

A New Approach to Post-Quantum Non-Malleability

Xiao Liang¹, Omkant Pandey^{2*}, and Takashi Yamakawa³

¹ Rice University, Houston, USA

xiao.crypto@gmail.com

² Stony Brook University, Stony Brook, USA

omkant@cs.stonybrook.edu

³ NTT Social Informatics Laboratories, Tokyo, Japan

takashi.yamakawa.ga@hco.ntt.co.jp

Abstract. We provide the first *constant-round* construction of post-quantum non-malleable commitments under the minimal assumption that *post-quantum one-way functions* exist. We achieve the standard notion of non-malleability with respect to commitments. Prior constructions required $\Omega(\log^* \lambda)$ rounds under the same assumption.

We achieve our results through a new technique for constant-round non-malleable commitments which is easier to use in the post-quantum setting. The technique also yields an almost elementary proof of security for constant-round non-malleable commitments in the classical setting, which may be of independent interest.

When combined with existing work, our results yield the first constant-round quantum-secure multiparty computation for both classical and quantum functionalities *in the plain model*, under the *polynomial* hardness of quantum fully-homomorphic encryption and quantum learning with errors.

Keywords: Non-Malleable · Post-Quantum · Constant-Round

* Supported in part by DARPA SIEVE Award HR00112020026, NSF CAREER Award 2144303, NSF grants 1907908, 2028920, 2106263, and 2128187. 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, DARPA, or NSF.

Table of Contents

Table of Contents	iii
1 Introduction	1
1.1 Challenge: Robust Simulatable Extraction	2
1.2 Our Results	4
2 Technical Overview	5
2.1 Small-Tag, One-Sided, Synchronous, Classical Setting: Construction	6
2.2 Proving Non-Malleability	7
2.3 Small-Tag, One-Sided, Synchronous, Post-Quantum Setting	14
2.4 Full-Fledged Post-Quantum Non-Malleable Commitments	16
2.5 Organization	18
3 Preliminaries	18
3.1 Basic Notations	18
3.2 Commitment Schemes	19
3.3 Non-Malleable Commitments	20
3.4 Extractable Commitments	21
3.5 Witness-Indistinguishable Proofs of Knowledge	22
4 Small-Tag, One-Sided, Synchronous, Classical Setting	24
4.1 Construction of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$	25
4.2 Proving Non-Malleability (Proof of Thm. 2)	26
4.3 Proof of Lem. 6	27
4.4 Proof of Lem. 7	29
4.5 Extractor \mathcal{K} (Proof of Lem. 9)	35
4.6 Proof of Lem. 10	38
4.7 Proof of Claim 10	42
4.8 Replacing the Injective OWF with Any OWF	43
4.9 Summarizing the Proof Strategy for Thm. 2	44
5 Small-Tag, Synchronous, Classical Setting	46
5.1 High-Level Idea	46
5.2 Construction of $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$	47
5.3 Proving Non-Malleability (Proof of Thm. 11)	49
5.3.1 The Case of $t < \tilde{t}$	49
5.3.2 The Case of $t > \tilde{t}$	50
6 Small-Tag, Asynchronous, Classical Setting	50
6.1 High-Level Idea	50
6.2 Construction of $\langle C, R \rangle_{\text{tg}}^{\text{async}}$	51
6.3 Adaptive Schedule vs. Predetermined Schedule	54
6.4 Proving Non-Malleability (Proof of Thm. 13)	55
6.5 Non-Malleability for Good Schedules	55
6.5.1 The Case of $t < \tilde{t}$	55
6.5.2 The Case of $t > \tilde{t}$	57
6.6 Non-Malleability for Bad Schedules	58
6.6.1 Analysis of Bad 1	59
6.6.2 Analysis of Bad 2	59
6.6.3 Analysis of Bad 3	61

6.6.4	Analysis of Bad 4	62
6.6.5	Analysis of Bad 5	63
7	Extract-and-Simulate Lemma	63
8	Post-Quantum Non-Malleable Commitments	64
8.1	Small-Tag, One-sided, Synchronous, Post-Quantum Setting	64
8.2	Small-Tag, Asynchronous, Post-Quantum Setting	64
8.3	Tag Amplification	67
9	Application: Quantum-Secure MPC in Constant Rounds from QLWE	68
	References	71
A	Proof of Extract-and-Simulate Lemma (Lem. 20)	72
A.1	Preparation	72
A.2	Proving Lem. 20	74
B	Small-Tag, One-sided, Synchronous, Post-Quantum Setting	79
B.1	Construction	79
B.2	Proving Non-Malleability	80
B.3	Proof of Lem. 28	81
B.4	Proof of Lem. 29	83
B.5	Extractor \mathcal{K} (Proof of Lem. 31)	86
B.6	Proof of Lem. 32	89
B.7	Proof of Claim 37	93

1 Introduction

Commitments are one of the central primitives in modern cryptography. They are two-party protocols that enable a sender (or committer) to commit to a message to a receiver. It is required that an efficient receiver should learn nothing about the committed message until later the committer chooses to open (or decommit to) it. However, this vanilla promise of “information hiding” does not rule out the following mauling or man-in-the-middle (MIM) attack: An adversary \mathcal{M} could play the role of a receiver in one instance of the commitment (referred to as the *left session*), while simultaneously playing as a committer in another (referred to as the *right session*). In this situation, \mathcal{M} can potentially make the value committed in the right session depend on the value in the left session, in a malicious manner that is to her advantage. Notice that this is not breaking the hiding property of the commitment scheme, as \mathcal{M} could conduct the above attack without explicitly learning the value committed in the left session.

To protect against such attacks, Dolev, Dwork, and Naor [DDN91] introduced the concept of *non-malleable commitments*. Such commitments capture the MIM attack by requiring that in the MIM execution, the *joint distribution* of \mathcal{M} 's final output *and* the value committed in the right session is computationally indistinguishable for any values committed by the honest committer in the left session. Of course, \mathcal{M} can always relay without any modifications the messages between the left honest committer and the right honest receiver. This is handled by augmenting the commitment scheme with a *tag*, and \mathcal{M} is considered winning the MIM game only if she does not use the same tag on both sides.⁴

Non-malleable commitments have found several applications in cryptography. They turn out to be a critical ingredient in resolving the exact round complexity of secure multiparty computation in the standard (plain) model [KOS03, PPV08, Wee10, Goy11, GMPP16, BGJ⁺18, CCG⁺19], as well as protecting such protocols against concurrent attacks [Can00, Can01, CLOS02, Pas03, PS04, MPR06, BDH⁺17, CLP20].

Since their introduction, the central question in this area has been the construction of *constant-round* non-malleable commitments under the minimal assumption that one-way functions (OWFs) exist. The original work [DDN91] presented a logarithmic-round construction assuming only OWFs. The works of Barak [Bar02], and Pass and Rosen [PR05] succeeded in obtaining constant-round constructions under the (stronger but standard) assumption of (polynomially-hard) collision-resistant hash functions; this assumption is inherited due to the use of non-black-box simulation techniques [Bar01]. After a long line of follow-up works [PPV08, LPV08, LP09, Wee10], constant-round non-malleable commitments assuming only OWFs were first constructed in independent and concurrent works of Goyal [Goy11] and Lin and Pass [LP11]. Since then, several constructions optimizing various aspects of this primitive, such as the *exact* round complexity [Pas13, GRRV14, GPR16, COSV17, LPS17, Khu17, KS17, BL18, GR19, Khu21] and black-box usage of primitives [Wee10, Goy11, GLOV12], have been proposed, thus achieving an almost complete understanding of this primitive in the *classical* setting where all parties, including the adversary, as well as the communication are classical.

Unfortunately, the results in the classical setting do not usually translate to the quantum setting where one or more parties may be quantum machines. Existing classical techniques often require the ability to rewind and copy the adversaries' code, both of which are not possible in the quantum setting due to the no-cloning theorem [WZ82] and quantum state disturbance [FP96].

These issues, together with the fact that quantum computers might be possible one day, have resulted in a significant push toward developing tools and techniques to reason about security in the

⁴ Copying the tag can be shown equivalent to copying the entire interaction assuming one-way functions.

presence of quantum adversaries. Some examples include zero-knowledge [Wat06, BS20], signatures [BZ13], pseudo-random functions [Zha12], secure computation [BCKM21b, ABG⁺21], and so on.

Post-Quantum Non-Malleable Commitments. A particularly appealing goal in this direction is the construction of the so-called *post-quantum* secure protocols, where the honest parties and their communication channels are entirely *classical* but the adversary is allowed to be a *quantum* (polynomial-time) algorithm. Such protocols have the feature that even if the adversary somehow gains early access to quantum computing capabilities, the honest parties are not required to catch up to remain protected against it.

Agarwal, Bartusek, Goyal, Khurana, and Malavolta [ABG⁺21]—in their pursuit of constant-round post-quantum secure multiparty computation—construct a constant-round post-quantum non-malleable commitments w.r.t. the special case of *synchronous schedules* where, upon receiving a message in the left session, \mathcal{M} must respond with the corresponding message of the right session immediately, in each round of the protocol. They achieve this assuming super-polynomial quantum hardness of the learning with errors (QLWE) problem.

In the general (i.e., asynchronous) setting, the first positive result was recently obtained by Bitansky, Lin, and Shmueli [BLS22] who, in the post-quantum setting, construct a $\log^*(\lambda)$ -round protocol assuming only post-quantum OWFs. They present a general compiler to convert any k -round (ε -simulatable) post-quantum extractable commitment to $k^{O(1)} \cdot \log^*(\lambda)$ -round post-quantum non-malleable commitments, and then rely on the very recent work of Chia, Chung, Liang, and Yamakawa [CCLY22] where such an extractable commitment protocol is constructed in constant rounds from only post-quantum OWFs (see also [CCLY22, Section 1.2]). This brings us tantalizingly close to constant rounds. However, the techniques in [BLS22] rely on scheduling and amplification techniques, which inherently require non-constant rounds to support (the standard requirement of) large tags or identities. It is unclear if these techniques can yield constant rounds.

The central question in this area thus still remains open:

Question 1: *Do constant-round post-quantum non-malleable commitments, assuming only post-quantum one-way functions, exist?*

1.1 Challenge: Robust Simulatable Extraction

Toward answering **Question 1**, let us first discuss about the challenges.

As mentioned above, proving non-malleability of a commitment scheme requires one to show that the *joint distribution* of the final state of the MIM \mathcal{M} (denoted by $\text{st}_{\mathcal{M}}$) and the value committed in the right session (denoted by \tilde{m}) are computationally indistinguishable when the left-session committed value changes from any m_0 to any $m_1 \neq m_0$. Typically, this is done by a careful design of a sequence of hybrids, where the left value is gradually changed from m_0 to m_1 ; And non-malleability will then be established by showing that the joint distribution of $(\text{st}_{\mathcal{M}}, \tilde{m})$ is indistinguishable between each pair of adjacent hybrids. As in many other cryptographic proofs, one usually needs to reduce the indistinguishability between adjacent hybrids to some computational hardness assumptions. But this step is especially hard for non-malleability proofs due to the following “inefficient testability” issue—The value \tilde{m} is hidden in the transcript of the right interaction, and no efficient machine could obtain it. Thus, the event that “the joint distribution of $(\text{st}_{\mathcal{M}}, \tilde{m})$ changes” is not efficiently testable, thus forbidding an efficient reduction to the underlying hardness assumptions.

The most common template to address the above issue is to (efficiently) extract the value \tilde{m} from some “extractable gadget” (e.g., extractable commitments or proofs of knowledge) in the right session. The hope is: if the extracted value, denoted by \tilde{m}' , is equal to \tilde{m} with good-enough

probability, then one can conduct the above reduction using $(\text{st}_{\mathcal{M}}, \tilde{m}')$ in place of $(\text{st}_{\mathcal{M}}, \tilde{m})$. That is, since the extraction of \tilde{m}' is efficient, it saves the reduction from the aforementioned inefficient testability issue.

To properly implement this template, it is crucial to maintain the following conditions (the combination of which we call *Robust Simulatable Extractability*.)

- **Extractability:** One can extract the committed message \tilde{m} in the right from some extractable gadget;
- **Simulatability:** One can simulate \mathcal{M} 's final (i.e., post-extraction) state while extracting \tilde{m} ;
- **Robustness:** The extraction of \tilde{m} (in the right session) does not harm the hiding property of the left session (or some left-session gadget on whose hiding property the security reduction relies).

Roughly speaking, previous designs of non-malleable protocols (in the classical setting) can be thought of as developing different (and better-and-better) techniques that enable the above robust simulatable extractability.

However, robust simulatable extractability turns out to be hard to obtain *in the post-quantum setting*. First, as mentioned earlier, special techniques are needed to perform simulation for quantum adversaries due to the no-cloning theorem; Extracting a desired value while simultaneously simulating the quantum adversaries' internal state is even harder. Fortunately, the recent works [CCLY22, LMS21] did provide a constant-round solution to simulatable extraction based solely on post-quantum OWFs.

However, the picture becomes unclear when we additionally require robustness. Known simulatable extraction techniques treat the adversary as a single reversible operation (i.e., unitary) and “rewind” it coherently. However, if the adversary talks in straight-line with the external left committer C (that cannot be rewound), those techniques seem inapplicable. Actually, this robustness issue already appeared in the classical setting. But we would like to point out that the quantum power of adversaries complicates it further: It is reasonable to expect that the extractor (to be constructed) may need to “read” the messages exchanged between the MIM adversary \mathcal{M} and the left committer C . However, known quantum rewinding strategies need to treat the adversary as a reversible operation. When we view (C, \mathcal{M}) as a joint adversary to perform quantum rewinding, it is unclear if the simulator can “read” (technically, *measure*) the messages exchanged between C and \mathcal{M} , because this may irreversibly collapse the internal quantum state of the joint (C, \mathcal{M}) adversary. Therefore, it is unclear if existing post-quantum rewinding techniques could be used in the MIM setting when robustness is a concern. This indeed represents the major difficulty when one tries to build constant-round post-quantum non-malleable commitments by quantizing the security proofs of the classical ones (e.g., [Goy11, LP11, GLOV12, GRRV14, GPR16, COSV17, GR19]).

Existing Techniques. The two recent works mentioned earlier demonstrated possible solutions to the robust simulatable extractability issue, if one is willing to make stronger hardness assumptions, *or* does not insist on constant rounds:

- [BLS22] took a similar approach as in [DDN91]. Roughly, the idea is to introduce enough rewinding opportunities for extraction (usually referred to as “slots”) in the construction such that there is always a “free slot”, namely, a slot that does not interleave with any messages exchange between \mathcal{M} and the left committer C (who will become the external challenger of the hardness-providing gadget in the security proof); Then, one can extract \tilde{m} from this free slot using known post-quantum rewinding strategies. This approach is unlikely to give a constant-round construction (even in the classical setting). The constructions from [BLS22] require at least $\log^*(\lambda)$ rounds.

- [ABG⁺21] took a similar approach as in [PW10, KS17], which essentially obtained robust simulatable extractability via *complexity leveraging*. At a high level, the idea is to hide the committed value in some computationally hard gadget, whose hardness depends on the tag. This hard gadget is cleverly designed to admit the following strategy: In the MIM execution, because of the asymmetry of the left and right tags (i.e., $t \neq \tilde{t}$), the reduction can extract \tilde{m} by breaking the *right* hard gadget via brute force in some slightly super-polynomial time $T(\lambda)$; However, the MIM adversary cannot break the *left* hard gadget in time $\text{poly}(T(\lambda))$. Since the extraction in this template is conducted via brute force, the robust simulatable extraction issue can be circumvented. However, this approach is unlikely to work without super-polynomial hardness assumptions due to the use of complexity leveraging.

1.2 Our Results

In this work, we answer **Question 1** affirmatively by providing the first constant-round construction of post-quantum non-malleable commitments assuming only post-quantum one-way functions.

As discussed above, major classical approaches to OWF-based constant-round non-malleable commitments seem not to be “quantum-friendly.” Therefore, we first propose a new classical construction whose security proof is quantum-friendly. That is, when designing the new protocol, we restrict ourselves to techniques that avoid state-cloning and the above robust simulatable extractability issue. We find this result already interesting as it improves the diversity of the approaches known in the classical setting.

Next, we show that our new classical construction can be made quantum-secure once its building blocks are replaced by their post-quantum counterparts. This is possible because our classical construction is deliberately designed to be quantum-friendly.

Theorem 1 (Informal). *Assuming the existence of post-quantum one-way functions, there exists a constant-round construction of post-quantum non-malleable commitments.*

Thm. 1 yields interesting corollaries w.r.t. *multi-party computation* (MPC) in the quantum era. For classical functionalities, the recent work [ABG⁺21] presents the first post-quantum MPC in constant rounds, from a mildly super-polynomial *quantum hardness of Learning with Errors* (QLWE) assumption and a QLWE-based circular security assumption. They need the super-polynomial hardness of QLWE (only) to build constant-round post-quantum non-malleable commitments, which serve as a building block to their MPC. Plugging our non-malleable commitment into their framework yields the following result:

Corollary 1 (Informal). *Assuming (polynomial) QLWE and the QLWE-based circular security assumption (as in [ABG⁺21]), there exists a constant-round construction of post-quantum MPC for classical functionalities, i.e., an MPC protocol secure against QPT adversaries where honest parties only need to perform classical computation.*

For quantum functionalities, the recent work [BCKM21a] presents a constant-round quantum-secure MPC for quantum functionalities *in the CRS model*, based on the hardness of QLWE. It is easy to see that **Corollary 1** provides a constant-round implementation for the CRS required by the [BCKM21a] protocol. This observation leads to the first constant-round quantum-secure MPC for quantum functionalities from polynomial hardness assumptions *without any trusted setup*.

Corollary 2 (Informal). *Assuming (polynomial) QLWE and the QLWE-based circular security assumption (as in [ABG⁺21]), there exists a constant-round construction of quantum-secure MPC for quantum functionalities.*

2 Technical Overview

We will first present a construction $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ that is non-malleable *in the classical setting* with the following restrictions:

- **Small-Tag:** It only supports tags from the polynomial-size space $[n] := \{1, \dots, n(\lambda)\}$, where $n(\lambda)$ is a polynomial on the security parameter λ ;
- **One-Sided:** Its non-malleability holds only if, in the MIM game, the left-session tag t is smaller than the right-session tag \tilde{t} .
- **Synchronous:** Its non-malleability holds only in the synchronous setting. This refers to the setting where upon receiving a message in the right (resp. left) session, the MIM adversary immediately responds with the corresponding message in the left (resp. right) session.

It is a common approach in the literature to first obtain a construction under the above conditions, and then convert it to a full-fledged non-malleable commitment. While there exist standard techniques that take care of the latter step, the former step (i.e., constructing $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$) is typically where difficulty lies.

In the classical setting, once we obtain $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$, we can apply known techniques to remove the **One-Sided** restriction, yielding a small-tag, synchronous protocol, which we denote as $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$. Then, known compilers (e.g., [Wee10]) can be used to remove the **Small-Tag** and **Synchronous** restrictions *at one stroke*. This leads to a full-fledged non-malleable commitment $\langle C, R \rangle_{\text{TG}}^{\text{async}}$ in the classical setting.

We emphasize that all the above protocols are non-malleable *only in the classical setting*. However, they are designed deliberately using quantum-friendly techniques, which makes it possible to quantize their security proofs. We elaborate on that in the sequel.

Post-Quantum Tag Amplification. Recall that our eventual goal is to achieve *post-quantum* non-malleability. The use of quantum-friendly techniques will allow us to prove *post-quantum* non-malleability of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$. Also, the aforementioned (classical) conversion from $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ to $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$ extends naturally to the post-quantum setting as well. However, the classical compiler from $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$ to $\langle C, R \rangle_{\text{TG}}^{\text{async}}$ seems not to extend to the post-quantum setting.

This has already been observed in [BLS22]. The authors of [BLS22] addressed this issue by constructing a new tag amplifier that converts a small-tag, asynchronous post-quantum non-malleable commitment to a large-tag (i.e, $t \in [2^\lambda]$), asynchronous post-quantum non-malleable commitment. But it is worth noting that this tag amplifier requires the small-tag protocol to be *asynchronously* secure; This is in contrast to the aforementioned classical compiler, which handles asynchronicity and tag-size amplification at one stroke.

To overcome this problem, we will first show that our post-quantum version of $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$ can be modified to achieve non-malleability in the asynchronous setting, yielding a protocol $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$, which is exactly a small-tag, asynchronously non-malleable commitment in the post-quantum setting. Now, the [BLS22] tag amplifier can be applied to $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$, leading to the full-fledged post-quantum non-malleable commitment $\langle C, R \rangle_{\text{TG}, \text{PQ}}^{\text{async}}$ we want.

Organization of Technical Overview. In Sec. 2.1 and 2.2, we overview the main idea behind our construction of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ in the classical setting. This is the most technically involved part where the main difficulty lies.

As mentioned earlier, protocol $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ is designed using only quantum-friendly techniques. Thus, the proof of its non-malleability extends to the post-quantum setting once we replace its building blocks by their post-quantum counterparts. (We provide a high-level explanation for that

on Page 12.) We then show how this can be done in Sec. 2.3, where we obtain the one-sided, small-tag, synchronous, *post-quantum* non-malleable commitment $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$.

Finally, we provide in Sec. 2.4 a brief overview on how to convert $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$ to the full-fledged post-quantum non-malleable commitment $\langle C, R \rangle_{\text{TG}, \text{PQ}}^{\text{async}}$, achieving our eventual goal.

2.1 Small-Tag, One-Sided, Synchronous, Classical Setting: Construction

Our construction of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ is easy to describe. It supports tags from the space $[n(\lambda)]$, where $n(\lambda)$ is any (fixed) polynomial on the security parameter λ .

To proceed with a tag $t \in [n]$, we first ask the committer C to commit to m using a statistically binding scheme $\text{com} = \text{Com}(m; r)$ (e.g., Naor’s commitment). Then, the receiver R sends a *non-interactive hard puzzle* that has exactly t distinct solutions; R also gives a witness-indistinguishable proof of knowledge⁵ (referred to as **WIPoK-1**) to prove that he knows one of the t solutions. Finally, C is required to prove using another WIPoK (referred to as **WIPoK-2**) that he knows *either* the value committed in com *or* one solution to R ’s hard puzzle.

We depict this construction (in the MIM setting) in Fig. 1, where the “hard puzzle” is implemented with the problem of “finding one of the preimages of the t OWF images $\{y_i\}_{i \in [t]}$ ”. Throughout this overview, we assume that the OWF f is injective (otherwise, this “hard puzzle” may have more than t solutions). But this is only to ease the presentation—We will describe in Sec. 4.8 a simple trick to remove this injectivity requirement.

The main intuition underlying this construction is best illustrated by the proof of its non-malleability, which we show next in Sec. 2.2.




Fig. 1: Man-in-the-Middle Execution of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$

⁵ Indeed, a WI *argument* of knowledge suffices (see Sec. 2.3).

2.2 Proving Non-Malleability

Instead of giving out the security proof directly, we start with a naive attempt that will fail. It will reveal the main difficulty in the proof of non-malleability and how our design addresses it.

The First Attempt. In the MIM execution shown in Fig. 1, imagine that we invoke the knowledge extractor⁶ for the right **WIPoK-2** to extract a witness w' . It is easy to see that the only possible values for w' is (\tilde{m}, \tilde{r}) or $(1, \tilde{x}_1)$, because they are the only information that \mathcal{M} could potentially possess assuming it cannot break the one-wayness of the OWF \tilde{f} in the right **Step-3**. Furthermore, we know for sure that w' can only be (\tilde{m}, \tilde{r}) because of the WI property of the right **WIPoK-1** and the following (standard) argument: Imagine that we switch to $(2, \tilde{x}_2)$ as the witness to finish the right **WIPoK-1**. By the WI property of the right **WIPoK-1**, we know that the distribution of the extracted w' cannot change (up to negligible probability). That is, *if w' could take the value $(1, \tilde{x}_1)$ before we switch to $(2, \tilde{x}_2)$ in the right **WIPoK-1**, it would keep taking the value $(1, \tilde{x}_1)$ when we are using $(2, \tilde{x}_2)$ in the right **WIPoK-1***. But this cannot happen as it breaks the one-wayness of \tilde{f} on the image \tilde{y}_1 using a standard reduction—Consider an external challenger for the one-wayness of \tilde{f} ; The challenger sends us a challenge y^* ; We run the game shown in Fig. 1, while using y^* in place of \tilde{y}_1 and using $(2, \tilde{x}_2)$ as the witness to perform the right **WIPoK-1**. If we finally extract $w' = (1, x_1)$, we find a preimage for y^* .

With the above observation, we *hope to* prove non-malleability in the following way: We change the left committed value from m_0 to m_1 , while extracting the right committed value $w' = (\tilde{m}, \tilde{r})$ using the knowledge extractor for the right **WIPoK-2**. If \mathcal{M} really makes \tilde{m} depend on the left committed value, then the reduction can detect this change by checking the extracted w' . This breaks the hiding property⁷ of the left execution when the left committed value changes from m_0 to m_1 . However, this naive idea will not work: When we invoke the knowledge extractor for the right **WIPoK-2**, \mathcal{M} also rewinds the left **WIPoK-2**. So, the left interaction is not “in straight-line” anymore. Thus, we could not build the reduction to the hiding property of the left interaction.

Reduction to Naor’s Commitment? Though the above attempt does not work, it inspires the following thoughts: What if we reduce non-malleability to the hiding property of Naor’s commitment in the left **Step-1** and **Step-2**? There is at least some hope as the left Naor’s commitment is not interleaved with the right **WIPoK-2** (recall that we are in the synchronous setting). That is, we can generate the **Step-1** and **Step-2** messages in the left interaction by forwarding messages between \mathcal{M} and an (external) challenger for the hiding property of Naor’s commitment; Meanwhile, we extract $w' = (\tilde{m}, \tilde{r})$ from the right **WIPoK-2** as before. Then, if \mathcal{M} makes the value \tilde{m} depend on the value committed in the first two steps in the left, we win the hiding game by checking the extracted w' .

To implement this idea, we first need to make the interaction happening after the left **Step-2** independent of m ; Otherwise, the above reduction will not work—Because the left Naor’s commitment is now coming from an external challenger so that the reduction does not possess the value m anymore, which is required to finish the remaining steps (in particular, the left **WIPoK-2**) of the left interaction. To address this issue, we consider an intermediate execution \mathcal{G}_1 shown in Fig. 2a—The only difference (shown in red) between \mathcal{G}_1 and the real MIM game (Fig. 1) is that \mathcal{G}_1 uses the knowledge extractor \mathcal{WE} in the left **WIPoK-1**, and uses the extracted witness (j, x_j)

⁶ We remark that there exist different definitions for *Proofs of Knowledge*. Throughout this paper, we use Lindell’s formalism called *Witness Extended Emulation* [Lin03]. This definition ensures the existence of a knowledge extractor \mathcal{WE} such that for any P^* that convinces V with probability p , \mathcal{WE}^{P^*} will extract a valid witness with probability $\geq p - \text{negl}(\lambda)$. See Sec. 3.5 for more details.

⁷ It is easy to see that our construction, in the stand-alone setting, is a computationally hiding commitment.

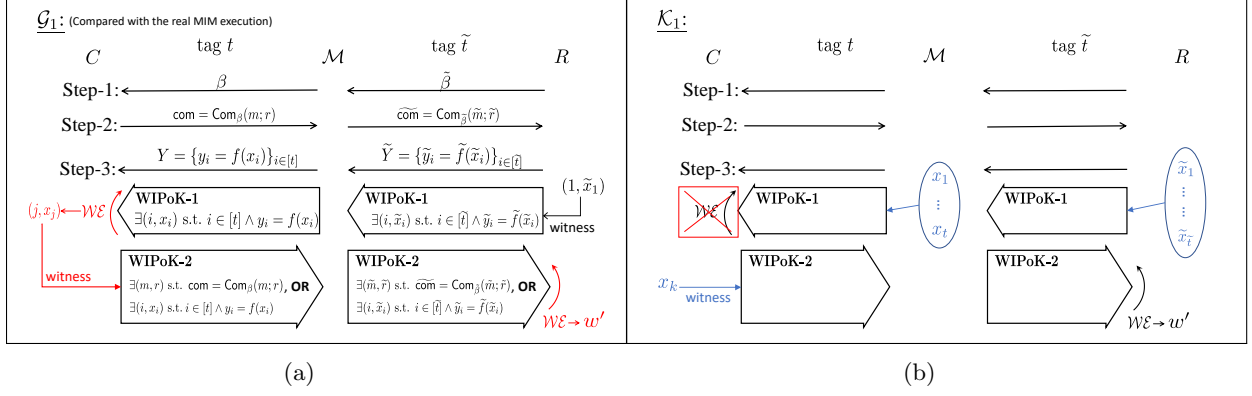


Fig. 2: Machines \mathcal{G}_1 and \mathcal{K}_1

as the witness to go through the left **WIPoK-2**. In this way, the interaction after the left **Step-2** does not depend on m anymore. However, this introduces new problems: We can no longer make use of the WI property of the right **WIPoK-1** as before, because this \mathcal{WE} needs to rewind the left **WIPoK-1**, resulting in rewindings of the right **WIPoK-1** as well (since we are in the **Synchronous** setting). Thus, we can no longer argue that the extracted w' must be (\tilde{m}, \tilde{r}) as in **The First Attempt** part.

Actually, there is a deeper reason why this naive attempt is bound to fail in establishing $w' = (\tilde{m}, \tilde{r})$. When we run \mathcal{WE} for the left **WIPoK-1**, \mathcal{M} could also learn the witness $(1, \tilde{x}_1)$ used by R in the right **WIPoK-1**. When we switch the witness from (m, r) to the extracted (j, x_j) in the left **WIPoK-2**, nothing stops \mathcal{M} from switching her witness from (\tilde{m}, \tilde{r}) to $(1, \tilde{x}_1)$. Therefore, we can no longer argue that the extracted w' must be (\tilde{m}, \tilde{r}) . We emphasize that the WI property of the left **WIPoK-2** does not help in ruling out this possibility—*WI does not protect against a man-in-the-middle adversary who is trying to make the right WIPoK instance depend on the left WIPoK instance*. Instead, this is a “non-malleability” type of requirement rather than (plain) witness indistinguishability. In that sense, the above argument did not really address the issue of non-malleability; Rather, it simply “pushed” the non-malleability requirement to the WIPoK used in the final **WIPoK-2** stage.

Thus, new ideas seem necessary to enforce non-malleability. This leads us to the real virtue of our design—In the sequel, we will demonstrate that the t -solution hard puzzle introduces “asymmetry” between the right and the left sessions as $t < \tilde{t}$; This will allow us to perform a pigeon-hole-style argument, which can be used to address the heart of the non-malleability problem.

A Pigeon-Hole Argument. To address the above issue, we consider another hybrid \mathcal{K}_1 shown in Fig. 2b. Hybrid \mathcal{K}_1 is similar to \mathcal{G}_1 but it does not rewind the left **WIPoK-1** for witness extraction (we will explain shortly which witness \mathcal{K}_1 will use to perform the left **WIPoK-2**).

We first make a crucial observation: There are \tilde{t} preimages $\{\tilde{x}_1, \dots, \tilde{x}_{\tilde{t}}\}$ that R could potentially use as the witness to perform the right **WIPoK-1**. In contrast, there are at most t preimages $\{x_1, \dots, x_t\}$ that \mathcal{M} could potentially use as the witness to perform the left **WIPoK-1**. Since the **WIPoK-1** stage itself is not “non-malleable”, it is possible that \mathcal{M} ’s witness used there depends on R ’s witness used in the right **WIPoK-1**.⁸ For example, if R uses \tilde{x}_1 to perform the right **WIPoK-1**, \mathcal{M} may use, say, x_3 in the left **WIPoK-1**; If R instead uses \tilde{x}_2 , \mathcal{M} may switch to x_5 in the left **WIPoK-1**.

⁸ Again, standard witness-indistinguishability does not rule out such dependency.

Now, recall that $t \leq \tilde{t} - 1$ (because we have the **One-Sided** condition). It then follows from the pigeon-hole principle that there must exist a distinct pair $\tilde{x}_i, \tilde{x}_j \in \{\tilde{x}_1, \dots, \tilde{x}_{\tilde{t}}\}$ and an $x_k \in \{x_1, \dots, x_t\}$ such that $(\tilde{x}_i, \tilde{x}_j, x_k)$ form the following correspondence: *No matter R uses \tilde{x}_i or \tilde{x}_j in the right **WIPoK-1**, \mathcal{M} always uses (the same) x_k in the left **WIPoK-1**.*

If we assume that \mathcal{K}_1 somehow “magically” knows this pigeon-hole tuple $(\tilde{x}_i, \tilde{x}_j, x_k)$, then we can prove $w' = \tilde{m}$ using the following argument. Consider two scenarios:

1. When R uses \tilde{x}_i in the right **WIPoK-1** and C (or \mathcal{K}_1) uses x_k in the left **WIPoK-2**, we can prove that the extracted w' can only take the values of \tilde{m} or \tilde{x}_i . Intuitively, this is because \tilde{m} and \tilde{x}_i are the only witnesses that \mathcal{M} could potentially use for the right **WIPoK-2**, assuming it cannot break the one-wayness of the right OWF \tilde{f} . We suppress the full proof in this informal discussion, as we will show a formal (and slightly different) argument in the next subsection.
2. When R uses \tilde{x}_j in the right **WIPoK-1** and C uses (the same) x_k in the left **WIPoK-2**, we can prove that the extracted w' can only take the values of \tilde{m} or \tilde{x}_j . This follows from a similar argument as in the above bullet.

Next, observe that \mathcal{K}_1 does *not* rewind the **WIPoK-1** stage. By the witness indistinguishability of the right **WIPoK-1**, it follows that the witness used by \mathcal{M} in the right **WIPoK-2** cannot change when R switches between \tilde{x}_i and \tilde{x}_j in the right **WIPoK-1** (otherwise, we can invoke the knowledge extractor of the right **WIPoK-2** to extract the used witness to detect this change). This, together with the above two bullets, implies that the extracted w' can only take the value \tilde{m} if R uses \tilde{x}_i (or \tilde{x}_j) in the right **WIPoK-1**. Notice that in this argument, we do not rely on the “non-malleability” (or even WI) of the final **WIPoK-2** stage, because in both scenarios described above, C uses the same x_k in the left **WIPoK-2**.

Although the above reasoning looks promising, it suffers from the following obstacles:

1. In the above, we assumed that \mathcal{K}_1 “magically” knows the pigeon-hole tuple $(\tilde{x}_i, \tilde{x}_j, x_k)$. However, it is unclear how \mathcal{K}_1 could (efficiently) learn this tuple. Note that \mathcal{K}_1 cannot try to learn x_k by rewinding (i.e., running the knowledge extractor for) the left **WIPoK-1**; Otherwise, we will be back to hybrid \mathcal{G}_1 and cannot rely on the WI property of the right **WIPoK-1** (as we did in the previous paragraph).
2. The above pigeon-hole correspondence is not accurate. In fact, it is possible that \mathcal{M} changes the *distributions* (instead of *concrete values*, as in the previous discussion) of the witness used in the left **WIPoK-1**. For example, when R uses \tilde{x}_1 in the right **WIPoK-1**, \mathcal{M} samples a witness for the left **WIPoK-1** uniformly at random from $\{x_1, \dots, x_t\}$; While R switches to \tilde{x}_2 , \mathcal{M} may decide to sample a witness according to the Gaussian distribution instead. In this case, even with unbounded computational power, it is unclear how to identify (or define) the tuple $(\tilde{x}_i, \tilde{x}_j, x_k)$.

Nevertheless, this hybrid \mathcal{K}_1 divulges the usefulness of the t -solution hard puzzle. It provides an approach to argue $w' = \tilde{m}$, thus having a potential to help us establish non-malleability. In the following, we show our main technical lemma which overcomes the above two obstacles, allowing us to formalize the above argument properly.

The Main Technical Lemma. Let us summarize what we have so far. We described two hybrids: Hybrid \mathcal{G}_1 seems to help us reduce non-malleability to the computational-hiding property of the left Naor’s commitment. But it actually pushes the non-malleability requirement to the final **WIPoK-2** stage (or more technically, we are not able to prove that the extracted w' is equal to the right-side committed value \tilde{m}). In contrast, hybrid \mathcal{K}_1 does not suffer from this issue. It instead relies on a pigeon-hole-style argument, which employs the structure of the t -solution hard puzzle to enforce non-malleability. However, to formalize this argument, one has to (i) figure out how \mathcal{K}_1 can learn

the pigeon-hole tuple (**Obstacle 1**), and (ii) perform the pigeon-hole argument in a “distributional” manner (**Obstacle 2**).

In the following, we present our main technical lemma. This lemma can be understood as a combination of \mathcal{G}_1 and \mathcal{K}_1 —In its proof, we will define analogs of these two hybrids (i.e., the $\{\mathcal{G}_i\}_{i \in [\tilde{t}]}$ and $\{\widehat{\mathcal{K}}_i\}_{i \in [\tilde{t}]}$ in **Fig. 3**) and “jump between” them to take advantage of both hybrids, while performing the pigeon-hole argument in a “distributional” fashion.

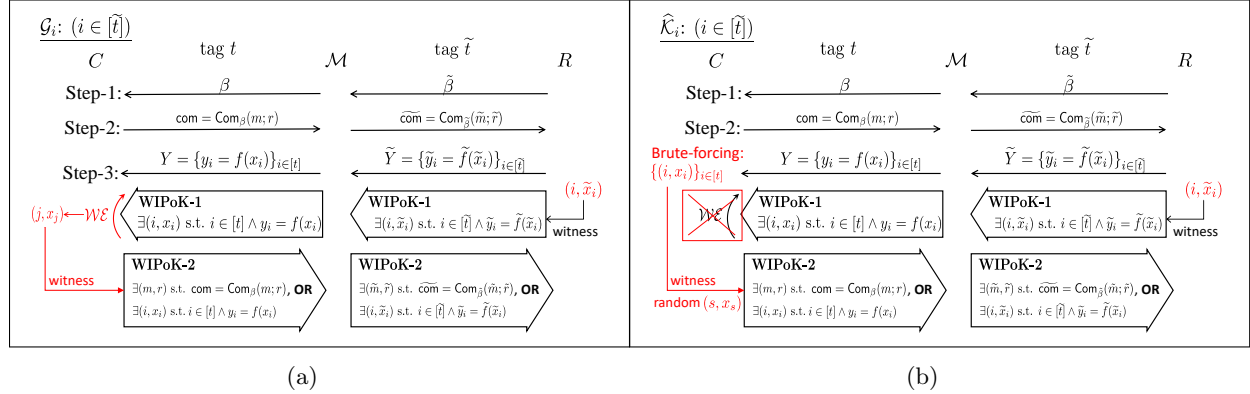


Fig. 3: Machines \mathcal{G}_i and $\widehat{\mathcal{K}}_i$ (Difference is highlighted in red color)

First, we define machines \mathcal{G}_i for each $i \in [\tilde{t}]$ (as shown in **Fig. 3a**). Machine \mathcal{G}_i for $i \neq 1$ behaves identically to the \mathcal{G}_1 described earlier (**Fig. 2a**), except that \mathcal{G}_i uses (i, \tilde{x}_i) as the witness when performing the right **WIPoK-1**. Again, we define the *extracted witness* w' to be the witness extracted by running the knowledge extractor for the right **WIPoK-2** in \mathcal{G}_i . Then, we show the main technical lemma, which we will use to establish non-malleability later.

Lemma 1 (Informal Version of Lem. 9). *Assume that in the real MIM execution, \mathcal{M} convinces the honest right receiver R with some noticeable probability $p(\lambda)$. Then, there exists some $i \in [\tilde{t}]$ such that the extracted witness w' in \mathcal{G}_i must be a valid opening (\tilde{m}, \tilde{r}) for $\widehat{\text{com}}$ with another noticeable probability $p'(\lambda)$.*

Proof of Lem. 1. We prove **Lem. 1** by contradiction. In the following, we assume for contradiction that for all $i \in [\tilde{t}]$, the extracted w' contains a valid opening to $\widehat{\text{com}}$ with at most *non-noticeable* probability. (For simplicity, we ignore the difference between “non-noticeable” and “negligible” and use the term “negligible” in place of “non-noticeable” in this proof.)

First, we claim that in game \mathcal{G}_i (for any $i \in [\tilde{t}]$), R must also be convinced with probability $p \pm \text{negl}(\lambda)$. This claim follows from the PoK property of the left **WIPoK-1** and the WI property of the right **WIPoK-1** and the left **WIPoK-2**. Since this proof is rather standard, we do not elaborate on it in this overview. More details can be found in **Sec. 4.5**.

Next, we observe that in \mathcal{G}_i , the extracted witness w' can only take the values (\tilde{m}, \tilde{r}) or (i, \tilde{x}_i) ; Otherwise, we break the one-wayness of the right OWF \tilde{f} (this follows from the same argument we made earlier). However, our assumption (for contradiction) says that in each \mathcal{G}_i , $w' \neq (\tilde{m}, \tilde{r})$ except for with negligible probability. Therefore, we obtain the following inequality:

$$\forall i \in [\tilde{t}], \quad \left| \Pr[w' = (i, \tilde{x}_i) \text{ in } \mathcal{G}_i] - p \right| \leq \text{negl}(\lambda). \quad (1)$$

Next, we define $\{\widehat{\mathcal{K}}_i\}_{i \in [\tilde{t}]}$ (shown in Fig. 3b)⁹. For each $i \in [\tilde{t}]$, $\widehat{\mathcal{K}}_i$ behaves identically as \mathcal{G}_i except that $\widehat{\mathcal{K}}_i$ does not invoke the knowledge extractor for the left **WIPoK-1**; Instead, it obtains all the preimages $\{x_i\}_{i \in [t]}$ for $\{y_i\}_{i \in [t]}$ in the left **Step-3** by *brute force*; It then picks a random index $s \stackrel{\$}{\leftarrow} [t]$ and uses preimage (s, x_s) as the witness to conduct the left **WIPoK-2**. First, observe that if the (s, x_s) picked by $\widehat{\mathcal{K}}_i$ hits the (j, x_j) extracted from the left **WIPoK-1** in \mathcal{G}_i (see Fig. 3a), then the games $\widehat{\mathcal{K}}_i$ and \mathcal{G}_i are identical (modulo that $\widehat{\mathcal{K}}_i$ is inefficient). Since $\widehat{\mathcal{K}}_i$ picks s uniformly at random from $[t]$, we know that $\widehat{\mathcal{K}}_i$ will be identical to \mathcal{G}_i with probability $\geq 1/t$. It then follows from Inequality (1) that

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) \text{ in } \widehat{\mathcal{K}}_i] \geq \frac{1}{t} \cdot p - \text{negl}(\lambda). \quad (2)$$

Next, observe that in $\widehat{\mathcal{K}}_i$, **WIPoK-1** is not rewound anymore. Therefore, by the (non-uniform, to compensate for the brute-forcing step) WI property of the right **WIPoK-1**, it follows that

$$\forall i \in [\tilde{t}], \left| \Pr[w' = (i, \tilde{x}_i) \text{ in } \widehat{\mathcal{K}}_i] - \Pr[w' = (i, \tilde{x}_i) \text{ in } \widehat{\mathcal{K}}_1] \right| \leq \text{negl}(\lambda). \quad (3)$$

Combining Inequalities (2) and (3), we have

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) \text{ in } \widehat{\mathcal{K}}_1] \geq \frac{1}{t} \cdot p - \text{negl}(\lambda). \quad (4)$$

Inequality (4) implies the following

$$\Pr[w' \text{ is a valid witness in } \widehat{\mathcal{K}}_1] \geq \tilde{t} \cdot \frac{1}{t} \cdot p - \text{negl}(\lambda) \geq p + \frac{p}{t} - \text{negl}(\lambda), \quad (5)$$

where “ w' is a valid witness” refers to the event that $w' = (\tilde{m}, \tilde{r}) \vee w' = (1, \tilde{x}_1) \vee \dots \vee w' = (\tilde{t}, \tilde{x}_{\tilde{t}})$, and the last “ \geq ” follows from the requirement that $t < \tilde{t}$ (i.e., we are in the **One-Sided** setting).

On the other hand, we claim that

$$\Pr[w' \text{ is a valid witness in } \widehat{\mathcal{K}}_1] \leq p + \text{negl}(\lambda). \quad (6)$$

This can be seen by comparing $\widehat{\mathcal{K}}_1$ with the real MIM execution shown in Fig. 1: Recall that p is the probability of R being convinced *in the real MIM execution*. The only difference between $\widehat{\mathcal{K}}_1$ and the real MIM execution is the witness used in the left **WIPoK-2**. By the (non-uniform) WI property of the left **WIPoK-2**, we know that R must be convinced in $\widehat{\mathcal{K}}_1$ with probability $\leq p + \text{negl}(\lambda)$. Then, by the proof of knowledge property, we must extract a valid witness in the right **WIPoK-2** in $\widehat{\mathcal{K}}_1$ with probability upper-bounded by $p + \text{negl}(\lambda)$ as well (this is exactly Inequality (6)).

Observe that Inequality (6) contradicts Inequality (5) because p/t is non-negligible. This gives us the desired contradiction, thus finishing the proof of Lem. 1.

Completing the Proof of Non-Malleability. With Lem. 1, non-malleability can be reduced to the computational-hiding property of Naor’s commitment as follows. Assuming that \mathcal{M} breaks the non-malleability of our scheme w.r.t. a distinguisher D , we construct an adversary $\mathcal{A}_{\text{hiding}}$ against the computational-hiding property of Naor’s commitment as follows:

⁹ For readers trying to find a correspondence between this technical overview and its main-body counterpart (i.e., Sec. 4), we would like to point out that the $\widehat{\mathcal{K}}_i$ defined here does not match any machine in Sec. 4. In the current overview, $\widehat{\mathcal{K}}_i$ actually serves the functionality of several machines defined in Sec. 4.5 to 4.7. But $\widehat{\mathcal{K}}_i$ is most similar to the \mathcal{K}_i'' on Page 40 (and depicted in Fig. 9b), which is essentially $\widehat{\mathcal{K}}_i$ but additionally runs a knowledge extractor for the right **WIPoK-2**.

- **Main-Thread Simulation:** $\mathcal{A}_{\text{hiding}}$ runs \mathcal{G}_1 where it embeds the instance of the hiding game of Naor’s commitment in the left **Step-1** and **Step-2**, and finishes other steps just as \mathcal{G}_1 . If R rejects, it sets $\tilde{m} := \perp$. (Notice that at the end of this step, \mathcal{M} will give an output that is computationally indistinguishable with the one from the real MIM execution. This is due to the similarity between \mathcal{G}_1 and the real MIM execution.)
- **Rewinding:** Unless R rejects in the above, $\mathcal{A}_{\text{hiding}}$ repeats the following N (to be specified later) times to extract the message \tilde{m} committed in the right session:
 - It rewinds \mathcal{M} to the point right after the completion of **Step-2**.
 - It picks a random index $i \xleftarrow{\$} [\tilde{t}]$, and finishes the execution in the same manner as \mathcal{G}_i (depicted in Fig. 3a).
 - Extract a witness \tilde{w} from the right **WIPoK-2** of the simulated execution.
 - If $\tilde{w} = (\tilde{m}, \tilde{r})$ is a valid opening to $\widetilde{\text{com}}$, it breaks the loop. Otherwise, it continues.

If it fails to extract \tilde{m} within N trials, it aborts and outputs a random guess.
- **Decision:** It runs the distinguisher D on \mathcal{M} ’s final output at the end of the **Main-Thread Simulation** step and \tilde{m} , and outputs whatever D outputs.

It is easy to see that the above reduction works unless it fails to extract \tilde{m} . We show that the failure probability can be made an arbitrarily small noticeable $\varepsilon(\lambda)$ if we take sufficiently large $N = \text{poly}(\lambda)$ depending on ε . We call the snapshot (including the transcript and \mathcal{M} ’s internal state) at the end of **Step-2** of the main thread a *prefix*. For each prefix **pref**, let p_{pref} be the probability that the right receiver accepts in \mathcal{G}_1 starting from **pref**. Then, [Lem. 1](#) ensures that¹⁰ for any fixed prefix **pref** such that $p_{\text{pref}} \geq \varepsilon(\lambda)$, if we run \mathcal{G}_i for random $i \xleftarrow{\$} [\tilde{t}]$ and extract \tilde{w} from the right **WIPoK-2** (as is done in each repetition described in the above **Rewinding** step), then we will extract $\tilde{w} = (\tilde{m}, \tilde{r})$ with at least another noticeable probability $\varepsilon'(\lambda)$ (More accurately, it should be the $\varepsilon'(\lambda)$ guaranteed by [Lem. 1](#) *divided by* \tilde{t} as we pick an i randomly from $[\tilde{t}]$. But this does not affect the argument here as \tilde{t} is bounded by a polynomial on λ —recall that we are in the **Small-Tag** setting.) Now, we consider the following two cases:

- The Case of $p_{\text{pref}} \geq \varepsilon(\lambda)$: In this case, if we set $N = \Omega(\varepsilon'(\lambda)^{-1} \cdot \lambda)$, the reduction algorithm fails with at most a negligible probability (i.e., $(1 - \varepsilon'(\lambda))^{\Omega(\varepsilon'(\lambda)^{-1} \cdot \lambda)}$).
- The Case of $p_{\text{pref}} < \varepsilon(\lambda)$: In this case, the right R rejects in the main-thread except for a probability upper-bounded by $\varepsilon(\lambda)$. Thus, the reduction fails with probability at most $\varepsilon(\lambda)$.

Overall, the reduction algorithm works with an additive loss of $\varepsilon(\lambda)$. Since we have the freedom to set $\varepsilon(\lambda)$ to an arbitrarily small noticeable function, the reduction from non-malleability to the computational hiding of Naor’s commitment can be done properly. We omit further details.

Why Our Proof Is “Quantum-Friendly”. Before going on, let us explain why the above proof of non-malleability in the classical setting could extend to the post-quantum setting. At a high-level, our security proof enjoys the following properties:

¹⁰ Strictly speaking, there are two minor differences from [Lem. 1](#). First, we are considering experiments for each fixed **pref** whereas [Lem. 1](#) considers the whole experiment. This is not an issue since the proof of [Lem. 1](#) does not touch messages before **Step-3** and thus works for any fixed **pref**. Another difference is that we define p_{pref} w.r.t. \mathcal{G}_1 whereas p is defined w.r.t. the real MIM experiment in [Lem. 1](#). This is not an issue since \mathcal{G}_1 is computationally indistinguishable from the real MIM experiment.

- The rewinding-extraction procedure described on [Page 12](#) never goes to touch the left Naor’s commitment, which is the “hiding gadget” that provides computational hardness for the reduction. This property saves us from the robust simulatable extractability issue described in [Sec. 1.1](#).
- Moreover, this extraction procedure is to first perform the “main-thread” execution as hybrid \mathcal{G}_1 ; If this main thread is accepted, then it starts N “rewinding threads” for extraction. This structure is very similar to the classical extractor for the canonical three-round extractable commitments in [\[PW09\]](#). Thus, it allows us to invoke (a generalization of) the quantum simulatable-extraction lemma from [\[CCLY22\]](#) (developed originally to quantize the [\[PW09\]](#) extractable commitments) to quantize the our proof of non-malleability. More details are provided in [Sec. 2.3](#).

In contrast, no previous constructions of (classical) constant-round non-malleable commitments achieved the above two properties *simultaneously*. This is why it is hard to quantize their security proofs *even with the [CCLY22] lemma in hand*.

On the Necessity of N Rewindings. Given that [Lem. 1](#) already guarantees a noticeable probability $p'(\lambda)$ for successful extraction, one may wonder why we need to perform N rewindings in the above reduction to the computational hiding of Naor’s commitment. That is, machine $\mathcal{A}_{\text{hiding}}$ seems to work as desired even if we set $N = 1$ —If it extracts \tilde{m} , then the reduction is done; Otherwise, simply ask the reduction to guess at random in the Naor’s hiding game. At the first glance, this seems to give a $\frac{1}{2} + p'(\lambda)$ advantage, which is sufficient for the security reduction.

Though this strategy looks appealing due to its simplicity, it does not work. To explain the reason, we need to refer to [Lem. 9](#) on page [Page 31](#) (i.e., the formal version of [Lem. 1](#)). The explanation (provided below) is rather technical and orthogonal to the current discussion. The reader may feel free to skip it and continue at [Sec. 2.3](#) directly.

Notice that [Lem. 9](#) refers to the prefix pref of the MIM execution, which consists of the first two rounds of the protocol. In particular, pref includes the left Naor’s commitment from the hiding challenger. [Lem. 9](#) says that for any noticeable probability ε , if a pref leads to an accepting MIM execution with probability greater than ε , then the extraction succeeds with a probability *lower-bounded* by ε'/t , a value polynomially related to (but smaller than) ε .

At a high level, [Lem. 9](#) is not directly applicable in our reduction because there may be a correlation between \mathcal{M} ’s view and the success of extraction. That is, conditioned on the success of extraction, the *joint distribution* of \mathcal{M} ’s view and the extracted right value may be distinguishable from those in the real MIM experiment. Consider the following distinguishing example: there are two prefixes pref_1 and pref_2 that lead to an accepting execution with the same probability ε , and they appear with the same probability in the real experiment. Suppose that the extraction success probability is ε'/t for pref_1 , but $\varepsilon (> \varepsilon'/t)$ for pref_2 . Then, if we run the reduction with [Lem. 9](#) directly, the distinguisher may get to see pref_2 much more often than pref_1 (conditioned on successful extraction). But they should happen with the same probability in the real experiment.

The N rewindings in our construction essentially “amplify” [Lem. 9](#) to the following stronger version: For any noticeable ε , if \mathcal{M} convinces R with probability greater than ε , then the extraction succeeds with overwhelming probability (i.e., $1 - \text{negl}(\lambda)$). This effectively ensures the following: Conditioned on a prefix that leads to acceptance with probability greater than ε *and* that R is indeed convinced, the joint distribution of \mathcal{M} ’s view and the extracted \tilde{m} are negligibly close to their distribution in the real MIM execution. (Other cases are easy to handle: If R is not convinced, we simply use \perp as the extracted value; Also, by an averaging argument, no more than ε fraction of prefixes can lead to acceptance with probability $< \varepsilon$.) Thus, our strategy eventually makes the

joint distribution from simulation to be ε -close to the real one. Since we have the freedom to pick any noticeable ε , our reduction can achieve a non-negligible advantage in Naor’s hiding game.

We refer the reader to [Sec. 4.4](#) for a formal treatment of this issue.

2.3 Small-Tag, One-Sided, Synchronous, Post-Quantum Setting

Next, we explain how to make the $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ described in [Sec. 2.1](#) non-malleable in the *post-quantum* setting.

For that, we simply instantiate $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ with post-quantum building blocks. (We denote this post-quantum version by $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$.) Here, we have to be careful about what security notion we should require for WIPoK. In the security proof in the classical setting, we often extract a witness from WIPoK while continuing the rest of the experiment. For this to work, we implicitly require the knowledge extractor to have a *simulation* property [[HSS11](#), [LN11](#)], i.e., it can simulate the prover’s internal state as in a real execution, while performing the extraction task. While this is almost trivial in the classical setting, it becomes hard in the post-quantum setting due to the no-cloning theorem. In particular, constant-round constructions for such “simulatable” post-quantum WIPoKs are not known from (polynomially-hard) post-quantum OWFs.¹¹

Fortunately, a recent work [[CCLY22](#)] shows that a constant-round construction from post-quantum OWFs is possible if we (i) relax the PoK to *argument* of knowledge (AoK), which only requires extractability against (quantum) polynomial-time adversaries, and (ii) relax the simulation requirement to the ε -close simulation property, where an (arbitrarily small) noticeable simulation error is allowed. Next, we show that such a ε -simulatable WIAoK suffices for our purpose. First, the relaxation from PoK to AoK is totally fine since we only apply the knowledge extractors for polynomial-time adversaries (possibly with non-uniform advice). Second, a noticeable error coming from extraction does not affect the proof of non-malleability if we take the noticeable error to be much smaller than the assumed MIM adversary’s advantage. Thus, we ignore the noticeable errors, and assume that the error is negligible in the rest of this overview.

By using such a post-quantum WIAoK (as well as the post-quantum version of Naor’s commitment Com and a post-quantum injective OWF f), we can see that the main technical part of the security proof in the classical setting (the proof of [Lem. 1](#)) can be migrated to the post-quantum setting almost immediately.

The only non-trivial issue is how to complete the reduction to the computational hiding of Naor’s commitment, assuming (a post-quantum analog of) [Lem. 1](#). In the classical setting, we rely on a rewinding argument that is not applicable anymore when \mathcal{M} is a quantum machine (again, due to the no-cloning theorem). Thus, we use a different argument here—We observe that the following lemma, which is a generalization of [[CCLY22](#), Lemma 4], suffices for this step.

Lemma 2 (Extract-and-Simulate Lemma (Informal)). *Let \mathcal{G} and \mathcal{K} be QPT algorithms that satisfy the following:¹²*

1. \mathcal{K} takes 1^λ and a quantum state ρ as input and outputs some unique s^* or otherwise \perp .
2. For any noticeable function $\gamma(\lambda)$, there exists a noticeable function $\delta(\lambda)$ such that for any polynomial-size quantum state ρ , if

$$\Pr \left[b = \top : (b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, \rho) \right] \geq \gamma(\lambda),$$

¹¹ A recent work by Lombardi, Ma, and Spooner [[LMS21](#)] gave such WIPoKs from super-polynomial hardness of post-quantum OWFs.

¹² In the formal version of this lemma ([Lem. 20](#)), \mathcal{G} and \mathcal{K} additionally take $1^{\gamma^{-1}}$ as input so that they can run extractors with simulation errors depending on γ . We ignore this in this overview since we assume negligibly-close simulation for simplicity.

then

$$\Pr[\mathcal{K}(1^\lambda, \rho) = s^*] \geq \delta(\lambda).$$

Then, there exists a QPT algorithm \mathcal{SE} such that for any polynomial-size quantum state ρ and noticeable function $\varepsilon = \varepsilon(\lambda)$,

$$\{\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, \rho)\}_\lambda \stackrel{s}{\approx}_\varepsilon \{(\rho_{\text{out}}, \Gamma_b(s^*)) : (b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, \rho)\}_\lambda \quad (7)$$

where

$$\Gamma_b(s^*) := \begin{cases} s^* & \text{if } b = \top \\ \perp & \text{otherwise} \end{cases}.$$

The above lemma is a generalization of [CCLY22, Lemma 4] and can be proven in a similar manner.¹³ We will explain the main intuition for this proof at the end of this subsection.

Assuming a post-quantum analog of [Lem. 1](#), we can complete the reduction to the computational hiding of Naor's commitment by using [Lem. 2](#) as follows: We let ρ be the state of the experiment \mathcal{G}_1 (as defined in [Fig. 3a](#), [Sec. 2.1](#)) right after finishing **Step-2**; Let \mathcal{G} be the experiment that runs the rest of \mathcal{G}_1 starting from the state ρ ;¹⁴ At the end, \mathcal{G} outputs the final state of \mathcal{M} as ρ_{out} , and outputs the final decision of the right receiver R as b . We define \mathcal{K} as follows:

- Upon receiving a state ρ , simulate \mathcal{G}_i for $i \stackrel{\$}{\leftarrow} [\tilde{t}]$.
- Extract a witness \tilde{w} from the right **WIPoK-2** of the simulated execution.
- If $\tilde{w} = (\tilde{m}, \tilde{r})$ is a valid opening to $\widetilde{\text{com}}$, output \tilde{m} . Otherwise, output \perp .

It is worth noting that \mathcal{K} corresponds to the repeated steps in the rewinding argument in [Sec. 2.1](#). By construction, it is clear that \mathcal{K} outputs \tilde{m} whenever it does not output \perp (for an overwhelming fraction of $\tilde{\beta}$ by the statistical binding property of Naor's commitment). Therefore, it satisfies [Property 1](#) in [Lem. 2](#) with $s^* := \tilde{m}$ (for each fixed prefix). (We stress that we set $s^* = \tilde{m}$ instead of $s^* = (\tilde{m}, \tilde{r})$ for ensuring the uniqueness of s^* by the statistical binding property of Naor's commitment.) Moreover, noting that \mathcal{G}_1 is computationally indistinguishable from the real MIM experiment, [Lem. 1](#) implies that \mathcal{K} also satisfies [Property 2](#) in [Lem. 2](#). Thus, there exists a machine \mathcal{SE} that satisfies [Eq. \(7\)](#) for $s^* = \tilde{m}$. Using this \mathcal{SE} , we can construct an adversary $\mathcal{A}_{\text{hiding}}$ against the computational-hiding property of Naor's commitment as follows.

Assuming that \mathcal{M} breaks the non-malleability of our scheme w.r.t. a distinguisher D , we construct an adversary $\mathcal{A}_{\text{hiding}}$ against the computational-hiding property of Naor's commitment as follows.

- **Prefix Generation:** $\mathcal{A}_{\text{hiding}}$ runs the real MIM experiment¹⁵ for an adversary \mathcal{M} until **Step-2** finishes, where it embeds the instance of the hiding game of Naor's commitment in the left **Step-2** while simulating other steps of the left C and the right R honestly. Let ρ be the state of the experiment at this point.
- **Extract-and-Simulate:** It runs $\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, \rho)$ to extract \tilde{m}^* while simulating the output ρ_{out} of the experiment \mathcal{G}_1 . We remark that it succeeds in extracting \tilde{m}^* whenever R accepts in the simulated experiment except for probability ε .

¹³ Indeed, such a generalized version was mentioned in the technical overview of [CCLY22].

¹⁴ Similar to the classical setting in [Sec. 2.1](#), we stipulate that ρ does *not* include the left committer C 's internal state ST_C at the end of **Step-2** (otherwise, we cannot reduce non-malleability to the hiding of the left commitment). This requirement is valid as the remaining steps of \mathcal{G}_1 , which will be executed by \mathcal{G} , do not depend on ST_C .

¹⁵ Remark that the real MIM experiment and \mathcal{G}_1 are identical until **Step-2**.

– **Decision:** It runs the distinguisher D on ρ_{out} and \tilde{m} , and outputs whatever D outputs.

By [Lem. 2](#), the $(\rho_{\text{out}}, \tilde{m})$ given to D is indistinguishable from that in \mathcal{G}_1 (up to an error ε), which in turn is indistinguishable from that of the real MIM experiment. Since we have the freedom to set ε to any arbitrarily small noticeable function, this completes the reduction.

Proof Idea of [Lem. 2](#). We assume that ρ is a pure state w.l.o.g., and denote it by $|\psi\rangle$. By a similar usage of Jordan’s lemma as in [[CCY21](#), [CCLY22](#)], for any noticeable function $\delta(\lambda)$, we can decompose $|\psi\rangle$ as $|\psi\rangle = |\psi_{<\delta}\rangle + |\psi_{\geq\delta}\rangle$ such that

$$\Pr\left[\mathcal{K}(1^\lambda, |\psi_{<\delta}\rangle) = s^*\right] < \delta(\lambda) \tag{8}$$

and

$$\Pr\left[\mathcal{K}(1^\lambda, |\psi_{\geq\delta}\rangle) = s^*\right] \geq \delta(\lambda). \tag{9}$$

By the contraposition of [Property 2](#) of [Lem. 2](#) and [Eq. \(8\)](#), it holds that

$$\Pr\left[b = \top : (b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, |\psi_{<\delta}\rangle)\right] < \gamma(\lambda),$$

where $\gamma(\lambda)$ is a noticeable function that can be made arbitrarily small by taking sufficiently small $\delta(\lambda)$. That is, \mathcal{G} simply “rejects” (i.e., outputs $b = \perp$) except for a small probability $\gamma(\lambda)$, in which case $I_b(s^*) = \perp$. The simulation of this case is easy because we do not need to extract s^* except for probability $\gamma(\lambda)$, which can be made arbitrarily small. This shows [Lem. 2](#) for the case where $|\psi\rangle$ only has $|\psi_{<\delta}\rangle$ component.

On the other hand, by [Eq. \(9\)](#), we can extract s^* from $|\psi_{\geq\delta}\rangle$ with an overwhelming probability if we can repeat $\mathcal{K}(1^\lambda, |\psi_{\geq\delta}\rangle)$ for $\Theta(\delta(\lambda)^{-1}\lambda)$ times. Though such a repetition is not possible in general, the way of defining $|\psi_{\geq\delta}\rangle$ (as explained in [[CCLY22](#)]) enables us to do so. As a result, we can extract s^* from $|\psi_{\geq\delta}\rangle$ almost without disturbing the state by the Almost-as-Good-as-New lemma [[Aar05](#), Lemma 2.2] (a.k.a. the gentle measurement lemma), which shows [Lem. 2](#) for the case where $|\psi\rangle$ only has $|\psi_{\geq\delta}\rangle$ component.

In the above, we separately analyze $|\psi_{<\delta}\rangle$ and $|\psi_{\geq\delta}\rangle$. In fact, we can see from the way we define them that they do not interfere with each other during the above described extraction procedure. As a result, the case of a superposition of $|\psi_{<\delta}\rangle$ and $|\psi_{\geq\delta}\rangle$ can be reduced to the above two cases.

See [Lem. 20](#) for the formal statement of the lemma and [Appx. A](#) for the full proof.

2.4 Full-Fledged Post-Quantum Non-Malleable Commitments

To make $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$ a full-fledged post-quantum non-malleable commitments, we need to remove the one-side, small-tag, and synchronous restrictions. Due to space constraints, we will only briefly describe how to do that in this overview. We will provide more detailed overviews in the corresponding sections in the main body.

Removing the One-Sided Restriction. Notice that the non-malleability proof of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ (and $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$) makes use of the one-sided condition in both **Step-3** and **WIPoK-1**. Let us refer to these two steps as a **Slot**. As demonstrated in [Sec. 2.1](#), it is the asymmetry condition $t < \tilde{t}$ for this **Slot** that allows us to argue that the extracted w' must contain the right committed value \tilde{m} with good-enough probability.

We can use the following (standard) approach to remove the one-sided restriction. We simply repeat the **Slot** twice, sequentially. The first repetition, referred to as **Slot-A**, is identical to the

original **Slot**; In particular, both parties use the tag t . The second repetition, referred to as **Slot-B**, will instead use $(n - t)$ as the tag.

Now, in a MIM execution, if $t < \tilde{t}$, we can perform the same non-malleable proof as in [Sec. 2.1](#) (and [Sec. 2.3](#)) with **Slot-A** playing the role of **Slot**. Otherwise (i.e., $t > \tilde{t}$), it must hold that $(n - t) < (n - \tilde{t})$; In this case, we can perform the non-malleable proof again, with **Slot-B** playing the role of **Slot**. (Note that by the definition of non-malleability, we do not need to consider the case $t = \tilde{t}$.)

Handling Asynchronous Schedules. First, we observe that the proof strategy used in [Sec. 2.1](#), which is for the synchronous setting, relies crucially on the following three conditions:

1. Both **Step-1** and **Step-2** in both the left and right sessions should finish before **Step-3** in either of the two sessions starts. This is because our security proof considers each fixed prefix, which is the state right after the end of **Step-2** (and before the start of **Step-3**).
2. The right **WIPoK-2** should not be interleaved with the left **WIPoK-1**. This is because \mathcal{G}_i rewinds the left **WIPoK-1**; Meanwhile, we need to rewind the right **WIPoK-2** to extract \tilde{m} . We do not want these two parts to rewind each other recursively.
3. The left **Step-3** should happen before the right **WIPoK-2**. This is because in machine $\hat{\mathcal{K}}_i$, brute-forcing is performed for the left **Step-3**; At the same time, we need to rely on the WI property of the right **WIPoK-2** to argue the similarity among all the $\hat{\mathcal{K}}_i$'s for different $i \in [\tilde{t}]$ (i.e., [Inequality \(3\)](#)).

To prove non-malleability in the asynchronous setting, our strategy is to introduce more gadgets into the protocol that enforce the above three conditions. As long as these conditions are satisfied, we can then rely on essentially the same proof of non-malleability as in the synchronous case.

For example, to enforce the third condition above, we can add an extractable commitment **ExtCom** between **Step-3** and **WIPoK-1**, where C commits to m again (and additionally proves in **WIPoK-2** that he did it as required). Then, in the MIM execution, if the left **Step-3** happens after (the first message of) the right **WIPoK-2**, then the right **ExtCom** must happen before the left **Step-3**. In this case, we can extract \tilde{m} from the right **ExtCom** to finish the proof of non-malleability, instead of performing the same proof as in [Sec. 2.1](#). We will use similar ideas to enforce all the conditions mentioned above.

Of course, in the real proof, we need to do the above argument for the two slot-version described above in the **Removing the One-Sided Restriction** part. Moreover, whenever we add one more gadget, cautions are needed to resolve new problems that it may cause by interacting with the “already-introduced” gadgets in some undesirable way. We will elaborate on this in [Sec. 6](#), where we also provide a more detailed overview (in [Sec. 6.1](#)).

Tag Amplification. After the above two steps, we have already obtained a small-tag, asynchronous post-quantum non-malleable commitment $\langle C, R \rangle_{\text{tg,PQ}}^{\text{async}}$. As mentioned earlier, now we can apply the [\[BLS22\]](#) tag amplifier. Here is one more caveat—The tag amplifier requires $\langle C, R \rangle_{\text{tg,PQ}}^{\text{async}}$ to additionally satisfy a condition called *robust extractability*, which means that the extractability holds even if the cheating committer interacts with an external bounded-round machine that cannot be rewound. We show that our $\langle C, R \rangle_{\text{tg,PQ}}^{\text{async}}$ already satisfies it, thanks to the added instances of **ExtCom** (which was introduced to handle asynchronous schedules originally). Thus, we can apply their tag amplifier to our $\langle C, R \rangle_{\text{tg,PQ}}^{\text{async}}$ protocol to obtain the final full-fledged post-quantum non-malleable commitment scheme.

2.5 Organization

In [Sec. 3](#), we present necessary notations and preliminaries.

In [Sec. 4](#), we show the small-tag, one-sided, synchronous non-malleable commitment $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ in the classical setting.

In [Sec. 5](#), we remove the one-sided restriction, obtaining the small-tag, synchronous non-malleable commitment $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$ (still in the classical setting).

In [Sec. 6](#), we show how to handle asynchronous schedules *in the classical setting*. This leads to the small-tag, synchronous non-malleable commitment $\langle C, R \rangle_{\text{tg}}^{\text{async}}$. As mentioned earlier, the way we handle asynchronous schedules is quantum-friendly. We could have done this part directly for *the quantum version of* $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$; But we find that it conveys our ideas better to first show the techniques in the classical setting, because it can avoid confusing the reader with quantum phenomena that are orthogonal to the discussion of asynchronous schedules. Once one understands our techniques for asynchronous schedules in the classical setting, it is easy to see why it extends to the post-quantum setting.

In [Sec. 7](#), we show a post-quantum extract-and-simulate lemma, which plays an important role later when we quantize our classical proof of non-malleability.

In [Sec. 8](#), we describe the full-fledged post-quantum non-malleable commitment. To do that, we first show in [Sec. 8.1](#) how to quantize the non-malleable proof of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$. This gives its post-quantum analog $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$ (i.e., a small-tag, one-sided, synchronous, post-quantum non-malleable commitment). Next, we describe in [Sec. 8.2](#) how to remove the one-sided and synchronous restrictions from $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$, to obtain the protocol $\langle C, R \rangle_{\text{tg,PQ}}^{\text{async}}$; This step mimics the techniques we used in the classical setting converting $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ to $\langle C, R \rangle_{\text{tg}}^{\text{async}}$. Finally, in [Sec. 8.3](#), we obtain the final full-fledged post-quantum non-malleable commitment by applying the [\[BLS22\]](#) tag amplifier to $\langle C, R \rangle_{\text{tg,PQ}}^{\text{async}}$.

In [Sec. 9](#), we show the application of our post-quantum commitments to post-quantum MPC.

3 Preliminaries

3.1 Basic Notations

Let $\lambda \in \mathbb{N}$ denote security parameter. For a positive integer n , let $[n]$ denote the set $\{1, 2, \dots, n\}$. For a finite set \mathcal{X} , $x \leftarrow \mathcal{X}$ means that x is uniformly chosen from \mathcal{X} .

A function $f : \mathbb{N} \rightarrow [0, 1]$ is said to be *negligible* if for all polynomial p and sufficiently large $\lambda \in \mathbb{N}$, we have $f(\lambda) < 1/p(\lambda)$; it is said to be *overwhelming* if $1 - f$ is negligible, and said to be *noticeable* if there is a polynomial p such that $f(\lambda) \geq 1/p(\lambda)$ for sufficiently large $\lambda \in \mathbb{N}$. We denote by **poly** an unspecified polynomial and by **negl** an unspecified negligible function.

Honest (classical) parties are modeled as interactive Turing machines (ITMs). We use PPT and QPT to denote (classical) probabilistic polynomial time and quantum polynomial time, respectively. For a classical probabilistic or quantum algorithm \mathcal{A} , $y \leftarrow \mathcal{A}(x)$ means that \mathcal{A} is run on input x and outputs y . An adversary (or malicious party) is modeled as a non-uniform PPT (or QPT in the case of post-quantum security) algorithm. When we consider a non-uniform QPT adversary, we often specify it by a sequence of polynomial-size quantum circuits with quantum advice $\{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$. In an execution with the security parameter λ , \mathcal{A} runs \mathcal{A}_λ taking ρ_λ as the advice. For simplicity, we often omit the index λ and just write $\mathcal{A}(\rho)$ to mean a non-uniform QPT algorithm specified by $\{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$.

For an **NP** language \mathcal{L} and a true statement in this language $x \in \mathcal{L}$, we use $\mathcal{R}_{\mathcal{L}}(x)$ (\mathcal{R} stands for “relation”) to denote the set of all witnesses for x . We will refer to the OR-composition of **NP** languages, which are defined in [Def. 1](#).

Definition 1 (OR-Composition of NP Languages). Let \mathcal{L}_1 and \mathcal{L}_2 be two NP languages. The OR-composition of them (dubbed $\mathcal{L}_1 \vee \mathcal{L}_2$) is the new NP language defined as follows:

$$\mathcal{L}_1 \vee \mathcal{L}_2 := \{(x_1, x_2) \mid x_1 \in \mathcal{L}_1 \vee x_2 \in \mathcal{L}_2\}.$$

Notations for Indistinguishability. We may consider random variables over bit strings or over quantum states. This will be clear from the context. We use the same notations for classical and quantum computational indistinguishability, but there should be no fear of confusion; It means computational indistinguishability against PPT (resp. QPT) distinguishers whenever we consider classical (resp. post-quantum) security. For ensembles of random variables $\mathcal{X} = \{X_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$ and $\mathcal{Y} = \{Y_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$ over the same set of indices $I = \bigcup_{\lambda \in \mathbb{N}} I_\lambda$ and a function δ , we use $\mathcal{X} \stackrel{c}{\approx}_\delta \mathcal{Y}$ to mean that for any non-uniform PPT (resp. QPT) algorithm \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $i \in I_\lambda$, we have

$$|\Pr[\mathcal{A}(X_i)] - \Pr[\mathcal{A}(Y_i)]| \leq \delta(\lambda) + \text{negl}(\lambda).$$

We say that \mathcal{X} and \mathcal{Y} are δ -computationally indistinguishable if the above holds. In particular, when the above holds for $\delta = 0$, we say that \mathcal{X} and \mathcal{Y} are computationally indistinguishable, and simply write $\mathcal{X} \stackrel{c}{\approx} \mathcal{Y}$.

Similarly, we use $\mathcal{X} \stackrel{s}{\approx}_\delta \mathcal{Y}$ to mean that for any unbounded time algorithm \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $i \in I_\lambda$, we have

$$|\Pr[\mathcal{A}(X_i)] - \Pr[\mathcal{A}(Y_i)]| \leq \delta(\lambda) + \text{negl}(\lambda).$$

In particular, when the above hold for $\delta = 0$, we say that \mathcal{X} and \mathcal{Y} are statistically indistinguishable, and simply write $\mathcal{X} \stackrel{s}{\approx} \mathcal{Y}$. Moreover, we write $\mathcal{X} \equiv \mathcal{Y}$ to mean that X_i and Y_i are distributed identically for all $i \in I$.

When we consider an ensemble \mathcal{X} that is only indexed by λ (i.e., $I_\lambda = \{\lambda\}$), we write $\mathcal{X} = \{X_\lambda\}_\lambda$ for simplicity.

3.2 Commitment Schemes

We define (classically-secure and post-quantum) commitments. The following definitions are based on those in [CCLY22].

Definition 2 (Commitment). A commitment scheme is a pair of PPT ITMs $\langle C, R \rangle$ that satisfies the following. Let $m \in \{0, 1\}^{\ell(\lambda)}$ (where $\ell(\cdot)$ is some polynomial) is a message that C wants to commit to. The protocol consists of the following stages:

- **Commit Stage:** $C(m)$ and R interact with each other to generate a transcript τ , a decommitment decom as C 's private output, and a decision value $b_{\text{com}} \in \{\top, \perp\}$ as R 's private output. We denote this execution by $(\tau, \text{decom}, b_{\text{com}}) \leftarrow \langle C(m), R \rangle(1^\lambda)$.
- **Decommit Stage:** C sends the committed message m and the decommitment decom to R , and R outputs decision $b_{\text{dec}} \in \{\top, \perp\}$. We denote this execution by $b_{\text{dec}} \leftarrow \text{Verify}(\tau, m, \text{decom})$.¹⁶ W.l.o.g., we assume that whenever R rejects at the end of the commit stage (i.e., $b_{\text{com}} = \perp$), R will reject at the end of decommit stage (i.e., $b_{\text{dec}} = \perp$). (Note that w.l.o.g., one can think that b_{com} is included in τ , because we can always modify the protocol to ask R to send b_{com} as the last-round message of the commit stage.)

¹⁶ We assume that there is no state information for R kept from the commit stage.

The scheme satisfies the following completeness requirement:

1. **Completeness.** For any polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$ and any $m \in \{0, 1\}^{\ell(\lambda)}$, it holds that

$$\Pr \left[b_{\text{com}} = b_{\text{dec}} = \top : \begin{array}{l} (\tau, \text{decom}, b_{\text{com}}) \leftarrow \langle C(m), R \rangle(1^\lambda) \\ b_{\text{dec}} \leftarrow \text{Verify}(\text{com}, m, \text{decom}) \end{array} \right] = 1.$$

Definition 3 (Computationally Hiding). A commitment scheme $\langle C, R \rangle$ is computationally hiding if for any non-uniform PPT receiver R^* and any polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$, the following holds:

$$\left\{ \text{OUT}_{R^*} \langle C(m_0), R^* \rangle(1^\lambda) \right\}_{\lambda \in \mathbb{N}, m_0, m_1 \in \{0, 1\}^{\ell(\lambda)}} \stackrel{c}{\approx} \left\{ \text{OUT}_{R^*} \langle C(m_1), R^* \rangle(1^\lambda) \right\}_{\lambda \in \mathbb{N}, m_0, m_1 \in \{0, 1\}^{\ell(\lambda)}},$$

where $\text{OUT}_{R^*} \langle C(m_b), R^* \rangle(1^\lambda)$ ($b \in \{0, 1\}$) denotes the output of R^* at the end of the commit stage.

We say that $\langle C, R \rangle$ is post-quantum computationally hiding if the above holds for all non-uniform QPT R^* .

Definition 4 (Statistically Binding). A commitment scheme $\langle C, R \rangle$ is statistically binding if for any unbounded-time committer C^* , the following holds:

$$\Pr \left[\exists m_0, m_1, \text{decom}_0, \text{decom}_1, \text{ s.t. } m_0 \neq m_1 \wedge \begin{array}{l} \text{Verify}(\tau, m_0, \text{decom}_0) = \text{Verify}(\tau, m_1, \text{decom}_1) = \top \\ (\tau, \text{decom}, b_{\text{com}}) \leftarrow \langle C^*, R \rangle(1^\lambda) \end{array} \right] = \text{negl}(\lambda).$$

Definition 5 (Committed Values). For a commitment scheme $\langle C, R \rangle$, we define the value function as follows:

$$\text{val}(\tau) := \begin{cases} m & \text{if } \exists \text{ unique } m \text{ s.t. } \exists \text{ decom}, \text{Verify}(\tau, m, \text{decom}) = 1 \\ \perp & \text{otherwise} \end{cases},$$

where τ and decom are defined in the commit stage in [Def. 2](#).

3.3 Non-Malleable Commitments

Classically Non-Malleable Commitments. We first define non-malleable commitments in the classical setting. This definition follows the formalization in [\[LPV08, GPR16\]](#). We consider a man-in-the-middle adversary \mathcal{M} interacting with a committer C in the *left*, and a receiver R in the *right*. We denote the relevant entities used in the right interaction as the “tilde’d” version of the corresponding entities on the left. In particular, suppose that C commits to m in the left interaction, and \mathcal{M} commits to \tilde{m} on the right, i.e., we set $\tilde{m} = \text{val}(\tilde{\tau})$ where $\tilde{\tau}$ is the transcript of the right session. Let $\text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m, z)$ denote the random variable that is the pair $(\text{OUT}_{\mathcal{M}}, \tilde{m})$, consisting of \mathcal{M} ’s output as well as the value committed to by \mathcal{M} on the right (assuming C commits to m on the left), where z is \mathcal{M} ’s non-uniform advice. We use a *tag-based* (or “identity-based”) specification, and ensure that \mathcal{M} uses a distinct tag \tilde{t} on the right from the tag t it uses on the left. This is done by stipulating that $\text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m, z)$ outputs a special value \perp_{tag} when \mathcal{M} uses the same tag in both the left and right executions. The reasoning is that this corresponds to the uninteresting case when \mathcal{M} is simply acting as a channel, forwarding messages from C on the left to R on the right and vice versa.

Definition 6 (Classical Non-Malleable Commitments). A commitment scheme $\langle C, R \rangle$ is said to be (classically) non-malleable if for every (non-uniform) PPT man-in-the-middle adversary \mathcal{M} and every polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$, it holds that

$$\left\{ \text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m_0, z) \right\}_{\lambda \in \mathbb{N}, m_0, m_1 \in \{0, 1\}^{\ell(\lambda)}, z \in \{0, 1\}^*} \stackrel{c}{\approx} \left\{ \text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m_1, z) \right\}_{\lambda \in \mathbb{N}, m_0, m_1 \in \{0, 1\}^{\ell(\lambda)}, z \in \{0, 1\}^*}.$$

Post-Quantum Non-Malleable Commitments. In the post-quantum setting, non-malleability can be defined similarly, except that the man-in-the-middle adversary can be QPT, instead of being PPT. (In particular, the final output of the adversary could be a quantum state.) To avoid using confusing notations, we put an over-line on top of the variable mim in the post-quantum setting.

Definition 7 (Post-Quantum Non-Malleable Commitments). *A commitment scheme $\langle C, R \rangle$ is said to be post-quantumly non-malleable if for every (non-uniform) QPT man-in-the-middle adversary $\mathcal{M} = \{\mathcal{M}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ and every polynomial $\ell : \mathbb{N} \rightarrow \mathbb{N}$, it holds that*

$$\{\overline{\text{mim}}_{\langle C, R \rangle}^{\mathcal{M}_\lambda}(\lambda, m_0, \rho_\lambda)\}_{\lambda \in \mathbb{N}, m_0, m_1 \in \{0,1\}^{\ell(\lambda)}} \stackrel{c}{\approx} \{\overline{\text{mim}}_{\langle C, R \rangle}^{\mathcal{M}_\lambda}(\lambda, m_1, \rho_\lambda)\}_{\lambda \in \mathbb{N}, m_0, m_1 \in \{0,1\}^{\ell(\lambda)}}.$$

where “ $\stackrel{c}{\approx}$ ” refers to computational indistinguishability against QPT distinguishers.

Remark 1 (On Entangled Auxiliary Information.) The above definition does not consider entanglement between \mathcal{M} ’s auxiliary input and distinguisher’s auxiliary input. However, [BLS22, Claim 3.1] shows that the above definition implies the version that considers such entanglement.

Synchronous Adversaries: This notion refers to man-in-the-middle adversaries who upon receiving a message in the left (reps. right) session, immediately respond with the corresponding message in the right (resp. left) session. An adversary is said to be *asynchronous* if it is not synchronous.

3.4 Extractable Commitments

Classically Extractable Commitments. A commitment scheme is extractable if there exists an efficient extractor such that, the committed value can be extracted. We present the definition in Def. 8, which is taken from [PW09]. Constant-round constructions of (classically) extractable commitments are known from OWFs [PW09].

Definition 8 (Classically Extractable Commitments). *A commitment $\text{ExtCom} = \langle C, R \rangle$ is extractable if there exists an expected polynomial-time probabilistic oracle machine (the extractor) \mathcal{E} that given oracle access to any PPT cheating committer C^* outputs a pair (τ, σ^*) such that:*

- **Simulation:** τ is identically distributed to the view of C^* at the end of interacting with an honest receiver R in commitment phase.
- **Extraction:** the probability that τ is accepting and $\sigma^* = \perp$ is negligible.
- **Binding:** if $\sigma^* \neq \perp$, then it is statistically impossible to open τ to any value other than σ^* .

Post-Quantum (Robust) Extractable Commitments. We define the post-quantum analog of extractable commitments. In the post-quantum setting, we often need an extractor that (almost) does not disturb the (potentially malicious) committer’s state during the extraction. However, it is not known that such a post-quantum extractable commitments exist from (polynomially hard) post-quantum OWFs.¹⁷ Fortunately, a recent work [CCLY22] showed that a constant-round construction from post-quantum OWFs is possible if we relax the extractability to allow an (arbitrarily small) noticeable simulation error. The following definitions are taken from [BLS22] with some notational adaptations.

¹⁷ A recent work of [LMS21] gave a constant-round construction of such extractable commitments from super-polynomial hardness of post-quantum OWFs.

Let $\langle C, R \rangle$ be a (possibly tag-based) commitment scheme. A sequential committed-value oracle $O^\infty[\langle C, R \rangle]$ acts as follows in interaction with a committer C^* : it interacts with C^* in many sequential sessions; in each session,

- it participates with C^* in the commit phase of $\langle C, R \rangle$ as the honest receiver R (using a tag chosen adaptively by C^*), obtaining a transcript τ , **and**
- if C^* is non-aborting in the commit phase and sends request **break**, it returns $\text{val}(\tau)$.

The single-session oracle $O^1[\langle C, R \rangle]$ is similar to $O^\infty[\langle C, R \rangle]$, except that it interacts with the adversary in a single session. When the commitment scheme is clear from the context, we write O^∞ (or O^1) for simplicity.

Then, the definition of post-quantum ε -simulatable commitment is given below.

Definition 9 (Post-Quantum ε -Simulatable Extractable Commitment). *A commitment scheme $\langle C, R \rangle$ is said to be a post-quantum ε -simulatable extractable commitment if it satisfies the following in addition to post-quantum computational hiding and statistical binding. There exists a QPT algorithm \mathcal{SE} such that for any noticeable $\varepsilon(\lambda)$ and any non-uniform QPT committer $\{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$,*

$$\{\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, C_\lambda^*, \rho_\lambda)\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx}_\varepsilon \{C_\lambda^{*O^\infty}(\rho_\lambda)\}_{\lambda \in \mathbb{N}},$$

where “ $\stackrel{c}{\approx}_\varepsilon$ ” refers to ε -close computational indistinguishability against QPT distinguishers.

Lemma 3 ([CCLY22, Lemma 10]). *Assuming the existence of post-quantum OWFs, there exist constant-round post-quantum ε -simulatable extractable commitments.*

Remark 2. The ε -simulatable extractability as defined in Def. 9 is called post-quantum *strong* ε -simulatable extractability in [CCLY22]. Though their definition looks different from Def. 9, they are actually equivalent—First, it is easy to see that Def. 9 is equivalent to a modified version of Def. 9 where we give O^1 instead of O^∞ . Second, [BLS22, Lemma 3.2] shows that the O^1 and O^∞ versions of Def. 9 are equivalent.

We also define the *robust* version of the above definition following [BLS22]. Roughly speaking, r -robust ε -simulatable extractability means that the ε -simulatable extractability holds even if the adversary interacts with an external r -round interactive machine that cannot be rewound by the extractor. The formal definition is given below.

Definition 10 (Post-Quantum r -Robust ε -Simulatable Extractable Commitment). *A commitment scheme $\langle C, R \rangle$ is said to be a post-quantum r -robust ε -simulatable extractable commitment if it satisfies the following in addition to post-quantum computational hiding and statistical binding. There exists a QPT algorithm \mathcal{SE} such that for any noticeable $\varepsilon(\lambda)$, non-uniform QPT committer $\{C_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, and r -round interactive machine B ,*

$$\{\text{OUT}_{\mathcal{SE}}\langle B(1^\lambda, z), \mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, C_\lambda^*, \rho_\lambda) \rangle\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \stackrel{c}{\approx}_\varepsilon \{\text{OUT}_{C_\lambda^*}\langle B(1^\lambda, z), C_\lambda^{*O^\infty}(\rho_\lambda) \rangle\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}$$

where “ $\stackrel{c}{\approx}_\varepsilon$ ” refers to ε -close computational indistinguishability against QPT distinguishers.

3.5 Witness-Indistinguishable Proofs of Knowledge

Classically Witness-Indistinguishable Proofs of Knowledge. First, we provide a definition from [Gol01, Lin03].

Definition 11 (Proofs of Knowledge [Gol01, Lin03]). Let \mathcal{L} be an NP language and $\kappa : \mathbb{N} \rightarrow [0, 1]$. We say that a pair of PPT ITMs $\langle P, V \rangle$ is an interactive proof of knowledge for \mathcal{L} with knowledge error κ if the following two conditions hold:

1. **Completeness.** For every true statement $x \in \mathcal{L}$ and every witness $w \in \mathcal{R}_{\mathcal{L}}(x)$, it holds that

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

2. **Validity (with error κ).** There exists a polynomial $q(\cdot)$ and a probabilistic oracle machine K such that for every P^* , and every $x, z, r \in \{0, 1\}^*$, machine K satisfies the following condition: Denote by $p(x, z, r)$ the probability that $V(x)$ accepts when interacting with the prover specified by $P^*(x, z; r)$. If $p(x, z, r) > \kappa(|x|)$, then, on input x and with oracle access to oracle $P^*(x, z; r)$, K outputs a solution $w \in \mathcal{R}_{\mathcal{L}}(x)$ within an expected number of steps bounded by $\frac{q(|x|)}{p(x, z, r) - \kappa(|x|)}$.

We present an alternative formalism of knowledge soundness called *witness-extended emulation* due to Lindell [Lin03]. This formalism turns out to be more suitable for our application.

Definition 12 (PoK via Witness-Extended Emulation [Lin03]). Let \mathcal{L} be an NP language and $\kappa : \mathbb{N} \rightarrow [0, 1]$. We say that a pair of ITMs $\langle P, V \rangle$ is an interactive proof of knowledge for \mathcal{L} with witness-extended emulation if it satisfies the same completeness requirement as in Def. 11, and the following requirement: There exists an expected PPT oracle machine WE, called the emulation-extractor, such that for any machine P^* the following conditions hold

1. $\{(\text{ST}'_{P^*}, b') : (\text{ST}'_{P^*}, b', w') \leftarrow \mathcal{WE}^{P^*(x, z; r)}(x)\}_{\lambda, x, z, r} \equiv \{(\text{ST}_{P^*}, b) : (\text{ST}_{P^*}, b) \leftarrow \langle P^*(x, z; r), V(x) \rangle\}_{\lambda, x, z, r}$, where in the RHS above, ST_{P^*} and b denote the state of P^* and output of V respectively, at the end of the execution $\langle P^*(x, z; r), V(x) \rangle$.
2. $\Pr[b' = 1 \wedge w' \notin \mathcal{R}_{\mathcal{L}}(x) : (\text{ST}'_{P^*}, b', w') \leftarrow \mathcal{WE}^{P^*(x, z; r)}(x)] = \text{negl}(|x|)$.

The following lemma connects Def. 12 with Def. 11.

Lemma 4 ([Lin03, Lemma 3.1]). Let $\langle P, V \rangle$ be a proof of knowledge for $\mathcal{L} \in \text{NP}$ with negligible knowledge error. Then, $\langle P, V \rangle$ is also a proof of knowledge for \mathcal{L} with witness-extended emulation.

Next, we define witness indistinguishable proof of knowledge:

Definition 13 (WIPoKs). A witness-indistinguishable proof of knowledge for an NP language \mathcal{L} is a proof of knowledge for \mathcal{L} with witness extended emulation (as per Def. 12) which additionally satisfies the following property

- **Witness-Indistinguishability.** For any non-uniform PPT machine V^* , we have

$$\{\text{OUT}_{V^*} \langle P(w_0), V^* \rangle(x)\}_{\lambda, x, w_0, w_1} \stackrel{c}{\approx} \{\text{OUT}_{V^*} \langle P(w_1), V^* \rangle(x)\}_{\lambda, x, w_0, w_1},$$

where $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, and $w_0, w_1 \in \mathcal{R}_{\mathcal{L}}(x)$.

It is well-known that constant-round witness-indistinguishable proofs of knowledge (as per Def. 13) can be built from one-way functions.

Post-Quantum Witness-Indistinguishable Arguments of Knowledge We define a post-quantum analog of the WIPoKs described in Def. 13. A natural way to do so is to just replace all adversaries with any QPT algorithms. More precisely, we require that the witness indistinguishability holds against QPT distinguishers and that the knowledge extractor extracts a witness without

disturbing the prover’s state. However, a constant-round construction for such WIPoKs is not known based on (polynomially hard) post-quantum OWFs.¹⁸ Fortunately, a recent work [CCLY22] showed that a constant-round construction from post-quantum OWFs is possible if we relax the proof of knowledge via witness-extended emulation to *argument* of knowledge via witness-extended ε -close emulation, where a cheating prover is limited to QPT and a (arbitrarily small) noticeable simulation error is allowed. A formal definition of such post-quantum WIAoKs is given below.

Definition 14 (Post-Quantum WIAoK with ε -Close Emulation). *Let \mathcal{L} be an NP language. We say that a pair of PPT ITMs $\langle P, V \rangle$ is a post-quantum witness-distinguishable arguments of knowledge with ε -close emulation for \mathcal{L} if the following three conditions hold:*

1. **Completeness.** *For every $(x, w) \in \mathcal{L}$, it holds that*

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

2. **AoK via Witness-Extended ε -Close Emulation.** *There exists a QPT oracle machine \mathcal{WE} , called the witness-extended emulator, such that for any non-uniform QPT machine $\{P_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ the following conditions hold: for any noticeable function $\varepsilon(\cdot)$ (referred to as the error parameter),*

$$(a) \{(\text{ST}'_{P^*}, b') : (\text{ST}'_{P^*}, b', w') \leftarrow \mathcal{WE}(1^{\varepsilon^{-1}}, P_\lambda^*, \rho_\lambda, x)\}_{\lambda, x} \stackrel{c}{\approx}_\varepsilon \{(\text{ST}_{P^*}, b) : (\text{ST}_{P^*}, b) \leftarrow \langle P_\lambda^*(\rho_\lambda), V(x) \rangle\}_{\lambda, x}$$

where $\lambda \in \mathbb{N}$, $x \in \{0, 1\}^\lambda$, and “ $\stackrel{c}{\approx}_\varepsilon$ ” refers to ε -computational indistinguishability against QPT distinguishers;

$$(b) \Pr[b' = 1 \wedge w' \notin \mathcal{R}_\mathcal{L}(x) : (\text{ST}'_{P^*}, b', w') \leftarrow \mathcal{WE}(1^{\varepsilon^{-1}}, P_\lambda^*, \rho_\lambda, x)] = 0,^{19}$$

where $(\text{ST}_{P^}, b) \leftarrow \langle P_\lambda^*(\rho_\lambda), V(x) \rangle$ means that ST_{P^*} is the final state of P_λ^* and b is the decision output by V .*

3. **Witness-Indistinguishability.** *For any non-uniform QPT machine $\{V_\lambda^*, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, we have*

$$\{\text{OUT}_{V_\lambda^*} \langle P(w_0), V_\lambda^*(\rho_\lambda) \rangle(x)\}_{\lambda, x, w_0, w_1} \stackrel{c}{\approx} \{\text{OUT}_{V_\lambda^*} \langle P(w_1), V_\lambda^*(\rho_\lambda) \rangle(x)\}_{\lambda, x, w_0, w_1},$$

where $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, $w_0, w_1 \in \mathcal{R}_\mathcal{L}(x)$, and “ $\stackrel{c}{\approx}$ ” refers to computational indistinguishability against QPT distinguishers.

The work of [CCLY22] gives a constant-round construction of post-quantum ε -zero-knowledge arguments of knowledge from post-quantum OWFs. Since ε -zero-knowledge implies witness indistinguishability, we have the following lemma.

Lemma 5 ([CCLY22, Corollary 2]). *Assuming the existence of post-quantum OWFs, there exist constant-round post-quantum witness-indistinguishable arguments of knowledge as per Def. 14.*

4 Small-Tag, One-Sided, Synchronous, Classical Setting

In this section, we show the small-tag, one-sided, synchronous non-malleable commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ in the classical setting. Though our final goal is to construct *post-quantum* non-malleable commitments, we give a classically-secure construction here since the security proof for its post-quantum version $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$ (in Appx. B) is very similar to that in the classical case except for one step where we rely on the extract-and-simulate lemma (Lem. 20).

¹⁸ A recent work of [LMS21] gave such WIPoKs from super-polynomial hardness of post-quantum OWFs.

¹⁹ We can assume the probability is 0 instead of $\text{negl}(|x|)$ without loss of generality. If the probability is $\text{negl}(|x|)$, we can modify \mathcal{WE} to output $b' = 0$ whenever $w' \notin \mathcal{R}_\mathcal{L}(x)$. This only negligibly affects the distribution in the LHS of Condition 2a, which can be absorbed into $\stackrel{c}{\approx}_\varepsilon$.

4.1 Construction of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$

An intuitive explanation of this construction is already provided in [Sec. 2.1](#). We now present the formal description in [Prot. 1](#). This construction is based on the following building blocks:

- An *injective* OWF f ; To convey the main idea better, we first present the construction assuming the existence of *injective* OWFs. We explain how to relax the assumption to the existence to *any* OWFs in [Sec. 4.8](#).
- Naor’s commitment; We use β to denote the first message for Naor’s commitment, and $\text{Com}_\beta(m; r)$ to denote the second message where a string m is committed using randomness r .
- A witness-indistinguishable proof of knowledge WIPoK.

Protocol 1: Small-Tag One-Sided Synchronous NMCCom $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$
<p>The tag space is defined to be $[n]$ where n is a polynomial on λ. Let $t \in [n]$ be the tag for the following interaction. Let m be the message to be committed to.</p> <p>Commit Stage:</p> <ol style="list-style-type: none"> 1. Receiver R samples and sends the first message β for Naor’s commitment; 2. Committer C commits to m using the second message of Naor’s commitment. Formally, C samples a random tape r and sends $\text{com} = \text{Com}_\beta(m; r)$; 3. R computes $\{y_i = f(x_i)\}_{i \in [t]}$ with $x_i \xleftarrow{\\$} \{0, 1\}^\lambda$ for each $i \in [t]$. R sends $Y = (y_1, \dots, y_t)$ to C; 4. (WIPoK-1.)^a R and C execute an instance of WIPoK where R proves to C that he “knows” the pre-image of some y_i contained in Y (defined in Step 3). Formally, R proves that $Y \in \mathcal{L}_f^t$, where $\mathcal{L}_f^t := \{(y_1, \dots, y_t) \mid \exists(i, x_i) \text{ s.t. } i \in [t] \wedge y_i = f(x_i)\}. \quad (10)$ Note that R uses $(1, x_1)$ as the witness when executing this WIPoK. 5. (WIPoK-2.) C and R execute an instance of WIPoK where C proves to R that he “knows” <i>either</i> the message committed in com (defined in Step 2), <i>or</i> the pre-image of some y_i contained in Y (defined in Step 3). Formally, C proves that $(\text{com}, Y) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t$, where $\mathcal{L}_\beta \vee \mathcal{L}_f^t$ denotes the OR-composed language (as per Def. 1), \mathcal{L}_f^t was defined in Language (10) and $\mathcal{L}_\beta := \{\text{com} \mid \exists(m, r) \text{ s.t. } \text{com} = \text{Com}_\beta(m; r)\}. \quad (11)$ Note that C uses the (m, r) defined in Step 2 as the witness when executing this WIPoK. <p>Decommit Stage: C sends (m, r). R accepts if $\text{com} = \text{Com}_\beta(m; r)$, and rejects otherwise.</p>
<p>^a We assume that the first round of this step goes to the opposite direction to Step 3. This is true for typical WIPoK constructions. See Rmk. 3 for more details.</p>

Remark 3 (Inserting an ACK round if necessary). In [Prot. 1](#), we emphasize that [Step 3](#) and the first message of [Step 4](#) should be *separated*, i.e., they cannot be combined into one round even if the first message of [Step 4](#) happens to go from R to C as well, which may be possible depending on the choice of WIPoK protocols instantiating [Step 4](#). In that case, one can simply insert a dummy round right after [Step 3](#), where C sends an “ACK” symbol to acknowledge the reception of [Step 3](#); Only then will R starts the execution of [Step 4](#). This is necessary because some steps in our security

proof need to non-uniformly fix the execution up to the closing of [Step 3](#), and make use of the (non-uniform) WI property of [Step 4](#).

Security. Completeness is straightforward from the description of [Prot. 1](#). The statistical binding property follows from that of Naor’s commitment. Computational-hiding property of any non-malleable commitment scheme follows directly from its non-malleability. So, we only need to prove the non-malleability of our protocol, which is established by the following theorem.

Theorem 2. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ in [Prot. 1](#) is non-malleable against one-sided synchronous PPT adversaries with tag space $[n]$, with n being any polynomial on λ .*

We prove [Thm. 2](#) in subsequent subsections.

Structure of the Proof of [Thm. 2](#). Since the proof of [Thm. 2](#) is lengthy, we provide a road map here. We prove [Thm. 2](#) in [Sec. 4.2](#) to [4.7](#). In each subsection, we introduce a lemma or claim whose proof is deferred to the next subsection. Specifically,

- In the rest of [Sec. 4.2](#), we prove [Thm. 2](#) assuming that [Lem. 6](#) is correct.
- In [Sec. 4.3](#), we prove [Lem. 6](#) assuming that [Lem. 7](#) is correct.
- In [Sec. 4.4](#), we prove [Lem. 7](#) assuming that [Lem. 9](#) is correct.
- In [Sec. 4.5](#), we prove [Lem. 9](#) assuming that [Lem. 10](#) is correct.
- In [Sec. 4.6](#), we prove [Lem. 10](#) assuming that [Claim 10](#) is correct.
- In [Sec. 4.7](#), we prove [Claim 10](#).

Putting everything together, we complete the proof of [Thm. 2](#).

Remark 4. We remark that the proof of [Lem. 7](#) is the only step that cannot be directly translated to the post-quantum setting. Looking ahead, the post-quantum counterpart of this step (i.e., [Lem. 29](#)) will make use of a new *post-quantum extract-and-simulate* lemma, which we state and prove in [Sec. 7](#).

Moreover, we provide a moderately-detailed summary of this proof in [Sec. 4.9](#), with the hope to further elucidate the logic flow behind the proof of [Thm. 2](#). It is recommended that the reader start reading [Sec. 4.9](#) after he/she has read [Sec. 4.2](#) to [4.8](#) (at least once), because [Sec. 4.9](#) uses several notations defined in preceding subsections.

4.2 Proving Non-Malleability (Proof of [Thm. 2](#))

In the rest of this section, we write $\langle C, R \rangle$ to mean $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ for notational simplicity. To prove [Thm. 2](#), let us first define the game that captures the man-in-the-middle execution corresponding to $\text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m, z)$ w.r.t. the $\langle C, R \rangle$ defined in [Prot. 1](#).

Game $H^{\mathcal{M}}(\lambda, m, z)$: This is the man-in-the-middle execution of the commit stage of the $\langle C, R \rangle$ defined in [Prot. 1](#), where the left committer commits to m and \mathcal{M} ’s non-uniform advice is z . The output of this game consists of the following three parts:

1. $\text{OUT}_{\mathcal{M}}$: this is the output of \mathcal{M} at the end of this game;
2. $\tilde{\tau}$: this is defined to be $\tilde{\tau} := (\tilde{\beta}, \widetilde{\text{com}})$, where $\tilde{\beta}$ and $\widetilde{\text{com}}$ are the [Step 1](#) and [Step 2](#) messages exchanged between \mathcal{M} and the honest receiver R (i.e., in the right session);
3. $b \in \{\top, \perp\}$: this is the output of the honest receiver R , indicating if the man-in-the-middle’s commitment (i.e., the right session) is accepted ($b = \top$) or not ($b = \perp$).

Notation. For any $(\text{OUT}_{\mathcal{M}}, \tilde{\tau}, b)$ in the support of $H^{\mathcal{M}}(\lambda, m, z)$, the following $\text{val}_b(\tilde{\tau})$ defines the value committed in the man-in-the-middle's commitment (i.e., the right session):

$$\text{val}_b(\tilde{\tau}) := \begin{cases} \text{val}(\tilde{\tau}) & b = \top \\ \perp & b = \perp \end{cases}, \quad (12)$$

where $\text{val}(\tilde{\tau})$ denote the value statistically-bound in [Steps 1 and 2](#) of the right session. (Recall that these two steps constitute a Naor's commitment.) Also, throughout this paper, we assume for simplicity that Naor's commitment is perfectly binding (instead of only statistically binding). This assumption only suppresses an *additive error of negligible amount* in relevant lemmas, which will not affect any of our results.

By definition, it is easy to see that:

$$\{\text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m, z)\} \stackrel{\text{i.d.}}{=} \{(\text{OUT}_{\mathcal{M}}, \text{val}_b(\tilde{\tau})) : (\text{OUT}_{\mathcal{M}}, \tilde{\tau}, b) \leftarrow H^{\mathcal{M}}(\lambda, m, z)\}, \quad (13)$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$, $m \in \{0, 1\}^{\ell(\lambda)}$, and $z \in \{0, 1\}^*$.

Therefore, to prove that [Prot. 1](#) satisfies [Def. 6](#), it suffices to show the following:

$$\begin{aligned} & \{(\text{OUT}_{\mathcal{M}}^0, \text{val}_{b^0}(\tilde{\tau}^0)) : (\text{OUT}_{\mathcal{M}}^0, \tilde{\tau}^0, b^0) \leftarrow H^{\mathcal{M}}(\lambda, m_0, z)\} \\ & \stackrel{c}{\approx} \{(\text{OUT}_{\mathcal{M}}^1, \text{val}_{b^1}(\tilde{\tau}^1)) : (\text{OUT}_{\mathcal{M}}^1, \tilde{\tau}^1, b^1) \leftarrow H^{\mathcal{M}}(\lambda, m_1, z)\}, \end{aligned} \quad (14)$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$, $(m_0, m_1) \in \{0, 1\}^{\ell(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$, and $z \in \{0, 1\}^*$.

Proof by Contradiction. Our high-level strategy is proof by contradiction. We assume for contradiction that there are a (possibly non-uniform) PPT machine \mathcal{D} and a function $\delta(\lambda) = 1/\text{poly}(\lambda)$ such that for infinitely many $\lambda \in \mathbb{N}$,

$$\left| \Pr\left[\mathcal{D}(1^\lambda, \text{OUT}_{\mathcal{M}}^0, \text{val}_{b^0}(\tilde{\tau}^0)) = 1\right] - \Pr\left[\mathcal{D}(1^\lambda, \text{OUT}_{\mathcal{M}}^1, \text{val}_{b^1}(\tilde{\tau}^1)) = 1\right] \right| \geq 3 \cdot \delta(\lambda), \quad (15)$$

where the first probability is taken over the random procedure $(\text{OUT}_{\mathcal{M}}^0, \tilde{\tau}^0, b^0) \leftarrow H^{\mathcal{M}}(\lambda, m_0, z)$ and the second probability is taken over the random procedure $(\text{OUT}_{\mathcal{M}}^1, \tilde{\tau}^1, b^1) \leftarrow H^{\mathcal{M}}(\lambda, m_1, z)$.

We then show the following lemma:

Lemma 6. *There exists a hybrid G such that for any PPT \mathcal{M} and the $\delta(\lambda)$ defined above, the following holds*

1. $\{(\text{OUT}^0, \text{Val}^0) : (\text{OUT}^0, \text{Val}^0) \leftarrow G^{\mathcal{M}}(\lambda, m_0, z)\} \stackrel{c}{\approx} \{(\text{OUT}^1, \text{Val}^1) : (\text{OUT}^1, \text{Val}^1) \leftarrow G^{\mathcal{M}}(\lambda, m_1, z)\}$, where both ensembles are indexed by $\lambda \in \mathbb{N}$, $(m_0, m_1) \in \{0, 1\}^{\ell(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$, and $z \in \{0, 1\}^*$.
2. $\{(\text{OUT}^G, \text{Val}^G) : (\text{OUT}^G, \text{Val}^G) \leftarrow G^{\mathcal{M}}(\lambda, m, z)\} \stackrel{c}{\approx}_{\delta(\lambda)} \{(\text{OUT}^H, \text{val}_{b^H}(\tilde{\tau}^H)) : (\text{OUT}^H, \tilde{\tau}^H, b^H) \leftarrow H^{\mathcal{M}}(\lambda, m, z)\}$, where both ensembles are indexed by $\lambda \in \mathbb{N}$, $m \in \{0, 1\}^{\ell(\lambda)}$, and $z \in \{0, 1\}^*$.

It is easy to see that if [Lem. 6](#) is true, it contradicts our assumption in [Inequality \(15\)](#). Thus, it will finish the proof of non-malleability. Indeed, this lemma is the most technically involved part. We prove it in [Sec. 4.3](#).

4.3 Proof of [Lem. 6](#)

First, we provide a new but equivalent interpretation of the game $H^{\mathcal{M}}(\lambda, m, z)$ in [Algo. 4.1](#). (We also provide a picture in [Fig. 4a](#) to illustrate it.)

Algorithm 4.1: Re-interpretation of Game $H^{\mathcal{M}}(\lambda, m, z)$

Game $H^{\mathcal{M}}(\lambda, m, z)$ can be split into the following stages:

1. **Prefix Generation:** First, execute [Steps 1](#) and [2](#) of the man-in-the-middle game of [Prot. 1](#). That is, it plays as the left honest committer committing to m and the right honest receiver, with $\mathcal{M}(1^\lambda, z)$ being the man-in-the-middle adversary.

Notation: Let $\text{st}_{\mathcal{M}}$ denote the state of \mathcal{M} at the end of [Step 2](#); Let st_C (resp. st_R) denote the state of the honest committer (resp. receiver) at the end of [Step 2](#); Let $\tilde{\tau} = (\tilde{\beta}, \widetilde{\text{com}})^a$ and $\tau = (\beta, \text{com})$. In terms of notation, we denote the execution of this stage by

$$(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau}) \leftarrow H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z). \quad (16)$$

We will call the tuple $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ the *prefix* and denote it by **pref**. It is worth noting that this **pref** contains all the information such that a PPT machine can “complete” the remaining execution of $H^{\mathcal{M}}(\lambda, m, z)$ starting from **pref**.

2. **The Remainder:** Next, it simply resumes from where the **Prefix Generation** stage stops, to finish the remaining steps of the man-in-the-middle execution $H^{\mathcal{M}}(\lambda, m, z)$.

Notation: We introduce the following notations to describe this stage. Define a PPT machine \mathcal{A} that takes as input $(\text{st}_{\mathcal{M}}, \tilde{\tau})$; Machine \mathcal{A} is supposed to run the residual strategy of \mathcal{M} starting from $\text{st}_{\mathcal{M}}$. Also, define a PPT machine \mathcal{B} that takes as input $(\text{st}_C, \text{st}_R, \tilde{\tau})$; Machine \mathcal{B} is supposed to run the residual strategies of the honest committer C and receiver R , starting from st_C and st_R respectively^b. With the above notations, we can denote the execution of the remaining steps of $H^{\mathcal{M}}(\lambda, m, z)$ by

$$(\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau}), \quad (17)$$

where $\text{OUT}_{\mathcal{A}}$ is the output of \mathcal{A} , and $b \in \{\perp, \top\}$ is the output of the honest receiver R (in the right), indicating if the man-in-the-middle’s commitment (i.e., the right session) is accepted ($b = \top$) or not ($b = \perp$). (We remark that $\text{OUT}_{\mathcal{A}}$ is nothing but the man-in-the-middle \mathcal{M} ’s final output.)

3. **Output:** It outputs the tuple $(\text{OUT}_{\mathcal{A}}, \tilde{\tau}, b)$.

^a Recall that $\tilde{\beta}$ and $\widetilde{\text{com}}$ are the [Steps 1](#) and [2](#) messages in the right session; They constitute an execution of Naor’s commitment.

^b Note that it is actually redundant to give $\tilde{\tau}$ as common input to these parties—It can be included in $\text{st}_{\mathcal{M}}$ and st_R . We choose to make $\tilde{\tau}$ explicit for notational convenience.

We claim the following lemma regarding [Algo. 4.1](#).

Lemma 7. *Let $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, \mathcal{A} , \mathcal{B} be as defined in [Algo. 4.1](#). There exists a PPT machine \mathcal{SE} (the simulation-extractor) such that for any $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, any noticeable $\varepsilon(\lambda)$, it holds that*

$$\begin{aligned} & \{(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}) : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx}_{\varepsilon(\lambda)} \{(\text{OUT}_{\mathcal{A}}, \text{val}_b(\tilde{\tau})) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \end{aligned}$$

Remark 5. The third output $b_{\mathcal{SE}}$ of \mathcal{SE} is redundant in the above lemma. We include it in the output of \mathcal{SE} for convenience in the proof of this lemma in [Sec. 4.4](#).

Lem. 7 is the main technical lemma for the current proof of **Lem. 6**. We present its proof in **Sec. 4.4**. In the following, we finish the proof of **Lem. 6** assuming that **Lem. 7** is true.

Now, we are ready to present the description of G as required by **Lem. 6**.

<p>Algorithm 4.2: Hybrid $G^{\mathcal{M}}(\lambda, m, z)$</p> <p>This hybrid proceeds as follows:</p> <ol style="list-style-type: none"> 1. Prefix Generation: This stage is identical to Stage 1 of $H^{\mathcal{M}}(\lambda, m, z)$. Formally, it executes $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau}) \leftarrow H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z),$ where $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ is defined in Stage 1 of Algo. 4.1. 2. The Remainder: Define \mathcal{A} in the same way as in Stage 2 of $H^{\mathcal{M}}(\lambda, m, z)$. With this \mathcal{A} and the $(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ from the previous stage, $G^{\mathcal{M}}(\lambda, m, z)$ invokes the \mathcal{SE} prescribed in Lem. 7. Formally, it executes the following procedure: $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\delta^{-1}}, \text{st}_R, \tau, \tilde{\tau}),$ where the δ is the statistical distance that we want to show for Property 2 of Lem. 6. <p><i>Remark 6.</i> We emphasize that in this stage, $G^{\mathcal{M}}(\lambda, m, z)$ does <i>not</i> make use of st_C.</p> 3. Output: It outputs $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}})$.^a <p>^a Note that $b_{\mathcal{SE}}$ is not included in the output of G. Indeed, $b_{\mathcal{SE}}$ is not important for the current machine $G^{\mathcal{M}}(\lambda, m, z)$. We choose to include it in the output of $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\delta^{-1}}, \text{st}_R, \tau, \tilde{\tau})$ only for notational convenience when we define/prove properties about \mathcal{SE} itself.</p>

Proving Property 1 of Lem. 6. Observe that hybrid $G^{\mathcal{M}}(\lambda, m, z)$ is an efficient machine because both $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ and $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\delta^{-1}}, \text{st}_R, \tau, \tilde{\tau})$ are efficient. Moreover, it does not rewind **Steps 1** and **2** of the man-in-the-middle execution, and $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\delta^{-1}}, \text{st}_R, \tau, \tilde{\tau})$ does *not* need to know st_C . Therefore, **Property 1** of **Lem. 6** follows immediately from the computational-hiding property of the left Naor’s commitment (i.e., **Steps 1** and **2** in the left session).

Proving Property 2 of Lem. 6. First, observe that the distribution of the prefix $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ is identical in $G^{\mathcal{M}}(\lambda, m, z)$ and $H^{\mathcal{M}}(\lambda, m, z)$. For each fixed prefix, **Lem. 7** implies that $G^{\mathcal{M}}(\lambda, m, z)$ and $H^{\mathcal{M}}(\lambda, m, z)$ are $\delta(\lambda)$ -computationally indistinguishable where we remark that $G^{\mathcal{M}}(\lambda, m, z)$ runs \mathcal{SE} with the second input $1^{\delta^{-1}}$. This immediately implies **Property 2** of **Lem. 6**.

4.4 Proof of Lem. 7

Now, we prove **Lem. 7** by constructing the required machine \mathcal{SE} . Our machine \mathcal{SE} (for **Lem. 7**) will make use of two machines: (1) a \mathcal{G}_1 that simulates the main-thread, and (2) an extractor \mathcal{K} that extracts the value committed in $\tilde{\tau}$. Thus, in the following, we will first define these two machines, and then present the description of \mathcal{SE} in **Algo. 4.3**.

First, we describe a machine \mathcal{G}_1 that simulates the real execution *without using* st_C .

Machine \mathcal{G}_1 : (Illustrated in **Fig. 4b**) For any prefix pref , $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to **Stage 2** of $H^{\mathcal{M}}(\lambda, m, z)$ shown in **Algo. 4.1** (and depicted in **Fig. 4a**), except for the following difference: Instead of executing the left **WIPoK-1** honestly, it uses the witness-extended emulator \mathcal{WE} (as per **Def. 12**) of the left **WIPoK-1** to extract a witness, and

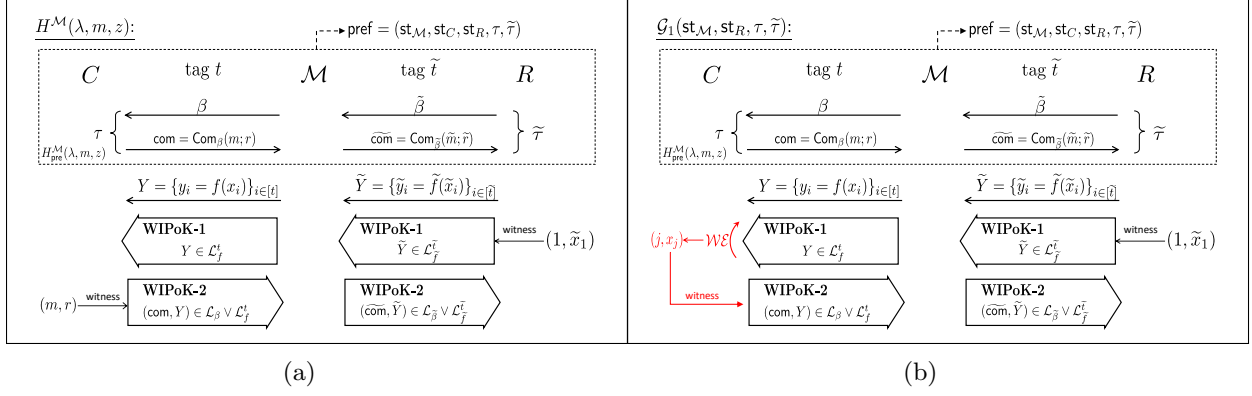


Fig. 4: Machines H^M and \mathcal{G}_1 (Difference is highlighted in red color)

- If the left committer accepts the left **WIPoK-1** and the extracted witness is valid (i.e., it is a (j, x_j) pair such that $y_j = f(x_j)$ for some $j \in [t]$), \mathcal{G}_1 uses (j, x_j) to finish the left **WIPoK-2**. Similarly to $H^M(\lambda, m, z)$, \mathcal{G}_1 eventually outputs \mathcal{M} 's final state and the right receiver's decision bit b ;
- If the left committer accepts the left **WIPoK-1** but the extracted witness is invalid, it aborts immediately and outputs (\perp, \perp) .
- If the left committer rejects the left **WIPoK-1**, it runs the rest of execution of $H^M(\lambda, m, z)$ to output \mathcal{M} 's final state and the right receiver's decision bit b . Note that it does not need to run the left **WIPoK-2** in this case since the left committer aborts after the left **WIPoK-1**.

We denote the above procedure by $(\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. It is worth noting that \mathcal{G}_1 does not need to know st_C .

Remark 7 (Precise Meaning of Extraction). In the above description of \mathcal{G}_1 , we say that “it uses the witness-extended emulator of the left **WIPoK-1** to extract a witness”. More accurately, this means the following: We consider a cheating prover against the left **WIPoK-1** that takes as advice the states of \mathcal{M} and the right receiver right before the start of the left **WIPoK-1**, interacts with the left committer by simulating \mathcal{M} and the right receiver, and outputs the states of \mathcal{M} and the right receiver at the end of the left **WIPoK-1**. Machine \mathcal{G}_1 runs the witness-extended emulator of the left **WIPoK-1** (as per Def. 12) w.r.t. the this cheating prover to simulate the states of \mathcal{M} and the right receiver right at the end of the left **WIPoK-1**, the decision bit of the left **WIPoK-1** of the left committer, while extracting a (candidate of) witness of \mathcal{L}_f^t . We will use similar convention many times throughout this paper.

Next, we prove a lemma that shows that \mathcal{G}_1 simulates the real MIM execution.

Lemma 8. For the $H_{\text{pre}}^M(\lambda, m, z)$, \mathcal{A} , and \mathcal{B} defined in Algo. 4.1, for any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^M(\lambda, m, z)$, it holds that

$$\begin{aligned} & \left\{ (\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau}) \right\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx} \left\{ (\text{OUT}_{\mathcal{A}}, b) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau}) \right\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

Proof. We first define a hybrid machine \mathcal{G}'_1 below.

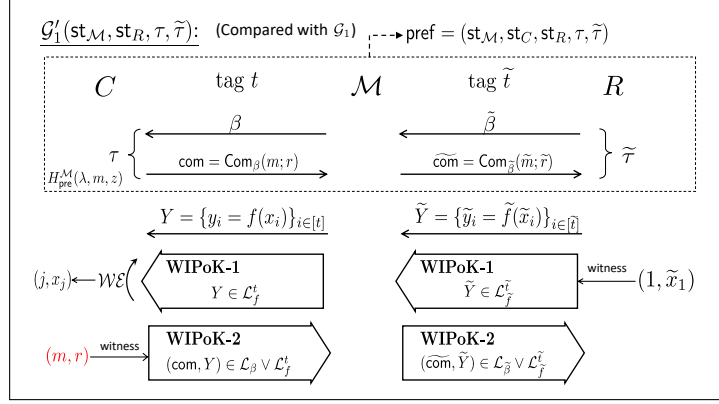


Fig. 5: Machine \mathcal{G}'_1 (Difference with \mathcal{G}_1 is highlighted in red color)

Machine \mathcal{G}'_1 : (Illustrated in Fig. 5) For any prefix pref , $\mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ except that it uses (m, r) as the witness in the left **WIPoK-2** even if it succeeds in extracting a valid witness from the left **WIPoK-1**.

By the WI property of the left **WIPoK-2**, it holds that

$$\begin{aligned} & \{(\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx} \{(\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned} \quad (18)$$

Note that the only difference between $\mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ and $\langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}) \rangle(1^\lambda, \tilde{\tau})$ is that the former runs the witness-extended emulator of the left **WIPoK-1** (but does not use the extracted witness at all). Thus, by the PoK property of the left **WIPoK-1**, it holds that

$$\begin{aligned} & \{(\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx} \{(\text{OUT}_{\mathcal{A}}, b) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}) \rangle(1^\lambda, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned} \quad (19)$$

By combining Eq. (18) and (19), we obtain Lem. 8. \square

Next, we define the probability of R being convinced in the execution of \mathcal{G}_1 . This value plays an important role later in our proof.

Definition 15. For any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, we define the following value $p_{\text{pref}}^{\text{Sim}}$:

$$p_{\text{pref}}^{\text{Sim}} := \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})].$$

Next, we show a technical lemma that gives an extractor \mathcal{K} without the simulation property.

Lemma 9. Let $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, \mathcal{A} , and \mathcal{B} be as defined in Algo. 4.1. There exists an expected PPT machine \mathcal{K} such that for any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ and any noticeable $\varepsilon(\lambda)$, the following holds:

1. **(Syntax.)** \mathcal{K} takes as input $(1^\lambda, \text{st}_{\mathcal{M}}, \tau, \tilde{\tau})$ and makes oracle access to $\mathcal{A}(\text{st}_{\mathcal{M}})$. It outputs a value $\text{Val}_{\mathcal{K}} \in \{0, 1\}^{\ell(\lambda)} \cup \{\perp\}$ such that $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$ whenever $\text{Val}_{\mathcal{K}} \neq \perp$. (Also see Rmk. 8 for an intuitive explanation.)

2. If $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then it holds that

$$\Pr\left[\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau}) : \text{Val}_{\mathcal{K}} \leftarrow \mathcal{K}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau})\right] \geq \frac{\varepsilon'(\lambda)}{\tilde{t}},$$

where $\varepsilon'(\lambda) := \frac{\varepsilon(\lambda)}{10t^2}$.

Remark 8 (On the Output of \mathcal{K}). The output $\text{Val}_{\mathcal{K}}$ is expected to be the value committed in $\tilde{\tau}$, i.e., $\text{val}(\tilde{\tau})$. If $\text{Val}_{\mathcal{K}} = \perp$, it indicates that \mathcal{K} did not extract the correct $\text{val}(\tilde{\tau})$. We explicitly include \perp in the range of \mathcal{K} for the following purpose: Looking ahead, the simulation-extractor \mathcal{SE} that we are going to build will invoke \mathcal{K} for several times, until the value $\text{val}(\tilde{\tau})$ is extracted. However, \mathcal{SE} cannot tell if the extracted value is indeed $\text{val}(\tilde{\tau})$. Thus, the case $\text{Val}_{\mathcal{K}} = \perp$ serves as an indicator, telling \mathcal{SE} if the extraction by \mathcal{K} succeeds (more accurately, fails).

As a vigilant reader may have already realized, there is an alternative formalism: simply ask \mathcal{K} to output both $\text{val}(\tau)$ and the decommitment information so that \mathcal{SE} can test by itself whether \mathcal{K} 's extraction is successful. We remark that this approach does work for the current proof *in classical setting*. However, it may not extend when we prove *post-quantum* non-malleability. In short, that is because to make this proof work in the quantum setting, our technique requires that the valid output of \mathcal{K} should be *unique*; Only in this way can we ask \mathcal{K} to “forget” other information, such that the extraction procedure can be “un-computed” to rewind \mathcal{M} back without much disturbance to its initial state. (This point will become clearer in the proof of [Lem. 20](#).) However, if we include the decommitment information in the output of \mathcal{K} , then the valid output may not be unique, even if the commitment scheme is perfect-binding—There could exist different ways to decommit to the *unique* committed value.

To make the proof consistent in both classical and quantum settings, we choose the current formalism in [Property 1](#).

We will prove [Lem. 9](#) in [Sec. 4.5](#). In the rest of this subsection, we finish the proof of [Lem. 7](#) assuming [Lem. 9](#) is true.

We present the description of \mathcal{SE} (for [Lem. 7](#)) in [Algo. 4.3](#).

Algorithm 4.3: Simulation-Extractor $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$

Let $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, \mathcal{A} , and \mathcal{B} be as defined in [Algo. 4.1](#). Let $\varepsilon(\lambda)$ be a noticeable function. For $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, machine $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$ proceeds as follows:

1. **Main-Thread Simulation.** It first uses the machine \mathcal{G}_1 to simulate [Stage 2](#) of $H^{\mathcal{M}}(\lambda, m, z)$. That is, it computes $(\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. It sets $\text{OUT}_{\mathcal{SE}} := \text{OUT}$ and $b_{\mathcal{SE}} := b$. Then,
 - if $b_{\mathcal{SE}} = \perp$, it sets $\text{Val}_{\mathcal{SE}} := \perp$ and jumps directly to [Stage 4](#);
 - otherwise, it goes to the next step.
2. **Rewinding:** $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$ then loops the following procedure for $\frac{\tilde{t}}{\varepsilon'(\lambda)} \cdot \lambda$ times, where $\varepsilon'(\lambda) := \frac{\varepsilon(\lambda)}{10t^2}$:
 - Execute $\text{Val}_{\mathcal{K}} \leftarrow \mathcal{K}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$, where \mathcal{K} is provided by [Lem. 9](#). If $\text{Val}_{\mathcal{K}} \neq \perp$,^a set $\text{Val}_{\mathcal{SE}} := \text{Val}_{\mathcal{K}}$ and jump to [Stage 4](#); Otherwise, go to the next loop.
3. It sets $\text{Val}_{\mathcal{SE}} = \perp$. (Note that if this stage is reached, it means that the number of the loop in the previous stage reached its upper bound $\frac{\tilde{t}}{\varepsilon'(\lambda)} \cdot \lambda$, but \mathcal{K} did not extract the message committed in $\tilde{\tau}$ yet.)

4. **Output:** it outputs $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}})$.

^a Note that by [Property 1](#) in [Lem. 9](#), this means $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$.

We first prove that the simulated main-thread is computationally indistinguishable from the real one.

Claim 3. Let $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$, \mathcal{A} , \mathcal{B} be as defined in [Algo. 4.1](#). For any $(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ and any noticeable $\varepsilon(\lambda)$, it holds that

$$\begin{aligned} & \{(\text{OUT}_{\mathcal{SE}}, \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau})) : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx} \{(\text{OUT}_{\mathcal{A}}, \text{val}_b(\tilde{\tau})) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}) \rangle(1^\lambda, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned} \quad (20)$$

Proof. Since we consider each fixed prefix, we can efficiently compute $\text{val}_{b_{\mathcal{SE}}}(\tilde{\tau})$ and $\text{val}_b(\tilde{\tau})$ from $b_{\mathcal{SE}}$ and b , respectively, by using $\text{val}(\tilde{\tau})$ as a non-uniform advice²⁰. Then, it suffices to prove

$$\begin{aligned} & \{(\text{OUT}_{\mathcal{SE}}, b_{\mathcal{SE}}) : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx} \{(\text{OUT}_{\mathcal{A}}, b) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}) \rangle(1^\lambda, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

Since \mathcal{SE} just runs \mathcal{G}_1 in the main thread, the above follows directly from [Lem. 8](#). □

Given [Claim 3](#), to prove [Lem. 7](#), it suffices to show the following inequality:

$$\Pr \left[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau}) \right] \leq \varepsilon(\lambda) + \text{negl}(\lambda). \quad (21)$$

To prove [Inequality \(21\)](#), let us first define two events:

- **Event $E_{\leq \varepsilon}$:**²¹ This is the event that the prefix $(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ lead to an accepting execution of \mathcal{SE} with probability at most ε . Formally, it denotes the event that $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ satisfies the following inequality:

$$\Pr \left[b_{\mathcal{SE}} = \top : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau}) \right] \leq \varepsilon(\lambda).$$

Remark 9. We emphasize that the above probability (on the LHS) is exactly the $p_{\text{pref}}^{\text{Sim}}$ defined in [Def. 15](#), because \mathcal{SE} obtains $b_{\mathcal{SE}}$ by running \mathcal{G}_1 (see [Stage 1](#) of [Algo. 4.3](#)).

- **Event E_{acc} :** This is the event that $b_{\mathcal{SE}} = \top$ at the end of the execution of $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$. Namely, it means that the honest receiver R accepts in the execution simulated by \mathcal{SE} . (The subscript “acc” stands for “accept”).

We first show that the the probability of $(E_{\text{acc}} \wedge E_{\leq \varepsilon})$ is upper-bounded by $\varepsilon(\lambda)$.

Claim 4. It holds that

$$\Pr \left[E_{\text{acc}} \wedge E_{\leq \varepsilon} : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau}) \right] \leq \varepsilon(\lambda).$$

²⁰ Note that this is possible because both the LHS and the RHS of [Eq. \(20\)](#) use the same *fixed* $\tilde{\tau}$, which does not depend on the randomness of the ensembles.

²¹ We remark that this event actually does not depend on the random procedure $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$. It is only about the property of the prefix $(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{C}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ given as input to machine \mathcal{SE} .

Proof. This is an immediate result of the definitions of these two events. In more details,

$$\Pr[E_{\text{acc}} \wedge E_{\leq \varepsilon}] = \Pr[E_{\text{acc}} \mid E_{\leq \varepsilon}] \cdot \Pr[E_{\leq \varepsilon}] \leq \Pr[E_{\text{acc}} \mid E_{\leq \varepsilon}] \cdot 1 = \varepsilon(\lambda) \cdot 1, \quad (22)$$

where all the probabilities are taken over $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$. \square

Next, we upper-bound the LHS of [Inequality \(21\)](#) by $\varepsilon(\lambda) + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \mid E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}]$:

Claim 5. *It holds that*

$$\Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau})] \leq \varepsilon(\lambda) + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \mid E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}],$$

where both probabilities are taken over $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$.

Proof. All the probabilities below are taken over $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$:

$$\begin{aligned} \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau})] &= \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge E_{\text{acc}}] + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge \neg E_{\text{acc}}] \\ &= \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge E_{\text{acc}}] \end{aligned} \quad (23)$$

$$\begin{aligned} &= \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge E_{\text{acc}} \wedge E_{\leq \varepsilon}] + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}] \\ &\leq \Pr[E_{\text{acc}} \wedge E_{\leq \varepsilon}] + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}] \\ &\leq \varepsilon(\lambda) + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \wedge E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}] \\ &\leq \varepsilon(\lambda) + \Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \mid E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}], \end{aligned} \quad (24)$$

where [Eq. \(23\)](#) follows from the fact that if $b_{\mathcal{SE}} = \perp$ then $\text{Val}_{\mathcal{SE}} = \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) = \perp$, and [Inequality \(24\)](#) follows from [Claim 4](#). \square

Recall that our goal is to establish [Inequality \(21\)](#). Due to [Claim 5](#), it now suffices to show

$$\Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \mid E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}] = \text{negl}(\lambda), \quad (25)$$

where the probability is taken over $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$.

Proof of Eq. (25). Given E_{acc} (i.e., $b_{\mathcal{SE}} = \top$), we know that $\mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$ must have entered [Stage 2](#) to execute the machine $\mathcal{K}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$ for at most $\frac{\tilde{t}}{\varepsilon'(\lambda)} \cdot \lambda$ times, with early termination only if the event $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$ happens (recall that $\text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) = \text{val}(\tilde{\tau})$ when $b_{\mathcal{SE}} = \top$).

Given $\neg E_{\leq \varepsilon}$ (thus, $p_{\text{pref}}^{\text{Sim}} > \varepsilon(\lambda)$, see [Rmk. 9](#)), it follows from [Property 2](#) in [Lem. 9](#) that each time \mathcal{K} is invoked, it outputs $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$ with probability $\geq \varepsilon'(\lambda)/\tilde{t}$.

Therefore, we have

$$\Pr[\text{Val}_{\mathcal{SE}} \neq \text{val}_{b_{\mathcal{SE}}}(\tilde{\tau}) \mid E_{\text{acc}} \wedge \neg E_{\leq \varepsilon}] = \left(1 - \frac{\varepsilon'(\lambda)}{\tilde{t}}\right)^{\frac{\tilde{t}}{\varepsilon'(\lambda)} \cdot \lambda} = 2^{-O(\lambda)} = \text{negl}(\lambda),$$

where the probability is taken over $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}, b_{\mathcal{SE}}) \leftarrow \mathcal{SE}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, 1^{\varepsilon^{-1}}, \text{st}_R, \tau, \tilde{\tau})$.

This finishes the proof of [Eq. \(25\)](#). \square

This eventually finishes the proof of [Lem. 7](#) (modulo the proof of [Lem. 9](#), which we present in [Sec. 4.5](#)).

4.5 Extractor \mathcal{K} (Proof of Lem. 9)

Machine \mathcal{G}_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 6.) Recall that we have already defined the machine $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ on Page 29. Now, for $i \in [\tilde{t}] \setminus \{1\}$, $\mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it uses (i, \tilde{x}_i) as the witness in the right **WIPoK-1**.

Claim 6. $\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda)$.

Proof. We define hybrid machines \mathcal{G}'_i and \mathcal{G}''_i as follows.

Machine \mathcal{G}'_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 7a.) Recall that we have already defined the machine $\mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ on Page 31. Now, for $i \in [\tilde{t}] \setminus \{1\}$, $\mathcal{G}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it uses (i, \tilde{x}_i) as the witness in the right **WIPoK-1**.

Machine \mathcal{G}''_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 7b) For any prefix pref , $\mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it honestly runs the left **WIPoK-1**, instead of running the witness-extended emulator \mathcal{WE} . In other words, $\mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_{\mathcal{C}}, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})$ except that it uses (i, \tilde{x}_i) as the witness in the right **WIPoK-1**. In particular, $\mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ is identical to $\langle \mathcal{A}(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_{\mathcal{C}}, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})$.

Then, Claim 6 follows from the following sequence of inequalities.

– By the WI property of the left **WIPoK-2** and the definition of $p_{\text{pref}}^{\text{Sim}}$ (Def. 15), it holds that:

$$\Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda).$$

– By the PoK property (per Def. 12) of the left **WIPoK-1** and the above inequality, it holds that:

$$\Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}''_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda).$$

– By the WI property of the right **WIPoK-1** and the above inequality, it holds that:

$$\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda).$$

– By the PoK property of the left **WIPoK-1** and the above inequality, it holds that:

$$\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda).$$

– By the WI property of the left **WIPoK-2** and the above inequality, it holds that:

$$\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda).$$

This finishes the proof of Claim 6. □

Machine \mathcal{K}_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 8a.) For a prefix pref , $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to the $\mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ depicted in Fig. 6, except for the following difference. Machine $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ uses the witness-extended emulator \mathcal{WE} (as per Def. 12) to finish the right **WIPoK-2**, instead of playing the role of the honest receiver.

\mathcal{K}_i 's Output: Let w' denote the third output of the witness-extended emulator \mathcal{WE} (see Def. 12), which is supposed to be the witness used by \mathcal{M} in the right **WIPoK-2** (for the statement $(\widetilde{\text{com}}, \tilde{Y})$ w.r.t. the language $\mathcal{L}_{\tilde{\beta}} \vee \mathcal{L}_{\tilde{\gamma}}$). Depending on the value of w' , we define a value $\text{Val} \in \{0, 1\}^{\ell(\lambda)} \cup \{\perp_{\tilde{\gamma}}, \perp_{\text{invalid}}\}$ as follows:

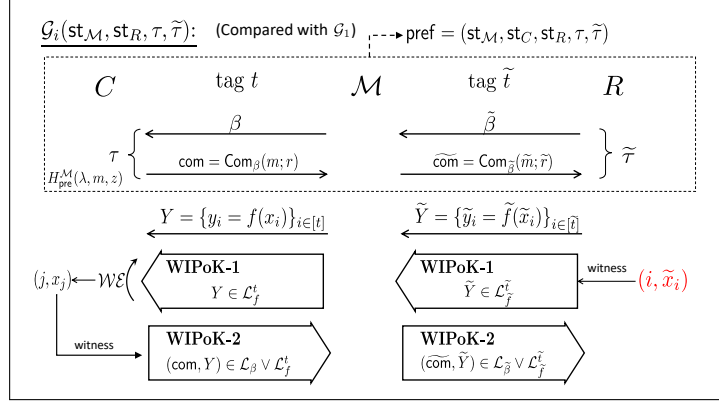


Fig. 6: Machine \mathcal{G}_i (Difference with \mathcal{G}_1 is highlighted in red color)

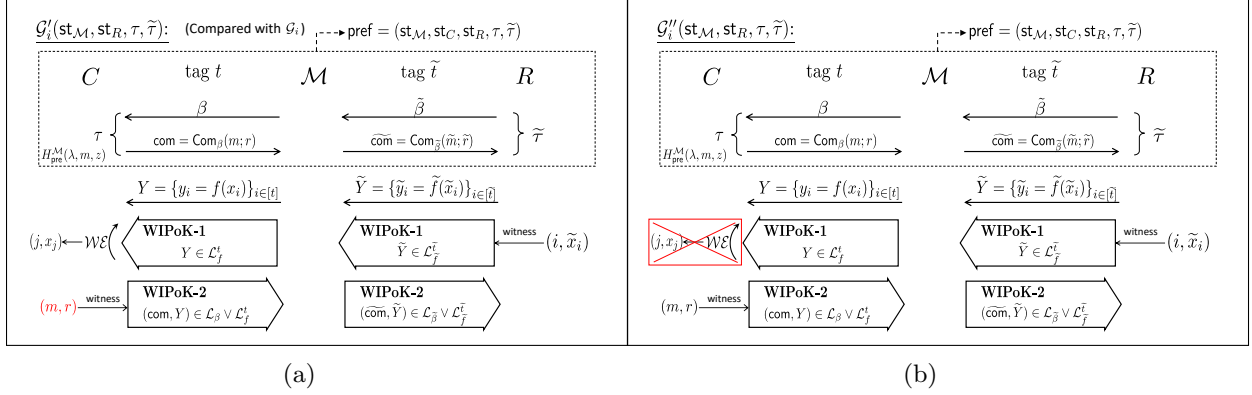


Fig. 7: Machines \mathcal{G}'_i and \mathcal{G}''_i (Difference is highlighted in red color)

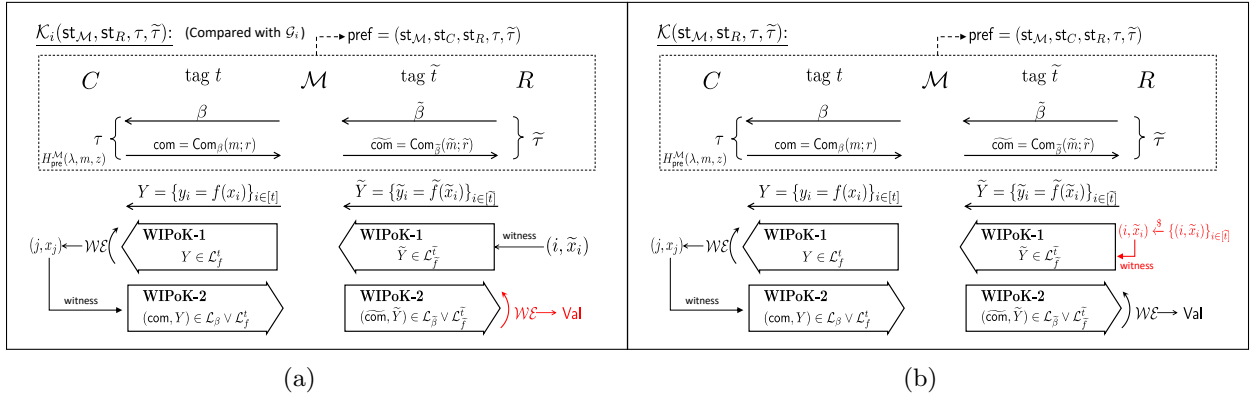


Fig. 8: Machines \mathcal{K}_i and \mathcal{K} (Difference is highlighted in red color)

1. If w' is a valid witness for $(\widetilde{\text{com}}, \tilde{Y}) \in \mathcal{L}_{\tilde{\beta}} \vee \mathcal{L}_{\tilde{f}}^{\tilde{t}}$, then there are two sub-cases:

(a) w' is a valid witness for $\widetilde{\text{com}} \in \mathcal{L}_{\widetilde{\beta}}$. In this case²², w' consists of the value $\text{val}(\widetilde{\tau})$, i.e., the value committed in $\widetilde{\tau} = (\widetilde{\beta}, \widetilde{\text{com}})$, and the randomness \widetilde{r} . We set $\text{Val} := \text{val}(\widetilde{\tau})$. Importantly, notice that we do *not* include the randomness \widetilde{r} in Val (as explained in [Rmk. 8.](#)).

(b) w' is a valid witness for $\widetilde{Y} \in \mathcal{L}_{\widetilde{f}}^{\widetilde{t}}$: In this case, we set $\text{Val} := \perp_{\widetilde{f}}$.

2. Otherwise, set $\text{Val} := \perp_{\text{invalid}}$.

The output of \mathcal{K}_i is defined to be the above Val . Notice that this is in contrast to all previous machines, for which the output is defined to be the man-in-the-middle \mathcal{M} 's output and the honest receiver's decision bit. We emphasize that such a Val satisfies the syntax requirement in [Property 1](#) of [Lem. 9](#). In particular, $\text{Val} = \text{val}(\widetilde{\tau})$ whenever $\text{Val} \neq \perp$.²³

Claim 7. $\forall i \in [\widetilde{t}], \Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \widetilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda)$.

Proof. This claim follows immediately from [Claim 6](#) and the PoK property (as per [Def. 12](#)) of the right **WIPoK-2**. □

Finally, we are ready to define the extractor \mathcal{K} as required by [Lem. 9](#). Intuitively, \mathcal{K} can be conceived as an average-case version of $\{\mathcal{K}_i\}_{i \in [\widetilde{t}]}$:

Extractor \mathcal{K} : (Illustrated in [Fig. 8b.](#)) For a prefix pref , $\mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \widetilde{\tau})$ samples uniformly at random an index $i \xleftarrow{\$} [\widetilde{t}]$, executes $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \widetilde{\tau})$, and outputs whatever $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \widetilde{\tau})$ outputs.

Remark 10 (On the Notation of \mathcal{K}). It is worth noting that in the above, we write machine \mathcal{K} as $\mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \widetilde{\tau})$, while it was written in [Lem. 9](#) as $\mathcal{K}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \widetilde{\tau})$. This is only a cosmetic difference.

Next, we show that the extractor \mathcal{K} satisfies the requirements in [Lem. 9](#).

Running Time of \mathcal{K} . Observe that for each $i \in [\widetilde{t}]$, $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \widetilde{\tau})$ differs from the real man-in-the-middle game only in the following places:

- (i, \widetilde{x}_i) is used in the right **WIPoK-1**;
- the witness-extended emulator \mathcal{WE} is used in the right **WIPoK-2** and the left **WIPoK-1**;
- the left **WIPoK-2** is done using the (j, x_j) extracted by \mathcal{WE} from the left **WIPoK-1**.

Since the witness-extended emulator \mathcal{WE} (as per [Def. 12](#)) runs in expected PPT, so does \mathcal{K}_i . Thus, \mathcal{K} is expected PPT.

Proving [Property 1](#) of [Lem. 9](#). It is straightforward to see that the \mathcal{K} defined above satisfies the syntax requirement in [Lem. 9](#) (see also [Rmk. 10](#)). In particular, we have $\text{Val}_{\mathcal{K}} = \text{val}(\widetilde{\tau})$ whenever $\text{Val}_{\mathcal{K}} \neq \perp$, because this is true for each \mathcal{K}_i by definition (see the paragraph for “ [\$\mathcal{K}_i\$'s Output](#)” on [Page 35](#)).

²² Technically, we should argue that \mathcal{K}_i is able to detect if this case is happening. This is easy to do as we explicitly give $\widetilde{\tau}$ as an input to \mathcal{K}_i —Upon obtaining w' , it can just re-execute Naor's commitment by itself to see if w' is consistent with $\widetilde{\tau}$.

²³ Note that here we defined two types of abortion: $\perp_{\widetilde{f}}$ and \perp_{invalid} , while [Property 1](#) of [Lem. 9](#) only allows a single abortion symbol \perp . We remark that this is only a cosmetic difference—it can be made consistent using the following rules: $\perp = \perp_{\widetilde{f}}$ and $\perp = \perp_{\text{invalid}}$ (i.e., $\text{Val} \neq \perp \Leftrightarrow (\text{Val} \neq \perp_{\widetilde{f}} \wedge \text{Val} \neq \perp_{\text{invalid}})$).

Proving Property 2 of Lem. 9. First, recall that [Property 2](#) requires us to show that for any pref in the support of $H_{\text{pref}}^{\mathcal{M}}(\lambda, m, z)$, if $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then it holds that

$$\Pr[\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau}) : \text{Val}_{\mathcal{K}} \leftarrow \mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \frac{\varepsilon'(\lambda)}{\tilde{t}}. \quad (26)$$

Also recall that $\mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ is defined to execute the machine $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ with i uniformly sampled from $[\tilde{t}]$. Therefore, [Inequality \(26\)](#) can be reduced to the following [Lem. 10](#). We will prove [Lem. 10](#) in [Sec. 4.6](#), which will eventually finish the current proof of [Lem. 9](#).

Lemma 10. *Let $\varepsilon(\lambda) = \frac{1}{\text{poly}(\lambda)}$ and $\varepsilon'(\lambda) = \frac{\varepsilon(\lambda)}{t^2}$. For any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$, if $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then there exists an $i \in [\tilde{t}]$ such that*

$$\Pr[\text{Val} = \text{val}(\tilde{\tau}) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \varepsilon'(\lambda).$$

4.6 Proof of [Lem. 10](#)

Notation. We highly recommend reviewing the “ [\$\mathcal{K}_i\$ ’s Output](#)” part on [Page 35](#) (in particular, the meanings of Val , $\perp_{\tilde{\gamma}}$, and \perp_{invalid}) before starting to read this subsection. Recall that \mathcal{K}_i ’s output Val is determined by the w' output by the witness-extended emulator of the right **WIPoK-2**. In this subsection, we will need to refer to this w' , though it is not explicitly included as a part of \mathcal{K}_i ’s output. Particularly, we will make use of the following notation: whenever we write an expression of the form

$$\Pr[\text{Some Event } E_{w'} \text{ about } w' : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})],$$

it should be understood as the probability of $E_{w'}$, where w' refers to the w' generated during the random procedure $\text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$, over which the probability is taken.

Using these notations, we can partition the event $\text{Val} = \perp_{\tilde{\gamma}}$ as the following *mutually exclusive* events: $w' = (1, \tilde{x}_1), \dots, w' = (\tilde{t}, \tilde{x}_{\tilde{t}})$, where $\tilde{y}_i = \tilde{f}(\tilde{x}_i)$ for each $i \in [\tilde{t}]$. Formally, we express this relation by

$$\Pr[\text{Val} = \perp_{\tilde{\gamma}} : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] = \sum_{i=1}^{\tilde{t}} \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \quad (27)$$

With the above notations, we prove [Lem. 10](#) in the following.

Proof for [Lem. 10](#). We assume for contradiction that for some pref satisfying $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, it holds that

$$\forall i \in [\tilde{t}], \Pr[\text{Val} = \text{val}(\tilde{\tau}) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] < \varepsilon'(\lambda). \quad (28)$$

Claim 8. *Under the assumption in [Inequality \(28\)](#), it holds that*

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda) - \text{negl}(\lambda).$$

Proof. In this proof, all the probabilities are taken over the random procedure $\text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$.

First, notice that

$$\begin{aligned} \forall i \in [\tilde{t}], \Pr[\text{Val} \neq \perp_{\text{invalid}}] &= \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \Pr[\text{Val} = \perp_{\tilde{\gamma}}] \\ &\text{(by Eq. (27))} = \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \Pr[w' = (i, \tilde{x}_i)] + \sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)]. \end{aligned} \quad (29)$$

Then, the following holds:

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i)] = \Pr[\text{Val} \neq \perp_{\text{invalid}}] - \Pr[\text{Val} = \text{val}(\tilde{\tau})] - \left(\sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)] \right) \quad (30)$$

$$\geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda) - \Pr[\text{Val} = \text{val}(\tilde{\tau})] - \left(\sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)] \right) \quad (31)$$

$$\geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda) - \varepsilon'(\lambda) - \left(\sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)] \right), \quad (32)$$

where Eq. (30) follows from Eq. (29), Inequality (31) follows from Claim 7, and Inequality (32) follows from Inequality (28).

Now, to prove Claim 8, it suffices to show that

$$\forall i \in [\tilde{t}], \forall j \in [\tilde{t}] \setminus \{i\}, \Pr[w' = (j, \tilde{x}_j) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] = \text{negl}(\lambda). \quad (33)$$

This can be reduced via standard techniques to the one-wayness of the OWF \tilde{f} in Step 3 of the right execution. In more details, we assume for contradiction that there exist $i^*, j^* \in [\tilde{t}]$ such that $i^* \neq j^*$ and that the probability $\Pr[w' = (j^*, \tilde{x}_{j^*}) : \text{Val} \leftarrow \mathcal{K}_{i^*}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})]$ is non-negligible, where, by definition, \tilde{x}_{j^*} is the preimage of \tilde{y}_{j^*} under the right OWF \tilde{f} . Then, we can build a PPT adversary \mathcal{A}_{OWF} breaking one-wayness in the following way: \mathcal{A}_{OWF} obtains the challenge y^* from the external one-wayness challenger; it then runs the machine $\mathcal{K}_{i^*}(\text{pref})$ internally, for which \mathcal{A}_{OWF} uses y^* in place of \tilde{y}_{j^*} when executing Step 3 in the right. Note that the internal execution of $\mathcal{K}_{i^*}(\text{pref})$ is identically to the real execution of \mathcal{K}_{i^*} , thus the extracted $w' = (j^*, \tilde{x}_{j^*})$ must satisfy $\tilde{f}(\tilde{x}_{j^*}) = \tilde{y}_{j^*} (= y^*)$ with non-negligible probability, breaking one-wayness.

This finishes the proof of Claim 8. □

Machine \mathcal{K}'_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 9a.) For a prefix pref , $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ proceeds as follows:

1. It first finishes Step 3 of the man-in-the-middle execution in the same way as $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. In particular, it will see in the left execution the values $Y = (y_1, \dots, y_t)$ sent from \mathcal{M} ;
2. It then recovers (x_1, \dots, x_t) by brute-force: Namely, for each $i \in [t]$, it inverts the OWF f to find x_i s.t. $f(x_i) = y_i$. It is possible that there exist some “bad” y_i ’s that are not in the range of f . For such bad i ’s, it sets $x_i = \perp$. If all the x_i ’s are bad, $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ halts and outputs Fail;
3. If this step is reached, we know that (x_1, \dots, x_t) cannot be all- \perp . $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ then picks an (s, x_s) uniformly at random from the good (i.e. non- \perp) x_i ’s.
4. Then, $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ continues to finish the execution in the same way as $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$, except that it uses (s, x_s) as the witness when executing the left **WIPoK-2**.

It is worth noting that the (j, x_j) extracted by the witness-extended emulator for the left **WIPoK-1** (inherited from $\mathcal{K}_i(\text{pref})$) is not used any more in $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$.

Obviously, if the (s, x_s) picked by $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ is equal to the (j, x_j) extracted in $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ from its left **WIPoK-1**, then $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ and $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ are identical. Notice that this step relies on the injectivity of f (see Rmk. 11). Since $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ samples (s, x_s) uniformly from all the good (i, x_i) ’s, it must hold with probability at least $1/t$ that $(s, x_s) = (j, x_j)$. Therefore, the following must hold:

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \frac{1}{t} \cdot \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})]. \quad (34)$$




Fig. 9: Machines \mathcal{K}'_i and \mathcal{K}''_i (Difference is highlighted in red color)

Remark 11 (On Injectivity of the OWF). We emphasize that throughout the proof of non-malleability, this is *the only* place where we rely on the injectivity of the OWF. In particular, we rely on the injectivity of the f in the left session to ensure that there is a unique preimage for each $\{y_i\}_{i \in [t]}$. Thus, if both the x_s with $s = j$ (sampled by \mathcal{K}'_i) and the extracted x_j (in \mathcal{K}_i) are a valid preimage for the same y_j , then the injectivity of f implies that $x_s = x_j$. We will show how to remove injectivity in [Sec. 4.8](#).

Machine \mathcal{K}''_i ($i \in [\tilde{t}]$): (Illustrated in [Fig. 9b](#).) For a prefix pref , $\mathcal{K}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it plays as the honest committer in the left **WIPoK-1**, instead of running the witness-extended emulator. Recall that starting from \mathcal{K}'_i , the witness (j, x_j) extracted by the witness-extended emulator from the left **WIPoK-1** is not used any more; Thus, machine \mathcal{K}''_i does not need to perform this witness-extended emulation.

By the PoK property of the left **WIPoK-1**, it holds that

$$\forall i \in [\tilde{t}], \left| \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] - \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda). \quad (35)$$

Next, by the (non-uniform, see [Rmk. 12](#)) WI property of the right **WIPoK-1**, it holds that

$$\forall i \in [\tilde{t}], \left| \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] - \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda). \quad (36)$$

Remark 12 (Power of Non-Uniform Reductions). Note that we can rely on the PoK and WI properties, although \mathcal{K}'_i and \mathcal{K}''_i perform brute-force to recover (x_1, \dots, x_t) . This is because the brute-forcing step is done before **WIPoK-1** or **WIPoK-2** starts; Thus, (x_1, \dots, x_t) can be treated as a non-uniform advice in the reductions. This non-uniform type of argument will be used again in this section later.

Then, we have the following claim:

Claim 9. $\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \frac{1}{\tilde{t}} \cdot (p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda)) - \text{negl}(\lambda).$

Proof. This claim follows from [Claim 8](#) and [Inequalities \(34\) to \(36\)](#). Formally, (in the following, we omit the input $(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ to $\mathcal{K}_i, \mathcal{K}'_i,$ and \mathcal{K}''_i)

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_1] \geq \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_i] - \text{negl}(\lambda) \quad (37)$$

$$\geq \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}'_i] - \text{negl}(\lambda) \quad (38)$$

$$\geq \frac{1}{t} \cdot \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i] - \text{negl}(\lambda) \quad (39)$$

$$\geq \frac{1}{t} \cdot (p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda)) - \text{negl}(\lambda), \quad (40)$$

where [Inequality \(37\)](#) follows from [Inequality \(36\)](#), [Inequality \(38\)](#) follows from [Inequality \(35\)](#), [Inequality \(39\)](#) follows from [Inequality \(34\)](#), and [Inequality \(40\)](#) follows from [Claim 8](#). \square

Now, we make the last claim which, together with [Claim 9](#), leads to the desired contradiction.

Claim 10. $\Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}''_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \leq p_{\text{pref}}^{\text{Sim}} + \text{negl}(\lambda)$.

Deriving the Contradiction. Before proving [Claim 10](#), we first show why [Claims 9](#) and [10](#) are contradictory (all the probabilities below are taken over $\text{Val} \leftarrow \mathcal{K}''_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$):

$$\begin{aligned} \Pr[\text{Val} \neq \perp_{\text{invalid}}] &= \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \Pr[\text{Val} = \perp_{\tilde{\gamma}}] \\ &= \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \sum_{i=1}^{\tilde{t}} \Pr[w' = (i, \tilde{x}_i)] \end{aligned} \quad (41)$$

$$\begin{aligned} &\geq \sum_{i=1}^{\tilde{t}} \Pr[w' = (i, \tilde{x}_i)] \\ &\geq \tilde{t} \cdot \frac{1}{t} \cdot (p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda)) - \text{negl}(\lambda) \end{aligned} \quad (42)$$

$$\geq (1 + \frac{1}{t}) \cdot (p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda)) - \text{negl}(\lambda) \quad (43)$$

$$\begin{aligned} &= p_{\text{pref}}^{\text{Sim}} + \left(\frac{p_{\text{pref}}^{\text{Sim}}}{t} - \varepsilon'(\lambda) - \frac{\varepsilon'(\lambda)}{t} \right) - \text{negl}(\lambda) \\ &\geq p_{\text{pref}}^{\text{Sim}} + \frac{10t^2 - t - 1}{10t^3} \cdot \varepsilon(\lambda) - \text{negl}(\lambda), \end{aligned} \quad (44)$$

where [Eq. \(41\)](#) follows from [Eq. \(27\)](#), [Inequality \(42\)](#) follows from [Claim 9](#), [Inequality \(43\)](#) follows from the assumption that $\tilde{t} \geq t+1$, and [Inequality \(44\)](#) follows from the assumption that $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$ and our parameter setting $\varepsilon'(\lambda) = \frac{\varepsilon(\lambda)}{10t^2}$.

Recall that t is the tag taking values from $[n]$ with n being a polynomial of λ . Also recall that $\varepsilon(\lambda)$ is an inverse polynomial on λ . Therefore, [Inequality \(44\)](#) can be written as:

$$\Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}''_1(\text{pref})] \geq p_{\text{pref}}^{\text{Sim}} + \frac{1}{\text{poly}(\lambda)} - \text{negl}(\lambda),$$

which contradicts [Claim 10](#).

This eventually finishes the proof of [Lem. 10](#) (modulo the proof of [Claim 10](#), which we show in [Sec. 4.7](#)).




Fig. 10: Machines \mathcal{K}_1^{**} and \mathcal{K}_1^* (Difference is highlighted in red color)

4.7 Proof of Claim 10

We first define two extra machines \mathcal{K}_1^{**} and \mathcal{K}_1^* .

Machine \mathcal{K}_1^{} :** (Illustrated in Fig. 10a.) For a prefix pref , $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ behaves identically as $\mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ except that $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ finishes the right **WIPoK-2** using the honest receiver's algorithm, instead of using the witness-extended emulator.

\mathcal{K}_1^{} 's Output.** We define the output of $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ to be the honest receiver's decision $b \in \{\top, \perp\}$. This is in contrast to previous hybrids \mathcal{K}_i' , \mathcal{K}_i'' , and \mathcal{K}_i , whose output is defined to be the value Val that depends on the value w' extracted by the \mathcal{WE} for the right **WIPoK-2**.

By the (non-uniform) PoK property (as per Def. 12) of the right **WIPoK-2**, it holds that

$$\left| \Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - \Pr[b = \top : b \leftarrow \mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda). \quad (45)$$

Machine \mathcal{K}_1^* : (Illustrated in Fig. 10b.) For a prefix pref , $\mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ behaves identically as $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ except the following difference: $\mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ uses the witness-extended emulator to extract a witness (j, x_j) from the left **WIPoK-1**, and if x_j is not a valid preimage for y_j , \mathcal{K}_1^* aborts.

By the (non-uniform) PoK property of the left **WIPoK-1**, it holds that

$$\left| \Pr[b = \top : b \leftarrow \mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda). \quad (46)$$

Compare \mathcal{K}_1^* with \mathcal{G}_1 . Now, let us compare \mathcal{K}_1^* with the \mathcal{G}_1 depicted in Fig. 6. They only differ in the witness used in the left **WIPoK-2** (and that \mathcal{G}_1 does not need to perform brute-forcing for Y , as it does not use those preimages). Therefore, by the (non-uniform) WI property of the left **WIPoK-2**, it holds that

$$\left| \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda).$$

Also, recall (from Def. 15) that $\Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})]$ is exactly the definition of $p_{\text{pref}}^{\text{Sim}}$. Thus, the above implies:

$$\left| \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - p_{\text{pref}}^{\text{Sim}} \right| \leq \text{negl}(\lambda). \quad (47)$$

Therefore, the following holds:

$$\Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \leq \Pr[b = \top : b \leftarrow \mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] + \text{negl}(\lambda) \quad (48)$$

$$\leq \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{pref})] + \text{negl}(\lambda) \quad (49)$$

$$\leq p_{\text{pref}}^{\text{Sim}} + \text{negl}(\lambda), \quad (50)$$

where [Inequality \(48\)](#) follows from [Inequality \(45\)](#), [Inequality \(49\)](#) follows from [Inequality \(46\)](#), and [Inequality \(50\)](#) follows from [Inequality \(47\)](#).

This finishes the proof of [Claim 10](#).

4.8 Replacing the Injective OWF with Any OWF

In this part, we show how to replace the injective OWF with any OWF in [Prot. 1](#).

As mentioned in [Rmk. 11](#), the injectivity is used only when we switch from \mathcal{K}_i to \mathcal{K}'_i (on [Page 39](#)). Let us first recall why injectivity is important there: \mathcal{K}'_i learns the preimages of each $\{y_i\}_{i \in [t]}$ via brute-forcing, while \mathcal{K}_i learns a preimage x_j of y_j for *some* $j \in [t]$ by extraction (from the left **WIPoK-1**). Our goal there is to show that if \mathcal{K}'_i picks a random preimage, then it will hit the (j, x_j) extracted by \mathcal{K}_i , except for a multiplicative $1/t$ loss on the probability (i.e., [Inequality \(34\)](#)). If the OWF f is not injective, there is no guarantee that the (j, x_j) extracted by \mathcal{K}_i is really contained in the set of all preimages learned by \mathcal{K}'_i via brute-forcing—Even if \mathcal{K}'_i guessed the index j correctly, it could have brute-forced a $x'_j \neq x_j$ (though it holds that $y_j = f(x_j) = f(x'_j)$). In this case, [Inequality \(34\)](#) may not hold anymore.

From the above discussion, it is clear that if we want to remove the requirement of injectivity, we only need to find an alternative approach to ensure that the preimage extracted by \mathcal{K}_i will fall in the set of the preimages learned by \mathcal{K}'_i via brute-forcing. It turns out that this can be achieved by the following simple modification to [Steps 3](#) and [4](#) of [Prot. 1](#):

- In [Step 3](#), ask R to *additionally* commit (using a statistically-binding scheme) to the preimages $\{x_i\}_{i \in [t]}$ that it used to generate $\{y_i\}_{i \in [t]}$. Formally, we ask R to send the following messages in [Step 3](#):

$$Y = \{y_i = f(x_i)\}_{i \in [t]}, \text{ and } \{\text{com}_i = \text{Com}_{\beta'}(x_i; r_i)\}_{i \in [t]},$$

where $\text{Com}_{\beta'}$ is Naor’s commitment with β' being the first Naor’s message from C (which could be sent in parallel with other messages in [Step 2](#)).

- In [Step 4](#), ask R to additionally prove that the Y and $\{\text{com}_i\}_{i \in [t]}$ are “consistent”. Formally, instead of proving $Y \in \mathcal{L}_f^t$ (defined in [Language \(10\)](#)), we ask R to prove $(Y, \{\text{com}_i\}_{i \in [t]}) \in \mathcal{L}_{f, \beta'}^t$ using the WIPoK, where

$$\mathcal{L}_{f, \beta'}^t := \{(y_1, \dots, y_t), (\text{com}_1, \dots, \text{com}_t) \mid \exists (i, x_i, r_i) \text{ s.t. } i \in [t] \wedge y_i = f(x_i) \wedge \text{com}_i = \text{Com}_{\beta'}(x_i; r_i)\}. \quad (51)$$

It is worth noting that we do *not* need to modify [Step 5](#). In particular, C does *not* need to prove that it knows one of the (x_i, r_i) ’s used by R to generate $\{\text{com}_i\}_{i \in [t]}$; C still only needs to prove $(\text{com}, Y) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t$.

Now, instead of performing brute-force over $\{y_i\}_{i \in [t]}$, \mathcal{K}'_i does that over $\{\text{com}_i\}_{i \in [t]}$ to learn the committed values, and then picks a random x_s satisfying $y_s = f(x_s)$ to use in the left **WIPoK-2**.²⁴ It is easy to see that the statistical-binding property of Com ensures that [Inequality \(34\)](#) still holds for an overwhelming fraction of β' .

²⁴ Note that \mathcal{K}'_i (resp. \mathcal{K}_i) will also learn the randomness used to generate the commitment via brute-forcing (resp. via extraction from **WIPoK-1**). They simply ignore the randomness—indeed, both machines only need to know a preimage (j, x_j) to finish **WIPoK-2**; the corresponding randomness r_i is irrelevant.

Remark 13 (On One-Wayness). Technically, we also need to argue that the above modification does not enable the man-in-the-middle \mathcal{M} to invert the one-way function \tilde{f} on $\tilde{Y} = \{\tilde{y}_i\}_{i \in [\tilde{\ell}]}$. (Recall that the one-wayness of \tilde{f} is used (only) in the proof of [Claim 8](#).) This is straightforward—Any OWF remains one-way even if a commitment to its preimage is given.

4.9 Summarizing the Proof Strategy for [Thm. 2](#)

Since the proof of [Thm. 2](#) is complicate, we summarize the logic flow in this subsection. We hope that it will help the reader obtain a better understanding of the whole picture, without being disturbed by details of secondary importance. But one could safely skip this subsection as it does not contain new information. In particular, it will not affect the understanding of succeeding sections.

At a high level, the proof of [Thm. 2](#) can be divided into the following five steps:

1. (**Proof by Contradiction.**) In the proof of [Thm. 2](#), we first denoted the MIM execution by $H^{\mathcal{M}}(\lambda, m, z)$ (on [Page 26](#)). We then assumed for contradiction that a (potentially non-uniform) PPT adversary can distinguish the outputs of $H^{\mathcal{M}}(\lambda, m_0, z)$ and $H^{\mathcal{M}}(\lambda, m_1, z)$ with advantage $3 \cdot \delta(\lambda)$, with $\delta(\lambda)$ being some polynomial (i.e., [Inequality \(15\)](#)). Next, we required the existence of a machine $G^{\mathcal{M}}$ (i.e., [Lem. 6](#)) such that

- (a) the outputs of $G^{\mathcal{M}}(\lambda, m_0, z)$ and $G^{\mathcal{M}}(\lambda, m_1, z)$ are computationally indistinguishable, and
- (b) $G^{\mathcal{M}}(\lambda, m_b, z)$ and $H^{\mathcal{M}}(\lambda, m_b, z)$ is (computationally) $\delta(\lambda)$ -close for both $b = 0$ and $b = 1$.

The existence of such a $G^{\mathcal{M}}$ finished the proof of [Thm. 2](#), because it contradicts the assumption that there exists some PPT distinguisher that tells the difference between $H^{\mathcal{M}}(\lambda, m_0, z)$ and $H^{\mathcal{M}}(\lambda, m_1, z)$ with advantage $3 \cdot \delta(\lambda)$. Thus, the remaining steps for the proof of [Thm. 2](#) were then devoted to constructing such a $G^{\mathcal{M}}$.

2. (**Constructing $G^{\mathcal{M}}$.**) To build the $G^{\mathcal{M}}$ required in [Step 1](#), we first broke $H^{\mathcal{M}}$ into two stages—the prefix part $H_{\text{pre}}^{\mathcal{M}}$ and the remainder $\langle \mathcal{A}, \mathcal{B} \rangle$ (i.e., [Algo. 4.1](#)). Here, we defined an important variable called “prefix”:

- An execution of $H_{\text{pre}}^{\mathcal{M}}$ fixes the prefix pref , which contains the internal state $\text{st}_C, \text{st}_{\mathcal{M}}, \text{st}_R$ of the corresponding parties at the end of [Step-2](#) in a MIM execution. Also, pref fixes the right Naor’s commitment $\tilde{\tau}$. Also note that the committed value in the right execution was defined to be $\text{val}_b(\tilde{\tau})$ (i.e., [Eq. \(12\)](#)).

Then, we constructed a simulation-extractor $\mathcal{SE}^{\mathcal{M}}$ (in [Lem. 7](#)) such that, when continuing from any prefix pref (Importantly, \mathcal{SE} cannot use the intermediate state st_C for the committer; Otherwise, we cannot show [Property 1a](#) by reducing it to the hiding property of Naor’s commitment), $\mathcal{SE}^{\mathcal{M}}$ will ε -simulate the remainder execution $\langle \mathcal{A}, \mathcal{B} \rangle$, while also extracting the committed value $\text{val}_b(\tilde{\tau})$ in the right execution.

With this $\mathcal{SE}^{\mathcal{M}}$, it is easy to construct $G^{\mathcal{M}}$ — $G^{\mathcal{M}}$ simply runs $H_{\text{pre}}^{\mathcal{M}}$ to obtain the prefix, and then runs \mathcal{SE} to simulate the remainder execution and also extract $\text{val}_b(\tilde{\tau})$. Thus, the remaining steps for the proof of [Thm. 2](#) were then devoted to constructing such a $\mathcal{SE}^{\mathcal{M}}$.

3. (**Constructing $\mathcal{SE}^{\mathcal{M}}$.**) To construct $\mathcal{SE}^{\mathcal{M}}$, we first built two machines:

- (a) a simulator \mathcal{G}_1 (in [Lem. 8](#)) that can simulate the execution $\langle \mathcal{A}, \mathcal{B} \rangle$ without knowing C ’s internal states st_C . This is to address the point that $\mathcal{SE}^{\mathcal{M}}$ is not allowed to use st_C . Here, we defined an important variable $p_{\text{pref}}^{\text{Sim}}$, which denotes the probability that the honest R is convinced in the execution simulated by \mathcal{G}_1 .

- (b) an extractor \mathcal{K} (in [Lem. 9](#)) such that for any prefix pref and noticeable $\varepsilon(\lambda)$, if $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then \mathcal{K} can extract the value committed in $\tilde{\tau}$ (contained in pref) with probability polynomially related to $\varepsilon(\lambda)$ (more accurately, the probability is $\frac{\varepsilon'(\lambda)}{t}$ where $\varepsilon'(\lambda) := \frac{\varepsilon(\lambda)}{10t^2}$).

With \mathcal{G}_1 and \mathcal{K} , $\mathcal{SE}^{\mathcal{M}}$ can be constructed as follows: on input a pref , it first runs \mathcal{G}_1 to simulate the main-tread execution of $\langle \mathcal{A}, \mathcal{B} \rangle$; If the honest receiver is not convinced in the main-thread, $\mathcal{SE}^{\mathcal{M}}$ simply sets the committed value in the right execution as \perp ; Otherwise, $\mathcal{SE}^{\mathcal{M}}$ will extract the value committed in $\tilde{\tau}$ (denoted as $\text{val}(\tilde{\tau})$) by repeating \mathcal{K} for polynomially many times (more accurately, $\frac{\tilde{t}}{\varepsilon'(\lambda)} \cdot \lambda$ times). As shown toward the end of [Sec. 4.4](#), such a $\mathcal{SE}^{\mathcal{M}}$ satisfied the requirements described in [Step 2](#). Thus, the remaining steps for the proof of [Thm. 2](#) were then devoted to constructing the extractor \mathcal{K} . (The construction and security of \mathcal{G}_1 is relatively straightforward, so we do not repeat it in this summary.)

4. (**Constructing \mathcal{K} .**) To construct \mathcal{K} , we first defined a sequence of machines $\{\mathcal{K}_i\}_{i \in [\tilde{t}]}$ as shown in [Fig. 8a](#). We then argued that there must exist an $i \in [\tilde{t}]$ (called “the good i ”) such that the Val output by \mathcal{K}_i must be $\text{val}(\tilde{\tau})$ with probability $\geq \varepsilon'(\lambda)$. Then, machine \mathcal{K} was defined to pick a random $i \xleftarrow{\$} [\tilde{t}]$ and run machine \mathcal{K}_i . Since such a \mathcal{K} will hit the good i with probability $1/\tilde{t}$, it will successfully extract $\text{val}(\tilde{\tau})$ with probability $\frac{\varepsilon'(\lambda)}{\tilde{t}}$, satisfying the requirement in [Step 3](#). Now, to finish the proof of [Thm. 2](#), the only thing left is to prove that $\{\mathcal{K}_i\}_{i \in [\tilde{t}]}$ satisfies the above requirement.
5. (**Existence of Good \mathcal{K}_i .**) Informally, our current goal is to show the following (see [Lem. 10](#) for the precise statement):

$$\exists i \in [\tilde{t}] \text{ s.t. } \Pr[\text{Val} = \text{val}(\tilde{\tau}) : \text{Val} \leftarrow \mathcal{K}_i] \geq \varepsilon'(\lambda).$$

We first assumed for contradiction that (i.e., [Inequality \(28\)](#)):

$$\forall i \in [\tilde{t}], \Pr[\text{Val} = \text{val}(\tilde{\tau}) : \text{Val} \leftarrow \mathcal{K}_i] \leq \varepsilon'(\lambda). \quad (52)$$

We then derived the contradiction in the following way:

- (a) We first showed that the Val output by \mathcal{K}_i must be a valid witness with probability $\geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda)$ (i.e., [Claim 7](#)). Next, observe that in \mathcal{K}_i , Val cannot be (j, \tilde{x}_j) for some $j \neq i$; Otherwise, it breaks the one-wayness of the right OWF \tilde{f} . Thus, when Val is a valid witness in \mathcal{K}_i , it can only take the values $\text{val}(\tilde{\tau})$ or (i, \tilde{x}_i) . This together with our assumption in [Inequality \(52\)](#) implies that

$$\forall i \in [\tilde{t}], \Pr[\text{Val} = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i] \geq p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda) - \text{negl}(\lambda). \quad (53)$$

(This is exactly [Claim 8](#).)

- (b) Next, we defined machine \mathcal{K}_i'' (see [Fig. 9b](#)), which differs from \mathcal{K}_i in the following way: \mathcal{K}_i'' does not rewind the left **WIPoK-1**. Instead, it extracts the preimages of the right $\{y_i\}_{i \in [t]}$ by brute force, and then picks a random preimage (s, x_s) to use in the left **WIPoK-2**. Since the (s, x_s) guessed by \mathcal{K}_i'' will hit the one used by \mathcal{K}_i , the following holds:

$$\forall i \in [\tilde{t}], \Pr[\text{Val} = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i''] \geq \frac{1}{t} \cdot \Pr[\text{Val} = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i]. \quad (54)$$

(This is essentially [Inequality \(34\)](#)—Note that in [Sec. 4.6](#), we also defined a machine \mathcal{K}_i' . But it is only a tool that helps us eventually connect \mathcal{K}_i'' and \mathcal{K}_i . So we ignore \mathcal{K}_i' in this relatively high-level summary.)

(c) Since in \mathcal{K}_i'' we do not rewind the left **WIPoK-1** anymore, we can now rely on the WI property of the right **WIPoK-1** to connect all the $\{\mathcal{K}_i''\}_{i \in [\tilde{t}]}$ to \mathcal{K}_1'' . Namely, we have

$$\forall i \in [\tilde{t}], \quad \left| \Pr[\text{Val} = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i''] - \Pr[\text{Val} = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_1''] \right| \leq \text{negl}(\lambda). \quad (55)$$

(This is exactly [Inequality \(36\)](#).) We want to remark that in this step, we took advantage of the synchronous schedule: in the synchronous setting, it is guaranteed that the brute-forcing for $\{y_i\}_{i \in [n]}$ happens *before* the right **WIPoK-1**; Otherwise, we will not be able to reduce [Inequality \(55\)](#) to the WI property of the right **WIPoK-1**.

Now, by [Inequalities \(53\)](#) to [\(55\)](#), the following holds (this is exactly [Claim 9](#))

$$\forall i \in [\tilde{t}], \quad \Pr[\text{Val} = (i, \tilde{x}_i) \text{ in } \mathcal{K}_1''] \geq \frac{1}{t} \cdot (p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda)) - \text{negl}(\lambda). \quad (56)$$

This implies

$$\begin{aligned} \Pr[\text{Val} = (1, \tilde{x}_1) \vee \dots \vee \text{Val} = (\tilde{t}, \tilde{x}_{\tilde{t}}) : \text{Val} \leftarrow \mathcal{K}_1''(\text{pref})] &\geq \tilde{t} \cdot \frac{1}{t} \cdot (p_{\text{pref}}^{\text{Sim}} - \varepsilon'(\lambda)) + \text{negl}(\lambda) \\ &\text{(By our parameter setting. See [Inequality \(44\)](#))} \geq p_{\text{pref}}^{\text{Sim}} + \frac{10t^2 - t - 1}{10t^3} \cdot \varepsilon(\lambda) - \text{negl}(\lambda) \end{aligned}$$

On the other hand, [Claim 10](#) implies that (we do not repeat the proof of [Claim 10](#) here as it is relatively straightforward)

$$\Pr[\text{Val} = (1, \tilde{x}_1) \vee \dots \vee \text{Val} = (\tilde{t}, \tilde{x}_{\tilde{t}}) : \text{Val} \leftarrow \mathcal{K}_1''(\text{pref})] \leq p_{\text{pref}} + \text{negl}(\lambda)$$

This gives us the desired contradiction because $\frac{10t^2 - t - 1}{10t^3} \cdot \varepsilon(\lambda)$ is an (inverse) polynomial on λ .

5 Small-Tag, Synchronous, Classical Setting

In this section, we show how to make $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ (i.e., [Prot. 1](#)) secure without the “one-sided” restriction. We emphasize that this section still focuses on the small-tag, synchronous setting against classical (i.e., non-uniform PPT) adversaries.

5.1 High-Level Idea

Our proof for the non-malleability of $\langle C, R \rangle_{\text{tg}}^{\text{OneSided}}$ (shown in [Prot. 1](#)) works only if $t < \tilde{t}$. However, this is not guaranteed in the real main-in-the-middle attack—the adversary can of course use a smaller tag in the right session. Fortunately, this problem can be addressed by the so-called “two-slot” technique proposed by Pass and Rosen [[PR05](#)]. The idea is to create a situation where no matter how the man-in-the-middle \mathcal{M} schedules the messages, there is always a “slot” for which the “ $t < \tilde{t}$ ” condition holds; As long as this is true, non-malleability can be proven using the same technique as in [Sec. 4](#).

To do that, first observe that the only place where [Prot. 1](#) makes use of the tag t is [Steps 3](#) and [4](#). Let us call these two messages a “**Slot**”. Our modification is to execute this **Slot** twice sequentially, where

- for the first execution (referred to as **Slot-A**), we ask R to use tag t as in the original construction;
- for the second execution (referred to as **Slot-B**), we ask R to use $n - t$ as the tag.

We also modify the language instance proven by **WIPoK-2** in the expected way—Now C needs to prove that it “knows” the (m, r) tuple, or the preimage of one of the y_i ’s sent in **Slot-A**, or the preimage of one of the y_i ’s sent in **Slot-B**.

By the above design, it is easy to see that one of the following case must happen for any man-in-the-middle execution (note that we are still in the synchronous setting):

1. $t = \tilde{t}$: This is the trivial case that is already ruled out by the definition of non-malleability.
2. $t < \tilde{t}$: In this case, non-malleability follows by applying the argument in [Sec. 4](#) to **Slot-A**. We sometimes refer to this case by saying that **Slot-A** is the good slot.
3. $t > \tilde{t}$: In this case, it holds that $(n - t) < (n - \tilde{t})$. In other words, the tag for the left **Slot-B** is smaller than the tag for the right **Slot-B**. Therefore, non-malleability follows by applying the argument in [Sec. 4](#) to **Slot-B**. We sometimes refer to this case by saying that **Slot-B** is the good slot.

Therefore, the modified protocol is non-malleable without the “one-sided” restriction.

5.2 Construction of $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$

We refer to the construction described in [Sec. 5.1](#) as $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$, and formally present it in [Prot. 2](#). This protocol makes use of the same building blocks as for [Prot. 1](#), namely:

- An OWF f , (Note that we do not require injectivity any more, because [Prot. 2](#) makes use of the technique discussed in [Sec. 4.8](#).)
- Naor’s commitment Com ;
- A witness-indistinguishable proof of knowledge WIPoK (as per [Def. 13](#)).

We remark that we further optimize the “two-slot” approach from [Sec. 5.1](#) by sending the first message of **Slot-A** and the first message of **Slot-B** together.

Protocol 2: Small-Tag Synchronous NMCCom $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$

The tag space is defined to be $[n]$ where n is a polynomial on λ . Let $t \in [n]$ be the tag for the following interaction. Let m be the message to be committed to.

Commit Stage:

1. Receiver R samples and sends the first message β for Naor’s commitment.
2. Committer C commits to m using the second message of Naor’s commitment. C also sends a β' that will be used as the first Naor’s message for R to generate a commitment in next step. Formally, C samples r and β' and sends the tuple $(\text{com} = \text{Com}_{\beta}(m; r), \beta')$.

Comment: The β' is to address the injectivity issue as discussed in [Sec. 4.8](#).

3. R performs the following computation:

- (a) R computes $\{y_i^A = f(x_i^A)\}_{i \in [t]}$ with $x_i^A \xleftarrow{\$} \{0, 1\}^\lambda$ for each $i \in [t]$, and sets $Y^A := (y_1^A, \dots, y_t^A)$. R also computes $\{\text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A)\}_{i \in [t]}$, where r_i^A is a random string for each $i \in [t]$.
- (b) R computes $\{y_i^B = f(x_i^B)\}_{i \in [n-t]}$ with $x_i^B \xleftarrow{\$} \{0, 1\}^\lambda$ for each $i \in [n-t]$, and sets $Y^B := (y_1^B, \dots, y_{n-t}^B)$. R also computes $\{\text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B)\}_{i \in [n-t]}$, where r_i^B is a random string for each $i \in [n-t]$.

R sends the tuple $(Y^A, \{\text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A)\}_{i \in [t]}, Y^B, \{\text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B)\}_{i \in [n-t]})$.

Comment: Compared with [Step 3](#) of [Prot. 1](#), the current [Step 3](#) has two differences: (1) R additionally commits to the preimages of y_i 's. This is to address the injectivity issue as discussed in [Sec. 4.8](#). (2) R sends *two copies* of the $(Y, \{\text{com}_i\}_i)$ tuple, indexed by superscripts A and B respectively. This is corresponding to the “two-slot” approach discussed in [Sec. 5.1](#) (with the first message of **Slot-B** being sent in parallel with the first message of **Slot-A**).

4. (**WIPoK-1-A.**) R and C execute an instance of WIPoK where R proves to C that he “knows” a pre-image of some y_i^A contained in Y^A , and y_i^A are consistent with com_i^A (as defined in [Step 3a](#)). Formally, R proves that $(Y^A, \{\text{com}_i^A\}_{i \in [t]}) \in \mathcal{L}_{f, \beta'}^t$, where

$$\mathcal{L}_{f, \beta'}^t := \left\{ (y_1^A, \dots, y_t^A), (\text{com}_1^A, \dots, \text{com}_t^A) \mid \exists (i, x_i^A, r_i^A) \text{ s.t. } \begin{array}{l} i \in [t] \wedge \\ y_i^A = f(x_i^A) \wedge \\ \text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A) \end{array} \right\}. \quad (57)$$

Note that R uses $(1, x_1^A, r_1^A)$ as the witness when executing this WIPoK.

5. (**WIPoK-1-B.**) R and C execute an instance of WIPoK where R proves to C that he “knows” a pre-image of some y_i^B contained in Y^B , and y_i^B are consistent with com_i^B (as defined in [Step 3b](#)). Formally, R proves that $(Y^B, \{\text{com}_i^B\}_{i \in [n-t]}) \in \mathcal{L}_{f, \beta'}^{n-t}$, where

$$\mathcal{L}_{f, \beta'}^{n-t} := \left\{ (y_1^B, \dots, y_{n-t}^B), (\text{com}_1^B, \dots, \text{com}_{n-t}^B) \mid \exists (i, x_i^B, r_i^B) \text{ s.t. } \begin{array}{l} i \in [n-t] \wedge \\ y_i^B = f(x_i^B) \wedge \\ \text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B) \end{array} \right\}. \quad (58)$$

Note that R uses $(1, x_1^B, r_1^B)$ as the witness when executing this WIPoK.

6. (**WIPoK-2.**) C and R execute an instance of WIPoK where C proves to R that he “knows” the message committed in com (defined in [Step 2](#)), *or* a pre-image of some y_i^A contained in Y^A (defined in [Step 3a](#)), *or* a pre-image of some y_i^B contained in Y^B (defined in [Step 3b](#)). Formally, C proves that $(\text{com}, Y^A, Y^B) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t \vee \mathcal{L}_f^{n-t}$, where

$$\mathcal{L}_\beta := \{\text{com} \mid \exists (m, r) \text{ s.t. } \text{com} = \text{Com}_\beta(m; r)\}, \quad (59)$$

$$\mathcal{L}_f^t := \{(y_1^A, \dots, y_t^A) \mid \exists (i, x_i^A) \text{ s.t. } i \in [t] \wedge y_i^A = f(x_i^A)\}, \quad (60)$$

$$\mathcal{L}_f^{n-t} := \{(y_1^B, \dots, y_{n-t}^B) \mid \exists (i, x_i^B) \text{ s.t. } i \in [n-t] \wedge y_i^B = f(x_i^B)\}. \quad (61)$$

Note that C uses the (m, r) defined in [Step 2](#) as the witness when executing this WIPoK.

Comment: Compared with [Step 5](#) of [Prot. 1](#), we add \mathcal{L}_f^{n-t} as an additional OR part to the language to be proven. This is corresponding to the new **Slot-B** as discussed in [Sec. 5.1](#).

Decommit Stage: C sends (m, r) . R accepts if $\text{com} = \text{Com}_\beta(m; r)$, and rejects otherwise.

Security. Completeness is straightforward from the description of [Prot. 2](#). The statistical binding property follows from that of Naor’s commitment. Computational hiding of any commitment scheme follows directly from non-malleability. So it remains for us to show that our commitment protocol is non-malleable, which we establish by the following theorem.

Theorem 11. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$ in [Prot. 2](#) is non-malleable against synchronous PPT adversaries with tag space $[n]$, with n being any polynomial on λ .*

5.3 Proving Non-Malleability (Proof of Thm. 11)

Since we only consider synchronous adversaries in this section, we can assume that the adversary declares t and \tilde{t} at the beginning w.l.o.g. Thus, we can analyze the three cases of $t = \tilde{t}$, $t < \tilde{t}$, and $t > \tilde{t}$ *separately*. Since the case of $t = \tilde{t}$ is already ruled out by the definition of non-malleability, we prove non-malleability for the other two cases below.

5.3.1 The Case of $t < \tilde{t}$

The proof for this case is almost identical to the proof of Prot. 1 in Sec. 4. To avoid re-doing the same work, we only focus on the places where modifications are needed. In the following, we assume that the reader has already read and understood the proof in Sec. 4.

First, by repeating the similar arguments in Sec. 4.2 to 4.4, we can see that it suffices to construct a machine \mathcal{K} that satisfies similar properties as required in Lem. 9. Roughly speaking, \mathcal{K} is required to extract the committed message \tilde{m} in the right session with a noticeable probability if it is given a prefix that leads to acceptance with a noticeable probability.²⁵ The construction of \mathcal{K} is identical to that in Sec. 4.5 except that it extracts a witness (j, x_j^A, r_j^A) from **WIPoK-1-A** and always uses the same witness $(1, \tilde{x}_1^B, \tilde{r}_1^B)$ for the right **WIPoK-1-B**. Specifically, \mathcal{K} honestly runs the man-in-the-middle experiment starting from the given prefix except for the following modifications:

1. Use the witness $(i, \tilde{x}_i^A, \tilde{r}_i^A)$ for the right **WIPoK-1-A** for $i \stackrel{\$}{\leftarrow} [\tilde{t}]$;
2. Extract a witness (j, x_j^A, r_j^A) from the left **WIPoK-1-A**;
3. Use the witness (j, x_j^A) in the left **WIPoK-2**;
4. Extract a witness \tilde{w} from the right **WIPoK-2**;
5. If \tilde{w} is a valid witness (\tilde{m}, \tilde{r}) for $\widehat{\text{com}} \in \mathcal{L}_{\tilde{\beta}}$, it outputs \tilde{m} . Otherwise, it outputs \perp .

We stress again that \mathcal{K} uses the same witness $(1, \tilde{x}_1^B, \tilde{r}_1^B)$ for the right **WIPoK-1-B** regardless of the i sampled in Item 1.²⁶ We also remark that \mathcal{K} extracts a witness from the left **WIPoK-1-A** but not from the left **WIPoK-1-B**.

The proof that the above \mathcal{K} satisfies similar properties as in Lem. 9 is almost identical to that in Sec. 4.5 to 4.7 in the one-sided setting. The only difference is the following. We define \mathcal{K}_i similarly to that in Sec. 4.5, i.e., it works similarly to \mathcal{K} with the random i (in Item 1) being fixed. We need to ensure that the witness \tilde{w} extracted from the right **WIPoK-2** by \mathcal{K}_i cannot be a witness for **Slot-B**, i.e., a witness for $Y^B \in \mathcal{L}_f^{n-t}$ with a non-negligible probability. This is needed in the proof of a counterpart of Claim 8 in Sec. 4.6. It can be shown as follows:

- It is easy to see that $\tilde{w} = (j, \tilde{x}_j^B)$ for $j \neq 1$ with a negligible probability: Since \mathcal{K}_i uses $(1, \tilde{x}_1^B, \tilde{r}_1^B)$ for the right **WIPoK-1-B**, extracting $(j, \tilde{x}_j^B, \tilde{r}_j^B)$ for $j \neq 1$ directly breaks the one-wayness of f or the computational hiding of Com.
- Suppose that $\tilde{w} = (1, \tilde{x}_1^B)$ with a non-negligible probability. The probability remains non-negligible even if we replace the witness used in the right **WIPoK-1-B** with $(2, \tilde{x}_2^B, \tilde{r}_2^B)$ by the WI property of the right **WIPoK-1-B**. Then, due to a similar reason to the above case, this also breaks the one-wayness of f or the computational hiding of Com. It is worth mentioning that this step relies on the fact that the left **WIPoK-1-A** and the right **WIPoK-1-B** do not interleave with each

²⁵ Strictly speaking, the input of \mathcal{K} is not a prefix, but a prefix *without* C 's state ST_C . We also remark that the assumption is not that the prefix leads to R 's acceptance with a noticeable probability in the real experiment, but in the *simulated* experiment defined similarly to \mathcal{G}_1 in Sec. 4.4.

²⁶ Changing the witness for **Slot-B** to $(i, \tilde{x}_i^B, \tilde{r}_i^B)$ by using the WI property also works, but that is redundant.

other by the synchronicity assumption. (Otherwise we have to rewind the right **WIPoK-1-B** when rewinding the left **WIPoK-1-A**, in which case we cannot rely on the WI property of the right **WIPoK-1-B**.) This point will become relevant when dealing with asynchronous adversaries in [Sec. 6](#).

Except for the above, the rest of the proof follows from a straightforward adaptation of that in [Sec. 4](#).

5.3.2 The Case of $t > \tilde{t}$

In this case, we have $n - t < n - \tilde{t}$. Thus, we can simply do the same analysis as the case of $t < \tilde{t}$ with the roles of **Slot-A** and **Slot-B** being swapped.

By combining the analysis for the above cases, we obtain [Thm. 11](#).

6 Small-Tag, Asynchronous, Classical Setting

In this section, we show how to make $\langle C, R \rangle_{\text{tg}}^{\text{sync}}$ (i.e., [Prot. 2](#)) secure against asynchronous adversaries, to obtain the protocol $\langle C, R \rangle_{\text{tg}}^{\text{async}}$. We emphasize that this section still focuses on the small-tag setting against classical (i.e., non-uniform PPT) adversaries.

6.1 High-Level Idea

By a careful inspection, the proof of non-malleability of [Prot. 2](#) (which is based on the proof of non-malleability of [Prot. 1](#)) relies on the following assumptions about the schedule.

- The prefix of both the left and right sessions should be generated before moving forward. That is, the left (resp. right) message of [Step 2](#) is sent before the right (resp. left) message of [Step 3](#) is sent. This is necessary because we need to analyze the protocol for each fixed prefix `pref` in the security proof for [Prot. 2](#).
- The left message of [Step 3](#) is sent before the right **WIPoK-1-A** starts. This is needed because we rely on brute-force search to break hiding of com_i^A or com_i^B in an intermediate hybrid, where we rely on (non-uniform) WI of the right **WIPoK-1-A**. This is possible only if we can do brute-force before the right **WIPoK-1-A** starts so that the result of the brute-force can be treated as a non-uniform advice in the reduction to WI. (See [Rmk. 12](#).)
- The left (resp. right) **WIPoK-1-A** and the right (resp. left) **WIPoK-1-B** do not interleave with each other. This is needed because we need to ensure that when rewinding the left **WIPoK-1-A** (resp. **WIPoK-1-B**), we do not rewind the right **WIPoK-1-B** (resp. **WIPoK-1-A**) as explained in [Sec. 5.3.1](#).
- The left **WIPoK-1-B** finishes before the right **WIPoK-2** starts. This is needed because we rewind both of them simultaneously in the construction of the extractor \mathcal{K} .²⁷

To deal with a schedule that does not satisfy some of the above conditions, we need to modify the protocol. Our modification is based on a combination of the following tricks.

1. Insert an extractable commitment to m by C between Steps X and $X + 1$. We also modify the language for **WIPoK-2** to ensure that the committed message in the extractable commitment is the correct one. This gives us an opportunity to extract the right message \tilde{m} without extracting

²⁷ In fact, it is also fine that the left **WIPoK-1-A** starts after the right **WIPoK-2** finishes, since they do not interleave with each other in this case either.

the left message m in the man-in-the-middle game if the right Step $X + 1$ starts before the left Step X finishes.

2. Insert a WIPoK proving some “trapdoor statement” where R plays the role of the prover between Steps X and $X + 1$. Here, the trapdoor is a witness for an NP language chosen by R . We also modify the language for **WIPoK-2** to allow C to use the trapdoor as a witness. Now, if the left Step $X + 1$ starts before the right Step X finishes, then we can extract the left trapdoor without extracting the right trapdoor in the man-in-the-middle game. Then, we can use the extracted trapdoor as a witness for the left **WIPoK-2**. As a result, we can simulate the honest committer in the left session without using the message m . This enables us to extract the right message \tilde{m} from the right **WIPoK-2** without using the left message m , and eventually to prove non-malleability.
3. Repeat Step X as many times as the total round complexity of all steps before Step X plus one. By the pigeonhole principle, this ensures that there is at least one execution of the left (resp. right) Step X that does not interleave with all steps before Step X in the right (resp. left). In particular, when we rewind that execution, it does not affect the security of any primitive on the other side.

6.2 Construction of $\langle C, R \rangle_{\text{tg}}^{\text{async}}$

We present our construction in [Prot. 3](#) where constants n_5, \dots, n_9 are specified later. This protocol makes use of the same building blocks as for [Prot. 2](#) plus an extractable commitment, namely:

- An OWF f ,
- Naor’s commitment Com ;
- A witness-indistinguishable proof of knowledge WIPoK (as per [Def. 13](#));
- An extractable commitment ExtCom (as per [Def. 8](#)). We denote by $\text{Verify}_{\text{ExtCom}}$ the verification algorithm of ExtCom in the decommit stage.

<p>Protocol 3: Small-Tag Asynchronous NMCCom $\langle C, R \rangle_{\text{tg}}^{\text{async}}$</p> <p>The tag space is defined to be $[n]$ where n is a polynomial on λ. Let $t \in [n]$ be the tag for the following interaction. Let m be the message to be committed to.</p> <p>Commit Stage:</p> <ol style="list-style-type: none"> 1. Receiver R samples and sends the first message β for Naor’s commitment. 2. Committer C commits to m using the second message of Naor’s commitment. C also sends a β' that will be used as the first Naor’s message for R to generate a commitment in next step. Formally, C samples r and β' and sends the tuple $(\text{com} = \text{Com}_\beta(m; r), \beta')$. 3. R and C do the following: <ol style="list-style-type: none"> (a) (TrapGen.) R computes $V_0 = f(v_0)$ and $V_1 = f(v_1)$ with $v_0, v_1 \xleftarrow{\\$} \{0, 1\}^\lambda$ and sends (V_0, V_1) to C (b) (WIPoK-Trap.) R and C execute an instance of WIPoK where R proves to C that he “knows” a pre-image of V_0 or V_1. Formally, R proves that $(V_0, V_1) \in \mathcal{L}_f^{\text{OR}}$, where $\mathcal{L}_f^{\text{OR}} := \{(V_0, V_1) \mid \exists v \text{ s.t. } f(v) = V_0 \vee f(v) = V_1\}. \quad (62)$ <p>Note that R uses v_0 as the witness when executing this WIPoK.</p>

Comment: This step is inserted for [Trick 2](#) in [Sec. 6.1](#) where a witness for $(V_0, V_1) \in \mathcal{L}_f^{\text{OR}}$ plays the role of a “trapdoor”.

4. R performs the following computation:

(a) R computes $\{y_i^A = f(x_i^A)\}_{i \in [t]}$ with $x_i^A \xleftarrow{\$} \{0, 1\}^\lambda$ for each $i \in [t]$, and sets $Y^A := (y_1^A, \dots, y_t^A)$. R also computes $\{\text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A)\}_{i \in [t]}$, where r_i^A is a random string for each $i \in [t]$.

(b) R computes $\{y_i^B = f(x_i^B)\}_{i \in [n-t]}$ with $x_i^B \xleftarrow{\$} \{0, 1\}^\lambda$ for each $i \in [n-t]$, and sets $Y^B := (y_1^B, \dots, y_{n-t}^B)$. R also computes $\{\text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B)\}_{i \in [n-t]}$, where r_i^B is a random string for each $i \in [n-t]$.

R sends the tuple $(Y^A, \{\text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A)\}_{i \in [t]}, Y^B, \{\text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B)\}_{i \in [n-t]})$.

5. (**ExtCom-1** $\times n_5$.) C and R sequentially execute n_5 instances of **ExtCom** where C commits to m . For $i \in [n_5]$, let τ_i^1 and decom_i^1 be the transcript and decommitment information (privately obtained by C) of the i -th execution.

Comment: This step is inserted for [Trick 1](#) and repeated for [Trick 3](#) in [Sec. 6.1](#).

6. (**WIPoK-1-A** $\times n_6$.) R and C sequentially execute n_6 instances of **WIPoK** where R proves to C that he “knows” a pre-image of some y_i^A contained in Y^A , and y_i^A are consistent with com_i^A (as defined in [Step 4a](#)). Formally, R proves that $(Y^A, \{\text{com}_i^A\}_{i \in [t]}) \in \mathcal{L}_{f, \beta'}^t$, where

$$\mathcal{L}_{f, \beta'}^t := \left\{ (y_1^A, \dots, y_t^A), (\text{com}_1^A, \dots, \text{com}_t^A) \mid \exists (i, x_i^A, r_i^A) \text{ s.t. } \begin{array}{l} i \in [t] \wedge \\ y_i^A = f(x_i^A) \wedge \\ \text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A) \end{array} \right\}. \quad (63)$$

Note that R uses $(1, x_1^A, r_1^A)$ as the witness when executing this **WIPoK**.

Comment: This step is repeated for [Trick 3](#) in [Sec. 6.1](#).

7. (**ExtCom-2** $\times n_7$.) C and R sequentially execute n_7 instances of **ExtCom** where C commits to m . Let τ_i^2 and decom_i^2 be the transcript and decommitment (privately obtained by C) of the i -th execution.

Comment: This step is inserted for [Trick 1](#) and repeated for [Trick 3](#) in [Sec. 6.1](#).

8. (**WIPoK-1-B** $\times n_8$.) R and C sequentially execute n_8 instances of **WIPoK** where R proves to C that he “knows” a pre-image of some y_i^B contained in Y^B , and y_i^B are consistent with com_i^B (as defined in [Step 4b](#)). Formally, R proves that $(Y^B, \{\text{com}_i^B\}_{i \in [n-t]}) \in \mathcal{L}_{f, \beta'}^{n-t}$, where

$$\mathcal{L}_{f, \beta'}^{n-t} := \left\{ (y_1^B, \dots, y_{n-t}^B), (\text{com}_1^B, \dots, \text{com}_{n-t}^B) \mid \exists (i, x_i^B, r_i^B) \text{ s.t. } \begin{array}{l} i \in [n-t] \wedge \\ y_i^B = f(x_i^B) \wedge \\ \text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B) \end{array} \right\}. \quad (64)$$

Note that R uses $(1, x_1^B, r_1^B)$ as the witness when executing this **WIPoK**.

Comment: This step is repeated for [Trick 3](#) in [Sec. 6.1](#).

9. (**ExtCom-3** $\times n_9$.) C and R sequentially execute n_9 instances of **ExtCom** where C commits to m . Let τ_i^2 and decom_i^2 be the transcript and decommitment (privately obtained by C) of the i -th execution.

Comment: This step is inserted for [Trick 1](#) and repeated for [Trick 3](#) in [Sec. 6.1](#).

10. (**WIPoK-2.**) C and R execute an instance of WIPoK where C proves to R that he “knows” the (same) message committed in com (defined in [Step 2](#)) and all commitments τ_i^c of ExtCom (in [Steps 5, 7](#) and [9](#)), or a pre-image of some y_i^A contained in Y^A (defined in [Step 4a](#)), or a pre-image of some y_i^B contained in Y^B (defined in [Step 4b](#)) or a preimage of either of V_0 or V_1 (defined in [Step 3a](#)). Formally, C proves that $(\text{com}, \{\tau_i^c\}_{(c,i) \in J}, Y^A, Y^B, V_0, V_1) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t \vee \mathcal{L}_f^{n-t} \vee \mathcal{L}_f^{\text{OR}}$, where $J = (\{1\} \times [n_5]) \cup (\{2\} \times [n_7]) \cup (\{3\} \times [n_9])$,

$$\mathcal{L}_\beta := \left\{ (\text{com}, \{\tau_i^c\}_{(c,i) \in J}) \left| \begin{array}{l} \exists (m, r, \{\text{decom}_i^c\}_{(c,i) \in J}) \text{ s.t.} \\ \text{com} = \text{Com}_\beta(m; r) \wedge \\ \forall (c, i) \in J, \text{Verify}_{\text{ExtCom}}(m, \tau_i^c, \text{decom}_i^c) = \top \end{array} \right. \right\} \quad (65)$$

$$\mathcal{L}_f^t := \{(y_1^A, \dots, y_t^A) \mid \exists (i, x_i^A) \text{ s.t. } i \in [t] \wedge y_i^A = f(x_i^A)\}, \quad (66)$$

$$\mathcal{L}_f^{n-t} := \{(y_1^B, \dots, y_{n-t}^B) \mid \exists (i, x_i^B) \text{ s.t. } i \in [n-t] \wedge y_i^B = f(x_i^B)\}, \quad (67)$$

and $\mathcal{L}_f^{\text{OR}}$ is defined in [Language \(62\)](#). Note that C uses the $(m, r, \{\text{decom}_i^c\}_{(c,i) \in J})$ defined in [Steps 2, 5, 7](#) and [9](#) as the witness when executing this WIPoK.

Comment: Compared with [Step 6](#) of [Prot. 2](#), we add $\mathcal{L}_f^{\text{OR}}$ as an additional OR part to the language to be proven for [Trick 2](#) in [Sec. 6.1](#). We also modify \mathcal{L}_β to prove that all commitments τ_i^c of ExtCom commit to the same message as com for [Trick 1](#) in [Sec. 6.1](#).

Decommit Stage: C sends (m, r) . R accepts if $\text{com} = \text{Com}_\beta(m; r)$, and rejects otherwise.

Choice of n_i . We recursively define n_5, \dots, n_9 so that n_i is larger than the total round complexity of Steps 1 to $i - 1$. We also require that n_5 is larger than the round complexities of WIPoK and ExtCom . It is easy to see that we can set n_i to be constant if both WIPoK and ExtCom have constant rounds. In particular, [Prot. 3](#) runs in constant rounds. We remark that the above choice of n_i is for simplifying the security analysis and we could significantly optimize the exact round complexity with more careful analysis. But we choose not to press the issue further.

Notation. For $c = 1, 2, 3$, we refer to the i -th execution of ExtCom in **ExtCom- c** as **ExtCom- c - i** . We also define **WIPoK-1-A- i** and **WIPoK-1-B- i** similarly.

Security. The completeness is easy to see. Statistical binding follows from that of Naor’s commitment. We prove computational hiding below.²⁸

Theorem 12. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{asynch}}$ in [Prot. 3](#) is computationally hiding.*

Proof. Let R^* be a non-uniform PPT adversary against the computational hiding property of [Prot. 3](#). We consider the following sequence of hybrids.

- Hybrid $H_1^{R^*}(\lambda, m)$: This is the real experiment that simulates the interaction between $C(m)$ and R^* . That is, it runs $\langle C(m), R^* \rangle_{\text{tg}}(1^\lambda)$ and outputs the final output of R^* . We have to prove

$$\left| \Pr \left[H_1^{R^*}(\lambda, m_0) = 1 \right] - \Pr \left[H_1^{R^*}(\lambda, m_1) = 1 \right] \right| = \text{negl}(\lambda). \quad (68)$$

for all m_0, m_1 .

²⁸ Though computational hiding follows from non-malleability, we prove it here because we rely on the computational hiding in the proof of non-malleability.

- Hybrid $H_2^{R^*}(\lambda, m)$: This is identical to $H_1^{R^*}(\lambda, m)$ except that it runs the emulation extractor for **WIPoK-1-A-1** to extract a witness (i, x_i^A, r_i^A) for $(Y^A, \{\text{com}_i^A\}_{i \in [t]}) \in \mathcal{L}_{f, \beta'}^t$ where Y^A and $\{\text{com}_i^A\}_{i \in [t]}$ are generated in [Step 4a](#). We note that this experiment extracts (i, x_i^A, r_i^A) but does not use it. By the PoK property (as per [Def. 12](#)) of **WIPoK-1-A-1**, we have

$$\left| \Pr \left[H_1^{R^*}(\lambda, m) = 1 \right] - \Pr \left[H_2^{R^*}(\lambda, m) = 1 \right] \right| = \text{negl}(\lambda) \quad (69)$$

for all m .

- Hybrid $H_3^{R^*}(\lambda, m)$: This is identical to $H_2^{R^*}(\lambda, m)$ except that it uses (i, x_i^A, r_i^A) extracted from **WIPoK-1-A-1** in **WIPoK-2**. By the WI property of **WIPoK-2**, we have

$$\left| \Pr \left[H_2^{R^*}(\lambda, m) = 1 \right] - \Pr \left[H_3^{R^*}(\lambda, m) = 1 \right] \right| = \text{negl}(\lambda) \quad (70)$$

for all m .

We observe that $H_3^{R^*}(\lambda, m_b)$ uses m only in the generation of com (in [Step 2](#)) and τ_i^c for $(c, i) \in J$ (in [Steps 5, 7](#) and [9](#)). Therefore, by the computational hiding property of Naor’s commitment and ExtCom, we have

$$\left| \Pr \left[H_3^{R^*}(\lambda, m_0) = 1 \right] - \Pr \left[H_3^{R^*}(\lambda, m_1) = 1 \right] \right| = \text{negl}(\lambda) \quad (71)$$

for all m_0, m_1 . Combining [Eq. \(69\)](#) to [\(71\)](#), we obtain [Eq. \(68\)](#).

This completes the proof of [Thm. 12](#). □

We establish non-malleability by [Thm. 13](#), whose proof will be presented in [Sec. 6.4](#).

Theorem 13. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{async}}$ in [Prot. 3](#) is non-malleable against asynchronous PPT adversaries with tag space $[n]$, with n being any polynomial on λ .*

6.3 Adaptive Schedule vs. Predetermined Schedule

Before proving the non-malleability of [Prot. 3](#), we prove a useful lemma that enable us to use *predetermined* schedules when proving non-malleability. The following [Lem. 11](#) should be attributed to [[BLS22](#), Lemma 5.1]: Although they only show this lemma for their specific construction, their proof technique is general enough. Our proof of [Lem. 11](#) is identical to theirs except for some notational changes customized to our application.

Lemma 11 ([\[BLS22, Lemma 5.1\]](#)). *Let \mathcal{S} be the set of all possible schedules of a MIM adversary against a commitment scheme $\langle C, R \rangle$. Let S_1, \dots, S_N be efficiently recognizable subsets of \mathcal{S} such that $\bigcup_{i \in [N]} S_i = \mathcal{S}$ for $N = \text{poly}(\lambda)$. If for every $i \in [N]$, $\langle C, R \rangle$ is non-malleable against non-uniform PPT MIM adversaries whose schedule belongs to S_i , then it is also non-malleable against arbitrary non-uniform PPT MIM adversaries.*

Proof. We can assume that S_1, \dots, S_N are disjoint w.l.o.g. (If they are not disjoint, we can define $S'_1 := S_1$, $S'_i := S_i \setminus \left(\bigcup_{j \in [i-1]} S_j \right)$ for $i > 1$ and consider disjoint sets S'_1, \dots, S'_N .)

Given an arbitrary non-uniform PPT MIM \mathcal{M} and non-uniform PPT distinguisher \mathcal{D} that break non-malleability for some values m_0, m_1 with advantage δ , we construct a new non-uniform PPT adversary with a schedule that belongs to S_i for a predetermined $i \in [n]$, which breaks the scheme with probability δ/N .

Consider an adversary \mathcal{M}' that first samples uniformly at random $i \xleftarrow{\$} [N]$. \mathcal{M}' then runs \mathcal{M} to complete the MIM attack. If the transcript generated in the execution belongs to S_i , \mathcal{M}' outputs whatever \mathcal{M} outputs, and otherwise outputs \perp .

Since every execution must belong to exactly one of S_1, \dots, S_N , \mathcal{M}' breaks non-malleability with probability exactly δ/N (with respect to the same distinguisher \mathcal{D} and m_0, m_1). Finally, by an averaging argument, we fix the choice of \mathcal{M}' for a schedule to be of some type S_i that maximizes \mathcal{D} 's distinguishing advantage. We obtain a corresponding MIM with a schedule that belongs to S_i for a predetermined $i \in [n]$ with at least the same advantage δ/N . □

6.4 Proving Non-Malleability (Proof of Thm. 13)

For proving Thm. 13, we consider the following “bad” schedules:

- Bad 1: The right message of Step 4 is sent before the left message of Step 2 is sent.
- Bad 2: The left message of Step 4 is sent before the right message of Step 2 is sent.
- Bad 3: The right **WIPoK-1-A** starts before the left message of Step 4 is sent.
- Bad 4: The right **WIPoK-1-B** starts before the left **WIPoK-1-A** finishes.
- Bad 5: The right **WIPoK-2** starts before the left **WIPoK-1-B** finishes.

In the above, when we say “**WIPoK-XX** starts”, this means that the first message of the first execution of **WIPoK** in **WIPoK-XX** is sent. Similarly, when we say “**WIPoK-XX** finishes”, this means that the final message of the final execution of **WIPoK** in **WIPoK-XX** is sent. We remark that the above bad cases are not disjoint.

We first prove non-malleability in the case where the schedule does not suffer from any of them (which we call a *good* schedule) in Sec. 6.5. Then, we prove it for the remaining schedules (which we call *bad* schedules) in Sec. 6.6. We remark that we can analyze them separately due to Lem. 11.

6.5 Non-Malleability for Good Schedules

Based on the observations made in Sec. 6.1, we can show the non-malleability analogously to the synchronous case if none of Bad 1-5 occurs.

Lemma 12. *The commitment scheme $(C, R)_{\text{tg}}^{\text{async}}$ in Prot. 3 is non-malleable against asynchronous PPT adversaries whose schedule does not suffer from any of Bad 1-5.*

Proof. Similarly to the proof of non-malleability of Prot. 2 in Sec. 5, we consider the three cases of $t = \tilde{t}$, $t < \tilde{t}$, and $\tilde{t} < t$ separately where t and \tilde{t} are the tags of the left and right sessions, respectively. We remark that we can assume that the reduction algorithm knows which of them happens from the beginning because of Lem. 11, even if the adversary adaptively determines the schedule.

Again, the case of $t = \tilde{t}$ is already ruled out by the definition of non-malleability. Next, we analyze the remaining two cases in Sec. 6.5.1 and 6.5.2, which will eventually complete the proof of Lem. 12. □

6.5.1 The Case of $t < \tilde{t}$.

For a PPT man-in-the-middle adversary \mathcal{M} with non-uniform advice z and a message m , we define the prefix generation algorithm $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ as follows similarly to that in Algo. 4.1: $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ runs the man-in-the-middle experiment $\text{mim}_{(C, R)_{\text{tg}}^{\text{async}}}^{\mathcal{M}}(\lambda, m, z)$ until \mathcal{M} receives the left message of

Step 2 and sends the right message of **Step 2**. Then, it outputs the prefix $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ where $\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R$ are the states of $\mathcal{M}, C,$ and R at this point, respectively; And τ (reps. $\tilde{\tau}$) is Naor's commitment in the right (resp. left) session generated in **Steps 1** and **2**. We denote by $\mathcal{A}(\text{st}_{\mathcal{M}})$ to mean the algorithm that works similarly to \mathcal{M} from the point where the prefix is generated. We remark that $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ does not reach to **Step 4** in either of the left or right session since we assume that neither of **Bad 1** or **Bad 2** occurs.

We define a machine \mathcal{G}_i similarly to that used in the proof of **Lem. 9** (recall it from **Fig. 6**) except that C commits to 0 instead of m in the left session. Formally, it works as follows:

Machine \mathcal{G}_i : For a prefix $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$, $\mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ runs the rest of $\text{mim}_{(C,R)_{\text{tg}}}^{\mathcal{M}}(\lambda, m, z)$ except for the following three differences:

1. C commits to 0 instead of m in the left **ExtCom- $\{1,2,3\}$** .
2. R uses (i, \tilde{x}_i^A) instead of $(1, \tilde{x}_1^A)$ in the right **WIPoK-1-A**.
3. It chooses i^* such that the left **WIPoK-1-A- i^*** does not interleave with the right **Step 1** to **5**, i.e., no message of the right **Step 1** to **5** is sent during the execution of the left **WIPoK-1-A- i^*** . Note that such i^* exists by the pigeonhole principle, and we call the smallest such i^* the *good* index. Then, instead of executing the left **WIPoK-1-A- i^*** honestly, it uses the witness-extended emulator \mathcal{WE} to extract a witness:
 - If the left committer accepts the left **WIPoK-1-A- i^*** and the extracted witness (j, x_j^A, r_j^A) is valid, \mathcal{G}_i uses (j, x_j^A) to finish the left **WIPoK-2** and outputs \mathcal{M} 's final state and the right receiver's decision bit b ;
 - If the left committer accepts the left **WIPoK-1-A- i^*** but the extracted witness is invalid, it aborts immediately and outputs (\perp, \perp) ;
 - If the left committer rejects the left **WIPoK-1-A- i^*** , it runs the rest of man-in-the-middle experiment to output \mathcal{M} 's final state and the right receiver's decision bit b . Note that it does not need to run the left **WIPoK-2** in this case since the left committer aborts after the left **WIPoK-1-A- i^*** .

We let $p_{\text{pref}}^{\text{Sim}}$ be the probability that $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ returns $b = \top$.

Remark 14. Strictly speaking, the above description of \mathcal{G}_i is not well-defined since the good index depends on the schedule, which is adaptively determined. However, since there are constant number of possibilities for the good index, we only have to prove non-malleability assuming that the good index is fixed at first by **Lem. 11**. A similar remark applies to later parts of the current proof. We will not repeat it as it should be clear from context.

By repeating similar arguments as in **Sec. 4.2** to **4.4**, we can see that it suffices to prove the following lemma, which is an analog of **Lem. 9**.

Lemma 13. *There exists an expected PPT machine \mathcal{K} such that for any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}}(\lambda, m, z)$ and any noticeable $\varepsilon(\lambda)$, the following holds:*

1. (**Syntax.**) \mathcal{K} takes as input $(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$ and makes oracle access to $\mathcal{A}(\text{st}_{\mathcal{M}})$. It output a value $\text{Val}_{\mathcal{K}} \in \{0, 1\}^{\ell(\lambda)} \cup \{\perp\}$ such that $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$ whenever $\text{Val}_{\mathcal{K}} \neq \perp$.
2. If $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then it holds that

$$\Pr \left[\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau}) : \text{Val}_{\mathcal{K}} \leftarrow \mathcal{K}^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau}) \right] \geq \frac{\varepsilon'(\lambda)}{t},$$

where $\varepsilon'(\lambda) := \frac{\varepsilon(\lambda)}{10t^2}$.

Proof of Lem. 13. We construct \mathcal{K}_i and \mathcal{K} similarly as in Sec. 5.3.1, which is in turn based on the technique in Sec. 4.5.

Machine \mathcal{K}_i ($i \in [\tilde{t}]$): On input $(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$, it behaves identically to $\mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except for the following difference. Machine $\mathcal{K}_i^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$ uses the witness-extended emulator \mathcal{WE} to finish the right **WIPoK-2**, instead of playing the role of the honest receiver. If it extracts a witness for $(\widetilde{\text{com}}, \{\tilde{\tau}_i^c\}_{(c,i) \in J}) \in \mathcal{L}_{\tilde{\beta}}$, then it sets Val to be the first component \tilde{m} of the witness (which is suppose to be the committed message in $\tilde{\tau}$) and outputs $\text{Val} = \tilde{m}$; Otherwise, it outputs \perp .

Extractor \mathcal{K} : On input $(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$, it samples uniformly at random an index $i \xleftarrow{\$} [\tilde{t}]$, executes $\mathcal{K}_i^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$, and outputs whatever $\mathcal{K}_i^{\mathcal{A}(\text{st}_{\mathcal{M}})}(1^\lambda, \text{st}_R, \tau, \tilde{\tau})$ outputs.

We can show that \mathcal{K} as described above satisfies Lem. 13 by similar analyses to those in Sec. 4.5 to 4.7 additionally relying on the computational hiding property of ExtCom along with the following observations:

1. By our choice of the good index i^* , the witness-extended emulator for the left **WIPoK-1-A- i^*** does not rewind the right **WIPoK-Trap**. Thus, \mathcal{K}_i does not extract a “trapdoor witness” (i.e., a witness for $(\tilde{V}_0, \tilde{V}_1) \in \mathcal{L}_f^{\text{OR}}$) except for a negligible probability by the one-wayness of f and the WI property of the right **WIPoK-Trap**.
2. Since we assume that Bad 4 does not occur, the witness-extended emulator for the left **WIPoK-1-A- i^*** does not rewind the right **WIPoK-1-B**. Thus, \mathcal{K}_i does not extract a witness for **Slot-B**, (i.e., a witness for $\tilde{Y}^B \in \mathcal{L}_f^{n-t}$) except for a negligible probability by the one-wayness of f and the WI property of the right **WIPoK-1-B**. This can be proven by a similar analysis to that in Sec. 5.3.1.
3. Since we assume that Bad 3 does not occur, we can rely on the WI property of the right **WIPoK-1-A** in a hybrid where we find the committed messages in $\{\text{com}_i^A\}_{i \in [t]}$ by brute-force. This is needed in a step corresponding to the proof of Inequality (36) in Sec. 4.6.
4. Since we assume that Bad 5 does not occur, the left **WIPoK-1-A** and the right **WIPoK-2** do not run simultaneously. Thus, there is no nesting of extractors in hybrids where we invoke extractors for both of them (e.g., \mathcal{K} and \mathcal{K}_i).

Remark 15. One might wonder how we rely on the computational hiding property of the left **ExtCom- $\{1,2,3\}$** even though the witness-extended emulator of the right **WIPoK-2** may rewind them. However, one can see by a close inspection of the proof that we only need to use computational hiding of the left **ExtCom- $\{1,2,3\}$** in hybrids that do not run the witness-extended emulator of the right **WIPoK-2**. This is similar to how we replace the witness used in the left **WIPoK-2** using the WI property in the proof for the one-sided synchronous case in Sec. 4.

This completes the proof of Lem. 13, which finishes the proof for the case of $t < \tilde{t}$. □

6.5.2 The Case of $t > \tilde{t}$

Similarly to the synchronous case (i.e., the proof of Thm. 11), the proof of this case can be done similarly to that of the case of $t < \tilde{t}$ except that we extract a witness from (the good index of) the left **WIPoK-1-B** instead of from **WIPoK-1-A**. In doing so, one has to be careful that Bad 4 is not symmetric for **Slot-A** and **Slot-B**. In particular, even if we assume that Bad 4 does not occur,

it may be still possible that the *left* **WIPoK-1-B** and the *right* **WIPoK-1-A** run concurrently. However, this is not an issue since we rewind the left **WIPoK-1-B** of the good index that does not interleave with *any* message of the right session before **WIPoK-1-B**. In particular, its rewinding does not rewind *any* slot of the right **WIPoK-1-A**. Except for the above remark, the proof of this case is almost identical to the proof for the case of $t < \tilde{t}$ with the roles of **Slot-A** and **Slot-B** being swapped.

6.6 Non-Malleability for Bad Schedules

Before analyzing bad cases, we show a lemma that is used many times in the analysis of bad cases.

Lemma 14. *Let C^* be a PPT cheating prover of the protocol in [Prot. 3](#) with advice z . For $(c, i) \in J$ (where $J = (\{1\} \times [n_5]) \cup (\{2\} \times [n_7]) \cup (\{3\} \times [n_9])$) as defined in [Step 10](#) of [Prot. 3](#)) we define experiments Exp and $\text{Exp}_{c,i}$ as follows.*

- Experiment $\text{Exp}(\lambda, z)$: *It simulates the interaction between C^* and R in the commit stage and outputs $(\text{OUT}_{C^*}, \text{val}_b(\tau))$ where OUT_{C^*} is the final output of C^* , b is R 's decision bit, τ is Naor's commitment generated in [Steps 1](#) and [2](#), and*

$$\text{val}_b(\tau) := \begin{cases} \text{val}(\tau) & b = \top \\ \perp & b = \perp \end{cases}.$$

- Experiment $\text{Exp}_{c,i}(\lambda, z)$: *This experiment works similarly to $\text{Exp}(\lambda, z)$ except that it runs the extractor of **ExtCom-c-i** to obtain the extracted message m^* and outputs $(\text{OUT}_{C^*}, \Gamma_b(m^*))$ where*

$$\Gamma_b(m^*) := \begin{cases} m^* & b = \top \\ \perp & b = \perp \end{cases}.$$

Then, for all $(c, i) \in J$, it holds that

$$\{\text{Exp}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*} \stackrel{s}{\approx} \{\text{Exp}_{c,i}(\lambda, z)\}_{\lambda \in \mathbb{N}, z \in \{0,1\}^*}.$$

Proof. First, we prove the following claim.

Claim 14.

$$\Pr[b = \top \wedge (\text{com}, \{\tau_i^c\}_{(c,i) \in J}) \notin \mathcal{L}_\beta] = \text{negl}(\lambda) \quad (72)$$

where the probability is taken over the execution of $\text{Exp}(\lambda, z)$.

Proof of Claim 14. We consider a modified experiment $\text{Exp}'(\lambda, z)$ that works similarly to $\text{Exp}(\lambda, z)$ except that it runs the witness-extended emulator for **WIPoK-2** instead of running it honestly. By the PoK property (as per [Def. 12](#)), it suffices to prove [Eq. \(72\)](#) holds in $\text{Exp}'(\lambda, z)$.

Let \tilde{w} be the extracted witness from **WIPoK-2**. By the PoK property (as per [Def. 12](#)), \tilde{w} is a valid witness for $(\text{com}, \{\tau_i^c\}_{(c,i) \in J}, Y^A, Y^B, V_0, V_1) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t \vee \mathcal{L}_f^{n-t} \vee \mathcal{L}_f^{\text{OR}}$ except for a negligible probability when $b = \top$.

We can show that \tilde{w} cannot be a witness of $Y^A \in \mathcal{L}_f^t$ except for a negligible probability by using the one-wayness of f and WI of **WIPoK-1-A**. Indeed, this can be seen as follows.

- In $\text{Exp}'(\lambda, z)$, the receiver uses $(1, x_1^A, r_1^A)$ as a witness for **WIPoK-1-A**. Therefore, if the probability that $\tilde{w} = (j, x_j^A)$ such that $y_j^A = f(x_j^A)$ for some $j \neq 1$ is non-negligible, it clearly breaks the one-wayness of f or the computational hiding property of **Com**.
- If the probability that $\tilde{w} = (1, x_1^A)$ such that $y_1^A = f(x_1^A)$ is non-negligible, by the WI property of **WIPoK-1-A**, a similar statement holds even if the receiver uses $(2, x_2^A, r_2^A)$ instead of $(1, x_1^A, r_1^A)$ as a witness for **WIPoK-1-A**. Then, by a similar argument to the above case, it also breaks the one-wayness of f or the computational hiding property of **Com**.

By similar arguments, \tilde{w} cannot be a witness of $Y^B \in \mathcal{L}_f^{n-t}$ or $(V_0, V_1) \in \mathcal{L}_f^{\text{OR}}$ except for a negligible probability. Thus, \tilde{w} is a witness for $(\text{com}, \{\tau_i^c\}_{(c,i) \in J}) \in \mathcal{L}_\beta$ except for a negligible probability when $b = \top$. This completes the proof of [Claim 14](#). □

[Claim 14](#) in particular means that we have $\text{val}(\tau) = \text{val}(\tau_c^i) \neq \perp$ whenever $b = \top$ except for a negligible probability. Then, the extractability of **ExtCom** immediately implies [Lem. 14](#). □

Next, we analyze the bad cases one by one.

6.6.1 Analysis of Bad 1

Lemma 15. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{async}}$ in [Prot. 3](#) is non-malleable against asynchronous PPT adversaries whose schedule satisfies **Bad 1**.*

Proof. The proof for this case is easy. When **Bad 1** occurs, the right commitment $\widetilde{\text{com}}$ (in [Step 2](#)) is sent before the left commitment **com** (in [Step 2](#)) is sent. Then, a straightforward reduction to the computational hiding of the left session works by extracting \tilde{m} from $\widetilde{\text{com}}$ by brute-force and treating it as non-uniform advice of an adversary against computational hiding. □

6.6.2 Analysis of Bad 2

Lemma 16. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{async}}$ in [Prot. 3](#) is non-malleable against asynchronous PPT adversaries whose schedule satisfies **Bad 2**.*

Proof. Recall that **Bad 2** means that the left message of [Step 4](#) is sent before the right message of [Step 2](#) is sent. We assume that \mathcal{M} receives the right message of [Step 1](#) at first w.l.o.g. In particular, no message in the right session is sent during the execution of the left **WIPoK-Trap**.

We consider the following sequence of hybrids.

- Hybrid $H_1^{\mathcal{M}}(\lambda, m, z)$: This hybrid is the real man-in-the-middle experiment $\text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m, z)$. The output of the experiment can be written as $(\text{OUT}_{\mathcal{M}}, \text{val}_b(\tilde{\tau}))$ where $\text{OUT}_{\mathcal{M}}$ is \mathcal{M} 's final output, b is the R 's decision bit, and $\tilde{\tau}$ is Naor's commitment generated in [Steps 1](#) and [2](#) in the right session, and

$$\text{val}_b(\tilde{\tau}) := \begin{cases} \text{val}(\tilde{\tau}) & b = \top \\ \perp & b = \perp \end{cases}.$$

- Hybrid $H_2^{\mathcal{M}}(\lambda, m, z)$: It is identical to $H_1^{\mathcal{M}}(\lambda, m, z)$ except that it runs the witness-extended emulator for the left **WIPoK-Trap** instead of running it honestly. Let v be the extracted witness. If C accepts the left **WIPoK-Trap** but v is not a valid witness for $(V_0, V_1) \in \mathcal{L}_f^{\text{OR}}$, it aborts with outputting (\perp, \perp) . Otherwise, it works similarly to $H_1^{\mathcal{M}}(\lambda, m, z)$. We remark that it does not use v though it extracts it.

- Hybrid $H_3^{\mathcal{M}}(\lambda, m, z)$: It is identical to $H_2^{\mathcal{M}}(\lambda, m, z)$ except that it uses v extracted from the left **WiPoK-Trap** as a witness in the left **WiPoK-2**.
- Hybrid $H_4^{\mathcal{M}}(\lambda, m, z)$: It is identical to $H_3^{\mathcal{M}}(\lambda, m, z)$ except that it commits to 0 instead of m in all slots of **ExtCom- $\{1,2,3\}$** .

We prove that each pair of neighboring hybrids is computationally indistinguishable.

Claim 15. *It holds that*

$$\{H_1^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{s}{\approx} \{H_2^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*}.$$

Proof of Claim 15. This follows from the PoK property (as per Def. 12) of the left **WiPoK-Trap** straightforwardly. □

Claim 16. *It holds that*

$$\{H_2^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{c}{\approx} \{H_3^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*}.$$

Proof of Claim 16. We introduce the following additional hybrids:

- Hybrid $\overline{H}_2^{\mathcal{M}}(\lambda, m, z)$: Choose the smallest $i^* \in [n_5]$ such that the right **ExtCom-1- i^*** does not interleave with the left **WiPoK-2**, i.e., no message of the left **WiPoK-2** is sent during the execution of the right **ExtCom-1- i^*** .²⁹ Such i^* exists by the pigeonhole principle since we assume that n_5 is larger than the round complexity of **WiPoK-2**. This experiment is identical to $H_2^{\mathcal{M}}(\lambda, m, z)$ except that it runs the extractor of the right **ExtCom-1- i^*** to extract \tilde{m}^* , and outputs $(\text{OUT}_{\mathcal{M}}, \Gamma_b(\tilde{m}^*))$ instead of $(\text{OUT}_{\mathcal{M}}, \text{val}_b(\tilde{\tau}))$ where

$$\Gamma_b(m^*) := \begin{cases} m^* & b = \top \\ \perp & b = \perp \end{cases}.$$

- Hybrid $\overline{H}_3^{\mathcal{M}}(\lambda, m, z)$: It is identical to $\overline{H}_2^{\mathcal{M}}(\lambda, m, z)$ except that it uses v extracted from the left **WiPoK-Trap** as a witness in the left **WiPoK-2**.

Note that both $\overline{H}_2^{\mathcal{M}}(\lambda, m, z)$ and $\overline{H}_3^{\mathcal{M}}(\lambda, m, z)$ are efficient since they no longer need to output $\text{val}_b(\tilde{\tau})$. Then, by the assumption that the right **ExtCom-1- i^*** does not interleave with the left **WiPoK-2**, the WI property of the left **WiPoK-2** gives

$$\{\overline{H}_2^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{c}{\approx} \{\overline{H}_3^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*}. \quad (73)$$

By the assumption that Bad 2 occurs and that \mathcal{M} receives the right message of Step 1 at first, hybrids $H_2^{\mathcal{M}}(\lambda, m, z)$ and $H_3^{\mathcal{M}}(\lambda, m, z)$ do not rewind the right session at all. Indeed, they only rewind the left **WiPoK-Trap**, but the right message of Step 1 is sent before the left **WiPoK-Trap** starts by our simplifying assumption and the rest of messages in the right session is sent after the left **WiPoK-Trap** finishes by the assumption that Bad 2 occurs. Thus, by considering the combination of \mathcal{M} and C as a single cheating committer, Lem. 14 gives

$$\{H_2^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{c}{\approx} \{\overline{H}_2^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \quad (74)$$

²⁹ Rmk. 14 applies here.

and

$$\{H_3^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{c}{\approx} \{\overline{H}_3^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*}. \quad (75)$$

By combining Eq. (73) to (75), we complete the proof of Claim 16. \square

Claim 17. *It holds that*

$$\{H_3^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{c}{\approx} \{H_4^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*}.$$

Proof of Claim 17. This follows from a similar argument to that in the proof of Claim 16. However, we remark that we cannot prove it at once since the *total* round complexity of **ExtCom- $\{1,2,3\}$** is larger than n_5 , in which case we cannot choose i^* such that the right **ExtCom- $1-i^*$** does not interleave with any message of the left **ExtCom- $\{1,2,3\}$** . To resolve this issue, we replace the committed message m with 0 for each execution of ExtCom in **ExtCom- $\{1,2,3\}$** one by one. Since we assume that n_5 is larger than the round complexity of ExtCom, a similar argument to the proof of Claim 16 works by using the computational indistinguishability of ExtCom. \square

Claim 18. *It holds that*

$$\{H_4^{\mathcal{M}}(\lambda, m_0, z)\}_{\lambda, m_0, m_1, z} \stackrel{c}{\approx} \{H_4^{\mathcal{M}}(\lambda, m_1, z)\}_{\lambda, m_0, m_1, z},$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$, $(m_0, m_1) \in \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{\ell(\lambda)}$, and $z \in \{0,1\}^*$.

Proof of Claim 18. This immediately follows from the computational hiding property of Naor's commitment noting that the left message m is only used for generating Naor's commitment in Step 2 in $H_4^{\mathcal{M}}(\lambda, m, z)$. \square

By combining Claims 15 to 18, we obtain

$$\{H_1^{\mathcal{M}}(\lambda, m_0, z)\}_{\lambda, m_0, m_1, z} \stackrel{c}{\approx} \{H_1^{\mathcal{M}}(\lambda, m_1, z)\}_{\lambda, m_0, m_1, z}$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$, $(m_0, m_1) \in \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{\ell(\lambda)}$, and $z \in \{0,1\}^*$. This completes the proof of Lem. 16. \square

6.6.3 Analysis of Bad 3

Lemma 17. *The commitment scheme $\langle C, R \rangle_{\text{tg}}^{\text{async}}$ in Prot. 3 is non-malleable against asynchronous PPT adversaries whose schedule satisfies Bad 3.*

Proof. Recall that Bad 3 means that the right **WIPoK-1-A** starts before the left message of Step 4 is sent.

We consider the following sequence of hybrids.

- Hybrid $H_1^{\mathcal{M}}(\lambda, m, z)$: This hybrid is the real man-in-the-middle experiment $\text{mim}_{\langle C, R \rangle}^{\mathcal{M}}(\lambda, m, z)$. The output of the experiment can be written as $(\text{OUT}_{\mathcal{M}}, \text{val}_b(\tilde{\tau}))$ where $\text{OUT}_{\mathcal{M}}$ is \mathcal{M} 's final output, b is the R 's decision bit, and $\tilde{\tau}$ is Naor's commitment generated in Steps 1 and 2 in the right session, and

$$\text{val}_b(\tilde{\tau}) := \begin{cases} \text{val}(\tilde{\tau}) & b = \top \\ \perp & b = \perp \end{cases}.$$

- Hybrid $H_2^{\mathcal{M}}(\lambda, m, z)$: Choose the smallest $i^* \in [n_5]$ such that the right **ExtCom-1- i^*** does not interleave with Steps 1 to 4 in the left, i.e., no message of the left **WiPoK-2** is sent during the execution of the right **ExtCom-1- i^*** .³⁰ Such i^* exists by the pigeonhole principle since we assume that n_5 is larger than the total round complexity of Steps 1 to 4. This experiment is identical to $H_1^{\mathcal{M}}(\lambda, m, z)$ except that it runs the extractor of the right **ExtCom-1- i^*** to extract \tilde{m}^* , and outputs $(\text{OUT}_{\mathcal{M}}, \Gamma_b(\tilde{m}^*))$ instead of $(\text{OUT}_{\mathcal{M}}, \text{val}_b(\tilde{\tau}))$ where

$$\Gamma_b(m^*) := \begin{cases} m^* & b = \top \\ \perp & b = \perp \end{cases}.$$

Claim 19. *It holds that*

$$\{H_1^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*} \stackrel{s}{\approx} \{H_2^{\mathcal{M}}(\lambda, m, z)\}_{\lambda \in \mathbb{N}, m \in \{0,1\}^{\ell(\lambda)}, z \in \{0,1\}^*}.$$

Proof of Claim 19. This immediately follows from [Lem. 14](#). □

Claim 20. *It holds that*

$$\{H_2^{\mathcal{M}}(\lambda, m_0, z)\}_{\lambda, m_0, m_1, z} \stackrel{s}{\approx} \{H_2^{\mathcal{M}}(\lambda, m_1, z)\}_{\lambda, m_0, m_1, z}$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$, $(m_0, m_1) \in \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{\ell(\lambda)}$, and $z \in \{0,1\}^*$.

Proof. By the choice of i^* and the assumption that **Bad 3** occurs, $H_2^{\mathcal{M}}(\lambda, m, z)$ does not rewind the left session at all. Therefore, by considering the combination of \mathcal{M} and R as a single cheating receiver, [Claim 20](#) follows from the computational hiding property of [Prot. 3](#). □

By combining [Claims 19](#) and [20](#), we obtain

$$\{H_1^{\mathcal{M}}(\lambda, m_0, z)\}_{\lambda, m_0, m_1, z} \stackrel{c}{\approx} \{H_1^{\mathcal{M}}(\lambda, m_1, z)\}_{\lambda, m_0, m_1, z}$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$, $m_0, m_1 \in \{0,1\}^{\ell(\lambda)}$, and $z \in \{0,1\}^*$. This completes the proof of [Lem. 17](#). □

6.6.4 Analysis of **Bad 4**

Lemma 18. *The commitment scheme $(C, R)_{\text{tg}}^{\text{async}}$ in [Prot. 3](#) is non-malleable against asynchronous PPT adversaries whose schedule satisfies **Bad 4**.*

Proof. Recall that **Bad 4** means that the right **WiPoK-1-B** starts before the left **WiPoK-1-A** finishes. [Lem. 18](#) can be proven similarly to [Lem. 17](#) except that we extract \tilde{m} from the right **ExtCom-2- i^*** that does not interleave with Steps 1 to 6 in the left. □

³⁰ [Rmk. 14](#) applies here.

6.6.5 Analysis of Bad 5

Lemma 19. *The commitment scheme $(C, R)_{\text{tg}}^{\text{async}}$ in [Prot. 3](#) is non-malleable against asynchronous PPT adversaries whose schedule satisfies Bad 5.*

Proof. Recall that Bad 5 means that the right **WIPoK-2** starts before the left **WIPoK-1-B** finishes. [Lem. 19](#) can be proven similarly to [Lem. 17](#) except that we extract \tilde{m} from the right **ExtCom-3- i^*** that does not interleave with Steps 1 to 10 in the left. □

By combining [Lem. 12](#) and [15](#) to [19](#) along with [Lem. 11](#), we complete the proof of [Thm. 13](#).

7 Extract-and-Simulate Lemma

In this section, we show a generalized version of the extract-and-simulate lemma given in [[CCLY22](#), Lemma 4].

Lemma 20 (Extract-and-Simulate Lemma). *Let \mathcal{G} be a QPT algorithm that takes the security parameter 1^λ , an error parameter $1^{\gamma^{-1}}$, a quantum state ρ , and a classical string z as input, and outputs $b \in \{\top, \perp\}$ and a quantum state ρ_{out} .*

Suppose that there exists a QPT algorithm \mathcal{K} (referred to as the simulation-less extractor) that takes as input the security parameter 1^λ , an error parameter $1^{\gamma^{-1}}$, a quantum state ρ , and a classical string z , and outputs $s \in \{0, 1\}^{\text{poly}(\lambda)} \cup \{\perp\}$ satisfying the following w.r.t. some sequence of classical strings $\{s_z^\}_{z \in \{0,1\}^*}$.*

1. *For any $\lambda, \gamma, \rho_\lambda$, and z_λ , the output s of $\mathcal{K}(1^\lambda, 1^{\gamma^{-1}}, \rho_\lambda, z_\lambda)$ is equal to $s_{z_\lambda}^*$ whenever $s \neq \perp$.*
2. *For any noticeable function $\gamma(\lambda)$, there exists a noticeable function $\delta(\lambda)$ that is efficiently computable from $\gamma(\lambda)$ and satisfies the following. For any sequence $\{\rho_\lambda, z_\lambda\}_{\lambda \in \mathbb{N}}$ of polynomial-size quantum states and classical strings, if*

$$\Pr \left[b = \top : (b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, \rho_\lambda, z_\lambda) \right] \geq \gamma(\lambda),$$

then

$$\Pr \left[\mathcal{K}(1^\lambda, 1^{\gamma^{-1}}, \rho_\lambda, z_\lambda) = s_{z_\lambda}^* \right] \geq \delta(\lambda).$$

Then, there exists a QPT algorithm \mathcal{SE} such that for any noticeable function $\varepsilon = \varepsilon(\lambda)$, there exists a noticeable function $\gamma = \gamma(\lambda) \leq \varepsilon(\lambda)$ that is efficiently computable from ε and satisfies the following: For any sequence $\{\rho_\lambda, z_\lambda\}_{\lambda \in \mathbb{N}}$ of polynomial-size quantum states and classical strings,

$$\{\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, \rho_\lambda, z_\lambda)\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_\varepsilon \{(\rho_{\text{out}}, \Gamma_b(s_{z_\lambda}^*)) : (b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, \rho_\lambda, z_\lambda)\}_{\lambda \in \mathbb{N}},$$

$$\text{where } \Gamma_b(s_{z_\lambda}^*) := \begin{cases} s_{z_\lambda}^* & \text{if } b = \top \\ \perp & \text{otherwise} \end{cases}.$$

In [[CCLY22](#)], the authors essentially proved a special case of the above lemma where the simulation-less extractor is fixed to Unruh's rewinding extractor [[Unr12](#)]. We observe that essentially the same proof works for general simulation-less extractors.³¹ We provide the full proof of [Lem. 20](#) in [Appx. A](#) for completeness.

³¹ Indeed, this observation is also mentioned in the technical overview of [[CCLY22](#)].

8 Post-Quantum Non-Malleable Commitments

In this section, we construct a full-fledged post-quantum non-malleable commitment scheme that supports an exponential-size tag space and is secure against QPT asynchronous MIM adversaries.

8.1 Small-Tag, One-sided, Synchronous, Post-Quantum Setting

First, we show that the [Prot. 1](#) given in [Sec. 4](#) is post-quantumly secure if we instantiate it with post-quantum building-blocks. Specifically, let $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$ be [Prot. 1](#) instantiated with post-quantum Naor’s commitment Com (which is constructed based on post-quantum OWFs), post-quantum injective OWF f , and a post-quantum witness-distinguishable arguments of knowledge with ε -close emulation as per [Def. 14](#) (see [Prot. 5](#) in [Appx. B](#) for the full description of $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$). Then, we can prove the following theorem:

Theorem 21. *The commitment scheme $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$ is non-malleable against one-sided synchronous QPT adversaries with tag space $[n]$, with n being any polynomial on λ .*

We remark that the injectivity assumption for f can be removed in exactly the same way as in the classical setting (as explained in [Sec. 4.8](#)).

The proof of [Thm. 21](#) is very similar to that of its classical counterpart ([Thm. 2](#)) except for one step as explained below. Thus, we only give a proof sketch below, and defer the full proof to [Appx. B](#).

Proof of [Thm. 21](#) (sketch). We observe that all steps of the proof of [Thm. 2](#), except for the proof of [Lem. 7](#) (in particular, [Lem. 9](#) in [Sec. 4.4](#)), can be translated into the post-quantum setting easily with the following remarks. The only difference in those steps (except for some superficial differences like PPT vs QPT) is that we only have WIAoK with ε -close emulation (as per [Def. 14](#)) instead of WIPoK (with negligibly-close emulation). First, the AoK property instead of the PoK property suffices because we run witness-extended emulators only for (possibly non-uniform) efficient cheating provers. Second, the ε -close emulation (in contrast to negligibly-close emulation) also suffices, because we can take a sufficiently small noticeable error parameter such that it is much smaller than the MIM adversary’s advantage. With these modifications, the whole proof will go through.

Thus, the only remaining issue is how to prove the post-quantum version of [Lem. 7](#) (using a post-quantum version of [Lem. 9](#)). Roughly, we have to construct a simulation-extractor \mathcal{SE} that extracts the right committed message $\text{val}(\tilde{\tau})$ while simulating \mathcal{M} ’s final state from a simulation-less extractor \mathcal{K} , which extracts $\text{val}(\tilde{\tau})$ with a noticeable probability *without simulating \mathcal{M} ’s final state*. This task is exactly what can be achieved by our new extract-and-simulate lemma ([Lem. 20](#)). Indeed, the requirements for \mathcal{K} in [Lem. 9](#) (more precisely, its post-quantum version [Lem. 31](#)) exactly correspond to the assumptions of [Lem. 20](#). Therefore, this step can be completed by reducing it to [Lem. 20](#). (We believe that readers with good understanding of the classical case can directly go to *Proof of [Lem. 29](#)* in [Page 86](#) after checking the statements of [Lem. 29](#) and [31](#) and relevant definitions to see how this step exactly works.)

□

8.2 Small-Tag, Asynchronous, Post-Quantum Setting

Next, we explain how to make $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$ (shown in [Prot. 5](#)) secure against two-sided asynchronous adversaries. We emphasize that this subsection still focuses on the small-tag setting. This step is essentially the same as its classical counterpart in [Sec. 5](#) and [6](#) except for that we rely on post-quantum building blocks.

We present our construction in [Prot. 4](#). This protocol makes use of the following building blocks.

- A post-quantum OWF f ;
- Naor’s commitment Com that is implemented with a post-quantum OWF;
- A post-quantum witness-indistinguishable argument of knowledge with ε -close emulation WIAoK (as per [Def. 14](#));
- A post-quantum ε -simulatable extractable commitment ExtCom (as per [Def. 9](#)). We denote by $\text{Verify}_{\text{ExtCom}}$ the verification algorithm of ExtCom in the decommit stage.

Similarly to [Sec. 6](#), we recursively define n_5, \dots, n_9 so that n_i is greater than the total round complexity of Steps 1 to $i - 1$. We also require that n_5 is greater than the round complexities of WIAoK and ExtCom . It is easy to see that we can set n_i to a constant if both WIPoK and ExtCom have constant rounds. In particular, [Prot. 4](#) runs in constant rounds.

Protocol 4: Post-Quantum Small-Tag Asynchronous NMCom $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$
<p>The tag space is defined to be $[n]$ where n is a polynomial on λ. Let $t \in [n]$ be the tag for the following interaction. Let m be the message to be committed to.</p> <p>Commit Stage:</p> <ol style="list-style-type: none"> 1. Receiver R samples and sends the first message β for Naor’s commitment. 2. Committer C commits to m using the second message of Naor’s commitment. C also sends a β' that will be used as the first Naor’s message for R to generate a commitment in next step. Formally, C samples r and β' and sends the tuple $(\text{com} = \text{Com}_\beta(m; r), \beta')$. 3. R and C do the following: <ol style="list-style-type: none"> (a) (TrapGen.) R computes $V_0 = f(v_0)$ and $V_1 = f(v_1)$ with $v_0, v_1 \xleftarrow{\\$} \{0, 1\}^\lambda$ and sends (V_0, V_1) to C (b) (WIAoK-Trap.) R and C execute an instance of WIAoK where R proves to C that he “knows” a pre-image of V_0 or V_1. Formally, R proves that $(V_0, V_1) \in \mathcal{L}_f^{\text{OR}}$, where $\mathcal{L}_f^{\text{OR}} := \{(V_0, V_1) \mid \exists v \text{ s.t. } f(v) = V_0 \vee f(v) = V_1\}. \quad (76)$ <p style="margin-left: 20px;">Note that R uses v_0 as the witness when executing this WIAoK.</p> 4. R performs the following computation: <ol style="list-style-type: none"> (a) R computes $\{y_i^A = f(x_i^A)\}_{i \in [t]}$ with $x_i^A \xleftarrow{\\$} \{0, 1\}^\lambda$ for each $i \in [t]$, and sets $Y^A := (y_1^A, \dots, y_t^A)$. R also computes $\{\text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A)\}_{i \in [t]}$, where r_i^A is a random string for each $i \in [t]$. (b) R computes $\{y_i^B = f(x_i^B)\}_{i \in [n-t]}$ with $x_i^B \xleftarrow{\\$} \{0, 1\}^\lambda$ for each $i \in [n-t]$, and sets $Y^B := (y_1^B, \dots, y_{n-t}^B)$. R also computes $\{\text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B)\}_{i \in [n-t]}$, where r_i^B is a random string for each $i \in [n-t]$. <p style="margin-left: 20px;">R sends the tuple $(Y^A, \{\text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A)\}_{i \in [t]}, Y^B, \{\text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B)\}_{i \in [n-t]})$.</p> 5. (ExtCom-1 $\times n_5$.) C and R sequentially execute n_5 instances of ExtCom where C commits to m. Let τ_i^1 and decom_i^1 be the transcript and decommitment (privately obtained by C) of the i-th execution. 6. (WIAoK-1-A $\times n_6$.) R and C sequentially execute n_6 instances of WIAoK where R proves to C that he “knows” a pre-image of some y_i^A contained in Y^A, and y_i^A are consistent with com_i^A

(as defined in [Step 4a](#)). Formally, R proves that $(Y^A, \{\text{com}_i^A\}_{i \in [t]}) \in \mathcal{L}_{f, \beta'}^t$, where

$$\mathcal{L}_{f, \beta'}^t := \left\{ (y_1^A, \dots, y_t^A), (\text{com}_1^A, \dots, \text{com}_t^A) \mid \exists (i, x_i^A, r_i^A) \text{ s.t. } \begin{array}{l} i \in [t] \wedge \\ y_i^A = f(x_i^A) \wedge \\ \text{com}_i^A = \text{Com}_{\beta'}(x_i^A; r_i^A) \end{array} \right\}. \quad (77)$$

Note that R uses $(1, x_1^A, r_1^A)$ as the witness when executing this WIAoK.

7. (**ExtCom-2** $\times n_7$.) C and R sequentially execute n_7 instances of ExtCom where C commits to m . Let τ_i^2 and decom_i^2 be the transcript and decommitment (privately obtained by C) of the i -th execution.
8. (**WIAoK-1-B** $\times n_8$) R and C sequentially execute n_8 instances of WIAoK where R proves to C that he “knows” a pre-image of some y_i^B contained in Y^B , and y_i^B are consistent with com_i^B (as defined in [Step 4b](#)). Formally, R proves that $(Y^B, \{\text{com}_i^B\}_{i \in [n-t]}) \in \mathcal{L}_{f, \beta'}^{n-t}$, where

$$\mathcal{L}_{f, \beta'}^{n-t} := \left\{ (y_1^B, \dots, y_{n-t}^B), (\text{com}_1^B, \dots, \text{com}_{n-t}^B) \mid \exists (i, x_i^B, r_i^B) \text{ s.t. } \begin{array}{l} i \in [n-t] \wedge \\ y_i^B = f(x_i^B) \wedge \\ \text{com}_i^B = \text{Com}_{\beta'}(x_i^B; r_i^B) \end{array} \right\}. \quad (78)$$

Note that R uses $(1, x_1^B, r_1^B)$ as the witness when executing this WIPoK.

9. (**ExtCom-3** $\times n_9$.) C and R sequentially execute n_9 instances of ExtCom where C commits to m . Let τ_i^2 and decom_i^2 be the transcript and decommitment (privately obtained by C) of the i -th execution.
10. (**WIAoK-2**.) C and R execute an instance of WIAoK where C proves to R that he “knows” the (same) message committed in com (defined in [Step 2](#)) and all commitments τ_i^c of ExtCom (in [Steps 5, 7](#) and [9](#)), or a pre-image of some y_i^A contained in Y^A (defined in [Step 4a](#)), or a pre-image of some y_i^B contained in Y^B (defined in [Step 4b](#)) or a preimage of either of V_0 or V_1 (defined in [Step 3a](#)). Formally, C proves that $(\text{com}, \{\tau_i^c\}_{(c,i) \in J}, Y^A, Y^B, V_0, V_1) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t \vee \mathcal{L}_f^{n-t} \vee \mathcal{L}_f^{\text{OR}}$, where $J = (\{1\} \times [n_5]) \cup (\{2\} \times [n_7]) \cup (\{3\} \times [n_9])$,

$$\mathcal{L}_\beta := \left\{ (\text{com}, \{\tau_i^c\}_{(c,i) \in J}) \mid \begin{array}{l} \exists (m, r, \{\text{decom}_i^c\}_{(c,i) \in J}) \text{ s.t.} \\ \text{com} = \text{Com}_\beta(m; r) \wedge \\ \forall (c, i) \in J, \text{Verify}_{\text{ExtCom}}(m, \tau_i^c, \text{decom}_i^c) = \top \end{array} \right\} \quad (79)$$

$$\mathcal{L}_f^t := \{(y_1^A, \dots, y_t^A) \mid \exists (i, x_i^A) \text{ s.t. } i \in [t] \wedge y_i^A = f(x_i^A)\}, \quad (80)$$

$$\mathcal{L}_f^{n-t} := \{(y_1^B, \dots, y_{n-t}^B) \mid \exists (i, x_i^B) \text{ s.t. } i \in [n-t] \wedge y_i^B = f(x_i^B)\}, \quad (81)$$

and $\mathcal{L}_f^{\text{OR}}$ is defined in [Language \(76\)](#). Note that C uses the $(m, r, \{\text{decom}_i^c\}_{(c,i) \in J})$ defined in [Steps 2, 5, 7](#) and [9](#) as the witness when executing this WIAoK.

Decommit Stage: C sends (m, r) . R accepts if $\text{com} = \text{Com}_\beta(m; r)$, and rejects otherwise.

We remark that [Prot. 4](#) is identical to [Prot. 3](#) except that we require post-quantum security for building blocks and use WIAoK instead of WIPoK.

Security. Completeness is straightforward. Statistical binding follows from that of Naor’s commitment.

For computational hiding and non-malleability, we observe that the proofs against classical adversaries in [Sec. 6](#) is already quantum-friendly, and they can be translated into proofs against quantum adversaries in a straightforward manner. The only difference is that we now have ε -close simulation for WIAoK and ExtCom though the classical counterparts have negligible simulation errors. Thus, whenever we invoke extractors of WIPoK or ExtCom in the original proof in the classical setting, a noticeable error occurs in the post-quantum setting. However, this is not a problem since the error can be an *arbitrarily small* noticeable function. Thus, by repeating the same proofs as those in the classical case, we can show that QPT MIM adversary's advantage is smaller than *any* noticeable function in the security parameter. This means that the advantage is negligible, which is nothing but non-malleability in the standard sense. Thus, we obtain the following theorems.

Theorem 22. *The commitment scheme $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ in [Prot. 4](#) is computationally hiding against QPT adversaries.*

Theorem 23. *The commitment scheme $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ in [Prot. 4](#) is non-malleable against asynchronous QPT adversaries with tag space $[n]$, with n being any polynomial on λ .*

8.3 Tag Amplification

Finally, we explain how to amplify the tag space of [Prot. 4](#). For this step, we rely on the following variant of a theorem shown in [\[BLS22\]](#).

Theorem 24 ([\[BLS22, Section 6\]](#)). *Assume the existence of*

- a post-quantum, k_1 -round, ε -simulatable extractable commitment (as per [Def. 9](#)), and
- a post-quantum, k_2 -round, witness-indistinguishable argument.

Let $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ be a post-quantumly non-malleable commitment scheme, where the subscript tg denotes the length of the maximum tag it can support. Assume $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ satisfies the following conditions:

1. $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ supports $\text{tg} \in \{3, 4, \dots, O(\log \lambda)\}$ bit tags, and
2. $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ is also a post-quantum $(k_1 + k_2)$ -robust ε -simulatable extractable commitment scheme (as per [Def. 10](#)).

Then, there exists a post-quantum non-malleable commitment scheme $\langle C, R \rangle_{\text{TG}, \text{PQ}}^{\text{async}}$ which supports tags of $\text{TG} = 2^{\text{tg}-1}$ bits. Moreover, $\langle C, R \rangle_{\text{TG}, \text{PQ}}^{\text{async}}$ has $k_1 + k_2 + O(1)$ rounds.

Remark 16. In the original version of the above theorem in [\[BLS22\]](#), they require negl -simulatability instead of ε -simulatability for the extractable commitment and the non-malleable commitment schemes (in [Condition 2](#)). However, we can see that the ε -simulatable version suffices due to a similar reason as explained in the end of the previous subsection. That is, if we use extractors with a noticeable error in their proof, we can show that the adversary's advantage is at most a noticeable function in λ . Since the noticeable error can be chosen to be arbitrarily small, this implies non-malleability in the standard sense.

Given the above [Thm. 24](#), we only need to prove that $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ in [Prot. 4](#) has $(k_1 + k_2)$ -robust ε -simulatable extractability where k_1 and k_2 are the round complexities of ExtCom and WIAoK used in the construction of $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$.³² We show this below.

³² For applying [Thm. 24](#), we only need to require computational soundness for the witness indistinguishable argument rather than the witness-extended emulation property. But since the latter is stronger than the former, we can simply instantiate it with any WIAoK that satisfies [Def. 14](#).

Theorem 25. *The commitment scheme $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{async}}$ in [Prot. 4](#) satisfies post-quantum $(k_1 + k_2)$ -robust ε -simulatable extractability (as per [Def. 10](#)) where k_1 and k_2 are round complexities of ExtCom and WIAoK.*

Proof. (sketch.) We first remark that we have $n_9 > k_1 + k_2$ since we assume that n_9 is larger than the total round complexity of Steps 1 to 8, which is clearly larger than $k_1 + k_2$ since those steps include repetitions of ExtCom and WIAoK. Then, by the pigeonhole principle, for each session with O^∞ , there is i^* such that **ExtCom-3- i^*** does not interleave with the interaction with the external $(k_1 + k_2)$ -round machine B in the definition of r -robust ε -simulatable extractability (in [Def. 10](#)). Then, we can extract the committed message by running the extractor for **ExtCom-3- i^*** without rewinding B .³³ We can show that the extracted message is the correct committed message in **com** (in [Step 1](#)) whenever the receiver accepts except for a negligible probability similarly to [Lem. 14](#). This enables us to simulate the oracle O^∞ without rewinding B . □

Finally, by combining [Thm. 23](#) to [25](#), we obtain the following theorem.

Theorem 26. *Assuming the existence of post-quantum OWFs, there exist constant-round post-quantum non-malleable commitments supporting tag space $[\Omega(2^\lambda)]$.*

9 Application: Quantum-Secure MPC in Constant Rounds from QLWE

MPC for Classical Functionalities. Agarwal et al. [[ABG⁺21](#)] constructed a constant-round post-quantum MPC protocol in the plain model assuming the super-polynomial hardness of QLWE and a QLWE-based circular security assumption. The only reason why they rely on the super-polynomial hardness of QLWE is for their construction of a constant-round post-quantum non-malleable commitment scheme.³⁴ Since we have constructed constant-round post-quantum non-malleable commitments from post-quantum OWFs ([Thm. 26](#)), we can weaken the assumption to the polynomial hardness of QLWE. Thus, we obtain the following theorem.

Theorem 27. *Assuming the polynomial hardness of QLWE and a QLWE-based circular security assumption (as in [[ABG⁺21](#)]), there exist constant-round constructions of post-quantum MPC for classical functionalities in the plain model.*

Recall that “post-quantum MPC” means an MPC protocol secure against QPT adversaries where honest parties only need to perform classical computation. This is the first construction of constant-round post-quantum MPC from polynomial hardness assumptions in the plain model.

Remark 17 (Synchronous Security Suffices.). For the construction of post-quantum MPC, we only need post-quantum non-malleable commitments secure against *synchronous* adversaries. It is simpler to obtain such non-malleable commitments than those against *asynchronous* adversaries. For example, this can be done by applying the tag amplification of [[ABG⁺21](#), Sec. 7.3] to our small-tag, synchronous construction (i.e., [Prot. 2](#) instantiated with post-quantum building blocks).

³³ We can assume that such i^* is known in advance w.l.o.g. by [Lem. 11](#).

³⁴ Actually, what they need is the so-called *many-to-one* non-malleable commitments. But it is known that (one-to-one) non-malleability as defined in [Def. 7](#) is equivalent to many-to-one non-malleability (even in the post-quantum setting, as noted in [[ABG⁺21](#), Lemma 7.3]).

MPC for Quantum Functionalities. Recently, Bartusek et al. [BCKM21a] built a constant-round quantum-secure MPC for quantum functionalities *in the CRS model*³⁵, based on the hardness of QLWE. It is easy to see that the post-quantum MPC from Thm. 27 can be used to instantiate the CRS required by the [BCKM21a] protocol. Since the protocol from Thm. 27 is also constant round, this leads to the first constant-round quantum-secure MPC for quantum functionalities from polynomial hardness assumptions *without any trusted setup*.

Theorem 28. *Assuming (polynomial) QLWE and the QLWE-based circular security assumption (as in [ABG⁺21]), there exists a constant-round construction of quantum-secure MPC for quantum functionalities in the plain model.*

References

- Aar05. Scott Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1(1):1–28, 2005. [16](#)
- ABG⁺21. Amit Agarwal, James Bartusek, Vipul Goyal, Dakshita Khurana, and Giulio Malavolta. Post-quantum multi-party computation. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 435–464. Springer, Heidelberg, October 2021. [2](#), [4](#), [68](#), [69](#)
- Bar01. Boaz Barak. How to go beyond the black-box simulation barrier. In *42nd FOCS*, pages 106–115. IEEE Computer Society Press, October 2001. [1](#)
- Bar02. Boaz Barak. Constant-round coin-tossing with a man in the middle or realizing the shared random string model. In *43rd FOCS*, pages 345–355. IEEE Computer Society Press, November 2002. [1](#)
- BCKM21a. James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. On the round complexity of secure quantum computation. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 406–435, Virtual Event, August 2021. Springer, Heidelberg. [4](#), [69](#)
- BCKM21b. James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 467–496, Virtual Event, August 2021. Springer, Heidelberg. [2](#)
- BDH⁺17. Brandon Broadnax, Nico Döttling, Gunnar Hartung, Jörn Müller-Quade, and Matthias Nagel. Concurrently composable security with shielded super-polynomial simulators. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 351–381. Springer, Heidelberg, April / May 2017. [1](#)
- BGJ⁺18. Saikrishna Badrinarayanan, Vipul Goyal, Abhishek Jain, Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Promise zero knowledge and its applications to round optimal MPC. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 459–487. Springer, Heidelberg, August 2018. [1](#)
- BL18. Nir Bitansky and Huijia Lin. One-message zero knowledge and non-malleable commitments. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part I*, volume 11239 of *LNCS*, pages 209–234. Springer, Heidelberg, November 2018. [1](#)
- BLS22. Nir Bitansky, Huijia Lin, and Omri Shmueli. Non-malleable commitments against quantum attacks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 519–550. Springer, 2022. [2](#), [3](#), [5](#), [17](#), [18](#), [21](#), [22](#), [54](#), [67](#)
- BS20. Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *52nd ACM STOC*, pages 269–279. ACM Press, June 2020. [2](#), [72](#)
- BZ13. Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 361–379. Springer, Heidelberg, August 2013. [2](#)
- Can00. Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, January 2000. [1](#)

³⁵ More accurately, [BCKM21a] presented their construction in the OT-hybrid model. But it is known that the OT functionality can be instantiated using the post-quantum straight-line simulatable construction in the CRS model from [PVW08], which is also constant-round and based on the hardness of QLWE.

- Can01. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001. [1](#)
- CCG⁺19. Arka Rai Choudhuri, Michele Ciampi, Vipul Goyal, Abhishek Jain, and Rafail Ostrovsky. Round optimal secure multiparty computation from minimal assumptions. Cryptology ePrint Archive, Report 2019/216, 2019. <https://eprint.iacr.org/2019/216>. [1](#)
- CCLY22. Nai-Hui Chia, Kai-Min Chung, Xiao Liang, and Takashi Yamakawa. Post-quantum simulatable extraction with minimal assumptions: Black-box and constant-round. In *Annual International Cryptology Conference*. Springer, 2022. [2](#), [3](#), [13](#), [14](#), [15](#), [16](#), [19](#), [21](#), [22](#), [24](#), [63](#), [72](#), [73](#)
- CCY21. Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. A black-box approach to post-quantum zero-knowledge in constant rounds. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 315–345, Virtual Event, August 2021. Springer, Heidelberg. [16](#), [73](#), [76](#), [77](#), [78](#)
- 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. ACM Press, May 2002. [1](#)
- CLP20. Rohit Chatterjee, Xiao Liang, and Omkant Pandey. Improved black-box constructions of composable secure computation. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *ICALP 2020*, volume 168 of *LIPICs*, pages 28:1–28:20. Schloss Dagstuhl, July 2020. [1](#)
- COSV17. Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 127–157. Springer, Heidelberg, August 2017. [1](#), [3](#)
- DDN91. Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography (extended abstract). In *23rd ACM STOC*, pages 542–552. ACM Press, May 1991. [1](#), [3](#)
- FP96. Christopher A Fuchs and Asher Peres. Quantum-state disturbance versus information gain: Uncertainty relations for quantum information. *Physical Review A*, 53(4):2038, 1996. [1](#)
- GLOV12. Vipul Goyal, Chen-Kuei Lee, Rafail Ostrovsky, and Ivan Visconti. Constructing non-malleable commitments: A black-box approach. In *53rd FOCS*, pages 51–60. IEEE Computer Society Press, October 2012. [1](#), [3](#)
- 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. Springer, Heidelberg, May 2016. [1](#)
- Gol01. Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001. [22](#), [23](#)
- 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. ACM Press, June 2011. [1](#), [3](#)
- GPR16. Vipul Goyal, Omkant Pandey, and Silas Richelson. Textbook non-malleable commitments. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 1128–1141. ACM Press, June 2016. [1](#), [3](#), [20](#)
- GR19. Vipul Goyal and Silas Richelson. Non-malleable commitments using Goldreich-Levin list decoding. In David Zuckerman, editor, *60th FOCS*, pages 686–699. IEEE Computer Society Press, November 2019. [1](#), [3](#)
- GRRV14. Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *55th FOCS*, pages 41–50. IEEE Computer Society Press, October 2014. [1](#), [3](#)
- HSS11. Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 411–428. Springer, Heidelberg, August 2011. [14](#)
- Khu17. Dakshita Khurana. Round optimal concurrent non-malleability from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part II*, volume 10678 of *LNCS*, pages 139–171. Springer, Heidelberg, November 2017. [1](#)
- Khu21. Dakshita Khurana. Non-interactive distributional indistinguishability (NIDI) and non-malleable commitments. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 186–215. Springer, Heidelberg, October 2021. [1](#)
- 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. Springer, Heidelberg, May 2003. [1](#)
- KS17. Dakshita Khurana and Amit Sahai. How to achieve non-malleability in one or two rounds. In Chris Umans, editor, *58th FOCS*, pages 564–575. IEEE Computer Society Press, October 2017. [1](#), [4](#)

- Lin03. Yehuda Lindell. Parallel coin-tossing and constant-round secure two-party computation. *Journal of Cryptology*, 16(3):143–184, June 2003. 7, 22, 23
- LMS21. Alex Lombardi, Fermi Ma, and Nicholas Spooner. Post-quantum zero knowledge, revisited (or: How to do quantum rewinding undetectably). Cryptology ePrint Archive, Report 2021/1543, 2021. <https://ia.cr/2021/1543>. 3, 14, 21, 24
- LN11. Carolin Lunemann and Jesper Buus Nielsen. Fully simulatable quantum-secure coin-flipping and applications. In Abderrahmane Nitaï and David Pointcheval, editors, *AFRICACRYPT 11*, volume 6737 of *LNCS*, pages 21–40. Springer, Heidelberg, July 2011. 14
- LP09. Huijia Lin and Rafael Pass. Non-malleability amplification. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 189–198. ACM Press, May / June 2009. 1
- 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. ACM Press, June 2011. 1, 3
- LPS17. Huijia Lin, Rafael Pass, and Pratik Soni. Two-round and non-interactive concurrent non-malleable commitments from time-lock puzzles. In Chris Umans, editor, *58th FOCS*, pages 576–587. IEEE Computer Society Press, October 2017. 1
- LPV08. Huijia Lin, Rafael Pass, and Muthuramakrishnan Venkitasubramaniam. Concurrent non-malleable commitments from any one-way function. In Ran Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 571–588. Springer, Heidelberg, March 2008. 1, 20
- MPR06. Silvio Micali, Rafael Pass, and Alon Rosen. Input-indistinguishable computation. In *47th FOCS*, pages 367–378. IEEE Computer Society Press, October 2006. 1
- Pas03. Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 160–176. Springer, Heidelberg, May 2003. 1
- Pas13. Rafael Pass. Unprovable security of perfect NIZK and non-interactive non-malleable commitments. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 334–354. Springer, Heidelberg, March 2013. 1
- PPV08. Omkant Pandey, Rafael Pass, and Vinod Vaikuntanathan. Adaptive one-way functions and applications. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 57–74. Springer, Heidelberg, August 2008. 1
- PR05. Rafael Pass and Alon Rosen. New and improved constructions of non-malleable cryptographic protocols. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 533–542. ACM Press, May 2005. 1, 46
- PS04. Manoj Prabhakaran and Amit Sahai. New notions of security: Achieving universal composability without trusted setup. In László Babai, editor, *36th ACM STOC*, pages 242–251. ACM Press, June 2004. 1
- 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. Springer, Heidelberg, August 2008. 69
- PW09. Rafael Pass and Hoeteck Wee. Black-box constructions of two-party protocols from one-way functions. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 403–418. Springer, Heidelberg, March 2009. 13, 21
- 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. Springer, Heidelberg, May / June 2010. 4
- Unr12. Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, Heidelberg, April 2012. 63
- Wat06. John Watrous. Zero-knowledge against quantum attacks. In Jon M. Kleinberg, editor, *38th ACM STOC*, pages 296–305. ACM Press, May 2006. 2
- Wat09. John Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009. 72
- Wee10. Hoeteck Wee. Black-box, round-efficient secure computation via non-malleability amplification. In *51st FOCS*, pages 531–540. IEEE Computer Society Press, October 2010. 1, 5
- WZ82. William K Wootters and Wojciech H Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982. 1
- Zha12. Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012. 2

Appendix

A Proof of Extract-and-Simulate Lemma (Lem. 20)

We prove Lem. 20. The proof follows similar techniques as in the proof of [CCLY22, Lemma 4].

A.1 Preparation

We prepare several lemmas that will be used in the proof of Lem. 20.

Watrous' Rewinding Lemma. The following is Watrous' rewinding lemma [Wat09] in the form of [BS20, Lemma 2.1].

Lemma 21 (Watrous' Rewinding Lemma [Wat09]). *There is a quantum algorithm R that gets as input the following:*

- A quantum circuit Q that takes n -input qubits in register Inp and outputs a classical bit b (in a register outside Inp) and an m -qubit output.
- An n -qubit state ρ in register Inp .
- A number $T \in \mathbb{N}$ in unary.

$R(1^T, Q, \rho)$ executes in time $T \cdot |Q|$ and outputs a distribution over m -qubit states $D_\rho := R(1^T, Q, \rho)$ with the following guarantees.

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

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

then for every n -qubit state ρ ,

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

Lemma 22 ([CCLY22, Lemma 8]). *Let $|\phi_b\rangle = |\phi_{b,0}\rangle + |\phi_{b,1}\rangle$ be a normalized quantum state in a Hilbert space \mathcal{H} . Let F be a quantum algorithm that takes a state in \mathcal{H} as input and outputs a quantum state (not necessarily in \mathcal{H}) or a classical failure symbol **Fail**. Suppose that we have*

$$\Pr \left[F \left(\frac{|\phi_{b,0}\rangle \langle \phi_{b,0}|}{\|\phi_{b,0}\|^2} \right) = \text{Fail} \right] \geq 1 - \gamma$$

for $b \in \{0, 1\}$ and $\|\phi_{1,1}\rangle - \phi_{1,0}\rangle\| \leq \delta$. Then for any distinguisher D , it holds that

$$|\Pr[D(F(|\phi_0\rangle \langle \phi_0|)) = 1] - \Pr[D(F(|\phi_1\rangle \langle \phi_1|)) = 1]| \leq (12\gamma^{1/2} + 2\delta)^{1/2}.$$

Lemma 23 ([CCY21, Lemma 3.2]). *Let $|\psi\rangle_{\mathbf{X}}$ be a (not necessarily normalized) state over register \mathbf{X} and U be a unitary over registers $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$. Suppose that a measurement of register \mathbf{Z} of $U|\psi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y},\mathbf{Z}}$ results in a deterministic value except for probability ν , i.e., there is z^* such that*

$$\|(I - |z^*\rangle\langle z^*|)_{\mathbf{Z}} U|\psi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y},\mathbf{Z}}\|^2 \leq \nu.$$

If we let $R := (|0\rangle\langle 0|)_{\mathbf{Y},\mathbf{Z}} U^\dagger (|z^*\rangle\langle z^*|)_{\mathbf{Z}} U$, then we have

$$\| |\psi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y},\mathbf{Z}} - R|\psi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y},\mathbf{Z}} \| \leq \sqrt{\nu}.$$

Lemma 24 (A variant of [CCLY22, Lemma 10]). *Let Π be a projection over a Hilbert space $\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}}$. For any noticeable function $\delta = \delta(\lambda)$, there exists an orthogonal decomposition $(S_{<\delta}, S_{\geq\delta})$ of $\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}}$ that satisfies the following:*

1. ($S_{<\delta}$ and $S_{\geq\delta}$ are invariant under Π and $(|0\rangle\langle 0|)_{\mathbf{Y}}$.) For any $|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in S_{<\delta}$, we have

$$\Pi|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in S_{<\delta}, \quad (I_{\mathbf{X}} \otimes (|0\rangle\langle 0|)_{\mathbf{Y}})|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in S_{<\delta}.$$

Similarly, for any $|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in S_{\geq\delta}$, we have

$$\Pi|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in S_{\geq\delta}, \quad (I_{\mathbf{X}} \otimes (|0\rangle\langle 0|)_{\mathbf{Y}})|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in S_{\geq\delta}.$$

2. (Π succeeds with probability $< \delta$ and $\geq \delta$ in $S_{<\delta}$ and $S_{\geq\delta}$.) For any quantum state $|\phi\rangle_{\mathbf{X}} \in \mathcal{H}_{\mathbf{X}}$ s.t. $|\phi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y}} \in S_{<\delta}$ we have

$$\|\Pi|\phi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y}}\|^2 < \delta.$$

Similarly, for any quantum state $|\phi\rangle_{\mathbf{X}} \in \mathcal{H}_{\mathbf{X}}$ s.t. $|\phi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y}} \in S_{\geq\delta}$ we have

$$\|\Pi|\phi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y}}\|^2 \geq \delta.$$

3. (**Unitary for amplification.**) For any $T \in \mathbb{N}$, there exists a unitary $U_{\text{amp},T}$ over $\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}} \otimes \mathcal{H}_{\mathbf{B}} \otimes \mathcal{H}_{\text{Anc}}$ where \mathbf{B} is a register to store a qubit and Anc is a register to store ancillary qubits with the following properties:

(a) (**Mapped onto $\Pi(I_{\mathbf{X}} \otimes (|0\rangle\langle 0|)_{\mathbf{Y}})$ when \mathbf{B} contains 1.**) For any quantum state $|\psi\rangle_{\mathbf{X},\mathbf{Y}} \in \mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}}$, we can write

$$|1\rangle\langle 1|_{\mathbf{B}} U_{\text{amp},T} |\psi\rangle_{\mathbf{X},\mathbf{Y}} |0\rangle_{\mathbf{B},\text{Anc}} = \sum_{\text{anc}} |\psi'_{\text{anc}}\rangle_{\mathbf{X},\mathbf{Y}} |1\rangle_{\mathbf{B}} |\text{anc}\rangle_{\text{Anc}}$$

by using sub-normalized states $|\psi'_{\text{anc}}\rangle_{\mathbf{X},\mathbf{Y}}$ that are in the span of $\Pi(I_{\mathbf{X}} \otimes (|0\rangle\langle 0|)_{\mathbf{Y}})$.

(b) (**Amplification of success probability in $S_{\geq\delta}$.**) For any noticeable function $\nu = \nu(\lambda)$, there is $T = \text{poly}(\lambda)$ such that for any quantum state $|\phi\rangle_{\mathbf{X}} \in \mathcal{H}_{\mathbf{X}}$ s.t. $|\phi\rangle_{\mathbf{X}}|0\rangle_{\mathbf{Y}} \in S_{\geq\delta}$, we have

$$\| |1\rangle\langle 1|_{\mathbf{B}} U_{\text{amp},T} |\phi\rangle_{\mathbf{X}} |0\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{B},\text{Anc}} \|^2 \geq 1 - \nu.$$

(c) ($S_{<\delta}$ and $S_{\geq\delta}$ are invariant under $U_{\text{amp},T}$.) For any quantum state $|\psi_{<\delta}\rangle_{\mathbf{X},\mathbf{Y}} \in S_{<\delta}$ and any b, anc , we can write

$$U_{\text{amp},T} |\psi_{<\delta}\rangle_{\mathbf{X},\mathbf{Y}} |b, \text{anc}\rangle_{\mathbf{B},\text{Anc}} = \sum_{b', \text{anc}'} |\psi'_{<\delta, b', \text{anc}'}\rangle_{\mathbf{X},\mathbf{Y}} |b', \text{anc}'\rangle_{\mathbf{B},\text{Anc}}$$

by using sub-normalized states $|\psi'_{<\delta,b',anc'}\rangle_{\mathbf{X},\mathbf{Y}} \in S_{<\delta}$.

Similarly, for any quantum state $|\psi_{\geq\delta}\rangle_{\mathbf{X},\mathbf{Y}} \in S_{\geq\delta}$ and any b, anc , we can write

$$U_{\text{amp},T} |\psi_{\geq\delta}\rangle_{\mathbf{X},\mathbf{Y}} |b, anc\rangle_{\mathbf{B},\mathbf{Anc}} = \sum_{b',anc'} |\psi'_{\geq\delta,b',anc'}\rangle_{\mathbf{X},\mathbf{Y}} |b', anc'\rangle_{\mathbf{B},\mathbf{Anc}}$$

by using sub-normalized states $|\psi'_{\geq\delta,b',anc'}\rangle_{\mathbf{X},\mathbf{Y}} \in S_{\geq\delta}$.

4. (**Efficient Implementation of $U_{\text{amp},T}$** .) There exists a QPT algorithm **Amp** (whose description is independent of Π) that takes as input 1^T , a description of quantum circuit that perform a measurement $(\Pi, I_{\mathbf{X},\mathbf{Y}} - \Pi)$, and a state $|\psi\rangle_{\mathbf{X},\mathbf{Y},\mathbf{B},\mathbf{Anc}}$, and outputs $U_{\text{amp},T} |\psi\rangle_{\mathbf{X},\mathbf{Y},\mathbf{B},\mathbf{Anc}}$. Moreover, **Amp** uses the measurement circuit for only implementing an oracle that apply unitary to write a measurement result in a designated register in **Anc**, and it acts on **X** only through the oracle access.

A.2 Proving Lem. 20

Since any mixed state can be seen as a distribution over pure states, we assume \mathcal{G} 's input ρ_λ is a pure state and denote it by $|\psi\rangle$, omitting the dependence on λ . Similarly, we simply write z to mean z_λ and s^* to mean $s_{z_\lambda}^*$ for simplicity.

For a noticeable function $\gamma(\lambda)$ and a quantum state $|\psi\rangle$, we define an experiment $\text{Exp}(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$ as follows

$\text{Exp}(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$: Run $(b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$.

- If $b = \top$, the experiment outputs (ρ_{out}, s^*) . We remark that this step may not be done efficiently since we do not assume that s^* can be computed from z efficiently.
- If $b = \perp$, the experiment outputs $(\rho_{\text{out}}, \perp)$.

What we have to do is to construct a QPT \mathcal{SE} such that for any polynomial-size $|\psi\rangle$ and noticeable $\varepsilon = \varepsilon(\lambda)$,

$$\{\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_\varepsilon \{(\text{Exp}(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z))\}_{\lambda \in \mathbb{N}}$$

for some noticeable γ .

Let $\text{Exp}_\perp(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$ and $\text{Exp}_\top(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$ be the same as $\text{Exp}(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$ except that they output a failure symbol **Fail** in the cases of $b = \top$ and $b = \perp$, respectively. That is, they work as follows.

$\text{Exp}_\perp(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$: Run $(b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$.

- If $b = \top$, the experiment outputs **Fail**.
- If $b = \perp$, the experiment outputs $(\rho_{\text{out}}, \perp)$.

$\text{Exp}_\top(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$: Run $(b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)$.

- If $b = \top$, the experiment outputs (ρ_{out}, s^*) .
- If $b = \perp$, the experiment outputs **Fail**.

First, we give simulation extractors for each of these experiments.

Lemma 25 (Extract-and-Simulate for the Case of $b = \perp$). *For any noticeable γ that is efficiently computable from ε , there is a QPT algorithm \mathcal{SE}_\perp such that for any polynomial-size quantum state $|\psi\rangle$ and noticeable ε ,*

$$\{\mathcal{SE}_\perp(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}} \equiv \{\text{Exp}_\perp(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}}.$$

Proof of Lem. 25. Since Exp_\perp is efficient (because it never outputs s^*), \mathcal{SE}_\perp just needs to run Exp_\perp . \square

Lemma 26 (Extract-and-Simulate for the Case of $b = \top$). *There is a QPT algorithm \mathcal{SE}_\top such that for any polynomial-size quantum state $|\psi\rangle$ and noticeable ε ,*

$$\{\mathcal{SE}_\top(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_\varepsilon \{\text{Exp}_\top(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}},$$

where $\gamma := \left(\frac{\varepsilon}{5}\right)^4$.

Proof of Lem. 26. Let δ be a noticeable function that satisfies Item 2 of Lem. 20 for $\gamma = \left(\frac{\varepsilon}{5}\right)^4$.

We apply Lem. 24 with respect to a projection corresponding to the success of \mathcal{K} . Let $U_{\mathcal{K}}$ be the unitary that represents $\mathcal{K}(1^\lambda, 1^{\gamma^{-1}}, \cdot, z)$. More precisely, we define $U_{\mathcal{K}}$ over registers the input register **Inp**, working register **W**, and output register **Out** so that $\mathcal{K}(1^\lambda, 1^{\gamma^{-1}}, \cdot, z)$ can be described as follows: $\mathcal{K}(1^\lambda, 1^{\gamma^{-1}}, \cdot, z)$: It takes a quantum state $|\psi\rangle$ in the register **Inp** and initializes registers **W** and **Out** to be $|0\rangle_{\mathbf{W}, \mathbf{Out}}$. Then it applies the unitary $U_{\mathcal{K}}$, measures the register **Out** in the standard basis to obtain s , and outputs s .

We define a projection Π over $(\text{Inp}, \mathbf{W}, \mathbf{Out})$ as

$$\Pi := U_{\mathcal{K}}^\dagger \left(\sum_{s \neq \text{Fail}} |s\rangle \langle s| \right)_{\mathbf{Out}} U_{\mathcal{K}}. \quad (82)$$

Then the following claim immediately follows from the assumption about \mathcal{K} (Item 1 of Lem. 20).

Claim 29. *Given any state in the span of $\Pi(I_{\text{Inp}} \otimes (|0\rangle \langle 0|)_{\mathbf{W}, \mathbf{Out}})$, if we apply $U_{\mathcal{K}}$ and then measure register **Out**, then the measurement outcome is always s^**

We apply Lem. 24 for the above Π where $\mathcal{H}_{\mathbf{X}} := \mathcal{H}_{\text{Inp}}$, $\mathcal{H}_{\mathbf{Y}} := \mathcal{H}_{\mathbf{W}} \otimes \mathcal{H}_{\mathbf{Out}}$, and $T = \text{poly}(\lambda)$ is chosen in such a way that Item 3b of Lem. 24 holds for $\nu := \left(\frac{\varepsilon}{2}\right)^4$. Then we have a decomposition $(S_{<\delta}, S_{\geq\delta})$ of $\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}}$ and a unitary $U_{\text{amp}, T}$ over $\mathcal{H}_{\mathbf{X}} \otimes \mathcal{H}_{\mathbf{Y}} \otimes \mathcal{H}_{\mathbf{B}} \otimes \mathcal{H}_{\mathbf{Anc}}$ that satisfy the requirements in Lem. 24. We denote by **Other** to mean the registers **W**, **Out**, **B**, and **Anc** for brevity. We construct the extractor \mathcal{SE}_\top for Lem. 26 as follows:

$\mathcal{SE}_\top(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$:

1. Set $|\psi\rangle$ in register **Inp** and initialize register **Other** to be $|0\rangle$.
2. Apply $U_{\text{amp}, T}$ by using the algorithm **Amp** in Item 4 of Lem. 24.
3. Measure register **B** and let b be the outcome. If $b = 0$, output **Fail** and immediately halt. Otherwise, proceed to the next step.
4. Apply $U_{\mathcal{K}}$, measure register **Out** to obtain an outcome s_{Ext} , and apply $U_{\mathcal{K}}^\dagger$.
5. Apply $U_{\text{amp}, T}^\dagger$ by using the algorithm **Amp** in Item 4 of Lem. 24.
6. Measure register **Other**. If the outcome is not the all 0's string, output **Fail** and immediately halt. Otherwise, let $|\psi'\rangle$ be the state in register **Inp** at this point, and proceed to the next step.

7. Run $(b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, |\psi'\rangle, z)$

– If $b = \top$, output $(\rho_{\text{out}}, s_{\text{Ext}})$.

– If $b = \perp$, output Fail.

We can easily see the following claim:

Claim 30. *Whenever Step 4 of \mathcal{SE}_\top is invoked, s_{Ext} obtained in the step is always equal to s^* . Moreover, the step does not change the state in registers **Inp** and **Other**, that is, the states before and after the step are identical.*

Proof of Claim 30. Whenever Step 4 is invoked, the bit b obtained in Step 3 is equal to 1. In this case, by Item 3a of Lem. 24, the state in registers **Inp**, **W**, and **Out** is in the span of $\Pi(I_{\text{Inp}} \otimes (|0\rangle\langle 0|)_{\text{W,Out}})$. Then, Claim 29 implies that s_{Ext} is always equal to s^* . Then the measurement of **Out** does not collapse the state and thus the step does not change the state. \square

The rest of the proof is similar to that of [CCY21, Claim 4.5]. Let R be an operator defined as follows:

$$R := (|0\rangle\langle 0|)_{\text{Other}} U_{\text{amp},T}^\dagger (|1\rangle\langle 1|)_{\text{B}} U_{\text{amp},T}.$$

Let $\Pi_{<\delta}$ and $\Pi_{\geq\delta}$ be projections onto $S_{<\delta}$ and $S_{\geq\delta}$, respectively. To apply Lemma 22, we define states $|\phi_0\rangle = |\phi_{0,0}\rangle + |\phi_{0,1}\rangle$ and $|\phi_1\rangle = |\phi_{1,0}\rangle + |\phi_{1,1}\rangle$ over $(\mathbf{D}, \text{Inp}, \text{Other})$ where \mathbf{D} is an additional one-qubit register as follows:

$$\begin{aligned} |\phi_0\rangle &:= |1\rangle_{\mathbf{D}} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}}, \\ |\phi_{0,0}\rangle &:= |1\rangle_{\mathbf{D}} \Pi_{<\delta} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}}, \\ |\phi_{0,1}\rangle &:= |1\rangle_{\mathbf{D}} \Pi_{\geq\delta} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}}, \\ |\phi_1\rangle &:= |1\rangle_{\mathbf{D}} R |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} + \alpha |0\rangle_{\mathbf{D}} |0\rangle_{\text{Inp}} |0\rangle_{\text{Other}}, \\ |\phi_{1,0}\rangle &:= |1\rangle_{\mathbf{D}} R \Pi_{<\delta} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} + \alpha |0\rangle_{\mathbf{D}} |0\rangle_{\text{Inp}} |0\rangle_{\text{Other}}, \\ |\phi_{1,1}\rangle &:= |1\rangle_{\mathbf{D}} R \Pi_{\geq\delta} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} \end{aligned}$$

for $\alpha := \sqrt{1 - \|R|\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}}\|^2}$ (so that $|\phi_1\rangle$ is a normalized state). Let F be a quantum algorithm that works as follows:

$F(|\phi\rangle_{\mathbf{D}, \text{Inp}, \text{Other}})$: It measures \mathbf{D} , and outputs Fail if the outcome is 0. Otherwise, for the state $|\psi_{\text{inp}}\rangle$ in register **Inp**, it runs $(b, \rho_{\text{out}}) \leftarrow \mathcal{G}(1^\lambda, 1^{\gamma^{-1}}, |\psi_{\text{inp}}\rangle, z)$ and outputs $(\rho_{\text{out}}, s_{\text{Ext}})$ if $b = \top$ and otherwise outputs Fail.

It is easy to see that

$$\text{Exp}_\top(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z) \equiv F(|\phi_0\rangle\langle\phi_0|).$$

Moreover, by the definition of $\text{Ext}_{\text{Sim}, \text{na}}$ and Claim 30, we can see that

$$\mathcal{SE}_\top(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z) \equiv F(|\phi_1\rangle\langle\phi_1|).$$

Thus, it suffices to prove that the distinguishing advantage between $F(|\phi_0\rangle\langle\phi_0|)$ and $F(|\phi_1\rangle\langle\phi_1|)$ is at most ε . To apply Lem. 22, we prove the following claim.

Claim 31. *The following hold:*

1. $\Pr \left[F \left(\frac{|\phi_{b,0}\rangle\langle\phi_{b,0}|}{\|\phi_{b,0}\|^2} \right) = \text{Fail} \right] \geq 1 - \gamma$ for $b \in \{0, 1\}$.
2. $\| |\phi_{1,1}\rangle - |\phi_{0,1}\rangle \| \leq \nu^{1/2}$.

Proof of Claim 31.

First item. We can write $\Pi_{<\delta} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} = |\psi_{<\delta}\rangle_{\text{Inp}} |0\rangle_{\text{Other}}$. By $|\psi_{<\delta}\rangle_{\text{Inp}} |0\rangle_{\text{Other}} \in S_{<\delta}$, [Item 2](#) of [Lem. 24](#), and [Item 1](#) of [Lem. 20](#), we have

$$\Pr \left[\mathcal{K} \left(1^\lambda, 1^{\gamma^{-1}} \frac{|\psi_{<\delta}\rangle}{\|\psi_{<\delta}\|}, z \right) = s_z^* \right] < \delta.$$

By the contraposition of [Item 2](#) of [Lem. 20](#), we have

$$\Pr \left[b = \top : (b, \rho_{\text{out}}) \leftarrow \mathcal{G} \left(1^\lambda, 1^{\gamma^{-1}} \frac{|\psi_{<\delta}\rangle}{\|\psi_{<\delta}\|}, z \right) \right] < \gamma.$$

Thus, we have

$$\Pr \left[F \left(\frac{|\phi_{0,0}\rangle\langle\phi_{0,0}|}{\|\phi_{0,0}\|^2} \right) \neq \text{Fail} \right] = \Pr \left[b = \top : (b, \rho_{\text{out}}) \leftarrow \mathcal{G} \left(1^\lambda, 1^{\gamma^{-1}} \frac{|\psi_{<\delta}\rangle}{\|\psi_{<\delta}\|}, z \right) \right] \leq \gamma.$$

This completes the proof of the first item for the case of $b = 0$. The case of $b = 1$ can be proven similarly noting that $R\Pi_{<\delta} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} \in S_{<\delta}$ by [Item 1](#) and [item 3c](#) of [Lemma 24](#).

Second Item. By [Item 3b](#) of [Lem. 24](#), we have

$$\| (|1\rangle\langle 1|)_{\mathbf{B}} U_{\text{amp}, T} \Pi_{\geq t} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} \|^2 \leq \nu.$$

Thus, [Lem. 23](#) implies

$$\| \Pi_{\geq t} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} - R\Pi_{\geq t} |\psi\rangle_{\text{Inp}} |0\rangle_{\text{Other}} \| \leq \nu^{1/2}.$$

□

By [Lem. 22](#) and [Claim 31](#) the distinguishing advantage between $F(|\phi_0\rangle\langle\phi_0|)$ and $F(|\phi_1\rangle\langle\phi_1|)$ is at most

$$\left(12\gamma^{1/2} + 2\nu^{1/2} \right)^{1/2}. \tag{83}$$

By using $\gamma = \left(\frac{\varepsilon}{5}\right)^4$ and $\nu = \left(\frac{\varepsilon}{2}\right)^4$, we can see that this is at most ε . This completes the proof of [Lem. 26](#).

We complete the proof of [Lem. 26](#).

□

Given [Lem. 25](#) and [Lem. 26](#), the rest of the proof of [Lem. 20](#) is very similar to the corresponding part of the ε -zero-knowledge property of the protocols in [[CCY21](#)]. We give the full proof for completeness.

Let $\mathcal{SE}_{\text{comb}}$ be an algorithm that works as follows:

$\mathcal{SE}_{\text{comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$:

1. Set $\varepsilon' := \frac{\varepsilon^2}{4 \log^4(\lambda)}$.

2. Choose $\text{mode} \leftarrow \{\top, \perp\}$.
3. Run and output $\mathcal{SE}_{\text{mode}}(1^\lambda, 1^{\varepsilon'^{-1}}, |\psi\rangle)$.

Lemma 27 ($\mathcal{SE}_{\text{comb}}$ Simulates Exp with Probability almost 1/2). *Let $p_{\text{comb}}^{\text{suc}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$ be the probability that $\mathcal{SE}_{\text{comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$ does not return Fail, and let*

$$D_{\text{ext,comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$$

be a conditional distribution of $\mathcal{SE}_{\text{comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$, conditioned on that it does not return Fail. Then we have

$$\left| p_{\text{comb}}^{\text{suc}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z) - 1/2 \right| \leq \varepsilon'/2 + \text{negl}(\lambda). \quad (84)$$

Moreover, we have

$$\{D_{\text{ext,comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_{4\varepsilon'} \{\text{Exp}(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}}, \quad (85)$$

where $\gamma := \left(\frac{\varepsilon'}{5}\right)^4$.

Proof. (sketch.) The intuition behind this proof is as follows. By [Lem. 25](#) and [26](#), \mathcal{SE}_\perp and \mathcal{SE}_\top almost simulate Exp conditioned on that $b = \perp$ and $b = \top$, respectively. Therefore, if we randomly guess b and runs either of \mathcal{SE}_\perp or \mathcal{SE}_\top that successfully works for the guessed case, the output distribution is close to the real output distribution of Exp conditioned on that the guess is correct, which happens with probability almost 1/2.

A formal proof can be obtained based on the above intuition and is exactly the same as the proof of [\[CCY21, Lemma 5.5\]](#) except for notational adaptations. □

Then, we convert $\mathcal{SE}_{\text{comb}}$ into a full-fledged simulator that does not return Fail by using Watrous' rewinding lemma ([Lem. 21](#)). Namely, we let \mathbf{Q} be a quantum algorithm that takes $|\psi\rangle$ as input and outputs $\mathcal{SE}_{\text{comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$ where $b := 0$ if and only if it does not return Fail, $p_0 := \frac{1}{4}$, $q := \frac{1}{2}$, $\alpha := \varepsilon'$, and $T := 2 \log(1/\varepsilon')$. Then it is easy to check that the conditions for [Lem. 21](#) is satisfied by [Inequality \(84\)](#) in [Lem. 27](#) (for sufficiently large λ). Then, by using [Lem. 21](#), we can see that $\mathbf{R}(1^T, \mathbf{Q}, |\psi\rangle)$ runs in time $T = \text{poly}(\lambda)$ and its output (seen as a mixed state) has a trace distance bounded by $4\sqrt{\alpha} \frac{\log(1/\alpha)}{p_0(1-p_0)}$ from $D_{\text{ext,comb}}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)$. Since we have $\alpha = \varepsilon' = \frac{\varepsilon^2}{4 \log^4(\lambda)} = 1/\text{poly}(\lambda)$, we have $4\sqrt{\alpha} \frac{\log(1/\alpha)}{p_0(1-p_0)} < \sqrt{\alpha} \log^2(\lambda) = \frac{\varepsilon}{2}$ for sufficiently large λ where we used $\log(1/\alpha) = \log(\text{poly}(\lambda)) = o(\log^2(\lambda))$ and $\frac{4}{p_0(1-p_0)} = O(1)$. Thus, by combining the above and [Eq. \(85\)](#) in [Lem. 27](#), if we define $\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z) := \mathbf{R}(1^T, \mathbf{Q}, |\psi\rangle)$, then we have

$$\{\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}} \stackrel{s}{\approx}_{\frac{\varepsilon}{2} + 4\varepsilon'} \{\text{Exp}(\lambda, 1^{\gamma^{-1}}, |\psi\rangle, z)\}_{\lambda \in \mathbb{N}}$$

where $\gamma = \left(\frac{\varepsilon'}{5}\right)^4 < \varepsilon' < \varepsilon$. We can conclude the proof of [Lem. 20](#) by noting that we have $\frac{\varepsilon}{2} + 4\varepsilon' < \varepsilon$ since we have $\varepsilon' = \frac{\varepsilon^2}{4 \log^4(\lambda)} < \frac{\varepsilon}{8}$ for sufficiently large λ .

B Small-Tag, One-sided, Synchronous, Post-Quantum Setting

B.1 Construction

In this section, we prove that [Prot. 1](#) is post-quantumly secure if we rely on post-quantum building-blocks. Similarly to the classical case in [Sec. 4](#), we present the construction assuming the existence of post-quantum *injective* OWFs. The assumption can be relaxed to *any* OWFs by an appropriate modification to the protocol similarly to the classical case. Since this is exactly the same as that in the classical setting (i.e., [Sec. 4.8](#)), we do not repeat it in this section.

Our construction is based on the following building blocks:

- A post-quantum *injective* OWF f ;
- Naor’s commitment Com that is implemented with a post-quantum OWF.³⁶;
- A post-quantum witness-indistinguishable argument of knowledge with ε -close emulation WIAoK (as per [Def. 14](#)).

Protocol 5: Small-Tag One-Sided Synchronous Post-Quantum NMCom $\langle C, R \rangle_{\text{tg,PQ}}^{\text{OneSided}}$
<p>The tag space is defined to be $[n]$ where n is a polynomial on λ. Let $t \in [n]$ be the tag for the following interaction. Let m be the message to be committed to.</p> <p>Commit Stage:</p> <ol style="list-style-type: none"> 1. Receiver R samples and sends the first message β for Naor’s commitment; 2. Committer C commits to m using the second message of Naor’s commitment. Formally, C samples a random tape r and sends $\text{com} = \text{Com}_\beta(m; r)$; 3. R computes $\{y_i = f(x_i)\}_{i \in [t]}$ with $x_i \xleftarrow{\\$} \{0, 1\}^\lambda$ for each $i \in [t]$. R sends $Y = (y_1, \dots, y_t)$ to C; 4. (WIAoK-1.) R and C execute an instance of WIAoK where R proves to C that he “knows” a pre-image of some y_i contained in Y (defined in Step 3). Formally, R proves that $Y \in \mathcal{L}_f^t$, where $\mathcal{L}_f^t := \{(y_1, \dots, y_t) \mid \exists(i, x_i) \text{ s.t. } i \in [t] \wedge y_i = f(x_i)\}. \quad (86)$ Note that R uses $(1, x_1)$ as the witness when executing this WIAoK. 5. (WIAoK-2.) C and R execute an instance of WIAoK where C proves to R that he “knows” <i>either</i> the message committed in com (defined in Step 2), <i>or</i> a pre-image of some y_i contained in Y (defined in Step 3). Formally, C proves that $(\text{com}, Y) \in \mathcal{L}_\beta \vee \mathcal{L}_f^t$, where $\mathcal{L}_\beta \vee \mathcal{L}_f^t$ denotes the OR-composed language (as per Def. 1), \mathcal{L}_f^t was defined in Language (86) and $\mathcal{L}_\beta := \{\text{com} \mid \exists(m, r) \text{ s.t. } \text{com} = \text{Com}_\beta(m; r)\}. \quad (87)$ Note that C uses the (m, r) defined in Step 2 as the witness when executing this WIAoK. <p>Decommit Stage: C sends (m, r). R accepts if $\text{com} = \text{Com}_\beta(m; r)$, and rejects otherwise.</p>

Remark that the above protocol is exactly the same as [Prot. 1](#) except that we use post-quantum WIAoK with ε -close emulation and assume post-quantum security for all other building-blocks.

³⁶ In this way, the Naor’s commitment is statistically binding and *post-quantumly* computationally hiding.

Security. Completeness is straightforward from the description of [Prot. 5](#). The statistical binding property follows from that of Naor’s commitment. Computational hiding of any non-malleable commitment follows from its non-malleability. So, we only need to show that [Prot. 5](#) is non-malleable. This is established by the following [Thm. 32](#), which we prove in subsequent subsections.

Theorem 32. *The commitment scheme $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$ in [Prot. 5](#) is non-malleable against one-sided synchronous QPT adversaries with tag space $[n]$, with n being any polynomial on λ .*

B.2 Proving Non-Malleability

We prove [Thm. 32](#) in [Appx. B.2](#) to [B.7](#). The proof follows the same template used in [Sec. 4.2](#). The only exception is the proof of [Lem. 29](#) (which is the post-quantum counterpart of [Lem. 7](#)). Recall that in [Lem. 7](#), we amplify the extractor \mathcal{K} to the simulation-extractor \mathcal{SE} by rewinding. In the current post-quantum setting, we cannot rewind quantum algorithms in general. Therefore, we rely on an alternative argument based on our new extract-and-simulate lemma ([Lem. 20](#)). The rest of the proof is almost the same as its classical counterpart; Thus, many parts of the proof are taken verbatim from there. Essentially, the only difference is that we have to deal with noticeable errors that come from witness-extended ε -close emulator (as per [Property 2](#) of [Def. 14](#)). We highlight differences from the classical counterpart in [purple color](#) throughout this section.

We use the same notation as in the proof of [Thm. 2](#), with the only difference that now the adversaries are non-uniform QPT (instead of PPT) machines. It is worth noting that the honest committer and receiver are still classical (i.e., non-uniform PPT) machines.

In the sequel, we write $\langle C, R \rangle$ to mean $\langle C, R \rangle_{\text{tg}, \text{PQ}}^{\text{OneSided}}$ for notational convenience.

Game $H^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)$: This game is identical to its classical counterpart defined on [Page 26](#), except that $\mathcal{M} = \{\mathcal{M}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$ now is a (non-uniform) QPT machine. That is, this is the man-in-the-middle execution of the commit stage of the $\langle C, R \rangle$ defined in [Prot. 5](#), where the left committer commits to m and \mathcal{M} ’s non-uniform advice is ρ_λ . The output of $H^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)$ is again defined to be $(\text{OUT}_{\mathcal{M}}, \tilde{\tau}, b)$, where $\text{OUT}_{\mathcal{M}}$ is the (quantum) output of $\mathcal{M}_\lambda(\rho_\lambda)$ at the end of this game, $\tilde{\tau}$ consists of the [Steps 1](#) and [2](#) messages exchanged in the right session, and $b \in \{\top, \perp\}$ is the honest receiver’s final decision.

Notation. Recall from [Def. 7](#) that the man-in-the-middle game for a QPT adversary \mathcal{M} is denoted by $\overline{\text{mim}}_{\langle C, R \rangle}^{\mathcal{M}_\lambda}(m, \rho_\lambda)$. Also, note that “ $\stackrel{c}{\approx}$ ” refers to *quantumly* computational indistinguishability throughout this section.

For any $(\text{OUT}_{\mathcal{M}}, \tilde{\tau}, b)$ in the support of $H^{\mathcal{M}_\lambda}(\lambda, m, z)$, we define $\text{val}_b(\tilde{\tau})$ similarly to the classical case, i.e.,

$$\text{val}_b(\tilde{\tau}) := \begin{cases} \text{val}(\tilde{\tau}) & b = \top \\ \perp & b = \perp \end{cases},$$

where $\text{val}(\tilde{\tau})$ denote the value statistically-bound in [Steps 1](#) and [2](#) of the right session.

Similarly to [Eq. \(13\)](#), we have

$$\{\overline{\text{mim}}_{\langle C, R \rangle}^{\mathcal{M}_\lambda}(m, \rho_\lambda)\} \stackrel{\text{i.d.}}{=} \{(\text{OUT}, \text{val}_b(\tilde{\tau})) : (\text{OUT}, \tilde{\tau}, b) \leftarrow H^{\mathcal{M}_\lambda}(m, \rho_\lambda)\}, \quad (88)$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$ and $m \in \{0, 1\}^{\ell(\lambda)}$.

Then, similarly to [Eq. \(14\)](#), the post-quantum non-malleability can be reduced to establishing the following equation:

$$\begin{aligned} & \{(\text{OUT}^0, \text{val}_{b^0}(\tilde{\tau}^0)) : (\text{OUT}^0, \tilde{\tau}^0, b^0) \leftarrow H^{\mathcal{M}_\lambda}(m_0, \rho_\lambda)\} \\ & \stackrel{c}{\approx} \{(\text{OUT}^1, \text{val}_{b^1}(\tilde{\tau}^1)) : (\text{OUT}^1, \tilde{\tau}^1, b^1) \leftarrow H^{\mathcal{M}_\lambda}(m_1, \rho_\lambda)\}, \end{aligned} \quad (89)$$

where both ensembles are indexed by $\lambda \in \mathbb{N}$ and $(m_0, m_1) \in \{0, 1\}^{\ell(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$.

Proof by Contradiction. Similarly to [Inequality \(15\)](#), we assume for contradiction that there are a (possibly non-uniform) QPT distinguisher $\mathcal{D} = \{\mathcal{D}_\lambda, \sigma_\lambda\}_{\lambda \in \mathbb{N}}$ and a function $\delta(\lambda) = 1/\text{poly}(\lambda)$ such that for infinitely many $\lambda \in \mathbb{N}$, it holds that

$$\left| \Pr[\mathcal{D}_\lambda(\text{OUT}^0, \text{val}_{b^0}(\tilde{\tau}^0); \sigma_\lambda) = 1] - \Pr[\mathcal{D}_\lambda(\text{OUT}^1, \text{val}_{b^1}(\tilde{\tau}^1); \sigma_\lambda) = 1] \right| \geq 3 \cdot \delta(\lambda), \quad (90)$$

where the first probability is taken over the random procedure $(\text{OUT}^0, \tilde{\tau}^0, b^0) \leftarrow H^{\mathcal{M}_\lambda}(m_0, \rho_\lambda)$, and the second probability is taken over the random procedure $(\text{OUT}^1, \tilde{\tau}^1, b^1) \leftarrow H^{\mathcal{M}_\lambda}(m_1, \rho_\lambda)$ (and the randomness due to the measurements performed by \mathcal{D}_λ (if any) for both probabilities).

Then, we show the following [Lem. 28](#), which should be understood as the post-quantum counterpart of [Lem. 6](#).

Lemma 28. *For the above $\delta(\lambda)$, there exists a hybrid G such that for any QPT $\mathcal{M} = \{\mathcal{M}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, the following holds*

1. $\{(\text{OUT}^0, \text{Val}^0) : (\text{OUT}^0, \text{Val}^0) \leftarrow G^{\mathcal{M}_\lambda}(\lambda, m_0, \rho_\lambda)\} \stackrel{c}{\approx} \{(\text{OUT}^1, \text{Val}^1) : (\text{OUT}^1, \text{Val}^1) \leftarrow G^{\mathcal{M}_\lambda}(\lambda, m_1, \rho_\lambda)\}$, where both ensembles are indexed by $\lambda \in \mathbb{N}$ and $(m_0, m_1) \in \{0, 1\}^{\ell(\lambda)} \times \{0, 1\}^{\ell(\lambda)}$.
2. $\{(\text{OUT}^G, \text{Val}^G) : (\text{OUT}^G, \text{Val}^G) \leftarrow G^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)\} \stackrel{c}{\approx}_{\delta(\lambda)} \{(\text{OUT}^H, \text{val}_{b^H}(\tilde{\tau}^H)) : (\text{OUT}^H, \tilde{\tau}^H, b^H) \leftarrow H^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)\}$, where both ensembles are indexed by $\lambda \in \mathbb{N}$ and $m \in \{0, 1\}^{\ell(\lambda)}$.

It is easy to see that if [Lem. 28](#) is true, it contradicts our assumption in [Inequality \(90\)](#). Therefore, it will finish the proof of non-malleability. Indeed, this lemma is the most technically involved part. We prove it in [Appx. B.3](#).

B.3 Proof of [Lem. 28](#)

Similarly to [Algo. 4.1](#), we provide a new but equivalent interpretation of the game $H^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)$ in [Algo. B.1](#). We also provide a picture in [Fig. 11a](#) to illustrate it.

Algorithm B.1: Re-interpretation of Game $H^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)$
<p>Game $H^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda)$ can be split into the following stages:</p> <ol style="list-style-type: none"> 1. Prefix Generation: First, execute Steps 1 and 2 of the man-in-the-middle game of Prot. 5. That is, it plays as the left honest committer committing to m and the right honest receiver, with $\mathcal{M}_\lambda(\rho_\lambda)$ being the man-in-the-middle adversary. <p><u>Notation:</u> Let $\text{st}_\mathcal{M}$ denote the state of \mathcal{M} at the end of Step 2; Let st_C (resp. st_R) denote the state of the honest committer (resp. receiver) at the end of Step 2; Let $\tilde{\tau}$ denote the tuple $(\tilde{\beta}, \widetilde{\text{com}})^a$. In terms of notation, we denote the execution of this stage by</p> $(\text{st}_\mathcal{M}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau}) \leftarrow H_{\text{pre}}^{\mathcal{M}_\lambda}(\lambda, m, \rho_\lambda). \quad (91)$

We will call the tuple $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ the *prefix* and denote it by pref . It is worth noting that this pref contains all the information such that a QPT machine can “complete” the remaining execution of $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ starting from pref .

2. **The Remainder:** Next, it simply resumes from where the **Prefix Generation** stage stops, to finish the remaining steps of the man-in-the-middle execution $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$.

Notation: We introduce the following notations to describe this stage. Define a QPT machine \mathcal{A} that takes as input $(\text{st}_{\mathcal{M}}, \tilde{\tau})$; Machine \mathcal{A}_λ is supposed to run the residual strategy of \mathcal{M}_λ starting from $\text{st}_{\mathcal{M}}$. Also, define a QPT machine \mathcal{B} that takes as input $(\text{st}_C, \text{st}_R, \tilde{\tau})$; Machine \mathcal{B} is supposed to run the residual strategies of the honest committer C and receiver R , starting from st_C and st_R respectively^b. With the above notations, we can denote the execution of the remaining steps of $H^{\mathcal{M}}(\lambda, m, z)$ by

$$(\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau}), \quad (92)$$

where $\text{OUT}_{\mathcal{A}}$ is the output of \mathcal{A}_λ , and $b \in \{\perp, \top\}$ is the output of the honest receiver R (in the right), indicating if the man-in-the-middle’s commitment (i.e., the right session) is accepted ($b = \top$) or not ($b = \perp$). (We remark that $\text{OUT}_{\mathcal{A}}$ is nothing but the man-in-the-middle \mathcal{M} ’s final output.)

3. **Output:** It outputs the tuple $(\text{OUT}_{\mathcal{A}}, \tilde{\tau}, b)$.

^a Recall that $\tilde{\beta}$ and $\widetilde{\text{com}}$ are the [Steps 1](#) and [2](#) messages in the right session; they constitutes an execution of Naor’s commitment.

^b Note that it is not necessary to give $\tilde{\tau}$ as common input to these parties; indeed, it can be included in their respective internal states. We choose to make $\tilde{\tau}$ explicit only to match the syntax of [Lem. 20](#).

We prove the following lemma.

Lemma 29. *Let $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, \mathcal{A}_λ , and \mathcal{B} be as defined in [Algo. B.1](#). There exists a QPT machine \mathcal{SE} (the simulation-extractor) such that for any $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, any noticeable $\varepsilon(\lambda)$, it holds that*

$$\begin{aligned} & \left\{ (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}) : (\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}) \leftarrow \mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, \mathcal{A}_\lambda, \text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau}) \right\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx}_{\varepsilon(\lambda)} \left\{ (\text{OUT}_{\mathcal{A}}, \text{val}_b(\tilde{\tau})) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau}) \right\}_{\lambda \in \mathbb{N}} \end{aligned}$$

Remark 18. Compared to the classical counterpart ([Lem. 7](#)), \mathcal{SE} takes \mathcal{A}_λ and $\text{st}_{\mathcal{M}}$ as part of its input instead of accessing $\mathcal{A}_\lambda(\text{st}_{\mathcal{M}})$. Though we can see that \mathcal{SE} actually makes only black-box use of $\mathcal{A}_\lambda(\text{st}_{\mathcal{M}})$ in a certain sense, we do not try to formally state it because black-box simulation is not our focus. Another minor difference is that we removed $b_{\mathcal{SE}}$ from the output of $b_{\mathcal{SE}}$. This is because $b_{\mathcal{SE}}$ was only used in the proof of the classical counterpart of [Lem. 29](#) (i.e., [Lem. 7](#)), which will be replaced with a different proof based on [Lem. 20](#).

[Lem. 29](#) is the main technical lemma for the current proof of [Lem. 28](#). We present its proof in [Appx. B.4](#). In the following, we finish the proof of [Lem. 28](#) assuming that [Lem. 29](#) is true.

With [Lem. 29](#), we are now ready to present the description of G .

Algorithm B.2: Hybrid $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$

This hybrid proceeds as follows:

1. **Prefix Generation:** This stage is identical to [Stage 1](#) of $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$. Formally, it executes

$$(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau}) \leftarrow H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda),$$

where $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ is defined in [Stage 1](#) of [Algo. B.1](#).

2. **The Remainder:** Define \mathcal{A}_λ in the same way as in [Stage 2](#) of $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$. With this \mathcal{A}_λ and the $(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ from the previous stage, $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ invokes the \mathcal{SE} prescribed in [Lem. 29](#). Formally, it executes the following procedure:

$$(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}}) \leftarrow \mathcal{SE}(1^\lambda, 1^{\delta^{-1}}, \mathcal{A}_\lambda, \text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau}),$$

where the δ is the statistical distance that we want to show for [Property 2](#) of [Lem. 28](#).

Remark 19. We emphasize that in this stage, $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ does *not* make use of st_C .

3. **Output:** It outputs $(\text{OUT}_{\mathcal{SE}}, \text{Val}_{\mathcal{SE}})$.

Proving [Property 1](#) of [Lem. 28](#). Observe that hybrid $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ is an efficient machine, since both $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ and $\mathcal{SE}(1^\lambda, 1^{\delta^{-1}}, \mathcal{A}_\lambda, \text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ are efficient. Moreover, it does not rewind [Steps 1](#) and [2](#) of the man-in-the-middle execution, and $\mathcal{SE}(1^\lambda, 1^{\delta^{-1}}, \mathcal{A}_\lambda, \text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ does *not* need to know st_C . Therefore, [Property 1](#) of [Lem. 28](#) follows immediately from the computational-hiding property of the left Naor's commitment (i.e., [Steps 1](#) and [2](#) in the left session).

Proving [Property 2](#) of [Lem. 28](#). First, observe that the distribution of the prefix is identical in $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ and $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$. For each fixed prefix, [Lem. 29](#) implies that $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ and $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ are $\delta(\lambda)$ -computationally indistinguishable (notice that $G^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ runs \mathcal{SE} with the second input $1^{\delta^{-1}}$). This immediately implies [Property 2](#) of [Lem. 28](#).

B.4 Proof of [Lem. 29](#)

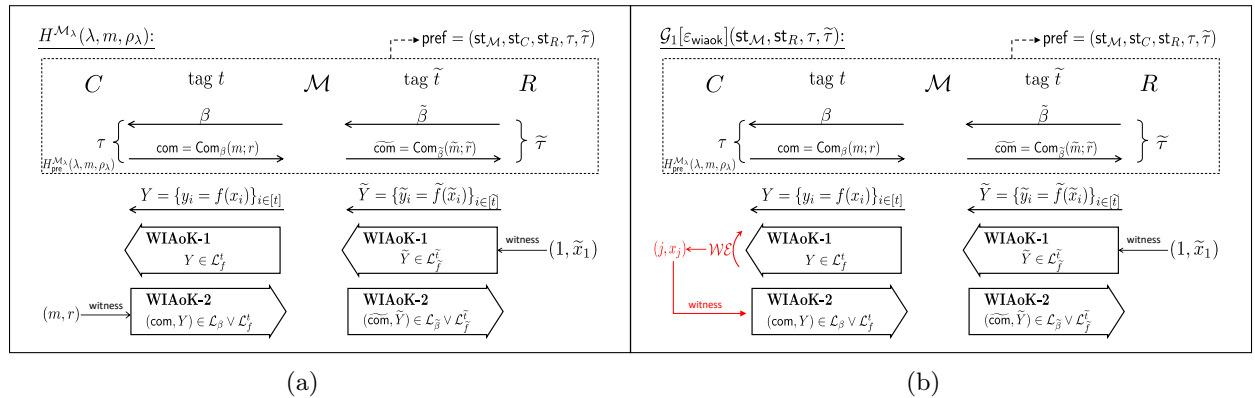


Fig. 11: Machines $H^{\mathcal{M}}$ and \mathcal{G}_1 (Difference is highlighted in red color)

First, we describe a machine \mathcal{G}_1 that simulates the real execution without using st_C . Unlike its classical counterpart, \mathcal{G}_1 is parameterized by a noticeable function $\epsilon_{\text{wiaok}}(\lambda)$, which is used as the error parameter for witness-extended emulators for WIAoK (as per [Property 2](#) of [Def. 14](#)).

Machine $\mathcal{G}_1[\varepsilon_{\text{wiaok}}]$: (Illustrated in Fig. 11b) For any prefix pref , $\mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to Stage 2 of $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ shown in Algo. B.1 (and depicted in Fig. 11a), except for the following difference: Instead of executing the left **WIAoK-1** honestly, it uses the witness-extended emulator (as per Property 2 of Def. 14) of the left **WIAoK-1** with the error parameter $\varepsilon_{\text{wiaok}}(\lambda)$ to extract a witness, and

- If the left committer accepts the left **WIAoK-1** and the extracted witness is valid (i.e., it is a (j, x_j) pair such that $y_j = f(x_j)$ for some $j \in [t]$), \mathcal{G}_1 uses (j, x_j) to finish the left **WIAoK-2**. Similarly to $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, \mathcal{G}_1 eventually outputs \mathcal{M} 's final state and the right receiver's decision bit b ;
- If the left committer accepts the left **WIAoK-1** but the extracted witness is invalid, it aborts immediately and outputs (\perp, \perp) .
- If the left committer rejects the left **WIAoK-1**, it runs the rest of execution of $H^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ to output \mathcal{M} 's final state and the right receiver's decision bit b . Note that it does not need to run the left **WIAoK-2** in this case since the left committer aborts after the left **WIAoK-1**.

We denote the above procedure by $(\text{OUT}, b) \leftarrow \mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. It is worth noting that \mathcal{G}_1 does *not* need to know st_C .

Next, we prove a lemma that shows that $\mathcal{G}_1[\varepsilon_{\text{wiaok}}]$ simulates the real execution up to an error $\varepsilon_{\text{wiaok}}$.

Lemma 30. *Let $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, \mathcal{A}_λ , and \mathcal{B} be as defined in Algo. B.1. For any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, it holds that*

$$\begin{aligned} & \left\{ (\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau}) \right\}_{\lambda \in \mathbb{N}} \\ & \stackrel{\approx}{\approx}_{\varepsilon_{\text{wiaok}}(\lambda)} \left\{ (\text{OUT}_{\mathcal{A}}, b) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau}) \right\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

Proof. We first define an intermediate machine $\mathcal{G}'_1[\varepsilon_{\text{wiaok}}]$ below.

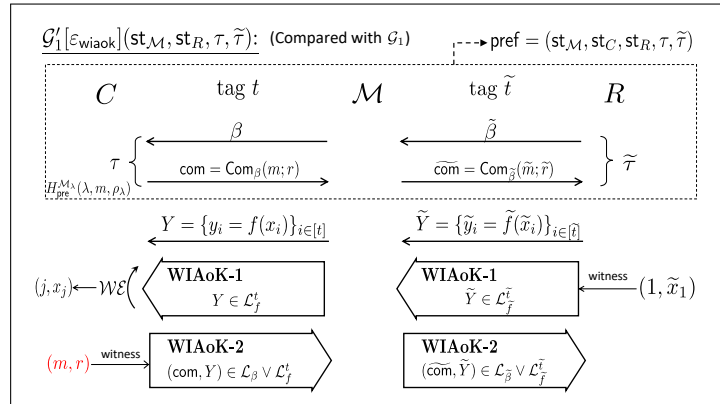


Fig. 12: Machine \mathcal{G}'_1 (Difference with \mathcal{G}_1 is highlighted in red color)

Machine $\mathcal{G}'_1[\varepsilon_{\text{wiaok}}]$: (Illustrated in Fig. 12) For any prefix pref , $\mathcal{G}'_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it uses (m, r) as the witness in the left **WIAoK-2** even if it succeeds in extracting a valid witness from the left **WIAoK-1**.

By the WI property of the left **WIAoK-2**, it holds that

$$\begin{aligned} & \{(\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx} \{(\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}'_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned} \quad (93)$$

Note that the only difference between $\mathcal{G}'_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ and $\langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})$ is that the former runs the witness-extended emulator of the left **WIAoK-1** with the error parameter $\varepsilon_{\text{wiaok}}(\lambda)$ (but does not use the extracted witness at all). Thus, by the AoK property of the left **WIAoK-1**, it holds that

$$\begin{aligned} & \{(\text{OUT}, b) : (\text{OUT}, b) \leftarrow \mathcal{G}'_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx}_{\varepsilon_{\text{wiaok}}(\lambda)} \{(\text{OUT}_{\mathcal{A}}, b) : (\text{OUT}_{\mathcal{A}}, b) \leftarrow \langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned} \quad (94)$$

By combining Eq. (93) and (94), we obtain Lem. 30. □

Next, we define the probability of R being convinced in the execution of $\mathcal{G}_1[\varepsilon_{\text{wiaok}}]$. This value plays an important role later in our proof.

Definition 16. For any $(\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$ in the support of $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$ and any noticeable $\varepsilon_{\text{wiaok}}$, we define the following value $p_{\text{pref}}^{\text{Sim}}[\varepsilon_{\text{wiaok}}]$:

$$p_{\text{pref}}^{\text{Sim}}[\varepsilon_{\text{wiaok}}] := \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})]$$

Next, we show a technical lemma that gives an extractor \mathcal{K} without the simulation property.

Lemma 31. Let $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, \mathcal{A}_λ , \mathcal{B} be as defined in Algo. B.1. There exists a QPT machine \mathcal{K} such that for any noticeable $\varepsilon(\lambda)$, there is a noticeable $\varepsilon_{\text{wiaok}}(\lambda) \leq \varepsilon(\lambda)$ that is efficiently computable from $\varepsilon(\lambda)$ such that the following holds for any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$:

1. (**Syntax.**) \mathcal{K} takes as input $(1^\lambda, 1^{\varepsilon^{-1}}, \mathcal{A}_\lambda, \text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. It outputs a value $\text{Val}_{\mathcal{K}} \in \{0, 1\}^{\ell(\lambda)} \cup \{\perp\}$ such that $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$ whenever $\text{Val}_{\mathcal{K}} \neq \perp$.
2. If $p_{\text{pref}}^{\text{Sim}}[\varepsilon_{\text{wiaok}}] \geq \varepsilon(\lambda)$, then it holds that

$$\Pr\left[\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau}) : \text{Val}_{\mathcal{K}} \leftarrow \mathcal{K}(1^\lambda, 1^{\varepsilon^{-1}}, \mathcal{A}_\lambda, \text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})\right] \geq \frac{\varepsilon'(\lambda)}{t},$$

$$\text{where } \varepsilon'(\lambda) := \frac{\varepsilon(\lambda)}{10t^2}.$$

Remark 20. Compared to the classical counterpart (Lem. 9), \mathcal{K} takes $1^{\varepsilon^{-1}}$ as an additional input. This is needed because we only assume the AoK property via witness-extended ε -close emulation (as per Property 2 of Def. 14). As a positive effect of the above difference, \mathcal{K} in Lem. 31 runs in *strict* QPT whereas it runs in *expected* PPT in the classical counterpart (Lem. 9). Also, we give \mathcal{A}_λ and $\text{st}_{\mathcal{M}}$ as part of input to \mathcal{K} instead of giving oracle access to $\mathcal{A}_\lambda(\text{st}_{\mathcal{M}})$ for a similar reason in Rmk. 18.

We will prove Lem. 31 in Appx. B.5. In the rest of this subsection, we finish the proof of Lem. 29 assuming Lem. 31 is true. This is the only part of the proof of Thm. 32 that significantly differs from its classical counterpart.

Proof of Lem. 29. We apply [Lem. 20](#) where $\mathcal{G}_1[\varepsilon_{\text{wiaok}}]$, \mathcal{K} , $(\mathcal{A}_\lambda, \text{st}_\mathcal{M}, \text{st}_R, \tau)$, $\tilde{\tau}$, and $\text{val}(\tilde{\tau})$ play the roles of \mathcal{G} , \mathcal{K} , ρ_λ , z_λ , and $s_{z_\lambda}^*$ in [Lem. 20](#), respectively. Then, [Lem. 31](#) ensures that the assumptions of [Lem. 20](#) are satisfied.³⁷ Thus, there exists a polynomial poly and a QPT machine \mathcal{SE}' such that for any noticeable function $\bar{\varepsilon}(\lambda)$,

$$\begin{aligned} & \{(\text{OUT}_{\mathcal{SE}'}, \text{Val}_{\mathcal{SE}'}) : (\text{OUT}_{\mathcal{SE}'}, \text{Val}_{\mathcal{SE}'}) \leftarrow \mathcal{SE}'(1^\lambda, 1^{\bar{\varepsilon}^{-1}}, \mathcal{A}_\lambda, \text{st}_\mathcal{M}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{s}{\approx}_{\bar{\varepsilon}(\lambda)} \{(\text{OUT}, \text{val}_b(\tilde{\tau})) : (\text{OUT}, b) \leftarrow \mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_\mathcal{M}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \end{aligned} \quad (95)$$

for a noticeable function $\varepsilon_{\text{wiaok}}(\lambda) \leq \bar{\varepsilon}(\lambda)$.

Moreover, since $\text{val}_b(\tilde{\tau})$ is determined by b for each fixed $\tilde{\tau}$, [Lem. 30](#) directly implies

$$\begin{aligned} & \{(\text{OUT}, \text{val}_b(\tilde{\tau})) : (\text{OUT}, b) \leftarrow \mathcal{G}_1[\varepsilon_{\text{wiaok}}](\text{st}_\mathcal{M}, \text{st}_R, \tau, \tilde{\tau})\}_{\lambda \in \mathbb{N}} \\ & \stackrel{c}{\approx}_{\varepsilon_{\text{wiaok}}(\lambda)} \{(\text{OUT}_\mathcal{A}, \text{val}_b(\tilde{\tau})) : (\text{OUT}_\mathcal{A}, b) \leftarrow \langle \mathcal{A}_\lambda(\text{st}_\mathcal{M}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})\}_{\lambda \in \mathbb{N}}. \end{aligned} \quad (96)$$

By setting $\bar{\varepsilon}(\lambda) := \varepsilon(\lambda)/2$ so that $\bar{\varepsilon}(\lambda) + \varepsilon_{\text{wiaok}}(\lambda) \leq \bar{\varepsilon}(\lambda) + \bar{\varepsilon}(\lambda) \leq \varepsilon(\lambda)$ and defining $\mathcal{SE}(1^\lambda, 1^{\varepsilon^{-1}}, \dots) := \mathcal{SE}'(1^\lambda, 1^{\bar{\varepsilon}^{-1}}, \dots)$, [Eq. \(95\)](#) and [\(96\)](#) give [Lem. 29](#). \square

B.5 Extractor \mathcal{K} (Proof of [Lem. 31](#))

In the following, we fix a noticeable function $\varepsilon(\lambda)$ for which we want to prove [Lem. 31](#). We show that it suffices to set $\varepsilon_{\text{wiaok}}(\lambda) := \frac{t+1}{t^2+6t+3} \cdot \varepsilon'(\lambda) = \frac{t+1}{10t^2(t^2+6t+3)} \cdot \varepsilon(\lambda)$. Since we fix ε and $\varepsilon_{\text{wiaok}}$, we omit the dependence on $\varepsilon_{\text{wiaok}}$ in our notation, i.e., we simply write $p_{\text{pref}}^{\text{Sim}}$ and \mathcal{G}_1 to mean $p_{\text{pref}}^{\text{Sim}}[\varepsilon_{\text{wiaok}}]$ and $\mathcal{G}_1[\varepsilon_{\text{wiaok}}]$. Similarly, machines introduced in the following also depend on $\varepsilon_{\text{wiaok}}$, but we do not explicitly write it in our notation. We also omit writing $(1^\lambda, 1^{\varepsilon^{-1}}, \mathcal{A}_\lambda)$ from inputs of those machines for notational simplicity.

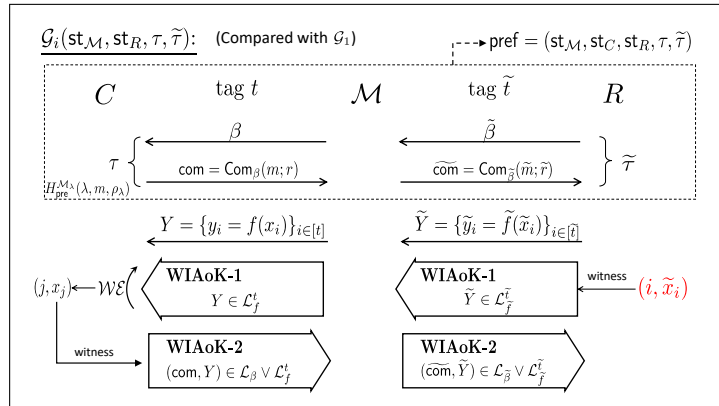


Fig. 13: Machine \mathcal{G}_i (Difference with \mathcal{G}_1 is highlighted in red color)

³⁷ Remark that $\varepsilon(\lambda)$ and $\frac{\varepsilon'(\lambda)}{t}$ in [Lem. 31](#) play the roles of $\gamma(\lambda)$ and $\delta(\lambda)$ in [Lem. 20](#), respectively.

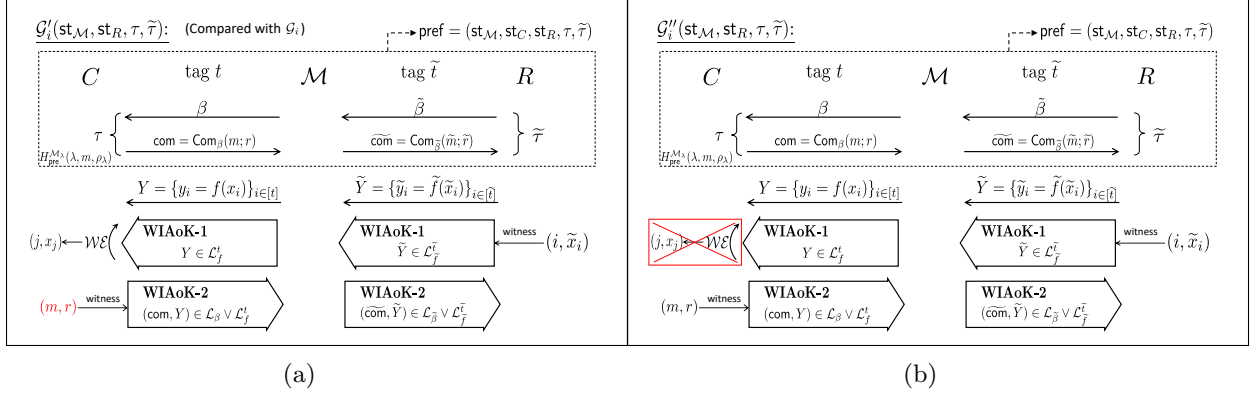


Fig. 14: Machines \mathcal{G}'_i and \mathcal{G}''_i (Difference is highlighted in red color)

Machine \mathcal{G}_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 13.) Recall that we have already defined the machine $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ on Page 84. Now, for $i \in [\tilde{t}] \setminus \{1\}$, $\mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it uses (i, \tilde{x}_i) as the witness in the right **WIAoK-1**.

Claim 33. $\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - 2\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda)$.

Proof. We define hybrid machines \mathcal{G}'_i and \mathcal{G}''_i as follows.

Machine \mathcal{G}'_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 14a.) Recall that we have already defined the machine $\mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ on Page 84. Now, for $i \in [\tilde{t}] \setminus \{1\}$, $\mathcal{G}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it uses (i, \tilde{x}_i) as the witness in the right **WIAoK-1**.

Machine \mathcal{G}''_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 14b) For any prefix pref , $\mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{G}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it honestly runs the left **WIAoK-1** instead of running the witness-extended emulator. In other words, $\mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})$ except that it uses (i, \tilde{x}_i) as the witness in the right **WIAoK-1**. In particular, $\mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ is identical to $\langle \mathcal{A}_\lambda(\text{st}_{\mathcal{M}}), \mathcal{B}(\text{st}_C, \text{st}_R) \rangle(1^\lambda, \tilde{\tau})$.

Then, Claim 33 follows from the following sequence of inequalities.

– By the WI property of the left **WIAoK-2** and the definition of $p_{\text{pref}}^{\text{Sim}}$, it holds that:

$$\Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}'_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \text{negl}(\lambda).$$

– By the AoK property of the left **WIAoK-1** and the above inequality, it holds that:

$$\Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}''_1(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda).$$

– By the WI property of the right **WIAoK-1** and the above inequality, it holds that:

$$\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - \varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda).$$

– By the AoK property of the left **WIAoK-1** and the above inequality, it holds that:

$$\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - 2\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda).$$

– By the WI property of the left **WIAoK-2** and the above inequality, it holds that:

$$\forall i \in [\tilde{t}], \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - 2\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda).$$

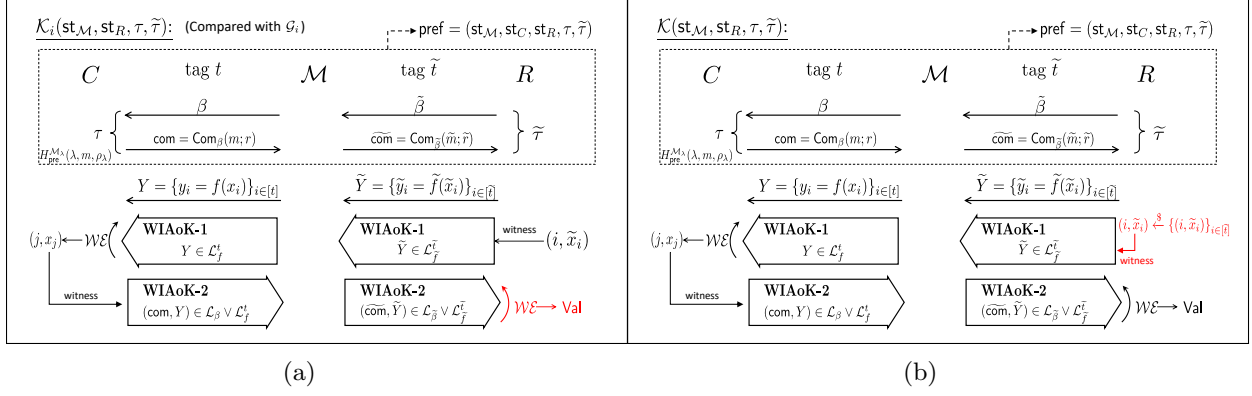


Fig. 15: Machines \mathcal{K}_i and \mathcal{K} (Difference is highlighted in red color)

□

Machine \mathcal{K}_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 15a.) For a prefix pref , $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to the $\mathcal{G}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ depicted in Fig. 13, except for the following difference. Machine $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ uses the witness-extended emulator (as per Property 2 of Def. 14) with the error parameter $\epsilon_{\text{wiaok}}(\lambda)$ to finish the right **WIAoK-2**, instead of playing the role of the honest receiver.

\mathcal{K}_i 's Output: Let w' denote the third output of the witness-extended emulator (see Property 2 of Def. 14), which is supposed to be the witness used by \mathcal{M} in the right **WIAoK-2** (for the statement $(\widetilde{\text{com}}, \tilde{Y})$ w.r.t. the language $\mathcal{L}_{\tilde{\beta}} \vee \mathcal{L}_{\tilde{f}}^t$). Depending on the value of w' , we define a value $\text{Val} \in \{0, 1\}^{\ell(\lambda)} \cup \{\perp_{\tilde{\gamma}}, \perp_{\text{invalid}}\}$ as follows:

1. If w' is a valid witness for $(\widetilde{\text{com}}, \tilde{Y}) \in \mathcal{L}_{\tilde{\beta}} \vee \mathcal{L}_{\tilde{f}}^t$. Then, there are two sub-cases:
 - (a) w' is a valid witness for $\widetilde{\text{com}} \in \mathcal{L}_{\tilde{\beta}}$. In this case, w' consists of the value $\text{val}(\tilde{\tau})$, i.e., the value committed in $\tilde{\tau} = (\tilde{\beta}, \widetilde{\text{com}})$, and the randomness \tilde{r} . We set $\text{Val} := \text{val}(\tilde{\tau})$. Importantly, notice that we do *not* include the randomness \tilde{r} in Val .
 - (b) w' is a valid witness for $\tilde{Y} \in \mathcal{L}_{\tilde{f}}^t$: In this case, we set $\text{Val} := \perp_{\tilde{\gamma}}$.
2. Otherwise, set $\text{Val} := \perp_{\text{invalid}}$.

The output of \mathcal{K}_i is defined to be the above Val . Notice that this is in contrast to all previous machines, for which the output is defined to be the main-in-the-middle's output and the honest receiver's decision bit. We emphasize that such a Val satisfies the syntax requirement in Property 1 of Lem. 31. In particular, $\text{Val} = \text{val}(\tilde{\tau})$ whenever $\text{Val} \neq \perp$.³⁸

Claim 34. $\forall i \in [\tilde{t}], \Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - 3\epsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda)$.

Proof. This claim follows from Claim 33 and the AoK property of the right **WIAoK-2**.

□

³⁸ Note that here we defined two types of abortion: $\perp_{\tilde{\gamma}}$ and \perp_{invalid} , while Property 1 of Lem. 31 only allows a single abortion symbol \perp . We remark that this is only a cosmetic difference—it can be made consistent using the following rules: $\perp = \perp_{\tilde{\gamma}}$ and $\perp = \perp_{\text{invalid}}$ (i.e., $\text{Val} \neq \perp \Leftrightarrow (\text{Val} \neq \perp_{\tilde{\gamma}} \wedge \text{Val} \neq \perp_{\text{invalid}})$).

Finally, we are ready to define the extractor \mathcal{K} . Intuitively, \mathcal{K} can be conceived as an average-case version of $\{\mathcal{K}_i\}_{i \in [\tilde{t}]}$:

Extractor \mathcal{K} : (Illustrated in Fig. 15b.) For a prefix pref , $\mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ samples uniformly at random an index $i \xleftarrow{\$} [\tilde{t}]$, executes $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$, and outputs whatever $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ outputs.

Next, we show that the extractor \mathcal{K} satisfies the requirements in Lem. 31.

Running Time of \mathcal{K} . Observe that for each $i \in [\tilde{t}]$, $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ differs from the real man-in-the-middle execution only in the following places:

- (i, \tilde{x}_i) is used in the right **WIAoK-1**;
- the witness-extended emulator is used in the right **WIAoK-2** and the left **WIAoK-1**.

Since the witness-extended emulator (as per Property 2 of Def. 14) runs in QPT, so does \mathcal{K}_i . Thus, \mathcal{K} is QPT.

Proving Property 1 of Lem. 31. It is straightforward to see that the \mathcal{K} defined above satisfies the syntax requirement in Lem. 31. In particular, we have $\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau})$ whenever $\text{Val}_{\mathcal{K}} \neq \perp$, because this is true for each \mathcal{K}_i by definition (see the paragraph for “ \mathcal{K}_i ’s Output” on Page 88).

Proving Property 2 of Lem. 31. First, recall that Property 2 requires that for any pref in the support of $H_{\text{pre}}^{\mathcal{M}\lambda}(\lambda, m, \rho_\lambda)$, if $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then it holds that

$$\Pr[\text{Val}_{\mathcal{K}} = \text{val}(\tilde{\tau}) : \text{Val}_{\mathcal{K}} \leftarrow \mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \frac{\varepsilon'(\lambda)}{t}. \quad (97)$$

Also recall that $\mathcal{K}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ is defined to execute the machine $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ with i uniformly sampled from $[\tilde{t}]$. Therefore, Inequality (97) can be reduced to the following Lem. 32. We will prove Lem. 32 in Appx. B.6, which will eventually finish the current proof of Lem. 31.

Lemma 32. *Let $\varepsilon(\lambda) = \frac{1}{\text{poly}(\lambda)}$ and $\varepsilon'(\lambda) = \frac{\varepsilon(\lambda)}{t^2}$. For any $\text{pref} = (\text{st}_{\mathcal{M}}, \text{st}_C, \text{st}_R, \tau, \tilde{\tau})$, if $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, then there exists an $i \in [\tilde{t}]$ such that*

$$\Pr[\text{Val} = \text{val}(\tilde{\tau}) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \varepsilon'(\lambda).$$

B.6 Proof of Lem. 32

Notation. We highly recommend reviewing the “ \mathcal{K}_i ’s Output” part on Page 88 (in particular, the meanings of Val , $\perp_{\tilde{\mathcal{Y}}}$, and \perp_{invalid}) before starting reading this subsection. Recall that \mathcal{K}_i ’s output Val is determined by the w' output by the witness-extended emulator of the right **WIAoK-2**. In this subsection, we will need to refer to this w' , although it is not included as a part of \mathcal{K}_i ’s output. Particularly, we will make use of the following notation: whenever we write an expression of the form

$$\Pr[\text{Some Event } E_{w'} \text{ about } w' : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})],$$

it should be understood as the probability of $E_{w'}$, where w' refers to the w' generated during the random procedure $\text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$, over which the probability is taken.

Using these notations, we can partition the event $\text{Val} = \perp_{\tilde{\mathcal{Y}}}$ as the following *mutually exclusive* events: $w' = (1, \tilde{x}_1)$, or, \dots , or $w' = (\tilde{t}, \tilde{x}_{\tilde{t}})$, where $\tilde{y}_i = \tilde{f}(\tilde{x}_i)$ for each $i \in [\tilde{t}]$. Formally, we express

this relation by

$$\Pr[\text{Val} = \perp_{\tilde{\gamma}} : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] = \sum_{i=1}^{\tilde{t}} \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \quad (98)$$

With the above notations, we prove [Lem. 32](#) in the following.

Proof of [Lem. 32](#). We assume for contradiction that for some pref satisfying $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$, it holds that

$$\forall i \in [\tilde{t}], \Pr[\text{Val} = \text{val}(\tilde{\tau}) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] < \varepsilon'(\lambda). \quad (99)$$

Claim 35. *Under the assumption in [Inequality \(99\)](#), it holds that*

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq p_{\text{pref}}^{\text{Sim}} - 3\varepsilon_{\text{wiaok}}(\lambda) - \varepsilon'(\lambda) - \text{negl}(\lambda).$$

Proof. In this proof, all the probabilities are taken over the random procedure $\text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. First, notice that

$$\begin{aligned} \forall i \in [\tilde{t}], \Pr[\text{Val} \neq \perp_{\text{invalid}}] &= \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \Pr[\text{Val} = \perp_{\tilde{\gamma}}] \\ &\text{(by [Eq. \(98\)](#))} = \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \Pr[w' = (i, \tilde{x}_i)] + \sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)]. \end{aligned} \quad (100)$$

Then, the following holds:

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i)] = \Pr[\text{Val} \neq \perp_{\text{invalid}}] - \Pr[\text{Val} = \text{val}(\tilde{\tau})] - \left(\sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)] \right) \quad (101)$$

$$\geq p_{\text{pref}}^{\text{Sim}} - 3\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda) - \Pr[\text{Val} = \text{val}(\tilde{\tau})] - \left(\sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)] \right) \quad (102)$$

$$\geq p_{\text{pref}}^{\text{Sim}} - 3\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda) - \varepsilon'(\lambda) - \left(\sum_{j \in [\tilde{t}] \setminus \{i\}} \Pr[w' = (j, \tilde{x}_j)] \right), \quad (103)$$

where [Eq. \(101\)](#) follows from [Eq. \(100\)](#), [Inequality \(102\)](#) follows from [Claim 34](#), and [Inequality \(103\)](#) follows from [Inequality \(99\)](#).

Now, to prove [Claim 35](#), it suffices to show that

$$\forall i \in [\tilde{t}], \forall j \in [\tilde{t}] \setminus \{i\}, \Pr[w' = (j, \tilde{x}_j) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] = \text{negl}(\lambda). \quad (104)$$

This can be reduced via standard techniques to the (post-quantum) one-wayness of the OWF \tilde{f} in [Step 3](#) of the right execution. In more details, we assume for contradiction that there exist $i^*, j^* \in [\tilde{t}]$ such that $i^* \neq j^*$ and that the probability $\Pr[w' = (j^*, \tilde{x}_{j^*}) : \text{Val} \leftarrow \mathcal{K}_{i^*}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})]$ is non-negligible, where, by definition, \tilde{x}_{j^*} is the preimage of \tilde{y}_{j^*} under the right OWF \tilde{f} . Then, we can build a QPT adversary \mathcal{A}_{OWF} breaking one-wayness in the following way: \mathcal{A}_{OWF} obtains the challenge y^* from the external one-wayness challenger; it then runs the machine $\mathcal{K}_{i^*}(\text{pref})$ internally, for which \mathcal{A}_{OWF} uses y^* in place of \tilde{y}_{j^*} when executing [Step 3](#) in the right. Note that the internal execution of $\mathcal{K}_{i^*}(\text{pref})$ is identically to the real execution of \mathcal{K}_{i^*} , thus the extracted $w' = (j^*, \tilde{x}_{j^*})$ must satisfy $\tilde{f}(\tilde{x}_{j^*}) = \tilde{y}_{j^*}$ ($= y^*$) with non-negligible probability, breaking one-wayness.

This finishes the proof of [Claim 35](#). □




Fig. 16: Machines \mathcal{K}'_i and \mathcal{K}''_i (Difference is highlighted in red color)

Machine \mathcal{K}'_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 16a.) For a prefix pref , $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ proceeds as follows:

1. It first finishes **Step 3** of the man-in-the-middle execution in the same way as $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$. In particular, it will see in the left execution the values $Y = (y_1, \dots, y_t)$ sent from \mathcal{M} ;
2. It then recovers (x_1, \dots, x_t) by brute-force: Namely, for each $i \in [t]$, it inverts the OWF f to find x_i s.t. $f(x_i) = y_i$. It is possible that there exist some “bad” y_i ’s that are not in the range of f . For such bad i ’s, it sets $x_i = \perp$. If all the x_i ’s are bad, $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ halts and outputs Fail;
3. If this step is reached, we know that (x_1, \dots, x_t) cannot be all- \perp . $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ then picks an (s, x_s) uniformly at random from the good (i.e. non- \perp) x_i ’s.
4. Then, $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ continues to finish the execution in the same way as $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$, except that it uses (s, x_s) as the witness when executing the left **WIAoK-2**.

It is worth noting that the (j, x_j) extracted by the witness-extended emulator for the left **WIAoK-1** (inherited from $\mathcal{K}_i(\text{pref})$) is not used any more by $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$.

It is easy to see that if the (s, x_s) picked by $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ is equal to the (j, x_j) extracted in $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ from its left **WIAoK-1**, then $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ and $\mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ are identical.³⁹ Since $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ samples (s, x_s) uniformly from all the good (i, x_i) ’s, it must hold with probability at least $1/t$ that $(s, x_s) = (j, x_j)$. Therefore, the following must hold:

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \frac{1}{t} \cdot \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})]. \quad (105)$$

Machine \mathcal{K}''_i ($i \in [\tilde{t}]$): (Illustrated in Fig. 16b.) For a prefix pref , $\mathcal{K}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically to $\mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that it plays as the honest committer in the left **WIAoK-1**, instead of running the witness-extended emulator. Recall that starting from \mathcal{K}'_i , the witness (j, x_j) extracted by the witness-extended emulator from the left **WIAoK-1** is not used any more. Thus, machine \mathcal{K}''_i does not need to perform this witness-extended emulation.

By the AoK property of the left **WIAoK-1**, it holds that

$$\forall i \in [\tilde{t}], \left| \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}''_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] - \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}'_i(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \right| \leq \epsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda). \quad (106)$$

³⁹ Similar to **Rmk. 11**, this is the only place we make use of the injectivity of f .

Next, by the (non-uniform) WI property of the right **WIAoK-1**, it holds that

$$\forall i \in [\tilde{t}], \left| \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] - \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda). \quad (107)$$

Remark 21 (Power of Non-Uniform Reductions). Note that we can rely on the AoK and WI properties even though \mathcal{K}_i' and \mathcal{K}_i'' perform brute-force to recover (x_1, \dots, x_t) . This is because the brute-forcing step is done before **WIAoK-1** or **WIAoK-2** starts; Thus, (x_1, \dots, x_t) can be treated as a non-uniform advice in the reductions. This non-uniform type of argument will be used again in this section later.

Then, we have the following claim:

Claim 36. *It holds that*

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \geq \frac{1}{t} \cdot (p_{\text{pref}}^{\text{Sim}} - (t+3)\varepsilon_{\text{wiaok}}(\lambda) - \varepsilon'(\lambda)) - \text{negl}(\lambda).$$

Proof. This claim follows from [Claim 35](#) and [Inequalities \(35\), \(36\)](#) and [\(105\)](#). Formally, (in the following, we omit the input $(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ to $\mathcal{K}_i, \mathcal{K}_i',$ and \mathcal{K}_i'')

$$\forall i \in [\tilde{t}], \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_1''] \geq \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i''] - \text{negl}(\lambda) \quad (108)$$

$$\geq \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i'] - \varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda) \quad (109)$$

$$\geq \frac{1}{t} \cdot \Pr[w' = (i, \tilde{x}_i) : \text{Val} \leftarrow \mathcal{K}_i] - \varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda) \quad (110)$$

$$\geq \frac{1}{t} \cdot (p_{\text{pref}}^{\text{Sim}} - 3\varepsilon_{\text{wiaok}}(\lambda) - \varepsilon'(\lambda)) - \varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda), \quad (111)$$

where [Inequality \(108\)](#) follows from [Inequality \(107\)](#), [Inequality \(109\)](#) follows from [Inequality \(106\)](#), [Inequality \(110\)](#) follows from [Inequality \(105\)](#), and [Inequality \(111\)](#) follows from [Claim 35](#). □

Now, we show the last claim which, together with [Claim 36](#), will lead to the desired contradiction.

Claim 37. *It holds that*

$$\Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})] \leq p_{\text{pref}}^{\text{Sim}} + 2\varepsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda).$$

Deriving the Contradiction. Before proving [Claim 37](#), we first show why [Claims 10](#) and [36](#) are contradictory. In the following, all the probabilities are taken over $\text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$:

$$\begin{aligned} \Pr[\text{Val} \neq \perp_{\text{invalid}}] &= \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \Pr[\text{Val} = \perp_{\tilde{\gamma}}] \\ &= \Pr[\text{Val} = \text{val}(\tilde{\tau})] + \sum_{i=1}^{\tilde{t}} \Pr[w' = (i, \tilde{x}_i)] \end{aligned} \quad (112)$$

$$\begin{aligned} &\geq \sum_{i=1}^{\tilde{t}} \Pr[w' = (i, \tilde{x}_i)] \\ &\geq \tilde{t} \cdot \frac{1}{t} \cdot (p_{\text{pref}} - (t+3)\varepsilon_{\text{wiaok}}(\lambda) - \varepsilon'(\lambda)) - \text{negl}(\lambda) \end{aligned} \quad (113)$$

$$\geq \left(1 + \frac{1}{t}\right) \cdot (p_{\text{pref}} - (t+3)\varepsilon_{\text{wiaok}}(\lambda) - \varepsilon'(\lambda)) - \text{negl}(\lambda) \quad (114)$$

$$\begin{aligned} &= \left(1 + \frac{1}{t}\right) \cdot \left(p_{\text{pref}} - \frac{t^2 + 6t + 3}{t+1} \cdot \varepsilon_{\text{wiaok}}(\lambda) - \varepsilon'(\lambda)\right) + 2\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda) \\ &= \left(1 + \frac{1}{t}\right) \cdot \left(p_{\text{pref}} - 2\varepsilon'(\lambda)\right) + 2\varepsilon_{\text{wiaok}}(\lambda) - \text{negl}(\lambda) \end{aligned} \quad (115)$$

$$\begin{aligned} &= p_{\text{pref}}^{\text{Sim}} + 2\varepsilon_{\text{wiaok}}(\lambda) + \left(\frac{p_{\text{pref}}^{\text{Sim}}}{t} - 2\varepsilon'(\lambda) - \frac{2\varepsilon'(\lambda)}{t}\right) - \text{negl}(\lambda) \\ &\geq p_{\text{pref}}^{\text{Sim}} + 2\varepsilon_{\text{wiaok}}(\lambda) + \frac{5t^2 - t - 1}{5t^3} \cdot \varepsilon(\lambda) - \text{negl}(\lambda) \end{aligned} \quad (116)$$

where [Eq. \(112\)](#) follows from [Eq. \(98\)](#), [Inequality \(113\)](#) follows from [Claim 36](#), [Inequality \(114\)](#) follows from the assumption that $\tilde{t} \geq t+1$, [Eq. \(115\)](#) follows from our parameter setting $\varepsilon_{\text{wiaok}}(\lambda) = \frac{t+1}{t^2+6t+3} \cdot \varepsilon'(\lambda)$, and [Inequality \(116\)](#) follows from the assumption that $p_{\text{pref}}^{\text{Sim}} \geq \varepsilon(\lambda)$ and our parameter setting $\varepsilon'(\lambda) = \frac{\varepsilon(\lambda)}{10t^2}$.

Recall that t is the tag taking values from $[n]$ with n being a polynomial of λ . Also recall that $\varepsilon(\lambda)$ is an inverse polynomial on λ . Therefore, [Inequality \(116\)](#) can be written as:

$$\Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_1''(\text{pref})] \geq p_{\text{pref}}^{\text{Sim}} + \frac{1}{\text{poly}(\lambda)} - \text{negl}(\lambda),$$

which contradicts [Claim 37](#).

This eventually finishes the proof of [Lem. 32](#) (modulo the proof of [Claim 37](#), which we show in [Appx. B.7](#)).

B.7 Proof of [Claim 37](#)

We first define two extra machines \mathcal{K}_1^{**} and \mathcal{K}_1^* .

Machine \mathcal{K}_1^{} :** (Illustrated in [Fig. 17a](#).) For a prefix pref , $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ behaves identically as $\mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ except that $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ finishes the right **WIAoK-2** using the honest receiver's algorithm, instead of using the witness-extended emulator.

\mathcal{K}_1^{} 's Output.** We define the output of $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_R, \tau, \tilde{\tau})$ to be the honest receiver's decision $b \in \{\top, \perp\}$. This is in contrast to previous hybrids \mathcal{K}_i' , \mathcal{K}_i'' , and \mathcal{K}_i , whose output is defined to be the value Val that depends on the value w' extracted by the \mathcal{SE} for the right **WIAoK-2**.




Fig. 17: Machines \mathcal{K}_1^{**} and \mathcal{K}_1^* (Difference is highlighted in red color)

By the AoK property of the right **WIAoK-2**, it holds that

$$\left| \Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - \Pr[b = \top : b \leftarrow \mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \right| \leq \varepsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda). \quad (117)$$

Machine \mathcal{K}_1^* : (Illustrated in Fig. 17b.) For a prefix pref , $\mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ behaves identically as $\mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ except the following difference: $\mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})$ uses the witness-extended emulator with the error parameter $\varepsilon_{\text{wiaok}}(\lambda)$ to extract a witness (j, x_j) from the left **WIAoK-1**, and if x_j is not a valid preimage for y_j , \mathcal{K}_1^* aborts.

By the AoK property of the left **WIAoK-1**, it holds that

$$\left| \Pr[b = \top : b \leftarrow \mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \right| \leq \varepsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda). \quad (118)$$

Compare \mathcal{K}_1^* with \mathcal{G}_1 . Now, let us compare \mathcal{K}_1^* with the \mathcal{G}_1 depicted in Fig. 13. They only differ in the witness used in the left **WIAoK-2** (and that \mathcal{G}_1 does not need to perform brute-forcing for Y , as it does not use those preimages). Therefore, by the (non-uniform) WI property of the left **WIAoK-2**, it holds that

$$\left| \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - \Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \right| \leq \text{negl}(\lambda).$$

Also, recall (from Def. 16) that $\Pr[b = \top : (\text{OUT}, b) \leftarrow \mathcal{G}_1(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})]$ is exactly the definition of $p_{\text{pref}}^{\text{Sim}}$. Thus, the above implies:

$$\left| \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] - p_{\text{pref}}^{\text{Sim}} \right| \leq \text{negl}(\lambda). \quad (119)$$

Therefore, the following holds:

$$\Pr[\text{Val} \neq \perp_{\text{invalid}} : \text{Val} \leftarrow \mathcal{K}_1''(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] \leq \Pr[b = \top : b \leftarrow \mathcal{K}_1^{**}(\text{st}_{\mathcal{M}}, \text{st}_{\mathcal{R}}, \tau, \tilde{\tau})] + \varepsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda) \quad (120)$$

$$\leq \Pr[b = \top : b \leftarrow \mathcal{K}_1^*(\text{pref})] + 2\varepsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda) \quad (121)$$

$$\leq p_{\text{pref}}^{\text{Sim}} + 2\varepsilon_{\text{wiaok}}(\lambda) + \text{negl}(\lambda), \quad (122)$$

where Inequality (120) follows from Inequality (117), Inequality (121) follows from Inequality (118), and Inequality (122) follows from Inequality (119).

This finishes the proof of Claim 37.