag tag'_α , where α denotes the smallest index such that $\text{tag}'_\alpha \neq \text{tag}_i$ for every $i \in [t/2]$.

\mathcal{H}_1 : In this hybrid, the challenger executes the simulator $\text{ZK.Sim}(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ on V_λ^* , which denotes a wrapper around the portion of the MIM that participates in Stage 2 of the protocol, and an instance-advice distribution $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ defined as follows:

- Set the state of MIM_λ to be ρ_λ .
- Execute Stages 0 and 1 of the protocol the same way as in the experiment $\text{VIEW}_{u,\perp}(\text{MIM})$, and set x, w, \mathcal{L} according to Protocol 5 on behalf of \mathcal{C} .
- Let $\sigma_\lambda^{(x_\lambda)}$ denote the joint distribution of the protocol transcript, the state of the MIM at the end of Stage 1, and the value v'_α committed by the MIM with tag tag'_α .

If there is an abort during sampling, or ZK.Sim causes V_λ^* to abort (this includes the MIM failing to provide an accepting proof), then output \perp . By Claim 6.2,

$$\text{VIEW}_{u,\perp}(\text{MIM}) \approx_c \mathcal{H}_1$$

²¹A more formal analysis of this can be found in [BS20, Lemma 3.2].

\mathcal{H}_2 : In this hybrid, the challenger behaves identically to \mathcal{H}_1 , except when generating $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$, it replaces the commitment to u with a commitment to v in the first parallel repetition, with tag tag_1 , of nmCom (while executing all other parallel repetitions the same way as \mathcal{H}_1). If there is an abort during sampling, or ZK.Sim causes V_λ^* to abort, then output \perp .

We prove in Claim C.2 that by one-to-one non-malleability of nmCom, for every $u, v \in \{0, 1\}^{p(\lambda)}$,

$$\mathcal{H}_1 \approx_c \mathcal{H}_2$$

\mathcal{H}_i for $i \in [3, (t/2 + 1)]$: In this hybrid, the challenger behaves identically to \mathcal{H}_{i-1} , except when generating $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$, it replaces the commitment to u with a commitment to v in the $(i - 1)^{\text{th}}$ parallel repetition, with tag tag_{i-1} , of nmCom (while executing all other parallel repetitions the same way as \mathcal{H}_{i-1}). If there is an abort during sampling, or ZK.Sim causes V^* to abort, then output \perp .

We prove in Claim C.2 that by one-to-one non-malleability of nmCom, for every $u, v \in \{0, 1\}^{p(\lambda)}$ and every $i \in [3, t/2 + 1]$,

$$\mathcal{H}_{i-1} \approx_c \mathcal{H}_i$$

Finally, by claim 6.2, we have that

$$\mathcal{H}_{(t/2+1)} \approx_c \text{VIEW}_{u, \perp}(\text{MIM})$$

Next, we state and prove Claim C.2.

Claim C.2. For all $u, v \in \{0, 1\}^{p(\lambda)}$ and all $i \in [2, t/2 + 1]$,

$$\mathcal{H}_i \approx_c \mathcal{H}_{i-1} \tag{19}$$

Proof. Suppose Claim C.2 is false. Then there exists values (u, v) , some $i \in [2, t/2 + 1]$ and a polynomial $\text{poly}(\cdot)$ such that for infinitely many $\lambda \in \mathbb{N}$,

$$|\Pr[\text{MIM}(\mathcal{H}_i) = 1] - \Pr[\text{MIM}(\mathcal{H}_{i-1}) = 1]| \geq \frac{1}{\text{poly}(\lambda)} \tag{20}$$

We will demonstrate an adversary MIM^β that contradicts the non-malleability of nmCom according to Definition 7.1, i.e. we will show that for infinitely many $\lambda \in \mathbb{N}$,

$$\left| \Pr[\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta} \langle C(v), \text{MIM}^\beta(\rho_\lambda), R \rangle(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1] - \Pr[\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta} \langle C(u), \text{MIM}^\beta(\rho_\lambda), R \rangle(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1] \right| \geq \frac{1}{\text{poly}(\lambda)} \tag{21}$$

where the two distributions $\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta} \langle C(v), \text{MIM}^\beta(\rho_\lambda), R \rangle(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1$ and $\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta} \langle C(u), \text{MIM}^\beta(\rho_\lambda), R \rangle(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1$ correspond to honest commitments to v and u respectively, for the small-tag commitment scheme.

MIM^β is defined as follows:

1. Obtain input values v, u , and begin an interaction with a challenger for nmCom.
2. Emulate the role of honest committer and honest receiver in an interaction with MIM executing Protocol 5. In more detail, in the role of a committer in a left session, participate in a session of Protocol 5 with MIM as receiver. At the same time, play the role of the receiver in a right session with MIM as committer. Recall that Protocol 5 contains k repetitions of nmCom and ZK.
3. In the left session, embed the challenger's messages in the $(i - 1)^{\text{th}}$ instance of nmCom, and forward the response of MIM corresponding to the $(i - 1)^{\text{th}}$ instance to the challenger. Execute remaining instances according to the strategy in \mathcal{H}_{i-1} .

4. In the right session, forward the message obtained from MIM in the α^{th} instance of nmCom to the challenger, and embed the challenger's response for that round as receiver message in the α^{th} instance. Use honest receiver strategy for all other instances of nmCom in the right session.
5. Obtain value v'_α from the challenger of the non-malleable commitment (representing the value in the commitment sent by MIM $^\beta$ to the challenger on the right).
6. Use the transcript, the obtained value v'_α and the state of MIM to define the instance-advice sample, and then execute $\text{ZK.Sim}(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$.
7. If an abort occurs at any point, output \perp .

We now analyze the probability that MIM $^\beta$ successfully contradicts Definition 7.1. To this end, we note that:

$$\Pr[\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta}(\mathcal{C}(u), \text{MIM}^\beta(\rho_\lambda), \mathbf{R})(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1] = \Pr[\text{MIM}(\mathcal{H}_{i-1}) = 1] \quad (22)$$

$$\Pr[\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta}(\mathcal{C}(v), \text{MIM}^\beta(\rho_\lambda), \mathbf{R})(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1] = \Pr[\text{MIM}(\mathcal{H}_i) = 1] \quad (23)$$

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

$$\begin{aligned} & \left| \Pr[\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta}(\mathcal{C}(u), \text{MIM}^\beta(\rho_\lambda), \mathbf{R})(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1] \right. \\ & \left. - \Pr[\text{MIM}^\beta(\text{VIEW}_{\text{MIM}_\lambda^\beta}(\mathcal{C}(v), \text{MIM}^\beta(\rho_\lambda), \mathbf{R})(1^\lambda, \text{tag}_{i-1}, \text{tag}'_\alpha)) = 1] \right| = \\ & \left| \Pr[\text{MIM}(\mathcal{H}_{i-1}) = 1] - \Pr[\text{MIM}(\mathcal{H}_i) = 1] \right| \geq \frac{1}{\text{poly}(\lambda)} \end{aligned}$$

which is a contradiction, as desired. □

This completes the proof of Claim C.1. □

Claim C.3.

$$\text{VIEW}_{u,\perp}^{\text{Abort}}(\text{MIM}) \approx_c \text{VIEW}_{v,\perp}^{\text{Abort}}(\text{MIM})$$

Proof. We prove this via a sequence of hybrids. We use \mathcal{H}_k to denote the joint distribution of the MIM's view (consisting of commitment and proof transcripts along with the MIM's state) in Hybrid k .

\mathcal{H}_1 : In this hybrid, the challenger executes one iteration of the simulator $\text{ZK.SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ on \mathbf{V}_λ^* , where \mathbf{V}_λ^* denotes the portion of the MIM that participates in Stage 2 of the protocol, and an instance-advice distribution $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$ defined as follows:

- Set the state of MIM $_\lambda$ to be ρ_λ .
- Execute Stage 1 of the protocol the same way as in the experiment $\text{real}(\text{MIM})$, and set x, w, \mathcal{L} according to Protocol 5 on behalf of \mathcal{C} .
- If an abort occurs, output the transcript and state of the MIM until the abort.
- Otherwise, let $\sigma_\lambda^{(x_\lambda)}$ denote the joint distribution of the protocol transcript and the state of the MIM at the end of Stage 1.

If $\text{ZK.SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$ outputs a non-aborting transcript and state, output \perp , otherwise return the output of $\text{ZK.SimAbort}_\perp(1^\lambda, x_\lambda, \mathbf{V}_\lambda^*, \sigma_\lambda^{(x_\lambda)})$. By Claim 6.3,

$$\text{VIEW}_{u,\perp}^{\text{Abort}}(\text{MIM}) \approx_c \mathcal{H}_1$$

\mathcal{H}_2 : In this hybrid, the challenger behaves identically to \mathcal{H}_1 , except when generating $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$, it replaces the commitment to u with a commitment to v in the first parallel repetition, with tag tag_1 , of nmCom (while executing all other parallel repetitions the same way as \mathcal{H}_1). If a non-aborting transcript is produced, then output \perp .

We prove in Claim C.4 that by hiding of nmCom , for every $u, v \in \{0, 1\}^{p(\lambda)}$,

$$\mathcal{H}_1 \approx_c \mathcal{H}_2$$

\mathcal{H}_i for $i \in [3, (t/2 + 1)]$: In this hybrid, the challenger behaves identically to \mathcal{H}_{i-1} , except when generating $(x_\lambda, \sigma_\lambda^{(x_\lambda)})$, it replaces the commitment to u with a commitment to v in the $(i - 1)^{\text{th}}$ parallel repetition, with tag tag_{i-1} , of nmCom (while executing all other parallel repetitions the same way as \mathcal{H}_{i-1}). If a non-aborting transcript is produced, then output \perp . We prove in Claim C.4 that by hiding of nmCom , for every $u, v \in \{0, 1\}^{p(\lambda)}$ and every $i \in [3, t/2 + 1]$,

$$\mathcal{H}_{i-1} \approx_c \mathcal{H}_i$$

Finally, by claim 6.3, we have that

$$\mathcal{H}_{(t/2+1)} \approx_c \text{VIEW}_{v, \perp}(\text{MIM})$$

Next, we state and prove Claim C.4.

Claim C.4. For every $u, v \in \{0, 1\}^{p(\lambda)}$ and all $i \in [2, t/2 + 1]$,

$$\mathcal{H}_i \approx_c \mathcal{H}_{i-1}$$

Proof. Suppose Claim C.4 is false. Then there exist $u, v \in \{0, 1\}^{p(\lambda)}$, some $i \in [2, t/2 + 1]$, a PPT distinguisher D and a polynomial $\text{poly}(\cdot)$ such that for infinitely many $\lambda \in \mathbb{N}$,

$$|\Pr[D(\mathcal{H}_i) = 1] - \Pr[D(\mathcal{H}_{i-1}) = 1]| \geq \frac{1}{\text{poly}(\lambda)} \quad (24)$$

We will demonstrate a receiver that contradicts the hiding property of nmCom , i.e. we will show that there exists R^* such that for infinitely many $\lambda \in \mathbb{N}$,

$$|\Pr[R^*(\text{nmCom}\langle C(u), R^* \rangle) = 1] - \Pr[R^*(\text{nmCom}\langle C(v), R^* \rangle) = 1]| \geq \frac{1}{\text{poly}(\lambda)} \quad (25)$$

R^* obtains input u, v , and begins an interaction with a challenger for the hiding of nmCom . It then emulates the role of honest committer and honest receiver in an interaction with MIM, executing Protocol 5. In the left session, it embeds the challenger's messages in the $(i - 1)^{\text{th}}$ instance of nmCom , and forwards the response of MIM corresponding to the $(i - 1)^{\text{th}}$ instance to the challenger. It executes the remaining instances in the left session, and all instances of the right session according to the strategy in \mathcal{H}_{i-1} . Next, it uses the transcript and state of the MIM to define the instance-advice sample, and executes one iteration of $\text{ZK.SimAbort}_\perp(1^\lambda, x_\lambda, V_\lambda^*, \sigma_\lambda^{(x_\lambda)})$.

If an abort occurs at some point, then R^* runs $D(\tau, \text{st})$ where τ and st denote the transcript and the state of the adversary until the point in the protocol where the abort occurs. If no abort occurs throughout the protocol, then R^* outputs 0.

We now analyze the probability that R^* successfully contradicts Definition 7.1. Here, we note that:

$$\Pr[R^*(\text{nmCom}\langle C(u), R^* \rangle) = 1] = \Pr[D(\mathcal{H}_{i-1}) = 1] \text{ and,}$$

$$\Pr[R^*(\text{nmCom}\langle C(v), R^* \rangle) = 1] = \Pr[D(\mathcal{H}_i) = 1]$$

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

$$\begin{aligned} & \left| \Pr[R^*(\text{nmCom}\langle C(u), R^* \rangle) = 1] - \Pr[R^*(\text{nmCom}\langle C(v), R^* \rangle) = 1] \right| \\ &= \left| \Pr[D(\mathcal{H}_{i-1}) = 1] - \Pr[D(\mathcal{H}_i) = 1] \right| \geq \frac{1}{\text{poly}(\lambda)} \end{aligned}$$

which is a contradiction, as desired. This completes the proof of Claim C.4. \square

This completes the proof of Claim C.3. □

Together, these claims complete the proof of Lemma 7.2.

D Multi-Committer Extractable Commitments against Arbitrary Distinguishers

Recall that in the setting of multi-committer extractable commitments 5.1, we only considered computational indistinguishability against any *compliant* non-uniform polynomial-size quantum distinguisher D . We will now demonstrate how to upgrade any multi-committer extractable commitment secure against any compliant non-uniform polynomial-size quantum distinguisher to one which is secure against any *arbitrary* non-uniform polynomial-size quantum distinguisher. The resulting commitment admits an extractor that makes use of the Quantum Rewinding lemma 3.1 to successfully generate both explainable and non-explainable transcripts.

Construction. Let $eCom$ denote any multi committer extractable commitment protocol which admits an extractor Ext . Consider a modified version of the protocol $eCom'$ which is identical to $eCom$ except that at the very end, each of the committers $\{C_i\}_{i \in [n]}$ sends a constant-round ZK argument to the receiver attesting to the fact that the committer messages were explainable. If the verification check passes for all the arguments, the receiver accepts all the commitments. Otherwise it rejects all the commitments, and the committed value is set to \perp . We denote this by $Reject$.

Analysis. Assuming $eCom$ admits an extractor Ext which satisfies the extractability property against compliant distinguishers, we will construct an extractor Ext' for $eCom'$ which satisfies the extractability property against arbitrary distinguishers. The extractor Ext' consists of a randomized extractor Ext_{comb} will consist of two sub-extractors, namely Ext_r and Ext_{nr} . The purpose of Ext_r is to simulate a transcript which generates a $Reject$ whereas the purpose of Ext_{nr} is to simulate a transcript which does not generate a $Reject$.

At a high level, Ext_{comb} will randomly call one of the two sub-extractors and try to produce a transcript which is indistinguishable from the real view. Looking ahead, this will result in the Ext_{comb} outputting a quantum state OUT that is indistinguishable from the real verifier output conditioned on $OUT \neq Fail$. Furthermore, $OUT \neq Fail$ will occur with probability negligibly close to $1/2$ (due to random choice of executing either Ext_r or Ext_{nr} and the computational indistinguishability of the view generated by Ext_r and Ext_{nr}). In other words, Ext_{comb} is going to succeed in extraction only with probability (negligibly close to) $1/2$. Once we have this, we can apply Watrous' quantum rewinding lemma 3.1 to amplify the success probability from $\approx 1/2$ to ≈ 1 .

We will now show the construction of Ext_{nr} , Ext_r , Ext_{comb} and finally Ext' .

$Ext_{nr}(1^\lambda, 1^n, I, C_\lambda^*, \rho)$:

1. Execute the extractor $Ext(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda)$ on the adversary C_λ^* (which controls a subset I of committers).
2. Participate as a honest verifier in $|I|$ ZK argument sessions with C_λ^* where C_λ^* sends messages on behalf of the prover.
3. Execute the verification algorithm on all $|I|$ argument transcripts. If verification check passes for all $|I|$ arguments, then accept all the commitments. Otherwise, say $Reject$ occurs.
4. If $Reject$ occurs, discard all information saved so far and output $Fail$. Otherwise output C_λ^* 's inner state and the extracted value.

$\text{Ext}_r(1^\lambda, 1^n, I, C_\lambda^*, \rho)$:

1. Interact with C_λ^* as an honest receiver of eCom' . If the verification check fails for some ZK argument transcript, count it as a Reject.
2. If Reject does not occur, discard all information saved so far and output Fail. Otherwise output $C_\lambda^{* \prime}$'s inner state, the transcript and \perp as the extracted value.

$\text{Ext}_{\text{comb}}(1^\lambda, 1^n, I, C_\lambda^*, \rho)$: Sample $b \xleftarrow{\$} \{r, nr\}$ and execute Ext_b .

$\text{Ext}'(1^\lambda, 1^n, I, C_\lambda^*, \rho)$:

1. Generate the circuit $\text{Ext}_{\text{comb}, C_\lambda^*}$ which is the circuit implementation of Ext_{comb} with hardwired input C_λ^* , that is, the only input to $\text{Ext}_{\text{comb}, C_\lambda^*}$ is the quantum state ρ .
2. Let R be the algorithm from Lemma 3.1. The output of the extractor is $R(\text{Ext}_{\text{comb}, C_\lambda^*}, \rho, \lambda)$

The following claim is similar to the definition of multi-committer extractability stated in Section 5.1 but generalized to handle *arbitrary* distinguishers (instead of just compliant ones).

Claim D.1. *For any arbitrary non-uniform polynomial-size quantum distinguisher $D' = \{D'_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$, such that for all adversaries $C^* = \{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset of n committers, namely, $\{C_i\}_{i \in I}$ for some set $I \subseteq [n]$, the following holds for all polynomial-size sequences of inputs $\{\{m_{i,\lambda}\}_{i \notin I}\}_{\lambda \in \mathbb{N}}$ and $\lambda \in \mathbb{N}$.*

$$\left| \Pr[D'_\lambda(\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i,\lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n)), \sigma_\lambda) = 1] \right. \\ \left. - \Pr[D'_\lambda(\text{Ext}'(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda), \sigma_\lambda) = 1] \right| \leq \mu(\lambda).$$

Here $\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i,\lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n))$ is defined to consist of the following:

- The view of C_λ^* on eCom' interaction with the honest receiver R and set $\{C_i(m_{i,\lambda})\}_{i \notin I}$ of honest parties; this view includes a set of transcripts $\{\tau_i\}_{i \in I}$ and a state st .
- A set of strings $\{m_i\}_{i \in I}$, where each m_i is defined relative to τ_i as follows. If there exists m'_i, r_i such that $R(1^\lambda, \tau_i, m'_i, r_i) = 1$, then $m_i = m'_i$, otherwise, $m_i = \perp$.

We will prove the above claim in several steps:

1. Simulating non Reject interactions using Ext_{nr}
2. Simulating Reject interactions using Ext_r
3. Applying Watrous rewinding lemma on the combined extractor Ext_{comb}

First, we introduce some notation:

- Let $\text{Ext}_{r,\perp}$ be the same distribution as $\text{Ext}_r(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda)$ except that whenever a Reject does not occur, the distribution output is \perp
- Let $\text{Ext}_{nr,\perp}$ be the same distribution as $\text{Ext}_{nr}(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda)$ except that whenever a Reject occurs, the distribution output is \perp
- Let $\text{VIEW}'_{r,\perp}$ be the same distribution as $\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i,\lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n))$ except that whenever a Reject does not occur, the distribution output is \perp
- Let $\text{VIEW}'_{nr,\perp}$ be the same distribution as $\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i,\lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n))$ except that whenever a Reject occurs, the distribution output is \perp

We use τ' to denote the input to the distinguisher D' where τ' can either be $\text{Ext}_{\{r, nr\}, \perp}$ or $\text{VIEW}'_{\{r, nr\}, \perp}$.

Claim D.2. For any arbitrary non-uniform polynomial-size quantum distinguisher $D' = \{D'_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$, such that for all adversaries $C^* = \{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset of n committers, namely, $\{C_i\}_{i \in I}$ for some set $I \subseteq [n]$, the following holds for all polynomial-size sequences of inputs $\{\{m_{i, \lambda}\}_{i \notin I}\}_{\lambda \in \mathbb{N}}$ and $\lambda \in \mathbb{N}$.

$$|\Pr[D'_\lambda(\text{VIEW}'_{nr, \perp}, \sigma_\lambda) = 1] - \Pr[D'_\lambda(\text{Ext}_{nr, \perp}, \sigma_\lambda) = 1]| \leq \mu(\lambda).$$

Proof. Conditioned on the event that Reject happens, both distributions ($\text{VIEW}'_{nr, \perp}$ and $\text{Ext}_{nr, \perp}$) output \perp by definition. Therefore, in such a case, these two distributions will be perfectly indistinguishable. Conditioned on the event that Reject does not happen, we can say that τ' is explainable with overwhelming probability. This holds due to the soundness of ZK. Having said that, we now prove that D' cannot distinguish between real and simulated τ' .

Suppose there exists an arbitrary non-uniform polynomial-size quantum distinguisher $D' = \{D'_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, a polynomial function $\text{poly}(\cdot)$, $C^* = \{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset of n committers, namely, $\{C_i\}_{i \in I}$ for some set $I \subseteq [n]$, s.t. the following holds for infinitely many polynomial-size sequence of input $\{\{m_{i, \lambda}\}_{i \notin I}\}_{\lambda \in \mathbb{N}}$ and $\lambda \in \mathbb{N}$.

$$|\Pr[D'_\lambda(\text{VIEW}'_{nr, \perp}, \sigma_\lambda) = 1 | \neg \text{Reject}] - \Pr[D'_\lambda(\text{Ext}_{nr, \perp}, \sigma_\lambda) = 1 | \neg \text{Reject}]| \geq 1/\text{poly}(\lambda).$$

We can use D' to build a compliant distinguisher D which contradicts the multi-committer extractability of eCom as per Definition 5.1. D first obtains a eCom transcript τ (and adversary's state) as a challenge. It then interacts as an honest verifier with C_λ^* in $|I|$ ZK argument sessions where C_λ^* proves that τ is explainable. D then verifies the ZK argument and outputs 0 if the argument rejects. Otherwise, it forwards the entire transcript along with C_λ^* 's internal state and the value inside commitment to D' . If D' returns 1, D returns 1. Otherwise D returns 0.

Note that since non-explainable transcripts that are not rejected occur with negligible probability (due to soundness of ZK), the probability that D outputs 1 on receiving a non-explainable transcript is negligible. Therefore, D is a compliant distinguisher. Moreover, the following holds:

$$\Pr[D_\lambda(\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i, \lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n)), \sigma_\lambda) = 1] = \Pr[D'_\lambda(\text{VIEW}'_{nr, \perp}, \sigma_\lambda) = 1 | \neg \text{Reject}] \text{ and}$$

$$\Pr[D_\lambda(\text{Ext}(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda), \sigma_\lambda) = 1] = \Pr[D'_\lambda(\text{Ext}_{nr, \perp}, \sigma_\lambda) = 1 | \neg \text{Reject}]$$

Therefore,

$$\begin{aligned} & |\Pr[D_\lambda(\text{VIEW}_{C_\lambda^*}^{\text{msg}}(\langle R, C_\lambda^*(\rho_\lambda), \{C_i(m_{i, \lambda})\}_{i \notin I}\rangle(1^\lambda, 1^n)), \sigma_\lambda) = 1] \\ & - \Pr[D_\lambda(\text{Ext}(1^\lambda, 1^n, I, C_\lambda^*, \rho_\lambda), \sigma_\lambda) = 1]| \geq 1/\text{poly}(\lambda). \end{aligned}$$

which gives a contradiction. \square

Claim D.3. For any arbitrary non-uniform polynomial-size quantum distinguisher $D' = \{D'_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$, such that for all adversaries $C^* = \{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ representing a subset of n committers, namely, $\{C_i\}_{i \in I}$ for some set $I \subseteq [n]$, the following holds for all polynomial-size sequences of inputs $\{\{m_{i, \lambda}\}_{i \notin I}\}_{\lambda \in \mathbb{N}}$ and $\lambda \in \mathbb{N}$.

$$|\Pr[D'_\lambda(\text{VIEW}'_{r, \perp}, \sigma_\lambda) = 1] - \Pr[D'_\lambda(\text{Ext}_{r, \perp}, \sigma_\lambda) = 1]| \leq \mu(\lambda).$$

Proof. Conditioned on the event that Reject does not happen, both distributions ($\text{VIEW}'_{r, \perp}$ and $\text{Ext}_{r, \perp}$) output \perp by definition. Therefore, in such a case, these two distributions will be perfectly indistinguishable. Conditioned on the event that Reject happens, the distribution produced by Ext_r is identical to the distribution produced by an honest receiver. This holds due to the construction of Ext_r . \square

To complete our proof, we will introduce some additional notation:

1. Let $\Pr_{\text{VIEW}'}^{\text{Reject}}$ be the probability that a Reject happens in VIEW'
2. Let $\Pr_{\text{Ext}_r}^{\text{Reject}}$ be the probability that a Reject happens in Ext_r
3. Let $\Pr_{\text{Ext}_{nr}}^{\text{Reject}}$ be the probability that a Reject happens in Ext_{nr}

Following [BS20], we now show that $\text{VIEW}' = \text{Ext}_{\text{comb}}$ via a sequence on hybrids. In particular, we show that:

$$\begin{aligned}
\text{VIEW}' &\stackrel{(1)}{=} (\text{VIEW}'|\text{Reject})\Pr_{\text{VIEW}'}^{\text{Reject}} + (\text{VIEW}'|\text{NoReject})(1 - \Pr_{\text{VIEW}'}^{\text{Reject}}) \\
&\stackrel{(2)}{\approx_c} (\text{VIEW}'|\text{Reject})\Pr_{\text{Ext}_r}^{\text{Reject}} + (\text{VIEW}'|\text{NoReject})(1 - \Pr_{\text{Ext}_{nr}}^{\text{Reject}}) \\
&\stackrel{(3)}{\approx_c} (\text{VIEW}'|\text{Reject})\Pr_{\text{Ext}_r}^{\text{Reject}} + (\text{Ext}_{nr}|\text{NoReject})(1 - \Pr_{\text{Ext}_{nr}}^{\text{Reject}}) \\
&\stackrel{(4)}{\approx_c} (\text{Ext}_r|\text{Reject})\Pr_{\text{Ext}_r}^{\text{Reject}} + (\text{Ext}_{nr}|\text{NoReject})(1 - \Pr_{\text{Ext}_{nr}}^{\text{Reject}}) \\
&\stackrel{(5)}{\approx_s} \text{Ext}_{\text{comb}}
\end{aligned}$$

where

1. The equality (1) follows by definition.
2. The indistinguishability (2) follows as Corollary of Claim D.3 and Claim D.2. Indeed, $\text{VIEW}'_{r,\perp} \approx_c \text{Ext}_{r,\perp}$ in particular implies that the difference in probability that Reject happens in the real interaction versus the simulated interaction is negligible, and likewise for $\text{VIEW}'_{nr,\perp} \approx_c \text{Ext}_{nr,\perp}$.
3. The indistinguishability (3) follows as a corollary of Claim D.2. This can be seen by considering two cases. First, if the probability that Reject happens in the real interaction is negligible, then $\text{VIEW}'|\text{NoReject} \approx_c \text{Ext}_{nr}|\text{NoReject}$ directly follows from Claim D.2, and the indistinguishability follows. Otherwise, this probability is non-negligible, meaning that $\text{VIEW}'|\text{Reject}$ is efficiently samplable. Thus, a reduction to D.2 can sample from the distribution $\text{VIEW}'|\text{Reject}$ whenever it receives \perp from the challenger.
4. The indistinguishability (4) follows as a corollary of Claim D.3 via a similar analysis as the last step.
5. The indistinguishability (5) follows from the definition of Ext_{comb}

Also, by an analysis similar to [BS20] Corollary 3.1, 3.2, we can say that the success probability of Ext_{comb} is negligibly close to $1/2$ and therefore the success probability of Ext_{comb} is *input-oblivious*.

Now we can apply the Quantum rewinding lemma 3.1 to amplify the success probability from $\approx 1/2$ to ≈ 1 following an analysis similar to [BS20]. Consider the quantum circuit $\text{Ext}_{\text{comb},C^*}$ which is the circuit implementation of Ext_{comb} with hardwired input C_λ^* , that is, the only input to $\text{Ext}_{\text{comb},C_\lambda^*}$ is the quantum state ρ . By denoting the success probability for input ρ by $p(\rho)$ and setting $\epsilon := \text{negl}(\lambda) + 2^{-\lambda \cdot \frac{3}{4}}$, $p_0 := 1/4$, and $q := 1/2$, we can satisfy all the conditions for Quantum Rewinding Lemma 3.1.

This implies that trace distance between $R(\text{Ext}_{\text{comb},C^*}, \rho, \lambda)$ and Ext_{comb} is bounded by a negligible function. Therefore, our final extractor $\text{Ext}'(1^\lambda, 1^n, I, C_\lambda^*, \rho) = R(\text{Ext}_{\text{comb},C_\lambda^*}, \rho, \lambda)$ completes the extraction successfully with probability negligibly close to 1.