On Non-Interactive Blind Signatures in the Plain Model Using Complexity Leveraging

Kazuki Yamamura[1], Tetsuya Okuda[1], and Eiichiro Fujisaki[2]

[1] NTT Social Informatics Laboratories, Tokyo, Japan, kazuki.yamamura@ntt.com
[2] Japan Advanced Institute of Science and Technology, Ishikawa, Japan

Abstract. Blind signatures, introduced by Chaum (Crypto'82), are a fundamental cryptographic primitive with various applications such as e-voting, e-cash, anonymous credentials, and more. Although blind signatures inherently require interaction between both parties, Hanzlik (Eurocrypt'23) introduced the notion of non-interactive blind signatures (NIBS), which allow signatures on random messages to be issued blindly without interaction. While Hanzlik's NIBS constructions are provably secure in the random oracle model, instantiating a provably secure NIBS in the plain model remains an open problem.

In this paper, we introduce a non-interactive blind signature scheme in the plain model based on complexity leveraging. The key to our construction is the use of the non-uniform reductions employed by Garg et al. (Crypto'11), which enables us to instantiate a provably secure NIBS without relying on a trusted setup.

Furthermore, we investigate whether our construction can avoid the use of complexity leveraging by applying the idea proposed by Kalai and Khurana (Crypto'19), wherein complexity leveraging can be replaced with classical and quantum assumptions. We introduce a weaker blindness notion called non-adaptive blindness and show that this property allows our construction to avoid using complexity leveraging.

Of independent interest, we provide separation results demonstrating the existence of a NIBS construction that satisfies nonce blindness but not recipient blindness, and vice versa. This result implies that any NIBS construction should be proven to satisfy both nonce blindness and recipient blindness.

Keywords: Blind Signatures · Non-Interactive · Plain Model · Standard Model · Complexity Leveraging

1 Introduction

Blind signature schemes, introduced by Chaum [13], are a fundamental cryptographic primitive with several applications, such as e-voting [10], e-cash [12, 13], and anonymous credentials [4]. Blind signature schemes can be viewed as an interactive protocol between a signer, holding a signing key, and a recipient, holding a message. Informally, the goal of blind signatures is for the user to obtain a signature from the signer without revealing the content of the message

being signed. As such, blind signatures must satisfy two key security properties: Blindness, meaning the signer cannot learn the signed message, and Unforgeability, meaning the user can only obtain valid signatures through interaction with the signer.

As described, all existing blind signatures require interaction. However, Hanzlik introduced a groundbreaking cryptographic protocol called non-interactive blind signatures (NIBS) for random messages [27]. NIBS eliminates the need for online interactions. In the NIBS setting, both the signer and the recipient (i.e., the user) possess a pair of public and private keys. To create a valid message-signature pair, the signer creates a presignature on a nonce and the recipient's public key[3], and the recipient uses their secret key to derive a signature on a random message from the presignature.

NIBS has the potential to replace interactive blind signatures in applications where the choice of message is not critical. For instance, Privacy Pass [14, 30] is a typical application of NIBS (for more details, see [27]). Other applications of NIBS include e-cash-like payment protocols [29]. Implementing NIBS allows for batch-issuing and offline issuance via hardware security modules (HSM) or cold wallets [27]. Given the security incidents at cryptocurrency exchanges, such as signing key leakage, utilizing NIBS can mitigate these risks through offline issuance in HSMs or cold wallets.

Hanzlik introduced two instantiations of NIBS schemes [27]: an efficient construction based on structure-preserving signatures of equivalence classes and a generic construction. Both are provably secure in the random oracle model. However, the question remains whether it is possible to design a NIBS scheme without relying on the random oracle model or the common reference string (CRS) model. Whether a cryptographic protocol requires a trusted setup assumption is a critical consideration since, in the trusted setup model, if an authority introduces a backdoor, security can no longer be guaranteed.

To explore NIBS schemes without any trusted setup, we first examine the existing round-optimal (interactive) blind signatures that do not require such setups. Fischlin and Schröder showed that under certain conditions, it is impossible to construct a three-move blind signature from falsifiable assumptions without any trusted setup [17]. However, Garg et al. demonstrated a round-optimal blind signature without any trusted setup by using complexity leveraging and non-uniform reductions [21][4]. Complexity leveraging refers to a type of reduction in which the reduction algorithm has more computational power than the adversary, and non-uniform reductions allow the algorithm to perform unbounded computation before it is given the problem instance (further details will be provided later).

In more recent work, Katsumata et al. demonstrated that it is possible to construct a round-optimal blind signature without relying on complexity leveraging,

[3] In the PKI model, the recipient never needs to send their public key, allowing them to create a valid message-signature pair without interaction.

[4] This does not contradict the impossibility result by Fischlin and Schröder [17], as it relies on complexity leveraging.

by employing both classical and quantum polynomial-time algorithms [32], building on the ideas of Kalai and Khurana [31]. Intuitively, their approach replaces super-polynomial security with quantum polynomial-time security, allowing the construction of round-optimal blind signatures in the plain model without using complexity leveraging.

Complexity Leveraging and Non-Uniform Reduction Complexity leveraging exploits the gap between the computational power of an adversary and that of the reduction algorithm in security proofs. To create this gap, the building blocks must provide super-polynomial-time security. This allows the reduction algorithm—operating in super-polynomial time—to extract information from the adversary, who is constrained to polynomial time, thereby reducing the security proof to a cryptographic hard problem. Using this approach, Pass [34] introduced a two-move zero-knowledge argument system, which surpasses the barrier posed by the impossibility result by Goldreich and Oren [23]. Using this approach, Pass [34] introduced a two-move zero-knowledge argument system that surpasses the barrier posed by the impossibility result of Goldreich and Oren [23]. Subsequently, several protocols have been instantiated using complexity leveraging [9, 11, 19–21, 31, 32].

Non-uniform reduction algorithms can be intuitively regarded as $\mathsf{P/poly}$-machines. More specifically, in this setting, reduction algorithms are treated as two-stage processes: the pre-computation phase and the online phase. In the pre-computation phase, the reduction algorithm takes the security parameter as input and, using unbounded computational power, computes a polynomial-length advice string. In the online phase, the reduction algorithm takes the problem instance along with the advice string as input and attempts to solve the problem in polynomial time. Both Garg et al. [21] and Katsumata et al. [32] constructed round-optimal blind signatures in the plain model using this non-uniform reduction approach.

1.1 Our Result

In this paper, we present a non-interactive blind signature (NIBS) scheme in the plain model. Our construction relies on complexity leveraging. Unlike previous schemes [3, 27], our construction does not require any setup assumptions, such as a common reference string or random oracle. Additionally, its security does not rely on interactive assumptions.

Furthermore, we explore whether our construction can avoid relying on complexity leveraging. Specifically, we investigate the possibility of constructing a NIBS scheme using a combination of classical and quantum hardness assumptions [31, 32].While complexity leveraging necessitates a gap between the computational power of the adversary and the reduction algorithm in the security proof—requiring large parameters for the underlying building blocks—using both classical and quantum hardness assumptions can improve the overall efficiency of our construction. However, we observe that Hanzlik's blindness def-

inition does not straightforwardly accommodate such techniques in our construction. On the other hand, a weaker notion of blindness, which we refer to as non-adaptive blindness, does allow these techniques to be applied. Further details are provided in Section 4.

Additionally, we analyze the distinction between two types of blindness definitions: recipient blindness and nonce blindness. A reader might wonder whether every scheme that satisfies recipient blindness also satisfies nonce blindness, or vice versa. To address this question, we present separation results demonstrating the existence of a NIBS scheme that satisfies nonce blindness but not recipient blindness, and vice versa. These findings suggest that any NIBS construction should be proven to satisfy both nonce blindness and recipient blindness. More details can be found in Section 2.1.

1.2 Technical Overview

Constructing NIBS in the plain model We provide an overview of our construction. Our main contribution is the demonstration of the NIBS scheme in the plain model. Unlike Hanzlik's construction [27], our scheme does not rely on trusted setup assumptions, such as the random oracle model or the CRS model. To better understand our construction, we revisit the generic NIBS scheme proposed by Hanzlik [27]. In the key generation phase, the signer generates a standard signature key pair $(\mathsf{spk}, \mathsf{ssk})$ while the recipient generates a verifiable random function (VRF) key pair $(\mathsf{vpk}, \mathsf{vsk})$. During the issuance phase, the signer creates a digital signature psig on both a nonce nonce and the recipient's public key vpk, resulting in a presignature. In the obtaining phase, the recipient evaluates the VRF on the nonce nonce to obtain a message m, and then creates a dual-mode witness indistinguishable proof (DMWI) [26] π, using a common reference string $\mathsf{H}_{\mathsf{crs}}(0)$, where $\mathsf{H}_{\mathsf{crs}}$ is treated as a random oracle. The proof π demonstrates one of the following:

- the presignature psig is a standard digital signature on $(\mathsf{nonce}, \mathsf{vpk})$ and the message m is a result of the VRF evaluation,

or

- the random oracle value $\mathsf{H}_{\mathsf{crs}}(1)$ acts as a common reference string in binding mode.

An honest recipient proves the first case, while in a security reduction, the latter case is demonstrated by programming the random oracle $\mathsf{H}_{\mathsf{crs}}$. The recipient thus obtains the NIBS message-signature pair (m, π). In the verification phase, the verifier checks whether the proof π is valid.

Next, we review the security proof strategy for Hanzlik's generic construction. To prove the one-more unforgeability of the scheme, we need to construct a reduction that uses an adversary against one-more unforgeability to compute a forgery for the underlying standard signature scheme. This reduction requires programming the random oracle $\mathsf{H}_{\mathsf{crs}}$ in such a way that it can extract the witness from the DMWI proofs computed by the adversary, which implies that the

common reference string $\mathsf{H}_{\mathsf{crs}}(0)$ needs to be in binding mode. On the other hand, to prove the blindness of the scheme, we need to simulate DMWI proofs without a witness used by an honest prover (i.e., the latter case mentioned above). This requires programming the common reference string $\mathsf{H}_{\mathsf{crs}}(1)$ to be in binding mode.

Thus, Hanzlik's generic construction relies on a 'trapdoor' to simulate the DMWI proof within the security proof, necessitating a trusted setup, such as the common random string model or the random oracle model. However, to achieve a NIBS scheme in the plain model, we cannot rely on these trusted setups. This means that we can no longer extract the witness for the proof (in the one-more unforgeability proof) or simulate the proof (in the blindness proof).

To address this challenge, we employ complexity leveraging. Specifically, the signer and recipient exchange a 'trapdoor' with one another. Then, reductions running in super-polynomial time extract the trapdoor, enabling us to prove the security of the underlying schemes. In other words, these trapdoors allow the reduction in the one-more unforgeability proof to extract signatures from the underlying scheme, and the reduction in the blindness proof to simulate message-signature pairs for the adversary.

Let us introduce the idea of our construction based on the above strategy. In the key generation phase, a singer samples a random value $y \leftarrow \{0,1\}^\lambda$ as well as a standard signature key pair $(\mathsf{spk}, \mathsf{ssk})$, setting $((\mathsf{spk}, y), \mathsf{ssk})$ as a signer's key pair. (Why we require the signer's public key to include this random value y is described later.) Meanwhile, a recipient samples a secret key $R \leftarrow \{0,1\}^\lambda$ for a pseudorandom function F and computes a commitment to R[5]; i.e., $\mathsf{rpk} \leftarrow \mathsf{Com}(R; \mathsf{r}_{\mathsf{rpk}})$ as a recipient's public key, where $\mathsf{r}_{\mathsf{rpk}}$ is a randomness.

In the issue phase, the signer creates a standard digital signature psig on a nonce nonce and recipient's public key rpk to receive a presignature as in Hanzlik's construction. In the obtaining phase, the recipient evaluates the pseudorandom function F on input nonce to receive a message m. It also computes the following two commitments:

1. $\mathsf{com}^{(0)} \leftarrow \mathsf{Com}'((\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig}); r^{(0)})$
2. $\mathsf{com}^{(1)} \leftarrow \mathsf{Com}'(0; r^{(1)})$,

where $r^{(0)}$ and $r^{(1)}$ are randomness for the commitments. Note that commitments Com and Com' are not identical, which reason we will describe later.

[5] Since VRF can be constructed in the plain model [24], even if we use a verifiable random function as a building block as in Hanzlik's construction then we can obtain NIBS in the plain model. However, we instead use a non-interactive commitment scheme and a pseudorandom function, which enables us to obtain a simpler and more general scheme.

Furthermore, it creates a non-interactive witness indistinguishable proof (NIWI)[6] π [6, 15], which proves one of the following:

- The recipient's public key rpk is committed to the secret key R for the pseudorandom function F,
- The presignature psig is a standard digital signature on $(\mathsf{nonce}, \mathsf{rpk})$,
- The message m is the result of evaluating F_R on input nonce with R, and
- $\mathsf{com}^{(0)}$ is committed to $(\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig})$,

or

- $\mathsf{com}^{(1)}$ is committed to a preimage a such that $y = f(a)$ holds.

An honest recipient proves the former case, while the reduction in the security proof for blindness proves the latter case. Specifically, in the one-more unforgeability proof, the reduction, which aims to break the unforgeability of the underlying signature scheme, extracts a valid message-signature pair $(\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig})$ from the commitment $\mathsf{com}^{(0)}$ as a forgery by running in super-polynomial time $T \cdot \mathsf{poly}(\lambda)$. Here, note that the hardness of the one-way function f ensures that the adversary running inside the reduction cannot create a proof for the latter case. In the blindness proof, the reduction, running in super-polynomial time $T' \cdot \mathsf{poly}(\lambda)$, extracts a value a from y such that $f(a) = y$, which allows us to create a NIWI proof proving the latter case.

However, this strategy presents a challenge. Suppose the one-way function f is super-polynomial-time $T \cdot \mathsf{poly}(\lambda)$ secure, and the commitment scheme Com' is super-polynomial $T' \cdot \mathsf{poly}(\lambda)$-time secure. We observe a contradiction: On one hand, T and T' must satisfy $T > T'$, as the one-more unforgeability proof requires the reduction to extract a preimage for f from the commitment $\mathsf{com}^{(1)}$. On the other hand, T and T' must satisfy $T < T'$, as the blindness proof requires the reduction to invert $y = f(a)$ from the signer's public key to simulate a NIWI proof.

To address this conflicting requirement, we turn to the non-uniform setting, as observed by [20, 21, 32]. In the non-uniform setting, reduction algorithms are viewed as two-stage processes. In the pre-computation phase, the reduction algorithm takes the security parameter as input and, using unbounded computational power, generates an advice string of polynomial length. Then, in the online phase, the reduction algorithm uses this advice string, along with the problem instance, to solve the problem in polynomial time.

Revisiting the strategy under the assumption that $T < T'$[7], the one-more unforgeability holds since $T < T'$. Meanwhile, in the blindness game, the pre-computation phase of the non-uniform reduction algorithm inverts $y = f(a)$ and

[6] NIWI can be constructed from NIZK proofs and derandomization assumptions [6, 15], from bilinear pairings [25] and indistinguishablity obfuscation [8]. NIWI is a powerful tool for constructing several cryptographic protocols in the plain model (e.g., verifiable random functions [24] and ring signatures [2]) as NIWI does not require a trusted setup assumption.

[7] Indeed, the commitment scheme Com' only needs to be $\mathsf{poly}(\lambda)$-secure.

outputs a as an advice string. In the online phase, given the problem instance and the advice string a, the algorithm can generate a NIWI proof demonstrating knowledge of the preimage a, meaning that the reduction can simulates the NIWI proof under the assumption that $T < T'$ (not $T > T'$).

Avoiding complexity leveraging As discussed earlier, complexity leveraging exploits a gap between computational power of an adversary and that of the reduction. To create this gap, certain building blocks of the protocol must be secure against super-polynomial-time adversaries. This requirement necessitates large parameters, which negatively impact overall efficiency. To address this issue, Kalai and Khurana [31] replace complexity leveraging with classical and quantum assumptions. This approach eliminates the need for super-polynomial-time secure primitives, significantly improving efficiency. Indeed, Katsumata et al. [32] applies this idea to a round optimal blind signature in the plain model from standard (polynomial-time) assumptions.

Can we simply replace complexity leveraging with classical and quantum assumptions in the NIBS scheme discussed above? Unfortunately, this direct replacement does not work, as our protocol requires three security levels of underlying primitives (since $T < T'$) whereas Kalai and Khurana's approach [31] requires only two security levels to achieve the replacement. To overcome this obstacle, we introduce a weaker blindness notion that allows the NIBS scheme to incorporate Kalai and Khurana's idea [31]. Intuitively, this weaker definition forces adversaries to fix the signer's public key before receiving the recipient's public key. As a result, the scheme no longer requires $T < T'$, meaning that it suffices to set $T' = \mathsf{poly}(\lambda)$. This adjustment enables us to replace complexity leverainging with classical and quantum assumptions.

Separation Results ntuitively, recipient blindness ensures that a given message-signature pair cannot be linked to any particular recipient. Meanwhile, nonce blindness ensures that a message-signature pair cannot be linked to any presignatures issued to a particular recipient, allowing the signer to issue multiple presignatures per user.

It is evident that every NIBS scheme should satisfy Both blindness property. However, an important question arises: we have one question: Does satisfying one blindness property necessarily imply satisfaction of the other? In other words, does a NIBS scheme with recipient blindness always satisfy nonce blindness, and vice versa? We answer this question by showing for each blindness property, there exists NIBS scheme which satisfies one but not the other (Theorem 1 and Theorem 2). Specifically, constructing a scheme with nonce blindness but not recipient blindness is straightforward—if a signature includes the recipient's public key, it can always be linked to the recipient, but this does not affect nonce blindness. However, constructing a scheme that satisfies recipient blindness but not nonce blindness requires more ingenuity. Our key idea is intuitively to incorporate a (one-time) ciphertext c that encrypts the nonce, defined as $c := \mathsf{rsk}+\mathsf{nonce}$, into the signature. This allows adversaries in the nonce blindness

game to obtain two (one-time) ciphertexts, from which they can extract the nonce information.

Meanwhile, this adjustment does not compromise the recipient blindness of the scheme, since different recipients typically have distinct secret keys. As a result, the (one-time) ciphertexts perfectly conceal nonce information from adversaries in the recipient blindness game.

1.3 Related Work

Non-interactive blind signatures Hanzlik [27] introduced two NIBS schemes in the random oracle model: an efficient construction based on structure-preserving signatures on equivalence classes (SPS-EQ) [18] and a generic construction.

Hanzlik's efficient construction is pairing-based, meaning it does not provide post-quantum security. In contrast, Baldimtsi et al. introduced an efficient lattice-based construction [3]. Additionally, they examined the definitional framework proposed by Hanzlik [27] and introduced a stronger blindness definition for NIBS, allowing an adversary to issue Obtain queries, which grant oracle access to the recipient's secret key. Both Hanzlik's schemes and Baldimtsi et al.'s construction are provably secure in the random oracle model. Currently, there is no NIBS scheme in the plain model.

Round-optimal blind signatures A round-optimal blind signature is a two-move protocol in which the user and signer each send one message to the other. Although many round-optimal blind signatures have been proposed [1, 7, 13, 16, 20, 28, 32, 33, 36], all existing schemes rely on random oracles, setup assumptions, interactive assumptions, or complexity leveraging. Indeed, several impossibility results exist for constructing round-optimal blind signatures in the plain model [5, 17, 35]. In particular, Fischlin and Schröder discovered a surprising result: three-move schemes cannot be constructed from non-interactive assumptions under certain conditions [17].

Since NIBS can be regarded as a two-move protocol, one might expect these impossibility results to apply to NIBS as well. However, unlike standard blind signatures, in NIBS, the message depends on the nonce chosen by the signer and the recipient's secret key. Therefore, as discussed in [27], whether it is possible to construct NIBS in the plain model remains unclear and requires further investigation.

2 Preliminaries

The security parameter is $\lambda \in \mathbb{Z}$. All algorithms receive λ implicitly as input. We denote the first N natural numbers by $[N] := \{1, \cdots, N\}$. For a finite set S, we write $x \leftarrow S$ if x is sampled uniformly at random from S. For a probabilistic algorithm $\mathcal{A}$, we write $y \leftarrow \mathcal{A}(x)$ if y is output from $\mathcal{A}$ on input x with uniformly sampled initialized randomness. We write $y \leftarrow \mathcal{A}(x; \rho)$ to make the initialized randomness ρ explicit, and $y \in \mathcal{A}(x)$ means that y is a possible output of $\mathcal{A}(x)$.

We also write $\mathcal{A}^{\mathcal{O}}(x)$ if $\mathcal{A}$ can access $\mathcal{O}$. We say that a function $f : \mathbb{N} \to \mathbb{R}_+$ is negligible in its input n, denoted by $f(n) = \mathsf{negl}(n)$, if $f \in n^{-\omega(1)}$. We use PPT and QPT to mean classical probabilistic polynomial time and quantum polynomial time.

In this section, we formally define non-interactive blind signatures (NIBS). Definitions of other cryptographic primitives can be found in Appendix A.

2.1 Non-Interactive Blind Signature

We introduce NIBS schemes as proposed by Hanzlik [27].

Definition 1. *A non-interactive blind signature* $\mathsf{NIBS} = (\mathsf{KeyGen}, \mathsf{RKeyGen}, \mathsf{Issue}, \mathsf{Obtain}, \mathsf{Verify})$ *consists of the following PPT algorithms:*

$\mathsf{KeyGen}(\lambda)$**:** *On input a security parameter* λ*, it outputs a signer's key pair* $(\mathsf{pk}, \mathsf{sk})$*. We assume that* pk *defines a nonce space* $\mathcal{N}_{\mathsf{pk}}$ *implicitly.*

$\mathsf{RKeyGen}(\lambda)$**:** *On input a security parameter* λ*, it outputs a recipient's key pair* $(\mathsf{rpk}, \mathsf{rsk})$*.*

$\mathsf{Issue}(\mathsf{pk}, \mathsf{sk}, \mathsf{rpk}, \mathsf{nonce})$**:** *On input a signer's key pair* $(\mathsf{pk}, \mathsf{sk})$ *a recipient's public key* rpk*, and a nonce* nonce*, it outputs a presignature* psig*.*

$\mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}, \mathsf{psig})$**:** *On input a signer's public key, a recipient's key pair* $(\mathsf{rpk}, \mathsf{rsk})$*, a nonce* nonce*, and a presignature* psig*, it outputs a valid message-signature pair* $(\mathsf{m}, \mathsf{sig})$ *or* $\perp$*.*

$\mathsf{Verify}(\mathsf{pk}, \mathsf{m}, \mathsf{sig})$**:** *On input a signer's public key* pk *and a message-signature pair* $(\mathsf{m}, \mathsf{sig})$*, it outputs a bit* b*.*

We require that NIBS *meets the following three properties: completeness, one-more unforgeability, and blindness:*

Completeness. *We say that* NIBS *satisfies completeness if for all* $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{KeyGen}(\lambda)$*, all* $(\mathsf{rpk}, \mathsf{rsk}) \in \mathsf{RKeyGen}(\lambda)$ *and all* $\mathsf{nonce} \in \mathcal{N}_{\mathsf{pk}}$*, it holds that*

$$\Pr\left[\mathsf{Verify}(\mathsf{pk}, \mathsf{m}, \mathsf{sig}) = 1 \;\middle|\; \begin{array}{c} \mathsf{psig} \leftarrow \mathsf{Issue}(\mathsf{pk}, \mathsf{sk}, \mathsf{rpk}, \mathsf{nonce}); \\ (\mathsf{m}, \mathsf{sig}) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}, \mathsf{psig}) \end{array}\right] = 1$$

One-More Unforgeability. *For an algorithm* $\mathcal{A}$*, we consider the following game* ℓ-$\mathsf{OMUF}^{\mathcal{A}}_{\mathsf{NIBS}}(\lambda)$*:*

1. *The game samples key pair* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(\lambda)$*.*
2. *Let* $\mathcal{O}$ *be an oracle inputting a recipient key* rpk *and a nonce* nonce *outputting presignature* $\mathsf{psig} \leftarrow \mathsf{Issue}(\mathsf{pk}, \mathsf{sk}, \mathsf{rpk}, \mathsf{nonce})$*. The game runs*

$$((\mathsf{m}_1, \mathsf{sig}_1), (\mathsf{m}_2, \mathsf{sig}_2), \cdots, (\mathsf{m}_k, \mathsf{sig}_k)) \leftarrow \mathcal{A}^{\mathcal{O}}(\mathsf{pk}),$$

where $\mathcal{A}$ *can query* $\mathcal{O}$ *in an arbitrarily interleaved way and complete at most* $\ell = \ell(\lambda)$ *of the interactions with* $\mathcal{O}$*.*

3. *The game outputs* 1 *if for all* $i \in [k]$, m_i *are distinct,* $\mathcal{A}$ *complete at most* $k-1$ *interactions with* $\mathcal{O}$ *and it holds that* $\mathsf{Verify}(\mathsf{pk}, \mathsf{m}_i, \mathsf{sig}_i) = 1$*; otherwise, it outputs* 0.

We say that NIBS *is* ℓ*-one-more unforgeable if for all PPT algorithm* $\mathcal{A}$*, the following advantage is negligible:*

$$\Pr\left[\ell\text{-}\mathsf{OMUF}^{\mathcal{A}}_{\mathsf{NIBS}}(\lambda) = 1\right].$$

Moreover, we say that NIBS *is one-more unforgeable if it is* ℓ*-one-more unforgeable for* $\ell = poly(\lambda)$.

Blindness. *We define two blindness properties: recipient blindness and nonce blindness.*

Recipient Blindness. *For an algorithms* $\mathcal{S}_{\mathsf{rbnd}}$*, we consider the following game* $\mathsf{RBND}^{\mathcal{S}_{\mathsf{rbnd}}}_{\mathsf{NIBS}}(\lambda)$:

1. *The game flips random bit* $b \leftarrow \{0,1\}$ *and samples* $(\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1) \leftarrow \mathsf{RKeyGen}(\lambda)$*. It then runs*

$$(\mathsf{pk}, \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1, \mathsf{St}) \leftarrow \mathcal{S}_{\mathsf{rbnd}}(\mathsf{issue}, \mathsf{rpk}_0, \mathsf{rpk}_1).$$

2. *The game computes*

$$(\mathsf{m}_0, \mathsf{sig}_0) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}_0, \mathsf{rsk}_0, \mathsf{nonce}_0, \mathsf{psig}_0) \text{ and}$$
$$(\mathsf{m}_1, \mathsf{sig}_1) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}_1, \mathsf{rsk}_1, \mathsf{nonce}_1, \mathsf{psig}_1).$$

If $(\mathsf{m}_0, \mathsf{sig}_0) \neq \perp$ *or* $(\mathsf{m}_1, \mathsf{sig}_1) \neq \perp$ *then it sets* $(\mathsf{m}_0, \mathsf{sig}_0) := \perp$ *and* $(\mathsf{m}_1, \mathsf{sig}_1) := \perp$

3. *The game runs*

$$b' \leftarrow \mathcal{S}_{\mathsf{rbnd}}(\mathsf{guess}, \mathsf{St}, \mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b}).$$

It finally outputs 1 *if* $b = b'$*;otherwise, it outputs* 0.

We say that NIBS *satisfies recipient blindness if for all PPT algorithm* $\mathcal{S}_{\mathsf{rbnd}}$ *the following advantage is negligible:*

$$2\Pr[\mathsf{RBND}^{\mathcal{S}_{\mathsf{rbnd}}}_{\mathsf{NIBS}}(\lambda) = 1] - 1.$$

Nonce Blindness. *For an algorithm* $\mathcal{S}_{\mathsf{nbnd}}$*, we consider the following game* $\mathsf{NBND}^{\mathcal{S}_{\mathsf{nbnd}}}_{\mathsf{NIBS}}(\lambda)$:

1. *The game flips random bit* $b \leftarrow \{0,1\}$ *and samples* $(\mathsf{rpk}, \mathsf{rsk}) \leftarrow \mathsf{RKeyGen}(\lambda)$*. It then runs*

$$(\mathsf{pk}, \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1, \mathsf{St}) \leftarrow \mathcal{S}_{\mathsf{nbnd}}(\mathsf{issue}, \mathsf{rpk}).$$

2. *The game computes*

$$(\mathsf{m}_0, \mathsf{sig}_0) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}_0, \mathsf{psig}_0) \text{ and}$$
$$(\mathsf{m}_1, \mathsf{sig}_1) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}_1, \mathsf{psig}_1).$$

If $(\mathsf{m}_0, \mathsf{sig}_0) \neq \perp$ *or* $(\mathsf{m}_1, \mathsf{sig}_1) \neq \perp$*. then it sets* $(\mathsf{m}_0, \mathsf{sig}_0) := \perp$ *and* $(\mathsf{m}_1, \mathsf{sig}_1) := \perp$

3. The game runs

$$b' \leftarrow \mathcal{S}_{\mathsf{nbnd}}(\mathsf{guess}, \mathsf{St}, \mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b}).$$

It finally outputs 1 *if* $b = b'$*; otherwise, it outputs* 0.

We say that NIBS *satisfies nonce blindness if for all PPT algorithm* $\mathcal{S}_{\mathsf{nbnd}}$*, the following advantage is negligible:*

$$2\Pr[\mathsf{NBND}^{\mathcal{S}_{\mathsf{nbnd}}}_{\mathsf{NIBS}}(\lambda) = 1] - 1.$$

In this paper, we provide separation results between recipient blindness and nonce blindness. Specifically, Theorem 1 demonstrates the existence of a NIBS scheme that satisfies nonce blindness but not recipient blindness, while Theorem 2 shows the existence of a scheme that satisfies recipient blindness but not nonce blindness. The full proofs can be found in Appendix B.

Theorem 1. *If there exists a* NIBS *scheme which satisfies nonce blindness and one-more unforgeability, then there exists a scheme which satisfies these same properties but which does not satisfies recipient blindness.*

Theorem 2. *If there exists a* NIBS *scheme which satisfies recipient blindness and one-more unforgeability, then there exists a scheme which satisfies these same properties but which does not satisfies nonce blindness.*

3 Our Construction

In this section, we provide a NIBS scheme in the plain model from complexity leveraging. Fix two super-polynomial functions T_1 and T_2 such that $T_1 < T_2$. Our construction relies on the following building blocks:

- $f : \{0,1\}^\lambda \to \{0,1\}^\lambda$ is an one-way permutation with T_1-one-wayness and T_2-invertible
- $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ is a digital signature scheme with T_1-EUF-CMA security.
- Com is a non-interactive perfect binding commitment with T_2-hiding.
- Com' is a non-interactive perfect binding commitment with hiding against non-uniform PPT adversaries and T_1-extractability.
- $F_R : \{0,1\}^\lambda \to \{0,1\}^\lambda$ is a pseudorandom function with pseudorandomness against non-uniform PPT adversaries.
- $\mathsf{NIWI} = (\mathsf{Prove}, \mathsf{Verify})$ is a non-interactive witness-indistinguishability proof system for the language $\mathcal{L}$ with witness-indistinguishability against non-uniform PPT adversaries.

Let us define the following NP language $\mathcal{L}$:

$$\mathcal{L} = \{(\mathsf{m}, \mathsf{pk} = (\mathsf{spk}, y), \mathsf{com}^{(0)}, \mathsf{com}^{(1)}) \mid (\mathsf{m}, \mathsf{spk}, \mathsf{com}^{(0)}) \in \mathcal{L}_0 \vee (y, \mathsf{com}^{(1)}) \in \mathcal{L}_1\},$$

where

$$\mathcal{L}_0 := \{(\mathsf{m}, \mathsf{spk}, \mathsf{com}^{(0)}) \,|\, \exists(\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig}, R, r_{\mathsf{rpk}}, r^{(0)}),\ \mathsf{SIG.Verify}(\mathsf{spk}, (\mathsf{nonce}, \mathsf{rpk}), \mathsf{psig}) = 1 \\ \wedge\, \mathsf{rpk} = \mathsf{Com}(R; r_{\mathsf{rpk}}) \\ \wedge\, \mathsf{m} = F_R(\mathsf{nonce}) \\ \wedge\, \mathsf{com}^{(0)} = \mathsf{Com}((\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig}); r^{(0)})\}$$

and

$$\mathcal{L}_1 := \left\{(y, \mathsf{com}^{(1)}) \;\middle|\; \exists(a, r^{(1)}),\ y = f(a) \wedge \mathsf{com}^{(1)} = \mathsf{Com}(a; r^{(1)})\right\}.$$

Then our scheme $\mathsf{NIBS} = (\mathsf{KeyGen}, \mathsf{RKeyGen}, \mathsf{Issue}, \mathsf{Obtain}, \mathsf{Verify})$ is described as follows:

$\mathsf{KeyGen}(\lambda)$**:**

1. $(\mathsf{spk}, \mathsf{ssk}) \leftarrow \mathsf{SIG.Gen}(1^\lambda)$;
2. $y \leftarrow \{0,1\}^\lambda$;
3. $\mathsf{pk} := (\mathsf{spk}, y)$; $\mathsf{sk} := (\mathsf{spk}, \mathsf{ssk})$;
4. Output $(\mathsf{pk}, \mathsf{sk})$;

$\mathsf{RKeyGen}(\lambda)$**:**

1. $R \leftarrow \{0,1\}^\lambda$; $r_{\mathsf{rpk}} \leftarrow \{0,1\}^\lambda$;
2. $\mathsf{rpk} \leftarrow \mathsf{Com}(R; r_{\mathsf{rpk}})$; $\mathsf{rsk} := (R, r_{\mathsf{rpk}})$;
3. Output $(\mathsf{rpk}, \mathsf{rsk})$;

$\mathsf{Issue}(\mathsf{sk}, \mathsf{rpk}, \mathsf{nonce})$**:**

1. $(\mathsf{spk}, \mathsf{ssk}) := \mathsf{sk}$;
2. $\mathsf{psig} \leftarrow \mathsf{SIG.Sign}(\mathsf{ssk}, (\mathsf{nonce}, \mathsf{rpk}))$;
3. Output psig;

$\mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}, \mathsf{psig})$**:**

1. $(R, r_{\mathsf{rpk}}) := \mathsf{rsk}$;
2. If $\mathsf{SIG.Verify}(\mathsf{spk}, (\mathsf{nonce}, \mathsf{rpk}), \mathsf{psig}) = 0$, then output $\perp$;
3. $\mathsf{m} \leftarrow F_R(\mathsf{nonce})$; $r^{(0)}, r^{(1)} \leftarrow \{0,1\}^\lambda$;
4. $\mathsf{com}^{(0)} \leftarrow \mathsf{Com}'((\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig}); r^{(0)})$; $\mathsf{com}^{(1)} \leftarrow \mathsf{Com}'(0; r^{(1)})$;
5. $x := (\mathsf{m}, \mathsf{pk}, \mathsf{com}^{(0)}, \mathsf{com}^{(1)})$; $w := (\mathsf{nonce}, \mathsf{rpk}, \mathsf{psig}, R, r_{\mathsf{rpk}}, r^{(0)})$;
6. $\pi \leftarrow \mathsf{NIWI.Prove}(x, w)$;
7. $\mathsf{sig} \leftarrow (\mathsf{com}^{(0)}, \mathsf{com}^{(1)}, \pi)$;
8. Output message-signature pair $(\mathsf{m}, \mathsf{sig})$;

$\mathsf{Verify}(\mathsf{pk}, (\mathsf{m}, \mathsf{sig}))$**:**

1. $(\mathsf{com}^{(0)}, \mathsf{com}^{(1)}, \pi) := \mathsf{sig}$;
2. $x := (\mathsf{m}, \mathsf{pk}, \mathsf{com}^{(0)}, \mathsf{com}^{(1)})$;
3. Output $\mathsf{NIWI.Verify}(x, \pi)$.

3.1 Security

We show that our construction satisfies one-more unforgeability, recipient blindness, and nonce blindness. Within all proofs, we assume that $T_1 < T_2$. The full proof can be found in Appendix C.

Theorem 3. *Our scheme satisfies one-more-unforgeability assuming that the signature scheme* SIG *satisfies* T_1-EUF-CMA *security, one-way permutation* f *satisfies* T_1*-one-wayness, non-interactive witness indistinguishable proof system* NIWI *satisfies perfect soundness, and the perfect binding commitment scheme* Com' *is* T_1-extractable.

Theorem 4. *Our scheme satisfies recipient blindness assuming the commitment scheme* Com *satisfies* T_2-hiding *and the non-interactive indistinguishable proof system* NIWI *satisfies witness-indistinguishability against non-uniform PPT adversaries, the one-way permutation* f *satisfies* T_2-invertible *and* F *is pseudorandom against non-uniform PPT adversaries.*

Theorem 5. *Our scheme satisfies nonce blindness assuming the commitment scheme* Com *satisfies* T_2-hiding *and the non-interactive indistinguishable proof system* NIWI *satisfies witness-indistinguishability against non-uniform PPT adversaries, the one-way permutation* f *satisfies* T_2-invertible *and* F *is pseudorandom against non-uniform PPT adversaries.*

4 Observations

Constructing NIBS from Classical and Quantum Hardness Assumptions Complexity leveraging requires a gap between the computational power of the adversary and the reduction algorithm in security proofs, necessitating large parameters for the building blocks. These large parameters reduce overall efficiency. To address this issue, Kalai and Khurana [31] suggest replacing complexity leveraging with quantum supremacy, converting super-polynomial hardness assumptions into quantum polynomial hardness. This approach allows the building blocks to avoid large parameters. Using this idea, Katsumata et al. [32] introduced a round-optimal blind signature in the plain model. This naturally raises the question:

Can the super-polynomial hardness assumptions in our construction be replaced with quantum polynomial hardness assumptions?

Unfortunately, this approach is not feasible in our case. Our security proof relies on complexity leveraging at two distinct points—T_1-hardness and T_2-hardness—requiring three levels of security for the underlying primitives. In contrast, combining classical and quantum polynomial hardness can provide only two security levels.

While prior work [32] encountered a similar obstacle, they skillfully avoided it using non-uniform reductions. Could this idea apply to our construction? Unfortunately, it does not. In the NIBS blindness game, adversaries receive the recipient's public key from the game before outputting any messages. This means that

non-uniform reductions cannot extract meaningful information during the pre-computation phase, as the recipient's public key, a problem instance, is required in the pre-computation phase (refer to **Claim7** in the proof of Theorem 4). In contrast, in the blindness game for standard blind signatures, the adversary must first output the signer's public key before receiving any messages, allowing non-uniform reductions to extract useful information in the pre-computation phase without needing the problem instance. This distinction allows the construction in [32] to achieve security in the plain model.

We now consider how Hanzlik's blindness definition needs to be weakened to achieve security in the plain model from classical and quantum hardness assumptions. The main obstacle in the case of NIBS is that the adversary in the blindness game receives the recipient's public key before outputting the signer's public key. A natural solution is to modify the game such that the adversary must first output the signer's public key before receiving the recipient's public key. Based on this idea, we introduce a weaker blindness definition called non-adaptive blindness, with the formal definition provided in Appendix D. This non-adaptive blindness property allows us to extract useful information from the signer's public key during the pre-computation phase (as referenced in **Claim7** in the non-adaptive recipient blindness proof of Theorem 7), enabling us to overcome the challenge above.

4.1 Construction

We present a NIBS scheme with non-adaptive blindness from classical and quantum hardness assumptions. The construction remains the same as the one introduced in Section 3, except for the underlying primitive, which is defined as follows:

- $f : \{0,1\}^\lambda \to \{0,1\}^\lambda$ is an one-way permutation against QPT adversaries.
- $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Verify})$ is a digital signature scheme with EUF-CMA security against QPT adversaries.
- Com is a non-interactive perfect binding commitment with computational hiding against non-uniform PPT adversaries.
- Com' is a non-interactive perfect binding commitment with computationally hiding against non-uniform PPT adversaries and QPT-extractability.
- $F_R : \{0,1\}^\lambda \to \{0,1\}^\lambda$ is a pseudorandom function with pseudorandomness against non-uniform PPT adversaries.
- $\mathsf{NIWI} = (\mathsf{Prove}, \mathsf{Verify})$ is a non-interactive witness-indistinguishability proof system for the language $\mathcal{L}$ with witness-indistinguishability against non-uniform PPT adversaries.

Security The full proof can be found in Appendix D.

Theorem 6. *Our scheme satisfies one-more-unforgeability assuming that the signature scheme* SIG *satisfies* EUF-CMA *against QPT adversaries, one-way permutation f satisfies one-wayness against QPT adversaries, non-interactive witness indistinguishable proof system* NIWI *satisfies perfect soundness, and the perfect binding commitment scheme* Com' *is* QPT-extractable.

Theorem 7. *Our scheme satisfies non-adaptive blindness assuming the commitment scheme* Com *satisfies computational hiding and the non-interactive indistinguishable proof system* NIWI *satisfies witness-indistinguishability against non-uniform PPT adversaries, the function f is a permutation and F is pseudorandom against non-uniform PPT adversaries.*

5 Conclusions and Open Problems

In this paper, we introduced a non-interactive blind signature (NIBS) scheme in the plain model. The core of our construction relies on complexity leveraging and non-uniform reductions [21], which allow us to extract useful information from adversaries without the need for a trusted setup. Additionally, we explored whether our construction could avoid complexity leveraging by utilizing both classical and quantum assumptions [31, 32]. To achieve this, we introduced a weaker blindness notion, called non-adaptive blindness, which enables our construction to bypass complexity leveraging.

Of independent interest, we provide separation results that demonstrate the existence of a NIBS construction satisfying nonce blindness but not recipient blindness, and vice versa. This result suggests that any NIBS construction should be proved to satisfy both nonce blindness and recipient blindness.

We leave open the question of whether it is possible to construct NIBS schemes with basic (not weaker) blindness, as introduced by Hanzlik [27], in the plain model using classical and quantum assumptions. Unlike standard blind signatures, the basic blindness of NIBS schemes requires adversaries to output their own signer's public key before receiving the recipient's public key, which prevents the approach used for constructing round-optimal blind signatures in the plain model [32] from being applied to NIBS. Overcoming this obstacle seems to be a significant challenge. We also leave as an open problem whether NIBS schemes with the stronger blindness introduced in [3] can be constructed in the plain model using complexity leveraging.

References

1. Agrawal, S., Kirshanova, E., Stehlé, D., Yadav, A.: Practical, round-optimal lattice-based blind signatures. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022. pp. 39–53. ACM (2022)
2. Backes, M., Döttling, N., Hanzlik, L., Kluczniak, K., Schneider, J.: Ring signatures: Logarithmic-size, no setup - from standard assumptions. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III. Lecture Notes in Computer Science, vol. 11478, pp. 281–311. Springer (2019)

3. Baldimtsi, F., Cheng, J., Goyal, R., Yadav, A.: Non-interactive blind signatures: Post-quantum and stronger security. In: Chung, K., Sasaki, Y. (eds.) Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part II. Lecture Notes in Computer Science, vol. 15485, pp. 70–104. Springer (2024)
4. Baldimtsi, F., Lysyanskaya, A.: Anonymous credentials light. In: Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security. pp. 1087–1098 (2013)
5. Baldimtsi, F., Lysyanskaya, A.: On the security of one-witness blind signature schemes. In: Sako, K., Sarkar, P. (eds.) Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part II. Lecture Notes in Computer Science, vol. 8270, pp. 82–99. Springer (2013)
6. Barak, B., Ong, S.J., Vadhan, S.P.: Derandomization in cryptography. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings. Lecture Notes in Computer Science, vol. 2729, pp. 299–315. Springer (2003)
7. Beullens, W., Lyubashevsky, V., Nguyen, N.K., Seiler, G.: Lattice-based blind signatures: Short, efficient, and round-optimal. In: Meng, W., Jensen, C.D., Cremers, C., Kirda, E. (eds.) Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS 2023, Copenhagen, Denmark, November 26-30, 2023. pp. 16–29. ACM (2023)
8. Bitansky, N., Paneth, O.: Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In: Dodis, Y., Nielsen, J.B. (eds.) Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9015, pp. 401–427. Springer (2015)
9. Branco, P., Döttling, N., Wohnig, S.: Universal ring signatures in the standard model. IACR Cryptol. ePrint Arch. p. 1265 (2022)
10. Canard, S., Gaud, M., Traoré, J.: Defeating malicious servers in a blind signatures based voting system. In: International Conference on Financial Cryptography and Data Security. pp. 148–153. Springer (2006)
11. Chardouvelis, O., Malavolta, G.: The round complexity of quantum zero-knowledge. In: Nissim, K., Waters, B. (eds.) Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part I. Lecture Notes in Computer Science, vol. 13042, pp. 121–148. Springer (2021)
12. Chaum, D.: Blind signatures for untraceable payments. In: CRYPTO. pp. 199–203 (1982)
13. Chaum, D.: Blind signature system. In: Advances in Cryptology: Proceedings of Crypto 83. pp. 153–153. Springer (1983)
14. Davidson, A., Goldberg, I., Sullivan, N., Tankersley, G., Valsorda, F.: Privacy pass: Bypassing internet challenges anonymously. Proceedings on Privacy Enhancing Technologies (2018)
15. Dwork, C., Naor, M.: Zaps and their applications. In: 41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA. pp. 283–293. IEEE Computer Society (2000)
16. Fischlin, M.: Round-optimal composable blind signatures in the common reference string model. In: Dwork, C. (ed.) Advances in Cryptology - CRYPTO 2006, 26th

Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4117, pp. 60–77. Springer (2006)
17. Fischlin, M., Schröder, D.: On the impossibility of three-move blind signature schemes. In: Gilbert, H. (ed.) Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6110, pp. 197–215. Springer (2010)
18. Fuchsbauer, G., Hanser, C., Slamanig, D.: Structure-preserving signatures on equivalence classes and constant-size anonymous credentials. J. Cryptol. **32**(2), 498–546 (2019)
19. Fuchsbauer, G., Konstantinov, M., Pietrzak, K., Rao, V.: Adaptive security of constrained prfs. In: Sarkar, P., Iwata, T. (eds.) Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II. Lecture Notes in Computer Science, vol. 8874, pp. 82–101. Springer (2014)
20. Garg, S., Gupta, D.: Efficient round optimal blind signatures. In: Nguyen, P.Q., Oswald, E. (eds.) Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8441, pp. 477–495. Springer (2014)
21. Garg, S., Rao, V., Sahai, A., Schröder, D., Unruh, D.: Round optimal blind signatures. In: Rogaway, P. (ed.) Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6841, pp. 630–648. Springer (2011)
22. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: Johnson, D.S. (ed.) Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA. pp. 25–32. ACM (1989)
23. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. J. Cryptol. **7**(1), 1–32 (1994)
24. Goyal, R., Hohenberger, S., Koppula, V., Waters, B.: A generic approach to constructing and proving verifiable random functions. In: Kalai, Y., Reyzin, L. (eds.) Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II. Lecture Notes in Computer Science, vol. 10678, pp. 537–566. Springer (2017)
25. Groth, J., Ostrovsky, R., Sahai, A.: Non-interactive zaps and new techniques for NIZK. In: Dwork, C. (ed.) Advances in Cryptology - CRYPTO 2006, 26th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4117, pp. 97–111. Springer (2006)
26. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Smart, N.P. (ed.) Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings. Lecture Notes in Computer Science, vol. 4965, pp. 415–432. Springer (2008)
27. Hanzlik, L.: Non-interactive blind signatures for random messages. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques,

Lyon, France, April 23-27, 2023, Proceedings, Part V. Lecture Notes in Computer Science, vol. 14008, pp. 722–752. Springer (2023)
28. Hanzlik, L., Loss, J., Wagner, B.: Rai-choo! evolving blind signatures to the next level. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 753–783. Springer (2023)
29. Heilman, E., Alshenibr, L., Baldimtsi, F., Scafuro, A., Goldberg, S.: Tumblebit: An untrusted bitcoin-compatible anonymous payment hub. In: Network and distributed system security symposium (2017)
30. IETF: Privacy Pass (privacypass), `https://datatracker.ietf.org/wg/privacypass/documents/`
31. Kalai, Y.T., Khurana, D.: Non-interactive non-malleability from quantum supremacy. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part III. Lecture Notes in Computer Science, vol. 11694, pp. 552–582. Springer (2019)
32. Katsumata, S., Nishimaki, R., Yamada, S., Yamakawa, T.: Round-optimal blind signatures in the plain model from classical and quantum standard assumptions. In: Canteaut, A., Standaert, F. (eds.) Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12696, pp. 404–434. Springer (2021)
33. Katsumata, S., Reichle, M., Sakai, Y.: Practical round-optimal blind signatures in the ROM from standard assumptions. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology - ASIACRYPT 2023 - 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4-8, 2023, Proceedings, Part II. Lecture Notes in Computer Science, vol. 14439, pp. 383–417. Springer (2023)
34. Pass, R.: Simulation in quasi-polynomial time, and its application to protocol composition. In: Biham, E. (ed.) Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings. Lecture Notes in Computer Science, vol. 2656, pp. 160–176. Springer (2003)
35. Pass, R.: Limits of provable security from standard assumptions. In: Fortnow, L., Vadhan, S.P. (eds.) Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011. pp. 109–118. ACM (2011)
36. del Pino, R., Katsumata, S.: A new framework for more efficient round-optimal lattice-based (partially) blind signature via trapdoor sampling. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II. Lecture Notes in Computer Science, vol. 13508, pp. 306–336. Springer (2022)

Appendix

A Formal Definitions of Cryptographic Primitives

A.1 One-Way Permutations

Definition 2. *We say that a permutation* $f : \{0,1\}^\lambda \to \{0,1\}^\lambda$ *is* one-way *if for all non-uniform PPT adversaries, the following advantage* $\mathbf{Adv}^{\mathsf{ow}}_{f,\mathcal{A}}(\lambda)$ *is negligible in* λ*:*

$$\Pr\left[f(x') = y \,|\, x \leftarrow \{0,1\}^\lambda; x' \leftarrow \mathcal{A}(f(x))\right] = \mathsf{negl}(\lambda).$$

Further, we say that f *is* T-one-way *if it is* one-way *against non-uniform adversaries running in time* $T \cdot \mathsf{poly}(\lambda)$*.*

Moreover, we require the permutation to be invertible after a certain time.

Definition 3. *A permutation* f *is* T-invertible *if there exists an PPT algorithm* $\mathcal{A}$ *running in time* $T \cdot \mathsf{poly}(\lambda)$ *such that for any* $x \in \{0,1\}^\lambda$*, it holds that*

$$\Pr[x = x' \,|\, x' \leftarrow \mathcal{A}(f(x))] = 1 - \mathsf{negl}(\lambda).$$

A.2 Pseudorandom Functions

Definition 4. *We say that a function* $F_R : \{0,1\}^\lambda \to \{0,1\}^\lambda$ *is* pseudorandom *if for all non-uniform PPT adversaries* $\mathcal{A}$*, the following advantage* $\mathbf{Adv}^{\mathsf{prf}}_{F,\mathcal{A}}(\lambda)$ *is negligible in* λ*:*

$$\left|\Pr[\mathcal{A}^{F_R(\cdot)}(\lambda) = 1] - \Pr[\mathcal{A}^{f(\cdot)}(\lambda) = 1]\right| = \mathsf{negl}(\lambda),$$

where $R \leftarrow \{0,1\}^\lambda$ *and* f *is chosen uniformly at random from the set of functions mapping* λ*-bit strings to* λ*-bit strings.*
Further, we say that F *is* T-invertible *if it is* inveritble *against adversaries running in time* $T \cdot \mathsf{poly}(\lambda)$*.*

A.3 Digital Signature Scheme

Definition 5. *A signature scheme* SIG *consists of the following PPT algorithms:*

$\mathsf{Gen}(\lambda)$**:** *Takes as input the security parameter* λ *and outputs a key pair of verification and signing keys* $(\mathsf{spk}, \mathsf{ssk})$*.*

$\mathsf{Sign}(\mathsf{ssk}, \mathsf{m})$**:** *Takes as input a signing key* ssk *and a message* m *and outputs a signature* σ*.*

$\mathsf{Verify}(\mathsf{spk}, \mathsf{m}, \sigma)$**:** *Takes as input a verification key* spk*, a message* m*, and a signature* σ *and outputs either* 0 *or* 1*.*

We require the following properties of a signature scheme.

Correctness: If for every security parameter κ, every message m, it hold that

$$\Pr[\mathsf{Verify}(\mathsf{spk}, \mathsf{m}, \sigma) = 1 \,|\, (\mathsf{spk}, \mathsf{ssk}) \leftarrow \mathsf{Gen}(\lambda); \sigma \leftarrow \mathsf{Sign}(\mathsf{ssk}, \mathsf{m})] = 1$$

EUF-CMA security: We say that a signature scheme SIG satisfies EUF-CMA security if for all non-uniform PPT (or QPT) adversaries $\mathcal{A}$, the advantage in the following game $\mathsf{EUF}^{\mathsf{CMA}}_{\mathsf{SIG},\mathcal{A}}(\lambda)$ is negligible in λ:

1. The game samples a pair of verification and signing keys $(\mathsf{spk}, \mathsf{ssk}) \leftarrow \mathsf{Gen}(\lambda)$ and provides spk to $\mathcal{A}$.
2. $\mathcal{A}$ is allowed to make signing queries of the message m, upon which the game returns $\sigma \leftarrow \mathsf{Sign}(\mathsf{ssk}, \mathsf{m})$.
3. Once $\mathcal{A}$ outputs as a forgery (m^*, σ^*), the game checks if m^* was not queried in a signing query and if it holds that $\mathsf{Verify}(\mathsf{spk}, \mathsf{m}, \sigma) = 1$. If so, it outputs 1, otherwise 0.

The advantage of $\mathcal{A}$ is defined as:

$$\mathbf{Adv}^{\mathsf{euf\text{-}cma}}_{\mathsf{SIG},\mathcal{A}}(\lambda) = \Pr[\mathsf{EUF}^{\mathsf{CMA}}_{\mathsf{SIG},\mathcal{A}}(\lambda) = 1].$$

Further, we say that SIG is T-EUF-CMA secure if it satisfies EUF-CMA security against adversaries running in time $T \cdot \mathsf{poly}(\lambda)$.

A.4 Non-interactive Commitment Schemes

Definition 6. *A non-interactive commitment scheme* Com *takes as input a message* $\mathsf{m} \in \{0,1\}^\lambda$ *and a randomness* $\mathsf{r} \in \{0,1\}^\lambda$*, outputting a commitment* com*; i.e.,* $\mathsf{com} \leftarrow \mathsf{Com}(\mathsf{m}; \mathsf{r})$.

We require that Com *meets the following properties.*

Perfect Binding: *There exists no values* $(\mathsf{m}, \mathsf{r}, \mathsf{m}', \mathsf{r}')$ *such that* $\mathsf{m} \neq \mathsf{m}'$ *and* $\mathsf{Com}(\mathsf{m}; \mathsf{r}) = \mathsf{Com}(\mathsf{m}'; \mathsf{r}')$

Computational Hiding: *For all non-uniform PPT (or QPT) adversaries* $\mathcal{A}$*, the following advantage* $\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com},\mathcal{A}}(\lambda)$ *is negligible in* λ*:*

$$2\Pr\left[b = b' \;\middle|\; \begin{array}{c} (\mathsf{m}_0, \mathsf{m}_1) \leftarrow \mathcal{A}(\lambda);\ b \leftarrow \{0,1\}; \\ \mathsf{r} \leftarrow \{0,1\}^\lambda;\ \mathsf{com} \leftarrow \mathsf{Com}(\mathsf{m}_b; \mathsf{r}) \\ b' \leftarrow \mathcal{A}(\mathsf{com}) \end{array}\right] - 1 = \mathsf{negl}(\lambda).$$

Further, we say that Com *satisfies* T-hiding *if it satisfies computational hiding against adversaries running in time* $T \cdot \mathsf{poly}(\lambda)$.

Moreover, we require the commitments to be extractable after a certain time. i.e., there exists an algorithm that takes as input a commitment and outputs the contained message. More formally:

Definition 7. *We say that a commitment scheme* Com *is* $T-\mathsf{extractable}$ *if there exists an PPT algorithm* $\mathcal{A}$ *running in time* $T \cdot \mathsf{poly}(\lambda)$ *such that for any message* $\mathsf{m} \in \{0,1\}^\lambda$ *and randomness* $\mathsf{r} \in \{0,1\}^\lambda$*,*

$$\Pr[\mathsf{m} = \mathsf{m}' \,|\, \mathsf{com} \leftarrow \mathsf{Com}(\mathsf{m}; \mathsf{r}); \mathsf{m}' \leftarrow \mathcal{A}(\mathsf{com})] = 1 - \mathsf{negl}(\lambda).$$

Additionally, we say that Com *is* $\mathsf{QPT\text{-}extractable}$ *if there exists such a QPT algorithm* $\mathcal{A}$.

Non-interactive prefect binding commitment schemes can be constructed from any injective one-way function via the Goldreich-Levin hardcore bit [22].

A.5 Non-Interactive Witness-Indistinguishability Proof Systems

Let $\mathcal{L}$ be an NP language with a relation R, where for $(x, w) \in R$ we call x a statement and w a witness, i.e., $\mathcal{L} = \{x \mid \exists w : (x, w) \in R\}$.

Definition 8. *Let $\mathcal{L}_{\mathcal{R}}$ be an NP language with associated relation $\mathcal{R}$. A non-interactive witness-indistinguishable proof system* NIWI *for language $\mathcal{L}_{\mathcal{R}}$ consists of the following PPT algorithm:*

$\mathsf{Prove}(\lambda, x, w)$**:** *takes as input a security parameter λ, a statement x, and a witness w and outputs a proof π.*

$\mathsf{Verify}(x, \pi)$**:** *takes as input a statement x and a proof π and outputs a bit b.*

We require the following properties of NIWI*.*

Perfect Completeness: *For all security parameter λ, all statement $x \in \mathcal{L}_{\mathcal{R}}$, and all statement-witness pair $(x, w) \in \mathcal{R}$, it holds that* $\mathsf{Verify}(x, \mathsf{Prove}(\lambda, x, w)) = 1$.

Perfect Soundness: *For all security parameter λ, all statement $x \notin \mathcal{L}_{\mathcal{R}}$, all proof π, it hold that* $\mathsf{Verify}(x, \pi) = 0$.

Witness Indistinguishability: *For any non-uniform PPT adversary $\mathcal{A}$, the advantage in the following game $\mathsf{WI}_{\mathsf{NIWI},\mathcal{A}}(\lambda)$ is negligible in λ :*

1. *$\mathcal{A}$ outputs (x, w_0, w_1) with $(x, w_0) \in \mathcal{R}$ and $(x, w_1) \in \mathcal{R}$ to the game.*
2. *The game select a random bit $b \leftarrow \{0, 1\}$, computes a proof $\pi_b \leftarrow \mathsf{Prove}(\lambda, x, w_b)$, and provides π_b to $\mathcal{A}$.*
3. *$\mathcal{A}$ outputs a guess b'. If $b' = b$, then the game outputs 1, otherwise 0.*

The advantage of $\mathcal{A}$ is defined as

$$\mathbf{Adv}^{\mathsf{wi}}_{\mathsf{NIWI},\mathcal{A}}(\lambda) = 2\Pr[\mathsf{WI}_{\mathsf{NIWI},\mathcal{A}}(\lambda) = 1] - 1.$$

B Separations between nonce blindness and recipient blindness

Proof. (Theorem 1) Let $\Pi = (\mathsf{KeyGen}, \mathsf{RKeyGen}, \mathsf{Issue}, \mathsf{Obtain}, \mathsf{Verify})$ be a NIBS scheme that satisfies recipient blindness and one-more unforgeability. We construct the following scheme $\Pi' = (\mathsf{KeyGen}', \mathsf{RKeyGen}', \mathsf{Issue}', \mathsf{Obtain}', \mathsf{Verify}')$: the signer's key generation algorithm KeyGen', the recipient's key generation algorithm $\mathsf{RKeyGen}'$, the issue algorithm Issue' are identical to the counterpart KeyGen, $\mathsf{RKeyGen}$, Issue, respectively. However, the message-signature pair generation algorithm Obtain' additionally outputs the recipient's public key rpk along with the message-signature pair $(\mathsf{m}, \mathsf{sig}) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}, \mathsf{psig})$. The verification algorithm Verify' discards the recipient's public key rpk and outputs a bit $b \leftarrow \mathsf{Verify}(\mathsf{pk}, \mathsf{m}, \mathsf{sig})$.

Clearly, the above scheme satisfies one-more unforgeability. Moreover, it satisfies nonce blindness, as adversaries in the nonce blindness game can obtain two message-signature pairs from a single recipient, where the recipient public key is the same in both pairs. However, it clearly does not satisfy recipient blindness, as adversaries in the recipient blindness game can obtain two message-signature pairs from two different recipients, where the recipient public keys differ. □

Proof. (Theorem 2) Let $\Pi = (\mathsf{KeyGen}, \mathsf{RKeyGen}, \mathsf{Issue}, \mathsf{Obtain}, \mathsf{Verify})$ be a NIBS scheme that satisfies recipient blindness and one-more unforgeability. Let q be an odd prime number such that $1/q = \mathsf{negl}(\lambda)$. We then construct the following scheme $\Pi' = (\mathsf{KeyGen}', \mathsf{RKeyGen}', \mathsf{Issue}', \mathsf{Obtain}', \mathsf{Verify}')$:

$\mathsf{KeyGen}'(\lambda)$**:** It samples $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(\lambda)$. It then sets $(\mathsf{pk}', \mathsf{sk}') := (\mathsf{pk}, \mathsf{sk})$ and outputs a key pair $(\mathsf{pk}', \mathsf{sk}')$.

$\mathsf{RKeyGen}'(\lambda)$**:** It samples $s \leftarrow \mathbb{Z}_q$ and $(\mathsf{rpk}, \mathsf{rsk}) \leftarrow \mathsf{RKeyGen}(\lambda)$. It then sets $(\mathsf{rpk}', \mathsf{rsk}') := (\mathsf{rpk}, (\mathsf{rsk}, s))$ and outputs a key pair $(\mathsf{rpk}', \mathsf{rsk}')$.

$\mathsf{Issue}'(\mathsf{pk}', \mathsf{sk}', \mathsf{rpk}', \mathsf{nonce})$**:** Let $(\mathsf{pk}, \mathsf{sk}) := (\mathsf{pk}', \mathsf{sk}')$ and $\mathsf{rpk} := \mathsf{rpk}'$. It then outputs (psig, r) as a presignature, where $\mathsf{psig} \leftarrow \mathsf{Issue}(\mathsf{pk}, \mathsf{sk}, \mathsf{rpk}, \mathsf{nonce})$ and $r \leftarrow \mathbb{Z}_q$.

$\mathsf{Obtain}'(\mathsf{pk}', \mathsf{rpk}', \mathsf{rsk}', \mathsf{nonce}, \mathsf{psig}')$**:** Let $\mathsf{pk} := \mathsf{pk}'$, $(\mathsf{rpk}, (\mathsf{rsk}, s)) := (\mathsf{rpk}', \mathsf{rsk}')$, and $(\mathsf{psig}, r) := \mathsf{psig}'$. It computes $(\mathsf{m}, \mathsf{sig}) \leftarrow \mathsf{Obtain}'(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}, \mathsf{psig})$. It then outputs $(\mathsf{m}, (\mathsf{sig}, s + r \ (mod\, p)))$ as a message-signature pair $(\mathsf{m}', \mathsf{sig}')$.

$\mathsf{Verify}'(\mathsf{pk}', \mathsf{m}', \mathsf{sig}')$**:** Let $\mathsf{pk} := \mathsf{pk}'$ and $(\mathsf{m}, (\mathsf{sig}, c)) := (\mathsf{m}', \mathsf{sig}')$. It discards c and outputs a bit $b \leftarrow \mathsf{Verify}(\mathsf{pk}, \mathsf{m}, \mathsf{sig})$.

Clearly, the above scheme satisfies one-more unforgeability as Π and Π' are identical except that Issue' and Obtain' additionally outputs a random number from $\mathbb{Z}_q$, and Verify' discards the value c included in a signature. Furthermore, it satisfies recipient blindness as the ciphertexts adversaries obtain from the two recipients in the recipient blindness game are independently masked using the one-time pad. However, it does not satisfies nonce blindness as there exists an adversary $\mathcal{A}$ against the scheme Π' with the following steps:

1. Given a recipient key rpk, $\mathcal{A}$ honestly computes $(\mathsf{pk}', \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$, where psig_0 and psig_0 include randomness r_0 and r_1, respectively. It then outputs these to the game.
2. $\mathcal{A}$ receives two message signature pairs $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ where sig_b and sig_{1-b} contain ciphertexts c_b and c_{1-b}, respectively. If $c_{1-b} + r_0 = c_b + r_1$, then it outputs a bit $b' = 0$. Otherwise, it outputs a bit $b' = 1$.

We demonstrate that the adversary $\mathcal{A}$ succeeds with all but negligible probability. Assume that r_0 and r_1 are distinct without loss of generality. If $b = 0$, then:

$$c_1 + r_0 = r_1 + s + r_0 = r_1 + c_0,$$

which implies that $\mathcal{A}$ outputs $b' = 0$. Conversely, if $b = 1$, then:

$$c_0 + r_0 = r_0 + s + r_0 \neq r_1 + s + r_1 = c_1 + r_1,$$

which implies that $\mathcal{A}$ outputs $b' = 1$. □

C Full Proofs for Our Construction from Complexity Leveraging

Proof. (Theorem 3) Let $\mathcal{A}$ be a PPT adversary against the one-more unforgeability of NIBS. To prove the statement via a sequence of games, we denote by G_i the event that the adversary $\mathcal{A}$ wins in game Game_i.

Game_0**:** The original one-more unforgeability game.

Game_1**:** The same as Game_0, except that if, given $\mathsf{pk} = (\mathsf{spk}, y)$, the adversary $\mathcal{A}$ eventually produces ℓ forgeries, then G_1 extracts the contents from each ℓ commitments $\mathsf{com}_1^{(0)}, \cdots, \mathsf{com}_\ell^{(0)}$ from the forgeries. Furthermore, if there exists a value a within these commitments such that $y = f(a)$, the game outputs $\perp$.

We claim that:

$$|\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| = \mathsf{negl}(\lambda), \tag{1}$$

assuming that the one-way permutation f satisfies T_1-one-wayness and the commitment scheme Com' is T_1-extractable. For simplicity, we assume the existence of a T_1-time algorithm that can extract content from the commitment of Com' with probability 1. Let us denote by Abort the event that Game_1 aborts. Clearly, $\mathsf{G}_1 = \mathsf{G}_0 \cap \overline{\mathsf{Abort}}$, so we have:

$$|\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| \leq \Pr[\mathsf{Abort}].$$

We now aim to bound the probability that the event Abort occurs. Specifically, we show that there exists a reduction $\mathcal{R}$ running in time $T_1 \cdot \mathsf{poly}(\lambda)$ such that

$$\mathrm{Adv}^{\mathsf{ow}}_{\mathcal{R}}(\lambda) = \Pr[\mathsf{Abort}]. \tag{2}$$

Let $\mathcal{A}$ be an adversary in the game Game_1. The reduction $\mathcal{R}$ is as follows:

1. It samples a key pair $(\mathsf{spk}, \mathsf{ssk}) \leftarrow \mathsf{SIG.Gen}(\lambda)$ and sets $\mathsf{pk} := (\mathsf{spk}, y)$ and $\mathsf{sk} := (\mathsf{pk}, \mathsf{ssk})$.
2. It runs $\mathcal{A}$, which takes pk as input. To issue presignatures for $\mathcal{A}$'s queries, it uses the singing key sk.
3. Suppose $\mathcal{A}$ outputs ℓ forgeries:

$$((\mathsf{m}_1, (\mathsf{com}_1^{(0)}, \mathsf{com}_1^{(1)}, \pi_1)), \cdots, (\mathsf{m}_\ell, (\mathsf{com}_\ell^{(0)}, \mathsf{com}_\ell^{(1)}, \pi_\ell)).$$

The reduction extracts the content a_j from each commitment $\mathsf{com}_j^{(1)}$ (for all $j \in [\ell]$), as the commitment scheme Com' is T_1-extractable. If there exists $j \in [\ell]$ such that $y = f(a_j)$, then it outputs a_j; otherwise, it outputs $\perp$.

The event Abort occurs if and only if the reduction $\mathcal{R}$ outputs the preimage of y under f, implying that equation (2) holds.

Since we assume that the one-way permutation f satisfies T_1-one-wayness, equation (1) follows.

Next, we claim that the probability of the event Abort occurring is negligible in λ, assuming that the signature scheme SIG satisfies T_1-$\mathsf{EUF\text{-}CMA}$ security, the non-interactive witness indistinguishable proof system NIWI satisfies perfect soundness, and the commitment scheme Com is perfectly binding. For simplicity, we assume the existence of a T_1-time algorithm that can extract content from the commitment with probability 1. We now demonstrate that there exists an adversary $\mathcal{R}$ in Game_1, running in time T_1, such that:

$$\Pr[\mathsf{G}_1] = \mathrm{Adv}_{\mathcal{R}}^{\mathsf{euf\text{-}cma}}(\lambda). \tag{3}$$

Let $\mathcal{A}$ be an adversary in Game_1. The reduction $\mathcal{R}$ proceeds as follows:

1. Given a public key spk from the $\mathsf{EUF\text{-}CMA}$ game, $\mathcal{R}$ samples $y \leftarrow \{0,1\}^\lambda$ and sets $\mathsf{pk} := (\mathsf{spk}, y)$.
2. It runs $\mathcal{A}$, which takes pk as input. To create presignatures in response to $\mathcal{A}$'s queries, it utilizes an external signing oracle provided by the game. Specifically, when $\mathcal{A}$ supplies a nonce and recipient public key $(\mathsf{nonce}, \mathsf{rpk})$, $\mathcal{R}$ returns a presignature $\mathsf{psig} \leftarrow \mathcal{O}(\mathsf{nonce}, \mathsf{rpk})$, where $\mathcal{O}$ is the external signing oracle.
3. Suppose $\mathcal{A}$ produces ℓ forgeries:

$$(\mathsf{m}_1, (\mathsf{com}_1^{(0)}, \mathsf{com}_1^{(1)}, \pi_1)), \cdots, (\mathsf{m}_\ell, (\mathsf{com}_\ell^{(0)}, \mathsf{com}_\ell^{(1)}, \pi_\ell)).$$

   It extracts the contents $((\mathsf{nonce}_j, \mathsf{rpk}_j), \mathsf{psig}_j)$ from each $\mathsf{com}_j^{(1)}$ for all $j \in [\ell]$. If there exists a valid message-signature pair $((\mathsf{nonce}^*, \mathsf{rpk}^*), \mathsf{psig}^*)$ such that the message $(\mathsf{nonce}^*, \mathsf{rpk}^*)$ was not queried by $\mathcal{A}$, it outputs $((\mathsf{nonce}^*, \mathsf{rpk}^*), \mathsf{psig}^*)$ as a forgery in the $\mathsf{EUF\text{-}CMA}$ game; otherwise, it outputs $\perp$.

The simulation by $\mathcal{R}$ mimics Game_1 perfectly. Since the commitment scheme Com' is perfectly binding and the NIWI system satisfies perfect soundness, if the event G_1 occurs, i.e., $\mathcal{A}$ outputs valid ℓ forgeries:

$$(\mathsf{m}_1, (\mathsf{com}_1^{(0)}, \mathsf{com}_1^{(1)}, \pi_1)), \cdots, (\mathsf{m}_\ell, (\mathsf{com}_\ell^{(0)}, \mathsf{com}_\ell^{(1)}, \pi_\ell))$$

such that

$$\mathsf{m}_i \neq \mathsf{m}_j \ \wedge \ \mathsf{Verify}(\mathsf{m}_i, (\mathsf{com}_i^{(0)}, \mathsf{com}_i^{(1)}, \pi_i)) = 1$$

and

$$a_i \neq f^{-1}(y),$$

where a_i is the content extracted from the commitment $\mathsf{com}_i^{(1)}$, for all distinct $i, j \in [\ell]$, then:

$$\exists (R_i, r_{\mathsf{rpk}_i}, r_i^{(0)}), \ ((\mathsf{m}_i, \mathsf{spk}, \mathsf{com}_i^{(0)}), (\mathsf{nonce}_i, \mathsf{rpk}_i, \mathsf{psig}_i, R_i, r_{\mathsf{rpk}_i}, r_i^{(0)})) \in \mathcal{R}_{\mathcal{L}_0},$$

by the perfect soundness of NIWI.

We now show that $(\mathsf{nonce}_i, \mathsf{rpk}_i) \neq (\mathsf{nonce}_j, \mathsf{rpk}_j)$ for all distinct $i, j \in [\ell]$, relying on the perfect binding property of Com'. Assuming $\mathsf{rpk}_i = \mathsf{rpk}_j$ (if not, $(\mathsf{nonce}_i, \mathsf{rpk}_i) \neq (\mathsf{nonce}_j, \mathsf{rpk}_j)$ is trivially true), we observe:

$$[R_i = R_j \ \wedge \ F_{R_i}(\mathsf{nonce}_i) = \mathsf{m}_i \neq \mathsf{m}_j = F_{R_j}(\mathsf{nonce}_j)] \Rightarrow \mathsf{nonce}_i \neq \mathsf{nonce}_j.$$

Hence, if G_1 occurs, $\mathcal{R}$ can obtain at least one forgery for the EUF-CMA security game, implying that equation (3) holds.

Putting all together, we conclude:

$$\begin{aligned}\mathrm{Adv}^{\mathsf{omuf}}_{\mathsf{NIBS},\mathcal{A}}(\lambda) &\leq |\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| + \Pr[\mathsf{G}_1] \\ &= \mathsf{negl}(\lambda).\end{aligned}$$

□

Proof. (Theorem 4) We will prove this theorem by a series of hybrid arguments. Let $\mathcal{A}$ be a non-uniform PPT adversary against the recipient blindness of NIBS and let G_i denote the event where the adversary $\mathcal{A}$ wins in game Game_i.

Game_0**:** The original recipient blindness game.

Game_1**:** The same as Game_0, except $\mathsf{com}_0^{(1)} \leftarrow \mathsf{Com}'(a; r_0^{(1)})$ where $f(a) = y$.

Game_2**:** Same as Game_1, except $\mathsf{com}_1^{(1)} \leftarrow \mathsf{Com}'(a; r_1^{(1)})$, where $f(a) = y$.

Game_3**:** Same as Game_2, except when creating the NIWI proof π_0 in challenge signature sig_0, we use witness $w := (a, r_0^{(1)})$ instead of $(\mathsf{nonce}_0, \mathsf{rpk}_0, \mathsf{psig}_0, R_0, r_{\mathsf{rpk}_0}, r_0^{(0)})$.

Game_4**:** Same as Game_3, except when creating the NIWI proof π_1 in challenge signature sig_1, we use witness $w := (a, r_1^{(1)})$ instead of $(\mathsf{nonce}_1, \mathsf{rpk}_1, \mathsf{psig}_1, R_1, r_{\mathsf{rpk}_1}, r_1^{(0)})$.

Game_5**:** Same as Game_4, except $\mathsf{com}_0^{(0)} \leftarrow \mathsf{Com}'(0; r_0^{(0)})$ instead of $\mathsf{com}_0^{(0)} \leftarrow \mathsf{Com}'(\mathsf{nonce}_0, \mathsf{rpk}_0, \mathsf{psig}_0, r_0^{(0)})$.

Game_6**:** Same as Game_5, except $\mathsf{com}_1^{(0)} \leftarrow \mathsf{Com}'(0; r_1^{(0)})$ instead of $\mathsf{com}_1^{(0)} \leftarrow \mathsf{Com}'(\mathsf{nonce}_1, \mathsf{rpk}_1, \mathsf{psig}_1, r_1^{(0)})$.

Game_7**:** Same as Game_6, except $\mathsf{rpk}_0 \leftarrow \mathsf{Com}(0; r_{\mathsf{rpk}_0})$ instead of $\mathsf{rpk}_0 \leftarrow \mathsf{Com}(R_0; r_{\mathsf{rpk}_0})$.

Game_8**:** Same as Game_7, except $\mathsf{rpk}_1 \leftarrow \mathsf{Com}(0; r_{\mathsf{rpk}_1})$ instead of $\mathsf{rpk}_1 \leftarrow \mathsf{Com}(R_1; r_{\mathsf{rpk}_1})$.

Game_9**:** Same as Game_8, except $\mathsf{m}_0 \leftarrow \{0,1\}^\lambda$ instead of $\mathsf{m}_0 \leftarrow f_{R_0}(\mathsf{nonce}_0)$.

Game_{10}**:** Same as Game_9, except $\mathsf{m}_1 \leftarrow \{0,1\}^\lambda$ instead of $\mathsf{m}_1 \leftarrow f_{R_1}(\mathsf{nonce}_1)$.

Claim1. We claim that

$$|\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| = \mathsf{negl}(\lambda), \tag{4}$$

assuming that the commitment scheme Com' satisfies non-uniformly computational hiding. Specifically, we show the existence of a non-uniform PPT algorithm $\mathcal{R} := (\mathcal{R}_0, \mathcal{R}_1)$ that breaks the hiding property of Com':

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$:

1. It honestly samples two recipient key pairs $(\mathsf{rpk}_0, \mathsf{rsk}_0)$ and $(\mathsf{rpk}_1, \mathsf{rsk}_1)$, and runs $\mathcal{A}$ with $(\mathsf{rpk}_0, \mathsf{rpk}_1)$.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses its unbounded power to extract a such that $f(a) = y$.
3. It outputs $\mathsf{st} := ((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0, 1\}$.
2. It parses st as $((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$ and outputs $(0, a)$ to the game.
3. Given com_b from the game, it honestly computes the challenge message-signature pairs:

$$(\mathsf{m}_0, \mathsf{sig}_0) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}_0, \mathsf{rsk}_0, \mathsf{nonce}_0, \mathsf{psig}_0) \text{ and}$$
$$(\mathsf{m}_1, \mathsf{sig}_1) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}_1, \mathsf{rsk}_1, \mathsf{nonce}_1, \mathsf{psig}_1).$$

   It replaces $\mathsf{com}_0^{(1)}$ in sig_0 with com_b and provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.
4. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

By construction, the simulation provided by $\mathcal{R}$ for $\mathcal{A}$ behaves exactly like Game_0 when $b = 0$ and like Game_1 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com}',\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]|,\end{aligned}$$

implying that equation (4) holds under the assumption that Com' satisfies non-uniformly computationally hiding.

Claim2. We claim that

$$|\Pr[\mathsf{G}_1] - \Pr[\mathsf{G}_2]| = \mathsf{negl}(\lambda).$$

This follows directly from the same argument as in **Claim1**, based on the non-uniformly computational hiding of Com'.

Claim3. We claim that

$$|\Pr[\mathsf{G}_2] - \Pr[\mathsf{G}_3]| = \mathsf{negl}(\lambda), \tag{5}$$

assuming that NIWI satisfies non-uniformly witness-indistinguishability. Specifically, we show the existence of the following non-uniform PPT algorithm $\mathcal{R} := (\mathcal{R}_0, \mathcal{R}_1)$ against this property:

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$:

1. It honestly samples two recipient key pairs $(\mathsf{rpk}_0, \mathsf{rsk}_0)$ and $(\mathsf{rpk}_1, \mathsf{rsk}_1)$, and runs $\mathcal{A}$ with $(\mathsf{rpk}_0, \mathsf{rpk}_1)$.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses its unbounded power to extract a such that $f(a) = y$.

3. It outputs $\mathsf{st} := ((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0,1\}$.
2. It parses st as $((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$, and outputs $(\mathsf{nonce}_0, \mathsf{rpk}_0, \mathsf{psig}_0, R_0, r_{\mathsf{rpk}_0}, r_0^{(0)})$ and $(a, r_0^{(1)})$ to the game.
3. Given π_b from the game, it computes the same two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ as in Game_2, but replaces π_0 in sig_0 with π_b.
4. It then provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.
5. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

By construction, $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_2 when $b = 0$ and like Game_3 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{wi}}_{\mathsf{NIWI},\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_2] - \Pr[\mathsf{G}_3]|,\end{aligned}$$

implying that equation (5) holds under the assumption that NIWI satisfies the non-uniformly witness-indistinguishability.

Claim4 We claim that

$$|\Pr[\mathsf{G}_3] - \Pr[\mathsf{G}_4]| = \mathsf{negl}(\lambda).$$

This follows from the same argument as in **Claim3**, based on the non-uniform computational witness indistinguishability of NIWI.

Claim5 We claim that

$$|\Pr[\mathsf{G}_4] - \Pr[\mathsf{G}_5]| = \mathsf{negl}(\lambda), \tag{6}$$

assuming that the commitment scheme Com' satisfies non-uniform computational hiding. Specifically, we demonstrate the existence of the following non-uniform PPT algorithm $\mathcal{R} := (\mathcal{R}_0, \mathcal{R}_1)$, which breaks the hiding of Com':

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$:

1. It honestly samples two recipient key pairs $(\mathsf{rpk}_0, \mathsf{rsk}_0)$ and $(\mathsf{rpk}_1, \mathsf{rsk}_1)$, and runs $\mathcal{A}$ with $(\mathsf{rpk}_0, \mathsf{rpk}_1)$.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses its unbounded power to extract a such that $f(a) = y$.
3. It outputs $\mathsf{st} := ((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0,1\}$.
2. It parses st as $((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$, and outputs $((\mathsf{nonce}_0, \mathsf{rpk}_0, \mathsf{psig}_0), 0)$ to the game.
3. Given com_b from the game, it computes two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ in the same manner as Game_4, except it replaces $\mathsf{com}_0^{(0)}$ in sig_0 with com_b. It subsequently provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.

4. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

We observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_4 when $b = 0$ and like Game_5 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com}',\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_4] - \Pr[\mathsf{G}_5]|,\end{aligned}$$

implying that equation (6) holds under the assumption that Com' satisfies non-uniform computational hiding.

Claim6 We claim that

$$|\Pr[\mathsf{G}_5] - \Pr[\mathsf{G}_6]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as in **Claim5**, relying on the non-uniform computational hiding property of Com'.

Claim7 We claim that

$$|\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]| = \mathsf{negl}(\lambda), \tag{7}$$

assuming that the commitment scheme Com satisfies T_2-hiding and the function f is T_1-$\mathsf{invertible}$. Specifically, we present the following $T_2 \cdot \mathsf{poly}(\lambda)$-time algorithm $\mathcal{R}$, which breaks the hiding of Com:

1. It samples a random bit $b' \leftarrow \{0, 1\}$.
2. It honestly samples two recipient key pairs $(\mathsf{rpk}_0, \mathsf{rsk}_0)$ and $(\mathsf{rpk}_1, \mathsf{rsk}_1)$, and outputs $(R_0, 0)$ to the game.
3. Given com_b from the game, it runs $\mathcal{A}$ with $(\mathsf{com}_b, \mathsf{rpk}_1)$.
4. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses the T_1-invertibility of f to extract a such that $f(a) = y$. (Note that $\mathcal{R}$ runs in $T_2 \cdot \mathsf{poly}(\lambda)$ time and $T_1 < T_2$.) It then computes two challenge message-signature pairs as in Game_6, and provides $(\mathsf{m}'_b, \mathsf{sig}'_b, \mathsf{m}_{1-b'}, \mathsf{sig}_{1-b'})$ to $\mathcal{A}$.
5. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

For simplicity, we assume the existence of a $T_2 \cdot \mathsf{poly}(\lambda)$-time algorithm that can extract every preimage of f with probability 1.

We observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_6 when $b = 0$ and like Game_7 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com},\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]|,\end{aligned}$$

implying that equation (7) holds under the assumption that Com satisfies T_2 hiding.

Claim8 We claim that

$$|\Pr[\mathsf{G}_7] - \Pr[\mathsf{G}_8]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as **Claim7**, relying on the T_2-hiding property of Com and the T_1-invertibility of f.

Claim9 We claim that

$$|\Pr[\mathsf{G}_8] - \Pr[\mathsf{G}_9]| = \mathsf{negl}(\lambda), \tag{8}$$

assuming that the function F satisfies non-uniform pseudorandomness. Specifically, we present the following non-uniform PPT algorithm $\mathcal{R} = (\mathcal{R}_0, \mathcal{R}_1)$, which breaks the pseudorandomness of F:

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$: ,

1. It samples two recipient key pairs $(\mathsf{rpk}_0, \mathsf{rsk}_0)$ and $(\mathsf{rpk}_1, \mathsf{rsk}_1)$ as in Game_8, running $\mathcal{A}$ with $(\mathsf{rpk}_0, \mathsf{rpk}_1)$.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from A, it extracts a such that $f(a) = y$ using its unbounded power.
3. It outputs the state $\mathsf{st} = ((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0, 1\}$.
2. It parses st as $((\mathsf{rpk}_0, \mathsf{rsk}_0), (\mathsf{rpk}_1, \mathsf{rsk}_1), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a) := \mathsf{st}$ and outputs nonce_0 to the game.
3. Given Y from the game, it computes two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ in the same manner as Game_8, except it sets $\mathsf{m}_0 := Y$. It then provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.
4. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, then it outputs 0 to the game; otherwise it outputs 1.

We can observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_8 when $b = 0$ and like Game_9 when $b = 1$. Therefore, we have:

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{prf}}_{F,\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R} | b = 0] - \Pr[0 \leftarrow \mathcal{R} | b = 1]| \\
&= |\Pr[b^* = b' | b = 0] - \Pr[b^* = b' | b = 1]| \\
&= |\Pr[\mathsf{G}_8] - \Pr[\mathsf{G}_9]|,
\end{aligned}$$

implying that equation (8) holds under the assumption that F satisfies the pseudorandomness.

Claim10 We claim that

$$|\Pr[\mathsf{G}_9] - \Pr[\mathsf{G}_{10}]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as **Claim9**, relying on the pseudorandomness of F.

We claim that the advantage of adversary $\mathcal{A}$ in Game Game_{10} is equal to 0, as the challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ are independent of the recipient's public key.

Putting all together, we can conclude:

$$\begin{aligned}2\Pr[\mathsf{RBND}^{\mathcal{S}_{\mathsf{rbnd}}}_{\mathsf{NIBS}}(\lambda)=1]-1 &= 2\Pr[\mathsf{G}_0]-1\\ &\le 2\sum_{i=0}^{9}|\Pr[\mathsf{G}_i]-\Pr[\mathsf{G}_{i+1}]|+|2\Pr[\mathsf{G}_{10}]-1|\\ &= \mathsf{negl}(\lambda).\end{aligned}$$

□

Proof. (Theorem 5) We will prove this theorem by a series of hybrid arguments. Let $\mathcal{A}$ be a non-uniform PPT adversary against the recipient blindness of NIBS and let G_i denote the event where the adversary $\mathcal{A}$ wins in game Game_i.

Game_0**:** The original nonce blindness game.

Game_1**:** The same as Game_0, except $\mathsf{com}_0^{(1)} \leftarrow \mathsf{Com}'(a; r_0^{(1)})$ where $f(a) = y$.

Game_2**:** Same as Game_1, except $\mathsf{com}_1^{(1)} \leftarrow \mathsf{Com}'(a; r_1^{(1)})$, where $f(a) = y$.

Game_3**:** Same as Game_2, except when creating the NIWI proof π_0 in challenge signature sig_0, we use witness $w := (a, r_0^{(1)})$ instead of $(\mathsf{nonce}_0, \mathsf{rpk}, \mathsf{psig}_0, R, r_{\mathsf{rpk}}, r_0^{(0)})$.

Game_4**:** Same as Game_3, except when creating the NIWI proof π_1 in challenge signature sig_1, we use witness $w := (a, r_1^{(1)})$ instead of $(\mathsf{nonce}_1, \mathsf{rpk}, \mathsf{psig}_1, R, r_{\mathsf{rpk}}, r_1^{(0)})$.

Game_5**:** Same as Game_4, except $\mathsf{com}_0^{(0)} \leftarrow \mathsf{Com}'(0; r_0^{(0)})$ instead of $\mathsf{com}_0^{(0)} \leftarrow \mathsf{Com}'(\mathsf{nonce}_0, \mathsf{rpk}, \mathsf{psig}_0, r_0^{(0)})$.

Game_6**:** Same as Game_5, except $\mathsf{com}_1^{(0)} \leftarrow \mathsf{Com}'(0; r_1^{(0)})$ instead of $\mathsf{com}_1^{(0)} \leftarrow \mathsf{Com}'(\mathsf{nonce}_1, \mathsf{rpk}, \mathsf{psig}_1, r_1^{(0)})$.

Game_7**:** Same as Game_6, except $\mathsf{rpk} \leftarrow \mathsf{Com}(0; r_{\mathsf{rpk}})$ instead of $\mathsf{rpk} \leftarrow \mathsf{Com}(R; r_{\mathsf{rpk}})$.

Game_8**:** Same as Game_7, except for the following modification: if $\mathsf{nonce}_0 = \mathsf{nonce}_1$, then $\mathsf{m}_0 = \mathsf{m}_1 \leftarrow \{0,1\}^\lambda$; otherwise, $\mathsf{m}_0, \mathsf{m}_1 \leftarrow \{0,1\}^\lambda$ (instead of setting $\mathsf{m}_0 \leftarrow F_R(\mathsf{nonce}_0)$ and $\mathsf{m}_1 \leftarrow F_R(\mathsf{nonce}_1)$).

Claim1. We claim that

$$|\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| = \mathsf{negl}(\lambda), \tag{9}$$

assuming that the commitment scheme Com' satisfies non-uniformly computational hiding. Specifically, we show the existence of a non-uniform PPT algorithm $\mathcal{R} := (\mathcal{R}_0, \mathcal{R}_1)$ that breaks the hiding property of Com':
Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$:

1. It honestly samples a recipient key pair $(\mathsf{rpk}, \mathsf{rsk})$, and runs $\mathcal{A}$ with rpk.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses its unbounded power to extract a such that $f(a) = y$.
3. It outputs $\mathsf{st} := ((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0, 1\}$.
2. It parses st as $((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$ and outputs $(0, a)$ to the game.
3.
4. Given com_b from the game, it honestly computes the challenge message-signature pairs:

$$(\mathsf{m}_0, \mathsf{sig}_0) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}_0, \mathsf{psig}_0) \text{ and}$$
$$(\mathsf{m}_1, \mathsf{sig}_1) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}_1, \mathsf{psig}_1).$$

   It replaces $\mathsf{com}_0^{(1)}$ in sig_0 with com_b and provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.
5. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

By construction, the simulation provided by $\mathcal{R}$ for $\mathcal{A}$ behaves exactly like Game_0 when $b = 0$ and like Game_1 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com}',\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R} | b = 0] - \Pr[0 \leftarrow \mathcal{R} | b = 1]| \\ &= |\Pr[b^* = b' | b = 0] - \Pr[b^* = b' | b = 1]| \\ &= |\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]|,\end{aligned}$$

implying that equation (9) holds under the assumption that Com' satisfies non-uniformly computationally hiding.

Claim2. We claim that

$$|\Pr[\mathsf{G}_1] - \Pr[\mathsf{G}_2]| = \mathsf{negl}(\lambda).$$

This follows directly from the same argument as in **Claim1**, based on the non-uniformly computational hiding of Com'.

Claim3. We claim that

$$|\Pr[\mathsf{G}_2] - \Pr[\mathsf{G}_3]| = \mathsf{negl}(\lambda), \tag{10}$$

assuming that NIWI satisfies non-uniformly witness-indistinguishability. Specifically, we show the existence of the following non-uniform PPT algorithm $\mathcal{R} := (\mathcal{R}_0, \mathcal{R}_1)$ against this property:

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$:

1. It honestly samples a recipient key pair $(\mathsf{rpk}, \mathsf{rsk})$, and runs $\mathcal{A}$ with rpk.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses its unbounded power to extract a such that $f(a) = y$.

3. It outputs $\mathsf{st} := ((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0,1\}$.
2. It parses st as $((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$, and outputs $(\mathsf{nonce}_0, \mathsf{rpk}, \mathsf{psig}_0, R, r_{\mathsf{rpk}}, r_0^{(0)})$ and $(a, r_0^{(1)})$ to the game.
3. Given π_b from the game, it computes the same two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ as in Game_2, but replaces π_0 in sig_0 with π_b.
4. It then provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.
5. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

By construction, $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_2 when $b = 0$ and like Game_3 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{wi}}_{\mathsf{NIWI},\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_2] - \Pr[\mathsf{G}_3]|,\end{aligned}$$

implying that equation (10) holds under the assumption that NIWI satisfies the non-uniformly witness-indistinguishability.

Claim4 We claim that

$$|\Pr[\mathsf{G}_3] - \Pr[\mathsf{G}_4]| = \mathsf{negl}(\lambda).$$

This follows from the same argument as in **Claim3**, based on the non-uniform computational witness indistinguishability of NIWI.

Claim5 We claim that

$$|\Pr[\mathsf{G}_4] - \Pr[\mathsf{G}_5]| = \mathsf{negl}(\lambda), \tag{11}$$

assuming that the commitment scheme Com' satisfies non-uniform computational hiding. Specifically, we demonstrate the existence of the following non-uniform PPT algorithm $\mathcal{R} := (\mathcal{R}_0, \mathcal{R}_1)$, which breaks the hiding of Com':

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$:

1. It honestly samples a recipient key pair $(\mathsf{rpk}, \mathsf{rsk})$, and runs $\mathcal{A}$ with rpk.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses its unbounded power to extract a such that $f(a) = y$.
3. It outputs $\mathsf{st} := ((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0,1\}$.
2. It parses st as $((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$, and outputs $((\mathsf{nonce}_0, \mathsf{rpk}, \mathsf{psig}_0), 0)$ to the game.
3. Given com_b from the game, it computes two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ in the same manner as Game_4, except it replaces $\mathsf{com}_0^{(0)}$ in sig_0 with com_b. It subsequently provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.

4. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

We observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_4 when $b = 0$ and like Game_5 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com}',\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_4] - \Pr[\mathsf{G}_5]|,\end{aligned}$$

implying that equation (11) holds under the assumption that Com' satisfies non-uniform computational hiding.

Claim6 We claim that

$$|\Pr[\mathsf{G}_5] - \Pr[\mathsf{G}_6]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as in **Claim5**, relying on the non-uniform computational hiding property of Com'.

Claim7 We claim that

$$|\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]| = \mathsf{negl}(\lambda), \tag{12}$$

assuming that the commitment scheme Com satisfies T_2-hiding and the function f is T_1-$\mathsf{invertible}$. Specifically, we present the following $T_2 \cdot \mathsf{poly}(\lambda)$-time algorithm $\mathcal{R}$, which breaks the hiding of Com:

1. It samples a random bit $b' \leftarrow \{0, 1\}$.
2. It honestly samples a recipient key pair $(\mathsf{rpk}, \mathsf{rsk})$, and outputs $(R, 0)$ to the game.
3. Given com_b from the game, it runs $\mathcal{A}$ with $(\mathsf{com}_b, \mathsf{rpk})$.
4. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it uses the T_1-invertibility of f to extract a such that $f(a) = y$. (Note that $\mathcal{R}$ runs in $T_2 \cdot \mathsf{poly}(\lambda)$ time and $T_1 < T_2$.) It then computes two challenge message-signature pairs as in Game_6, and provides $(\mathsf{m}'_b, \mathsf{sig}'_b, \mathsf{m}_{1-b'}, \mathsf{sig}_{1-b'})$ to $\mathcal{A}$.
5. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, it outputs 0 to the game; otherwise, it outputs 1.

For simplicity, we assume the existence of a $T_2 \cdot \mathsf{poly}(\lambda)$-time algorithm that can extract every preimage of f with probability 1.

We observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_6 when $b = 0$ and like Game_7 when $b = 1$. Therefore, we have:

$$\begin{aligned}\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com},\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\ &= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\ &= |\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]|,\end{aligned}$$

implying that equation (12) holds under the assumption that Com satisfies T_2 hiding.

Claim8 We claim that

$$|\Pr[\mathsf{G}_7] - \Pr[\mathsf{G}_8]| = \mathsf{negl}(\lambda), \tag{13}$$

assuming that the function F satisfies non-uniform pseudorandomness. Specifically, we present the following non-uniform PPT algorithm $\mathcal{R} = (\mathcal{R}_0, \mathcal{R}_1)$, which breaks the pseudorandomness of F:

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$: ,

1. It samples a recipient key pair $(\mathsf{rpk}, \mathsf{rsk})$ as in Game_8, running $\mathcal{A}$ with $(\mathsf{rpk}, \mathsf{rsk})$.
2. Given $(\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it extracts a such that $f(a) = y$ using its unbounded power.
3. It outputs the state $\mathsf{st} = ((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0, 1\}$.
2. It parses st as $((\mathsf{rpk}, \mathsf{rsk}), (\mathsf{pk} = (\mathsf{spk}, y), \mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1), a) := \mathsf{st}$ and outputs nonce_0 and nonce_1 to the game.
3. Given Y_0 and Y_1 from the game, it computes two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ in the same manner as Game_7, except it sets $\mathsf{m}_0 := Y_0$ and $\mathsf{m}_1 := Y_1$. It then provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$.
4. Given a bit b^* from $\mathcal{A}$, if $b' = b^*$, then it outputs 0 to the game; otherwise it outputs 1.

We can observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_7 when $b = 0$ and like Game_8 when $b = 1$. Therefore, we have:

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{prf}}_{F,\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R} | b = 0] - \Pr[0 \leftarrow \mathcal{R} | b = 1]| \\
&= |\Pr[b^* = b' | b = 0] - \Pr[b^* = b' | b = 1]| \\
&= |\Pr[\mathsf{G}_7] - \Pr[\mathsf{G}_8]|,
\end{aligned}$$

implying that equation (13) holds under the assumption that F satisfies the pseudorandomness.

We claim that the advantage of adversary $\mathcal{A}$ in Game Game_8 is equal to 0, as the challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ are independent of the nonces nonce_0 and nonce_1.

Putting all together, we can conclude:

$$\begin{aligned}
2\Pr[\mathsf{NBND}^{\mathcal{S}_{\mathsf{nbnd}}}_{\mathsf{NIBS}}(\lambda) = 1] - 1 &= 2\Pr[\mathsf{G}_0] - 1 \\
&\leq 2\sum_{i=0}^{7} |\Pr[\mathsf{G}_i] - \Pr[\mathsf{G}_{i+1}]| + |2\Pr[\mathsf{G}_8] - 1| \\
&= \mathsf{negl}(\lambda).
\end{aligned}$$

□

D Full Proofs for Our Construction from Classical and Quantum Hardness Assumptions

Definition 9. *(Non-Adaptive Blindness) We say that* NIBS *satisfies non-adaptive blindness if for any PPT adversary* $\mathcal{A}$*, the probability that* $\mathcal{A}$ *wins in the following two game is equal to* $1/2 + \mathsf{negl}(\lambda)$:

Non-Adaptive Recipient Blindness *For an algorithm* $\mathsf{S}_{\mathsf{rbnd}}$*, we consider the following game* $\mathsf{NARBND}^{\mathcal{S}_{\mathsf{rbnd}}}_{\mathsf{NIBS}}(\lambda)$:

1. *The game flips random bit* $b \leftarrow \{0,1\}$ *and then runs*

$$(\mathsf{pk}, \mathsf{St}) \leftarrow \mathcal{S}_{\mathsf{rbnd}}(\mathsf{init}).$$

2. *It samples*

$$(\mathsf{rpk}_0, \mathsf{rsk}_0) \leftarrow \mathsf{RKeyGen}(\lambda) \text{ and } (\mathsf{rpk}_1, \mathsf{rsk}_1) \leftarrow \mathsf{RKeyGen}(\lambda)$$

It also runs

$$(\mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1, \mathsf{St}') \leftarrow \mathcal{S}_{\mathsf{rbnd}}(\mathsf{issue}, \mathsf{St}, \mathsf{rpk}_0, \mathsf{rpk}_1).$$

3. *The game computes*

$$(\mathsf{m}_0, \mathsf{sig}_0) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}_0, \mathsf{rsk}_0, \mathsf{nonce}_0, \mathsf{psig}_0) \text{ and}$$
$$(\mathsf{m}_1, \mathsf{sig}_1) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}_1, \mathsf{rsk}_1, \mathsf{nonce}_1, \mathsf{psig}_1).$$

If $(\mathsf{m}_0, \mathsf{sig}_0) \neq \perp$ *or* $(\mathsf{m}_1, \mathsf{sig}_1) \neq \perp$*. then it sets* $(\mathsf{m}_0, \mathsf{sig}_0) := \perp$ *and* $(\mathsf{m}_1, \mathsf{sig}_1) := \perp$

4. *The game runs*

$$b' \leftarrow \mathcal{S}_{\mathsf{rbnd}}(\mathsf{guess}, \mathsf{St}', \mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b}).$$

It finally outputs 1 *if* $b = b'$*; otherwise, it outputs* 0*.*

Non-Adaptive Nonce Blindness *For an algorithm* $\mathsf{S}_{\mathsf{nbnd}}$*, we consider the following game* $\mathsf{NANBND}^{\mathcal{S}_{\mathsf{nbnd}}}_{\mathsf{NIBS}}(\lambda)$:

1. *The game flips random bit* $b \leftarrow \{0,1\}$ *and then runs*

$$(\mathsf{pk}, \mathsf{St}) \leftarrow \mathcal{S}_{\mathsf{nbnd}}(\mathsf{init}).$$

2. *It samples* $(\mathsf{rpk}, \mathsf{rsk}) \leftarrow \mathsf{RKeyGen}(\lambda)$ *and then runs*

$$(\mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1, \mathsf{St}') \leftarrow \mathcal{S}_{\mathsf{nbnd}}(\mathsf{issue}, \mathsf{St}, \mathsf{rpk}).$$

3. *The game computes*

$$(\mathsf{m}_0, \mathsf{sig}_0) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}_0, \mathsf{psig}_0) \text{ and}$$
$$(\mathsf{m}_1, \mathsf{sig}_1) \leftarrow \mathsf{Obtain}(\mathsf{pk}, \mathsf{rpk}, \mathsf{rsk}, \mathsf{nonce}_1, \mathsf{psig}_1).$$

If $(\mathsf{m}_0, \mathsf{sig}_0) \neq \perp$ *or* $(\mathsf{m}_1, \mathsf{sig}_1) \neq \perp$*. then it sets* $(\mathsf{m}_0, \mathsf{sig}_0) := \perp$ *and* $(\mathsf{m}_1, \mathsf{sig}_1) := \perp$

4. The game runs

$$b' \leftarrow \mathcal{S}_{\mathsf{nbnd}}(\mathsf{guess}, \mathsf{St}', \mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b}).$$

It finally outputs 1 *if* $b = b'$*; otherwise, it outputs* 0.

Proof. (Theorem 6) We prove the statement via a sequence of the same games as Therorem 3. The proof follows the same structure as the proof of Theorem 3, except that the T_1-time hardness assumption is replaced with the quantum hardness assumption. □

Proof. (Theorem 7)

Non-adaptive recipient blindness We will prove this theorem by a series of hybrid arguments, following the same 11-game sequence as in the proof of Theorem 4. Let $\mathcal{A}$ be a non-uniform PPT adversary against the non-adaptive recipient blindness of NIBS and let G_i denote the event where the adversary $\mathcal{A}$ wins in game Game_i.

Claim1. We claim that

$$|\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| = \mathsf{negl}(\lambda),$$

assuming that the commitment scheme Com' satisfies non-uniformly computational hiding. This claim follows for the same reasons as **Claim1** in the proof of Theorem 4.

Claim2. We claim that

$$|\Pr[\mathsf{G}_1] - \Pr[\mathsf{G}_2]| = \mathsf{negl}(\lambda).$$

This follows directly from the same argument as in **Claim1**, based on the non-uniformly computational hiding of Com'.

Claim3. We claim that

$$|\Pr[\mathsf{G}_2] - \Pr[\mathsf{G}_3]| = \mathsf{negl}(\lambda),$$

assuming that NIWI satisfies non-uniformly witness-indistinguishability. This claim follows for the same reasons as **Claim3** in the proof of Theorem 4.

Claim4 We claim that

$$|\Pr[\mathsf{G}_3] - \Pr[\mathsf{G}_4]| = \mathsf{negl}(\lambda).$$

This follows from the same argument as in **Claim3**, based on the non-uniform computational witness indistinguishability of NIWI.

Claim5 We claim that

$$|\Pr[\mathsf{G}_4] - \Pr[\mathsf{G}_5]| = \mathsf{negl}(\lambda),$$

assuming that the commitment scheme Com' satisfies non-uniform computational hiding. This claim follows for the same reasons as **Claim5** in the proof of Theorem 4.

Claim6 We claim that

$$|\Pr[\mathsf{G}_5] - \Pr[\mathsf{G}_6]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as in **Claim5**, relying on the non-uniform computational hiding property of Com'.

Claim7 We claim that

$$|\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]| = \mathsf{negl}(\lambda), \tag{14}$$

assuming that the commitment scheme Com satisfies computational hiding. Specifically, we present the following non-uniform PPT algorithm $\mathcal{R} = (\mathcal{R}_0, \mathcal{R}_1)$, which breaks the computational hiding property of Com:

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$: ,

1. It runs $\mathcal{A}$ with init.
2. Given $\mathsf{pk} = (\mathsf{spk}, y)$ from $\mathcal{A}$, it extracts a such that $f(a) = y$ using its unbounded power.
3. It outputs the state $\mathsf{st} := (\mathsf{pk}, a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0,1\}$.
2. It parses st as $(\mathsf{pk} = (\mathsf{spk}, y), a) := \mathsf{st}$ and outputs $(R, 0)$ to the game, where $R \leftarrow \{0,1\}^\lambda$.
3. Given com_b from the game, it sets $\mathsf{rpk}_0 := \mathsf{com}_b$ and honestly samples $(\mathsf{rpk}_1, \mathsf{rsk}_1)$.
4. It runs $\mathcal{A}$ with $(\mathsf{rpk}_0, \mathsf{rpk}_1)$. Given $(\mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it computes two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ in the same manner as Game_6.
5. It then provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$ to obtain a bit b^*. If $b' = b^*$, then it outputs 0 to the game. otherwise, it outputs 1.

We can observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_6 when $b = 0$ and like Game_7 when $b = 1$. Therefore, we have:

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com},\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R} | b = 0] - \Pr[0 \leftarrow \mathcal{R} | b = 1]| \\
&= |\Pr[b^* = b' | b = 0] - \Pr[b^* = b' | b = 1]| \\
&= |\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]|,
\end{aligned}$$

implying that equation (14) holds under the assumption that Com satisfies computational hiding.

Claim8 We claim that

$$|\Pr[\mathsf{G}_7] - \Pr[\mathsf{G}_8]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as **Claim7**, relying on the computational hiding property of Com.

Claim9 We claim that

$$|\Pr[\mathsf{G}_8] - \Pr[\mathsf{G}_9]| = \mathsf{negl}(\lambda),$$

assuming that the function F satisfies non-uniform pseudorandomness. This claim follows for the same reasons as **Claim9** in the proof of Theorem 4.

Claim10 We claim that

$$|\Pr[\mathsf{G}_9] - \Pr[\mathsf{G}_{10}]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as **Claim9**, relying on the pseudorandomness of F.

We claim that the advantage of adversary $\mathcal{A}$ in Game Game_{10} is equal to 0, as the challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ are independent of the recipient's public key.

Putting all together, we can conclude:

$$\begin{aligned} 2\Pr[\mathsf{NARBND}^{\mathcal{S}_{\mathsf{rbnd}}}_{\mathsf{NIBS}}(\lambda) = 1] - 1 &= 2\Pr[\mathsf{G}_0] - 1 \\ &\leq 2\sum_{i=0}^{9} |\Pr[\mathsf{G}_i] - \Pr[\mathsf{G}_{i+1}]| + |2\Pr[\mathsf{G}_{10}] - 1| \\ &= \mathsf{negl}(\lambda). \end{aligned}$$

Non-adaptive nonce blindness We will prove this theorem by a series of hybrid arguments, following the same 9-game sequence as in the proof of Theorem 5. Let $\mathcal{A}$ be a non-uniform PPT adversary against the non-adaptive nonce blindness of NIBS and let G_i denote the event where the adversary $\mathcal{A}$ wins in game Game_i.

Claim1. We claim that

$$|\Pr[\mathsf{G}_0] - \Pr[\mathsf{G}_1]| = \mathsf{negl}(\lambda),$$

assuming that the commitment scheme Com' satisfies non-uniformly computational hiding. This claim follows for the same reasons as **Claim1** in the proof of Theorem 5.

Claim2. We claim that

$$|\Pr[\mathsf{G}_1] - \Pr[\mathsf{G}_2]| = \mathsf{negl}(\lambda).$$

This follows directly from the same argument as in **Claim1**, based on the non-uniformly computational hiding of Com'.

Claim3. We claim that

$$|\Pr[\mathsf{G}_2] - \Pr[\mathsf{G}_3]| = \mathsf{negl}(\lambda),$$

assuming that NIWI satisfies non-uniformly witness-indistinguishability. This claim follows for the same reasons as **Claim3** in the proof of Theorem 5.

Claim4 We claim that

$$|\Pr[\mathsf{G}_3] - \Pr[\mathsf{G}_4]| = \mathsf{negl}(\lambda).$$

This follows from the same argument as in **Claim3**, based on the non-uniform computational witness indistinguishability of NIWI.

Claim5 We claim that

$$|\Pr[\mathsf{G}_4] - \Pr[\mathsf{G}_5]| = \mathsf{negl}(\lambda),$$

assuming that the commitment scheme Com' satisfies non-uniform computational hiding. This claim follows for the same reasons as **Claim5** in the proof of Theorem 5.

Claim6 We claim that

$$|\Pr[\mathsf{G}_5] - \Pr[\mathsf{G}_6]| = \mathsf{negl}(\lambda).$$

This claim follows the same reasoning as in **Claim5**, relying on the non-uniform computational hiding property of Com'.

Claim7 We claim that

$$|\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]| = \mathsf{negl}(\lambda), \tag{15}$$

assuming that the commitment scheme Com satisfies computational hiding. Specifically, we present the following non-uniform PPT algorithm $\mathcal{R} = (\mathcal{R}_0, \mathcal{R}_1)$, which breaks the computational hiding property of Com:

Unbounded-time pre-computation phase $\mathcal{R}_0(\lambda)$: ,

1. It runs $\mathcal{A}$ with init.
2. Given $\mathsf{pk} = (\mathsf{spk}, y)$ from $\mathcal{A}$, it extracts a such that $f(a) = y$ using its unbounded power.
3. It outputs the state $\mathsf{st} := (\mathsf{pk}, a)$.

Online phase $\mathcal{R}_1(\mathsf{st})$:

1. It samples a random bit $b' \leftarrow \{0,1\}$.
2. It parses st as $(\mathsf{pk} = (\mathsf{spk}, y), a) := \mathsf{st}$ and outputs $(R, 0)$ to the game, where $R \leftarrow \{0,1\}^\lambda$.
3. Given com_b from the game, it sets $\mathsf{rpk} := \mathsf{com}_b$.
4. It runs $\mathcal{A}$ with rpk. Given $(\mathsf{nonce}_0, \mathsf{psig}_0, \mathsf{nonce}_1, \mathsf{psig}_1)$ from $\mathcal{A}$, it computes two challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ in the same manner as Game_6.
5. It then provides $(\mathsf{m}_b, \mathsf{sig}_b, \mathsf{m}_{1-b}, \mathsf{sig}_{1-b})$ to $\mathcal{A}$ to obtain a bit b^*. If $b' = b^*$, then it outputs 0 to the game. otherwise, it outputs 1.

We can observe that $\mathcal{R}$'s simulation for $\mathcal{A}$ behaves like Game_6 when $b = 0$ and like Game_7 when $b = 1$. Therefore, we have:

$$\begin{aligned}
\mathbf{Adv}^{\mathsf{hiding}}_{\mathsf{Com},\mathcal{R}}(\lambda) &= |\Pr[0 \leftarrow \mathcal{R}|b = 0] - \Pr[0 \leftarrow \mathcal{R}|b = 1]| \\
&= |\Pr[b^* = b'|b = 0] - \Pr[b^* = b'|b = 1]| \\
&= |\Pr[\mathsf{G}_6] - \Pr[\mathsf{G}_7]|,
\end{aligned}$$

implying that equation (15) holds under the assumption that Com satisfies computational hiding.

Claim8 We claim that

$$|\Pr[\mathsf{G}_7] - \Pr[\mathsf{G}_8]| = \mathsf{negl}(\lambda),$$

assuming that the function F satisfies non-uniform pseudorandomness. This claim follows for the same reasons as **Claim8** in the proof of Theorem 5.

We claim that the advantage of adversary $\mathcal{A}$ in Game Game_8 is equal to 0, as the challenge message-signature pairs $(\mathsf{m}_0, \mathsf{sig}_0)$ and $(\mathsf{m}_1, \mathsf{sig}_1)$ are independent of the nonces nonce_0 and nonce_1.
Putting all together, we can conclude:

$$\begin{aligned} 2\Pr[\mathsf{NANBND}^{\mathcal{S}_{\mathsf{nbnd}}}_{\mathsf{NIBS}}(\lambda) = 1] - 1 &= 2\Pr[\mathsf{G}_0] - 1 \\ &\leq 2\sum_{i=0}^{7} |\Pr[\mathsf{G}_i] - \Pr[\mathsf{G}_{i+1}]| + |2\Pr[\mathsf{G}_8] - 1| \\ &= \mathsf{negl}(\lambda). \end{aligned}$$

□