

RSA Blind Signatures with Public Metadata

Ghous Amjad

Kevin Yeo

Moti Yung

Abstract

Anonymous tokens are, essentially, digital signature schemes that enable issuers to provide users with signatures without learning the user inputs or the final signatures. These primitives allow applications to propagate trust while simultaneously protecting the user identity. They have become a core component for improving the privacy of several real-world applications including ad measurements, authorization protocols, spam detection, and VPNs.

In certain applications, it is natural to associate signatures with specific public metadata, ensuring that trust is only propagated with respect to only a certain set of users and scenarios. To solve this, we study the notion of anonymous tokens with public metadata. We present a variant of RSA blind signatures with public metadata where issuers may only generate signatures that verify for a certain choice of public metadata a modification of a scheme by Abe and Fujisaki [9]. Our protocol exclusively uses standard cryptography with widely available implementations. We prove security from the one-more RSA assumptions with multiple exponents that we introduce. Furthermore, we provide evidence that the concrete security bounds should be nearly identical to standard RSA blind signatures. We show that our protocol incurs minimal overhead over standard RSA blind signatures and report anonymous telemetry for a real-world deployment to showcase its scalability. Moreover, the protocol in this paper has been proposed as a technical specification in an IRTF internet draft [12].

1 Introduction

Anonymous tokens are a powerful primitive that have been studied for decades under various names including anonymous credentials and blind signatures. They were first introduced by Chaum [23, 24] in the context of untraceable electronic cash. At a high level, the idea was that any bank/treasury would be able to attest the validity of money with the guarantee that no one would ever be able to trace the usage of the money even if the bank knows the identity of the user that was originally granted the money. This is just one example that exhibits the usefulness of anonymous tokens. In recent years, anonymous tokens have become a core component of many real-world applications such as private click measurement [5], privacy-preserving telemetry [39], fraud detection [3], avoiding repeated CAPTCHA solving [31] and modern VPNs [4, 6].

There are typically three different roles for participants in anonymous token schemes. The user (also commonly referred to as the signature recipient or receiver) is the party that is requesting and receiving the signature. The signer (or issuer) is the party that receives a blind signing request and is responsible for issuing a response that enables the user to recover the final signature. Finally, the last party is the verifier whose responsibility is to check whether a signature is one that was correctly signed by the signer. Anonymous token schemes may be split into two types: designated or public verifiability. In the designated setting, verification requires the usage of the private key and certain parties must be explicitly assigned the role of verifier by being given the private key. In the public verifiability case, anyone can perform verification given they have access to the public key. In our work, we will exclusively focus on public verification. We choose to focus on public verification due to the flexibility that it offers when used in protocols where any party can verify signatures without needing the secret key. In particular, parties may act as verifiers without the ability to also issue signatures. We note that a party may take on more than one role in various applications. For example, a party may be responsible for both issuing signatures as well as verifying their validity.

Google, {gamjad,kwlyeo,moti}@google.com.

Anonymous tokens are required to satisfy three important properties: correctness, unforgeability and unlinkability. For correctness, it must be the case that a signature will be successfully accepted by a verifier as long as all parties involved were acting honestly (that is, following the algorithm correctly). Unforgeability describes the property that an adversary should not be able to create valid signatures without interacting with a signer. For example, an adversarial user that performs ℓ blind signing protocols with a signer should not be able to create $\ell + 1$ distinct pairs of message and valid signature of the corresponding message. Both correctness and unforgeability are properties required from any digital signature scheme. The last property of unlinkability formally describes the notion of anonymity with respect to an untrusted signer. In particular, the signer should be unable to link any blind signing request with any final signature. Suppose that the signer answers ℓ blind signing requests. Afterwards, the signer receives the ℓ resulting signatures that are randomly permuted. Then, the signer is unable to link any request with a specific signature with probability better than $1/\ell$. In other words, a signature remains anonymous in the view of the adversarial signer that may have maliciously chosen the system’s parameters.

In this work, we study a variant of anonymous tokens with public metadata (also known as partially blind signatures [10, 8, 9]) where each signature will be tied to a specific piece of public metadata. In standard anonymous tokens (without public metadata), the verifier only receives the signature and plaintext message. Furthermore, the signer issuing the signature never sees the plaintext message to maintain anonymity. In other words, there is no way for a signer to properly check and embed metadata into a signature that can be later used by the verifier. By enabling public metadata support, we allow additional information in the protocol that will be integral during verification. This feature ends up being quite useful in multiple applications as we will show later. Anonymous tokens with public metadata allow the user and signer to jointly agree on metadata explicitly encoded into each token. This metadata will be public to all parties (the input message remains hidden from the signer for anonymity). Additionally, the verifier will check signatures against their specific choices of public metadata. The primitives must guarantee that a signature will only be accepted if the verifier uses the correct public metadata used in signing.

The ability to augment anonymous tokens with public metadata serves as a way to propagate trust for a specific subset of settings. In most applications using anonymous tokens, the signer is typically responsible for checking that the user is permitted to receive tokens. During verification, the verifier no longer has any information about the user and, thus, may only check that tokens are valid. Public metadata enables propagating this permission in a fine-grained manner. The verifier can check this public metadata with trust that the signer embedded the public metadata. We present two real-world example applications where public metadata is critical.

Application 1: Geolocality-based VPN. As a real-world application, we can consider VPN where each user purchases access for a specific geolocality such as a country. Whenever a user accesses the VPN, it is critical that the user’s anonymity is maintained. However, the VPN provider also needs to check that the accessing user has paid for access in the specified country. In fact, this is the exact problem encountered by GoogleOne VPN [6]. We show that this can be solved with public metadata for anonymous tokens. To access the VPN, a user is first required to authenticate to prove that they have purchased access for a specific country. Afterwards, the user will receive several anonymous tokens where the public metadata is the country for authorized access. When the user accesses the VPN, the user will redeem anonymous tokens. The VPN provider can verify that each token was created for the correct country before permitting VPN access.

Application 2: Short-Lived Anonymous Tokens. As another example, we consider the general setting where anonymous tokens should have short expiration times or may be redeemed within a short time period (for example, 1 hour). This could be useful for many applications such as timed event access or promotions with expiration. A trivial approach is to utilize different keys for each expiration time or valid time period (such as one key per hour). However, managing a large number of keys that is a very challenging problem in real-world deployments [51, 32, 21]. For example, it requires a large number of key rotations such as once every hour. Typically, key rotations are performed after a certain number of operations to avoid security degradation or mitigating accidental key exposure. In this case, we are abusing key rotation to enable expirations. The large number of key rotations may introduce additional problems (such as missing keys, exposure, etc). It is advantageous to keep the key set small to avoid these potential issues. Instead, we can

use public metadata to avoid unnecessary key rotations. When issuing a token, the signer finds the correct expiration time (or valid time period) for each user and embeds it as public metadata. Then, the verifier will only permit access to the user if the anonymous token is redeemed at the right time period.

1.1 Our Contributions

RSA Blind Signatures with Public Metadata. As our main contribution, we present a blind signature with public metadata based on the RSA assumption that is a modification of a scheme introduced by Abe and Fujisaki [9]. The construction uses only standard cryptography widely available across most platforms while avoiding more complex and less supported operations (such as pairings). We formally prove unforgeability from a variant of the one-more RSA assumption as well as unconditional unlinkability. Additionally, we derive concrete security bounds to enable picking parameters in real-world deployments. We note that the prior work [9] did not have formal definitions or security proofs. One of this paper’s main contributions is proving that an adjusted version of their scheme is secure. Our protocol underwent a security review and audit by the NCC group that also verified our paper’s claim in their public report [61]. Finally, we show that our protocol incurs only slight overhead over standard RSA blind signatures. We note that there are prior schemes for public metadata that rely on cryptography with more limited production availability such as pairings [57]. In Appendix G, we compare the availability of the necessary RSA operations for our scheme to pairings in production cryptography libraries.

IRTF Draft. The protocol in this work has been proposed in a CFRG specification in the IRTF [12] due to the interest of practitioners for real-world applications. By constructing practical anonymous tokens with public metadata using widely available production cryptography, our work unblocks a significant number of new applications. For example, there are already work streams to build real-world architectures relying on our construction in the IETF Privacy Pass working group (such as [37]). One can view this work as the accompanying academic paper proving formal cryptographic guarantees for the usage of these protocols in real-world applications.

Experimental Evaluation and Deployment Telemetry. We perform experimental evaluation to show our protocol incurs minimal overhead over standard RSA blind signatures. Additionally, our protocol has been deployed in an application for Google VPN. We report on anonymous real-world telemetry to showcase the scalability of our protocol to millions of users.

1.2 Technical Overview

Before we present our protocol for RSA blind signatures with public metadata, we first show some failed approaches that provide insights into the design choices of our final algorithm. In each of these approaches, we take an underlying anonymous token scheme and aim to enable public metadata.

One Key for Each Public Metadata. The first idea is to provide isolation between the different choices of public metadata using different keys. Formally, for each potential option of public metadata D , we can generate a new RSA key pair (pk_D, sk_D) . This is equivalent to having a separate anonymous token system for each choice of public metadata. As the public metadata D is known to the signer and users, each party can use the appropriate keys corresponding to D to performing blind signing and verification.

While the above approach would work, the main downside of this approach is that key management becomes significantly more challenging with large key sets. As mentioned earlier, key management is one of the most difficult aspects in real-world deployments [51, 32, 21]. Instead, we want a protocol where the number of key pairs does not grow with the public metadata universe. The other disadvantage of this approach is that if the universe of public metadata changes, all parties in the system must also update their key sets accordingly to include the new public and private keys associated with the new public metadata. Given the above, we want an anonymous token with public metadata scheme that does not require multiple keys while still providing the flexibility of a potentially changing set of public metadata.

Failed Solution: Embedding into Message. Another approach is to append the public metadata D to the message M to obtain message $M' = M \parallel D$. Afterwards, the anonymous token scheme remains identical using new message M' . At first, this seems to achieve the desired goal as the privacy of the signature and message follows directly from the underlying anonymous token scheme. Furthermore, if one attempts to use the wrong public metadata, verification will clearly fail as the input message is incorrect. Unfortunately, this approach would only work in the case when we are considering signature schemes that do not provide anonymity. In particular, the metadata must be publicly known to the signer during the signature process. If we wish to maintain anonymity, then the message $M' = M \parallel D$ cannot be sent to the signer in plaintext. Therefore, the signer can no longer verify that they are signing with respect to a specific piece of public metadata D . As an example, consider an application that requires associating each signature to a specific area using public metadata. Even if the signer wishes to only create signatures for a specific public metadata D (such as a specific country), the signer is unable to check the public metadata for the specific blind signing request. Therefore, this approach fails to satisfy the necessary unforgeability requirements for adversarial users.

Existing Solutions. There are other potential solutions to support public metadata in anonymous tokens. Silde and Strand [57] present a pairing-based protocol based on BLS signatures [19]. In general, pairings are computationally expensive and not widely available in production libraries. Several works [40, 58] consider pairing-free solutions, but require three moves (resulting in multiple roundtrips during blind signing). There are other options such as Albrecht *et al.* [11] using torus-based fully homomorphic encryption (FHE), which is even more expensive than pairings. We consider two-move protocols using efficient cryptography to enable easier adoption.

Our Protocol. Using the above approaches as guidance, we can determine requirements necessary for our construction. First, we should only use a single key independent of the public metadata universe. Additionally, we still need the public metadata to be viewable by the signer during blind signing. Finally, we should avoid any complex cryptography that is not widely supported.

We first revisit the original RSA blind signature protocol. Recall that standard RSA blind signatures utilize a modulus that is the product of two primes $N = pq$ along with a public exponent e and a private exponent d such that $d = e^{-1} \pmod{\phi(N)}$ ¹. The signer’s private key consists of d while the public key is (N, e) . In the rest of the description, all operations are done in \mathbb{Z}_N^* . To perform blind signing for a message M , the user computes $A = H_{\mathcal{M}}(M) \cdot R^e$ where $H_{\mathcal{M}}(M)$ hashes the message M to an element of \mathbb{Z}_N^* and R is a uniformly random element of \mathbb{Z}_N^* . This is sent to the signer who signs by computing $B = A^d = H_{\mathcal{M}}(M)^d \cdot R^{ed} = H_{\mathcal{M}}(M)^d \cdot R$. Finally, the user computes $B \cdot R^{-1} = H_{\mathcal{M}}(M)^d \cdot R \cdot R^{-1} = H_{\mathcal{M}}(M)^d$ that is the final signature. To verify any signature S for message M using public key (e, N) , one simply computes S^e and checks if it equals to $H_{\mathcal{M}}(M)$.²

We take insight from the first failed approach that used different keys for each metadata along with prior work of Abe and Fujisaki [9]. Rather than using different keys, we can try to use a hash function to enable generating keys for each metadata. For RSA blind signatures, we will use a hash function H_{MD} to generate public and private exponents that are specific to some public metadata D . In more detail, we will set $e_{\text{MD}} = H_{\text{MD}}(D)$ as the public exponent as done in the scheme presented in [9]. In our scheme, we use a salt in H_{MD} but omit it here for simplicity. Blind signing for the user remains similar using e_{MD} as opposed to e to obtain $A = H_{\mathcal{M}}(M \parallel D) \cdot R^{e_{\text{MD}}}$ with the exception that we also put the public metadata D into the message hash now. Now, the signer receives both the blind signing request A as well as the public metadata D . First, the signer computes $e_{\text{MD}} = H_{\text{MD}}(D)$. We augment the private key to contain $\phi(N)$ allowing the signer to compute the inverse $d_{\text{MD}} = (e_{\text{MD}})^{-1} \pmod{\phi(N)}$. Afterwards, the signer uses d_{MD} as the private exponent to return $B = H_{\mathcal{M}}(M \parallel D)^{d_{\text{MD}}} \cdot R$ to the user. Finally, the user removes R to obtain $H_{\mathcal{M}}(M \parallel D)^{d_{\text{MD}}}$ that is the final signature.

¹Euler’s totient function $\phi(N)$ counts the positive integers up to the given integer N that are relatively prime to N and thus invertible modulo N . In other words, it outputs the order of the group \mathbb{Z}_N^* .

²For sake of simple exposition, we assumed that $H_{\mathcal{M}}$ is deterministic. However, there exists other message encodings that are randomized with more complex verification algorithms. See Section 5.1 for more details.

Unfortunately, the above construction has a slight problem. In particular, we assumed that the output $e_{\text{MD}} = H_{\text{MD}}(D)$ is always invertible modulo $\phi(N)$ and, thus, co-prime to $\phi(N)$. For standard RSA modulus $N = pq$ where p and q are distinct primes, it is not necessarily the case that a random element would be invertible modulo $\phi(N) = (p - 1) \cdot (q - 1)$. To solve this problem, we can use strong RSA moduli where we require that $N = pq$ such that both p and q are safe primes meaning that $p = 2p' + 1$ and $q = 2q' + 1$ such that both p' and q' are also prime numbers. If N is a strong RSA modulus, then we know that $\phi(N) = (p - 1) \cdot (q - 1) = 4p'q'$. Therefore, an element $x \in \mathbb{Z}_{\phi(N)}$ has an inverse as long as x is odd and not divisible by p' and q' . If we assume that both p and q are κ -bit prime numbers, then we can guarantee that the output of H_{MD} is always invertible in $\mathbb{Z}_{\phi(N)}$ using the following modifications. First, we ensure that H_{MD} always outputs an odd number. Next, we make sure that H_{MD} always outputs elements of length at most $\kappa - 2$ bits. As p and q are κ bits, we know that both p' and q' are at least $\kappa - 1$ bits. As a result, we can guarantee that H_{MD} is always invertible modulo $\phi(N)$ as it is odd and always smaller than both p' and q' . We note that similar ideas were previously presented in [9, 8]. However, to our knowledge, no formal security proofs have been provided for these constructions.

One-More RSA Inversion with Multiple Exponents. Before proving the security of our protocol, we first need to define an extension of the RSA assumption using the “one-more” style techniques of Bellare, Namprempre, Pointcheval and Semanko [15]. We explore several natural ways to define RSA assumptions with respect to multiple exponents before arriving at the weakest form that may be used to prove security of our protocol. To our knowledge, this is the first exploration of one-more RSA inversion assumptions with multiple exponents. Finally, we also explore connections with various natural definitions of one-more strong³ RSA type assumptions.

Security of Our Protocol. Finally, we prove the security of our new protocol with public metadata. We start with proving concrete security of the non-blind variant from the RSA assumption. In particular, we show that our new protocol has similar concrete security guarantees for unforgeability as standard RSA signatures. We note that our subtle modification where the underlying hash $H_{\mathcal{M}}(M || D)$ includes both the message M and the public metadata D is integral in our security proof.

Next, we move onto proving the security of our main protocol of RSA blind signatures with public metadata protocol. Using the one-more multi-exponent RSA inversion assumptions, we are able to prove the unforgeability of our protocol. For anonymity, we adapt the techniques introduced by Lysyanskaya [42] to ensure and prove that our protocols satisfy unlinkability even in the presence of keys that are generated by a malicious party.

Underlying Cryptography Operations. One benefit of starting from RSA blind signatures is that all the underlying cryptographic operations are widely supported. The only additional functionalities required by our protocol is the ability to hash strings to random numbers for H_{MD} and generating random safe primes (only needed for key generation). These are cryptographic operations that are widely supported for production usage. In contrast, pairings (relied upon by prior works such as [57]) are not yet available in most widely used production cryptography libraries. Although, recent interest in pairing-based cryptography (such as IRTF draft [18]) may lead to pairings being more widely available in the future. See Appendix G for availability of operations in production libraries.

1.3 Related Work

BLS with Public Metadata. We note that a prior work of Silde and Strand [57] also studied the notion of anonymous tokens with public metadata where several constructions were presented. For the setting of public verifiability, the authors presented a construction from pairings based on the Boneh, Lynn and Shacham (BLS) signature scheme [19] along with ideas from other works including [64, 31]. Comparing with [57], our protocol uses RSA-based cryptography as opposed to pairing-based cryptography. RSA-based cryptography is more readily available on all platforms compared to pairing-based cryptography enabling

³In the strong RSA assumption, the adversary gets to pick its own public exponent (Section 4.3). This is unrelated to a strong RSA modulus that is the product of two safe primes.

easier adoption. We note that recent IRTF work [18] may lead to pairings being more widely available for production usage in the future. See Appendix G for comparison of production availability of pairings compared to the RSA operations for our scheme.

Partial OPRF with Public Metadata. Several works have studied partially oblivious pseudo-random functions (POPRF) including [62, 57, 64] that essentially allows public metadata for OPRF evaluations. We note that there was a proof flaw in [64] that was fixed in [62] and all schemes are secure.

RSA Blind Signatures with Strong Moduli. We note that similar variants of RSA blind signatures has appeared in the past by Abe and Fujisaki [9] as well as Abe and Camenisch [8]. However, the security was proven heuristically and predated the standard definitions used in anonymous tokens. One can view the main contribution of our work as proving a modified variant of the scheme in [9] to be secure along with concrete security bounds. Furthermore, we combine recent work [42] to obtain protection against maliciously generated keys. Other works also used strong RSA modulus such as threshold signatures [55] and integer commitments [29].

Anonymous Tokens with Private Metadata. Recent work [41] studied anonymous tokens augmented with a single private metadata bit along with publicly verifiable variants [17, 47]. The metadata bit is explicitly specified only by the signer and observable by the verifier while hidden from the user. Another recent work [22] studied augmenting private metadata into anonymous tokens by using algebraic MACs as building blocks for their constructions instead of pseudo-random functions. They also showed how their constructions can be extended to support public metadata. Moreover, by using techniques similar to [57] they were able to achieve public verifiability although without private metadata support. To our knowledge, blind signatures cannot be used to instantiate anonymous tokens with private metadata. We leave it as future work to support private metadata bits using RSA-style signatures.

2 Anonymous Tokens

We present the formal definitions for anonymous token schemes. Our definitions will work for both anonymous and non-anonymous protocols with small modifications. We will solely focus on the settings of public metadata and public verifiability (that is, verification only requires the public key). Our definitions are interchangeable with the notion of partially blind signature schemes [8, 10] and we explicitly focus on two-move (message-response structure) blind signatures. Non-two-move signatures exist such as [10] but due to the the extra communication round-trip we do not focus on them in this work. Other variants of anonymous token schemes exist such as those only allowing verification using the private key (such as [31, 57]) as well as enabling private metadata (see [41]).

Definition 1. *A token scheme with public metadata Tok that is publicly verifiable is a tuple of efficient algorithms $\text{Tok} = (\text{Setup}, \text{Blind}, \text{Sign}, \text{Finalize}, \text{Verify})$.*

1. $(\text{pk}, \text{sk}) \leftarrow \text{Tok.Setup}(1^\lambda)$: *The setup algorithm receives the security parameter λ as input and outputs a pair of public and private keys (pk, sk) .*
2. $(\text{st}, B_M) \leftarrow \text{Tok.Blind}(M, D, \text{pk})$: *The blinding protocol is run by the user who receives the plaintext message M , public metadata D and public key pk . The output B_M is a blinded version of the message M under public metadata D and some state st .*
3. $S' \leftarrow \text{Tok.Sign}(B_M, D, \text{pk}, \text{sk})$: *The signing protocol is run by the signer who receives the blinded message B_M , public metadata D and both public and private keys (pk, sk) . The output S' is a signature on the blinded message B_M under public metadata D which needs to be finalized by the user.*
4. $S \leftarrow \text{Tok.Finalize}(\text{st}, S', M, D, \text{pk})$: *The user runs the Finalize protocol on the user state st , the signer's response S' , plaintext message M , public metadata D and public key pk . The output S is a signature of the message M under public metadata D .*

5. $b \leftarrow \text{Tok.Verify}(S, M, D, \text{pk})$: The verification algorithm receives the signature S , the plaintext message M , public metadata D and public key pk and outputs a bit $b \in \{0, 1\}$.

The token scheme Tok satisfies the correctness properties if, for all choices of messages M and public metadata D , the following holds:

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Tok.Setup}(1^\lambda) \\ (\text{st}, B_M) \leftarrow \text{Tok.Blind}(M, D, \text{pk}) \\ S' \leftarrow \text{Tok.Sign}(B_M, D, \text{pk}, \text{sk}) \\ S \leftarrow \text{Tok.Finalize}(\text{st}, S', M, D, \text{pk}) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

We note that the $\text{Blind}, \text{Sign}, \text{Finalize}$ protocols can be viewed as a single round-trip protocol between the user and signer. One could generalize the above definition to encompass multiple round, interactive protocols. However, the structure of the $\text{Blind}, \text{Sign}, \text{Finalize}$ protocol will be useful for proving unforgeability of our schemes. Therefore, we only focus on this structure throughout our work.

Definition for Non-Anonymous Tokens. To obtain the standard definition of non-anonymous tokens, we note that we can simply restrict the functionality of Blind and Finalize to perform no operations. We will use this when evaluating the concrete security of RSA signatures with public metadata.

1. $\text{Tok.Blind}(M, D, \text{pk})$:
 - (a) Return $(\text{st} \leftarrow \perp, B_M \leftarrow M)$.
2. $\text{Tok.Finalize}(\text{st}, S', M, D, \text{pk})$:
 - (a) Return $S \leftarrow S'$.

Relation with Blind Signatures. In our work, we treat anonymous tokens synonymously with blind signatures. In general, blind signatures may be used to instantiate anonymous tokens (as we do in our work). However, anonymous tokens may be more general and may allow for designated (secret key) verifiers as studied in [37]. As we focus exclusively on public verification with public metadata, Definition 1 is equivalent to partially blind signatures [10, 8]. We also note that it is unknown whether standard blind signatures may be used to instantiate private metadata [41].

2.1 Unforgeability

We start with the first cryptographic guarantee provided by tokens known as unforgeability. At a high level, unforgeability guarantees that no party is able to generate signatures that will correctly verify unless they have access to the private signing key. For our work, we will consider a modification of the standard definition of unforgeability for blind signatures introduced by Schröder and Unruh [53]. This standard definition is known as *strong one-more unforgeability* as it was a strengthening of previous definitions such as those introduced in [49, 15].

We modify the definition of unforgeability to enable public metadata in the following way. In the prior definitions (without public metadata), the adversary wins if it can construct $\ell + 1$ signatures using at most ℓ signing oracle queries. In our game to incorporate public metadata, the adversary succeeds if it is able to construct at least $\text{cnt}_D + 1$ signatures for any choice of public metadata D using at most cnt_D signing oracle queries with public metadata D . For example, this covers the case even when an adversarial user can forge a signature for some metadata D without ever sending a signing oracle query with D .

The security game can be found in Figure 1 and we present the formal definition below. We present our definition using concrete parameters for adversarial advantage ϵ , adversary running time t and number of oracle queries ℓ . This enables us to prove concrete security bounds throughout our paper.

Game $G_{\text{Tok}, \mathcal{A}}^{\text{SOMUF}}(\lambda)$:	Oracle $\mathcal{O}^{\text{Sign}}(M, D)$:
$(\text{pk}, \text{sk}) \leftarrow \text{Tok.Setup}(1^\lambda)$ Initialize $\text{cnt}_D \leftarrow 0$ for all choices of public metadata D . $D, (S_i, M_i)_{i \in [x]} \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Sign}}}(\text{pk})$ Return 1 if and only if all the following hold: - $\text{cnt}_D < x$ - $\forall i \neq j \in [x], S_i \neq S_j \vee M_i \neq M_j$ - $\forall i \in [x], \text{Tok.Verify}(S_i, M_i, D, \text{pk}) = 1$	$\text{cnt}_D \leftarrow \text{cnt}_D + 1$ Return $\text{Tok.Sign}(M, D, \text{pk}, \text{sk})$.

Figure 1: Strong One-More Unforgeability (SOMUF) Game with Public Metadata.

Definition 2 ((ϵ, t, ℓ) -Strong One-More Unforgeability). *Let λ be the security parameter and consider the game $G_{\text{Tok}, \mathcal{A}}^{\text{SOMUF}}$ in Figure 1. A token scheme Tok is (ϵ, t, ℓ) -strong one-more unforgeable if, for any adversary \mathcal{A} that runs in probabilistic time t and makes at most ℓ signing queries, the following holds*

$$\Pr[G_{\text{Tok}, \mathcal{A}}^{\text{SOMUF}}(\lambda) = 1] \leq \epsilon.$$

Difference with Previous Definition. A slightly different definition of unforgeability for public metadata was introduced in [57]. For this game, the adversary is able to make at most ℓ signing queries for any metadata. Then, the adversary wins the game if they pick some metadata D and can create $\ell + 1$ valid signatures for D . In our work, we instead bound the total number of signing queries by ℓ and the adversary wins the game if they are able to generate $\text{cnt}_D + 1$ valid signatures for any metadata D while making at most cnt_D oracle queries for metadata D . The reason we do this is to be able to derive concrete security bounds with respect to ℓ . Using the prior definitions with public metadata in [57], we are unable to derive an upper bound on the number of signing queries beyond the running time of the adversary as the adversary may continue to make ℓ oracle queries for new choices of public metadata D . In other words, we can only guarantee that the adversary executes a probabilistically polynomial number of queries when considering probabilistically polynomial time adversaries.

Furthermore, we note that our definition is no weaker than the prior definition. If any adversary can produce $\ell + 1$ signatures using at most ℓ oracle queries, then there must exist some public metadata D such that the adversary has more valid signatures than oracle signing queries. We point readers to Appendix B for more details and showing the equivalence between the prior definition and our current definition.

Multiple Signers. The above definition of unforgeability only considers a single signer whereas general settings would consider multiple signers. Therefore, it may be natural and important to extend the definition for multiple signers. However, it turns out that the single and multiple signers are equivalent up to some factors that depend only on the number of signers. Furthermore, the strong one-more unforgeability is identical in either the single or multiple signer settings for two-move protocols (as studied in this paper). So, it is sufficient to focus security proofs with respect to a single signer. For more details, we point readers to [42].

Anonymous vs. Non-Anonymous Tokens. We will use the same definition for both anonymous and non-anonymous tokens and slightly overload notation for convenience. For non-anonymous tokens, we will assume that the input to the signing oracle will be the plaintext message that needs to be signed. In the case of anonymous tokens, the input to the signing oracle will be the blinded (i.e., encrypted) message where the signing oracle will perform the signer's actions to manipulate the blinded message such that the querier can retrieve the signed message.

Verification Oracles. In past definitions, an oracle for verification was also provided to the adversary. Although, this is only needed for the designated verifier setting where only certain parties can perform verification. For the case of RSA (blind) signatures, verification may be performed by anyone with access to public keys (that is, RSA signatures enable public verifiability). In other words, any adversary can simply execute the verification algorithm without the usage of an oracle. Therefore, there is no need for a verification oracle when considering publicly verifiable schemes.

<p>Game $G_{\text{Tok}, \mathcal{A}}^{\text{UNLINK}}(\lambda, b)$:</p> <p>Adversary \mathcal{A} outputs $(\text{pk}, M_0, M_1, D) \leftarrow \mathcal{A}(1^\lambda)$.</p> <p>For $i \in \{0, 1\}$:</p> <p>Challenger \mathcal{C} computes $(B_i, \text{st}_i) \leftarrow \text{Tok.Blind}(\text{pk}, M_i, D)$.</p> <p>$\mathcal{C}$ samples uniformly random bit $b \leftarrow_R \{0, 1\}$.</p> <p>\mathcal{A} receives B_b, B_{1-b} and computes S'_b, S'_{1-b}.</p> <p>For $i \in \{0, 1\}$:</p> <p>\mathcal{C} computes $S_i \leftarrow \text{Tok.Finalize}(\text{pk}, \text{st}_i, S'_i, M_i, D)$.</p> <p>If $\text{Tok.Verify}(M_0, S_0, D, \text{pk}) = 1 \wedge \text{Tok.Verify}(M_1, S_1, D, \text{pk}) = 1$:</p> <p>$\mathcal{C}$ sends $(M_0, S_0), (M_1, S_1)$ to \mathcal{A}.</p> <p>Else:</p> <p>\mathcal{C} sends \perp to \mathcal{A}.</p> <p>\mathcal{A} outputs a bit b'.</p>
--

Figure 2: Unlinkability (UNLINK) Game.

2.2 Unlinkability

The second important cryptographic guarantee is anonymity that is typically referred to formally as unlinkability. In an anonymous token scheme, it must be impossible to determine the blind signing request that was utilized to create any signature even from the view of a signer that views the signing interactions, the final signatures as well as the corresponding input message of each signature.

We use the definition of unlinkability following the work from [7] for single-round schemes that we focus on throughout our work. We modify the definitions to encompass public metadata. In this game-based definition, the adversary picks the public key, two messages M_0 and M_1 as well as its choice of public metadata D . The challenger will first blind both messages and randomly permute the blinded messages. The blinded messages are submitted to the adversary in the randomly permuted order that will return signatures on the blinded messages. Finally, the challenger finalizes the adversary’s responses to obtain the final signatures for both messages that are given to the adversary according to the original message order. The goal of the adversary is to guess the correct permutation of the original blinded messages. This security game is formally presented in Figure 2 and we present our definition of unlinkability below:

Definition 3 ((ϵ, t) -Unlinkability). *Let λ be the security parameter and consider the game $G_{\text{Tok}, \mathcal{A}, q}^{\text{UNLINK}}$ in Figure 2. A token scheme Tok satisfies (ϵ, t) -unlinkability if, for any adversary \mathcal{A} running in time t , the following holds:*

$$|\Pr[G_{\text{Tok}, \mathcal{A}}^{\text{UNLINK}}(\lambda, 0) = 1] - \Pr[G_{\text{Tok}, \mathcal{A}}^{\text{UNLINK}}(\lambda, 1) = 1]| \leq \epsilon.$$

Discussion about Public Metadata. In our definition, the adversary is able to choose any choice of public metadata. However, the same public metadata must be used for signing both messages. This is necessary as if the adversary may pick different public metadata for each message then it can trivially win the game just by attempting to verify signatures with the two different public metadata. In practice, this implies that anonymity or unlinkability only applies to all groups of messages signed with the same public metadata. For real-world applications, it is important that the set of possible choices of public metadata is not too large to ensure anonymity is maintained. In an extreme case, if each message is signed with a unique choice of public metadata, then there are no anonymity guarantees anymore.

Adversarial Signer. Our choice of unlinkability definitions requires anonymity even against adversarial signers. In particular, we note that the adversary is able to choose the public key arbitrarily. Furthermore, the adversary is free to compute signatures on the blinded messages arbitrarily. Therefore, this encompasses all malicious signers that try to compromise anonymity.

Necessity of Successful Verification. In our definition, we note that the adversary is only permitted to see the resulting signatures if both signatures successfully verified. This is necessary to rule out one naive

Game $\text{G}_{\text{Gen}, \mathcal{A}}^{\text{RSA}}(\lambda)$:
$(N, \mathcal{D}_N, (p, q)) \leftarrow_R \text{Gen}(1^\lambda)$
$e \leftarrow_R \mathcal{D}_N$
$X \leftarrow_R \mathbb{Z}_N^*$
$Y \leftarrow \mathcal{A}(N, e, X)$
Return 1 if and only if $Y^e = X \pmod N$.

Figure 3: RSA Game.

strategy for the adversary where it would issue a valid signature in response to one of the queries but not the other. In order to rule out this strategy, the challenger allows the adversary to see the resulting signatures only if both of them verify. If one (or both) of the signatures do not verify, the adversary will have to guess the ordering without seeing the finalized signatures.

3 Prior RSA Assumptions

In this section, we outline prior RSA assumptions that are relevant to our work. In particular, we present the standard RSA assumption as well as one-more RSA inversion assumptions.

3.1 Strong RSA Modulus

We start by discussing the RSA group structure. In our work, we will focus on a special subset of RSA modulus that are strong RSA modulus. All strong RSA modulus $N = pq$ are the product of two safe primes that enables structure to the multiplicative group modulo $\phi(N)$.

Definition 4 (Strong RSA Modulus). *An integer N is a strong RSA modulus $N = p \cdot q$ where each of p and q are distinct safe primes. In other words, $p = 2p' + 1$ and $q = 2q' + 1$ where all of p, p', q, q' are distinct prime numbers. Therefore, $\phi(N) = 4p'q'$.*

As a side note, a key property of strong RSA modulus is that it provides significant structure to the group $\mathbb{Z}_{\phi(N)}^*$. One observation that we will rely upon throughout our work is that a random odd number from the set $[3, 2^{\kappa-2}]$ will be co-prime to $\phi(N)$. This is because $\phi(N) = 4p'q'$. Therefore, any odd number strictly smaller than both p' and q' is co-prime to $\phi(N)$. Since we will assume that each of p' and q' are at least $\kappa - 1$ bits and p and q are κ bits, any odd number no larger than $2^{\kappa-2}$ is co-prime to $\phi(N)$.

3.2 RSA Assumption

We will denote generating a random modulus, the public exponent distribution and secret parameters by $(N, \mathcal{D}_N, p, q) \leftarrow_R \text{Gen}(1^\lambda)$. The prime bit-length κ is chosen to obtain λ bits of security. We move onto defining the RSA assumption. At a high level, the RSA assumption assumes that any probabilistically polynomial time (PPT) adversary \mathcal{A} given a RSA modulus N and public exponent e coprime to $\phi(N)$ is unable to compute the e -th root modulo N for a random element $X \leftarrow_R \mathbb{Z}_N^*$ drawn from the multiplicative group modulo N . In other words, the PPT adversary \mathcal{A} is unable to compute $X^d \pmod N$ where $d = e^{-1} \pmod{\phi(N)}$. In general, the RSA assumption is a class of assumptions that is parameterized by some distribution \mathcal{D}_N determining how to pick the public exponent e . Thus, the Gen algorithm also outputs a distribution \mathcal{D}_N to sample e . Furthermore, for concrete security bounds that may be used for guiding practical implementations, we consider a more fine-grained version of the RSA assumption that is additionally parameterized by the running time t and advantage ϵ of the adversary. Throughout the rest of this paper, we will present all security with respect to concrete parameters. As a result, all our definitions will always have fine-grained parameters.

Definition 5 ((ϵ, t) -RSA Assumption). *Let λ be the security parameter and consider the game $\mathbf{G}_{\text{Gen}, \mathcal{A}}^{\text{RSA}}(\lambda)$ in Figure 3. The (ϵ, t) -RSA assumption is true for Gen if for any PPT adversary \mathcal{A} that runs in time t ,*

$$\Pr[\mathbf{G}_{\text{Gen}, \mathcal{A}}^{\text{RSA}}(\lambda) = 1] \leq \epsilon.$$

In our work, we will focus on the setting when Gen always outputs strong RSA modulus N that is the product of two safe primes. We do not use the safe primes for reasons related to strengthening the RSA assumption. Instead, we use safe primes to utilize the additional structure in the multiplicative group $\mathbb{Z}_{\phi(N)}^*$. In particular, we show that safe primes enable an efficient and simple method of hashing public metadata to a large subset of elements in $\mathbb{Z}_{\phi(N)}^*$ without requiring knowledge of $\phi(N)$. We point readers to Section 5.2 for more details on our usage of the strong RSA modulus. We present our definition of the RSA assumption restricted to strong RSA modulus below.

Definition 6 ((ϵ, t) -RSA Assumption for Strong Modulus). *Let λ be the security parameter and consider the game $\mathbf{G}_{\text{Gen}, \mathcal{A}}^{\text{RSA}}(\lambda)$ in Figure 3. The (ϵ, t) -RSA assumption is true for Gen such that Gen only outputs strong RSA modulus and, if for any PPT adversary \mathcal{A} that runs in time t ,*

$$\Pr[\mathbf{G}_{\text{Gen}, \mathcal{A}}^{\text{RSA}}(\lambda) = 1] \leq \epsilon.$$

As the RSA assumption for strong modulus considers only a subset of possible modulus, it seems natural that it is no weaker than the standard RSA assumption (Theorem 1). We formally prove this relation in the next theorem that utilizes prime number densities and the conjectured densities of safe prime numbers. We only assume that Gen generates N as the product of two uniformly random κ -bit primes, which is common practice (see BoringSSL [2] for example). See Appendix A for the proof of Theorem 1 and further details.

Theorem 1. *If the (ϵ, t) -RSA Assumption (Definition 5) is true and Gen generates modulus N as the product of two uniformly random κ -bit primes, then the $(O(\epsilon \cdot \kappa^2), t)$ -RSA Assumption for Strong Modulus (Definition 6) is also true.*

In our proofs, we will prove the unforgeability of our RSA (non-blind) signatures with public metadata using the RSA assumption for strong modulus. Using Theorem 1, one can then re-interpret all our results as assuming the standard RSA assumption.

Finally, we note that the above reduction works for all of the assumptions in our paper. In particular, we can interchangeably use various RSA assumptions for arbitrary modulus as well as strong modulus. Throughout the rest of our work, we will use these assumptions interchangeably at the costs of an $O(\epsilon \kappa^2)$ multiplicative factor to the security parameter.

Discussion about \mathcal{D}_N . In practical implementations, the distribution \mathcal{D}_N is typically a singleton fixed for all modulus N . The most common public exponents used in RSA implementations are 3 and 65537. On the other hand, theoretical works will typically consider \mathcal{D}_N to be non-trivial distributions such as uniform from $\mathbb{Z}_{\phi(N)}^*$ or some subset of $\mathbb{Z}_{\phi(N)}^*$ (such as [27]). We will follow the same approach in our work and assume \mathcal{D}_N outputs random odd γ -bit integers where γ is chosen later in our constructions.

Discussion about Gen Trapdoor. In our definition of Gen , the trapdoor output is essentially two prime factors p and q of the RSA modulus N . Prior works defined Gen to return the decryption exponent $d = e^{-1} \bmod \phi(N)$ as opposed to the prime factorization of N . We note the two outputs are equivalent. Given the prime factors p and q of N , one can easily compute $\phi(N)$ and thus invert e modulo $\phi(N)$. In fact, we could also choose to define Gen to simply return $\phi(N)$ that is sufficient for our work. On the other hand, it is well known that one can also compute the prime factors of N using only d (see [27] for example). Therefore, the definitions are equivalent and we choose to return p and q for convenience.

3.3 One-More RSA Inversion Assumption

Security of RSA blind signature schemes have been proven by using of a class of computational problems known as one-more RSA inversion problems introduced by Bellare, Namprempre, Pointcheval and

Game $\mathsf{G}_{\text{Gen}, \mathcal{A}}^{\text{CT-RSA}}(\lambda)$:	Oracle $\mathcal{O}^{\text{RSA}}(X)$:	Oracle $\mathcal{O}^{\mathcal{X}}()$:
$(N, \mathcal{D}_N, (p, q)) \leftarrow_R \text{Gen}(1^\lambda)$ $e \leftarrow_R \mathcal{D}_N$ $\phi(N) \leftarrow (p-1)(q-1)$ $\text{cnt} \leftarrow 0, \mathcal{S}_{\mathcal{X}} \leftarrow \emptyset$ $(X_i, Y_i)_{i \in [\text{cnt}+1]} \leftarrow \mathcal{A}^{\mathcal{O}^{\text{RSA}}, \mathcal{O}^{\mathcal{X}}}(N, e)$ Return 1 if and only if all the following hold: <ul style="list-style-type: none"> - $\forall i \neq j \in [\text{cnt}+1], X_i \neq X_j$ - $\forall i \in [\text{cnt}+1], X_i \in \mathcal{S}_{\mathcal{X}}$ - $\forall i \in [\text{cnt}+1], Y_i^e = X_i \bmod N$ 	$\text{cnt} \leftarrow \text{cnt} + 1$ $d \leftarrow e^{-1} \bmod \phi(N)$ $Y \leftarrow X^d \bmod N$ Return Y	$X \leftarrow_R \mathbb{Z}_N^*$ $\mathcal{S}_{\mathcal{X}} \leftarrow \mathcal{S}_{\mathcal{X}} \cup \{X\}$ Return X

Figure 4: Chosen-Target RSA Inversion Game.

Semanko [15]. In these problems, the adversary is given access to a decryption oracle \mathcal{O}^{RSA} and a challenge oracle $\mathcal{O}^{\mathcal{X}}$. The decryption oracle takes as input $X \in \mathbb{Z}_N^*$ and returns $\mathcal{O}^{\text{RSA}}(X) = Y$ where $Y^e = X \bmod N$. The challenge oracle $\mathcal{O}^{\mathcal{X}}$ will return random elements from \mathbb{Z}_N^* as random challenge targets. The adversary only succeeds if it inverts $\ell + 1$ challenge targets while making at most ℓ queries to the decryption oracle \mathcal{O}^{RSA} . That is, it computes at least one more RSA decryption than the number of oracle queries. We define chosen-target RSA inversion game in Figure 4 along with the following definition that were both introduced in [15]. In this definition, *chosen-target* implies that the adversary may choose to invert any of the targets output by the oracle $\mathcal{O}^{\mathcal{X}}$.

Definition 7 ((ϵ, t, ℓ) -Chosen-Target RSA Inversion Assumption). *Let λ be the security parameter and consider the game $\mathsf{G}_{\text{Gen}, \mathcal{A}}^{\text{CT-RSA}}(\lambda)$ in Figure 4. The (ϵ, t, ℓ) -chosen-target RSA inversion assumption is true for Gen if, for any adversary \mathcal{A} that runs in time t and makes at most ℓ decryption queries, the following holds*

$$\Pr[\mathsf{G}_{\text{Gen}, \mathcal{A}}^{\text{CT-RSA}}(\lambda) = 1] \leq \epsilon.$$

Discussion about Challenge Oracle. In the formulation of the chosen-target RSA inversion game, the adversary has access to a target oracle $\mathcal{O}^{\mathcal{X}}$ that it can use to get a set of potential targets to invert. Another definition may be to suppose that the set of possible targets is provided to the adversary as input instead. It turns out that the two definitions are equivalent as shown in [15]. Throughout the rest of the work, we will consider this formulation where the adversary will have access to a challenge oracle.

4 RSA Assumptions with Multiple Exponents

In this section, we introduce and explore various generalizations of the RSA assumption with multiple exponents. Afterwards, we further extend them to one-more variants (following the definitional extensions for the RSA assumption in [15]). To our knowledge, we are unaware of prior works that formally defined multi-exponent variants of the RSA assumption. In our work, we will explore various natural generalizations and analyze them carefully.

4.1 Multi-Exponent RSA Assumption

We start by extending the standard RSA assumption from one challenge exponent e to a set of ℓ exponents, e_1, \dots, e_ℓ . The adversary wins the game if it can find the e_i -th root of a random target X where e_i can be any of the ℓ exponents provided in the challenge. We formally present this game in Figure 5 that we denote as the multi-exponent RSA game.

Definition 8 ((ϵ, t, ℓ) -Multi-Exponent RSA Assumption). *Let λ be the security parameter and consider the game $\mathsf{G}_{\text{Gen}, \mathcal{A}, \ell}^{\text{ME-RSA}}(\lambda)$ in Figure 5. The (ϵ, t, ℓ) -multi-exponent RSA assumption is true for Gen , if for any*

Game $G_{\text{Gen}, \mathcal{A}, \ell}^{\text{ME-RSA}}(\lambda)$: $(N, \mathcal{D}_N, (p, q)) \leftarrow_R \text{Gen}(1^\lambda)$ $(e_1, \dots, e_\ell) \leftarrow_R \mathcal{D}_N$ $X \leftarrow_R \mathbb{Z}_N^*$ $(i, Y) \leftarrow \mathcal{A}(N, e_1, \dots, e_\ell, X)$ Return 1 if and only if $i \in [\ell]$ and $Y^{e_i} = X \pmod N$.
--

Figure 5: Multi-Exponent RSA Game. All differences with the RSA game, $G_{\text{Gen}, \mathcal{A}}^{\text{RSA}}(\lambda)$, are highlighted in blue.

adversary \mathcal{A} that runs in time t , the following holds

$$\Pr[G_{\text{Gen}, \mathcal{A}, \ell}^{\text{ME-RSA}}(\lambda) = 1] \leq \epsilon.$$

As this provides more flexibility for the adversary, it is a stronger assumption than the standard RSA assumption with a single exponent. However, we can show that the two assumptions differ by at most a multiplicative factor in ℓ .

Theorem 2. *If the (ϵ, t) -RSA assumption for strong modulus (Definition 6) is true with \mathcal{D}_N choosing uniformly random exponents from $[3, e_{\max}]$ where $e_{\max} \leq 2^{\kappa-2}$, then the $(\epsilon \cdot \ell, t - O(\ell), \ell)$ -multi-exponent RSA assumption (Definition 8) is true with \mathcal{D}_N .*

Proof. Towards a contradiction, suppose that there exists an adversary \mathcal{A}' that wins the multi-exponent RSA game with probability ϵ' with running time t' when given ℓ challenge exponents. We build an adversary \mathcal{A} for the RSA game that wins with probability $\epsilon = \epsilon'/\ell$ with running time $t = t' + O(\ell)$.

Recall that \mathcal{A} is given an exponent e from $[3, e_{\max}]$ and random target X and must produce Y satisfying $Y^e = X \pmod N$. To do this, \mathcal{A} will choose ℓ challenge exponents e_1, \dots, e_ℓ as follows. \mathcal{A} first picks a uniformly random number from $z \in [\ell]$. Then, \mathcal{A} will set $e_z = e$. For the remaining $i \in [\ell] \setminus \{z\}$, \mathcal{A} sets e_i to a random odd number from the set $[3, e_{\max}]$. As we assume N is a strong RSA modulus and the product of two κ -bit primes, we know that $\phi(N) = 4p'q'$ for some primes p' and q' of length at least $\kappa - 1$. Therefore, odd numbers from $[3, 2^{\kappa-2}]$ will be co-prime with $\phi(N)$ and, thus, valid exponents. As each e_i are chosen such that $e_i \leq e_{\max} \leq 2^{\kappa-2}$, we know that each e_i is a valid exponent. Finally, \mathcal{A} passes the modulus N , the ℓ exponents as well as the random target X to $\mathcal{A}'(N, e_1, \dots, e_\ell, X)$.

First, we note that the distribution of inputs seen by \mathcal{A}' is identical to the multi-exponent RSA game (see Figure 5). \mathcal{A}' successfully forges and, thus, returns i as well as a valid decryption of X under e_i with probability ϵ' . In other words, \mathcal{A}' outputs Y satisfying $Y = X^{1/e_i}$. Therefore, \mathcal{A} can output Y to win the RSA game as long as $i = z$. So, \mathcal{A} wins this game with probability ϵ'/ℓ .

For the running time, note that \mathcal{A} required an additional $O(\ell)$ time to generate random exponents. If \mathcal{A}' runs in time t' , then \mathcal{A} runs in time $t' + O(\ell)$. \square

We build upon the above ideas to obtain our one-more extension of the multi-exponent RSA assumption that we will use to prove security of our blind signature protocol.

Tighter Reductions under Different Exponent Distributions. We show that one can obtain a tighter reduction if we assume a slightly different distribution of exponents in the multi-exponent RSA assumption (that also appeared in [9]). The above reduction loses a multiplicative ℓ factor. Assuming a slightly different exponent distribution for the multi-exponent RSA assumption, we can instead show a different reduction with no security loss. We point readers to Appendix C.2 for more details.

4.2 One-More Multi-Exponent RSA Assumption

We start from the chosen-target variant of the one-more RSA assumption from [15] (we refer readers back to Figure 4 for full details). We will modify this game to obtain a one-more variant of the multi-exponent

Game $G_{\text{Gen}, \mathcal{A}}^{\text{CT-ARE-RSA}}(\lambda)$:	Oracle $\mathcal{O}^{\text{RSA}}(X, e)$:	Oracle $\mathcal{O}^{\mathcal{X}}()$:	Oracle $\mathcal{O}^{\text{exp}}()$:
$(N, \mathcal{D}_N, (p, q)) \leftarrow_R \text{Gen}(1^\lambda)$ $\mathcal{S}_{\mathcal{X}} \leftarrow \emptyset, \mathcal{S}_{\text{exp}} \leftarrow \emptyset$ $\phi(N) \leftarrow (p-1)(q-1)$ $\text{cnt}_e \leftarrow 0, \forall e \in \mathbb{Z}_{\phi(N)}^*$ $e, (X_i, Y_i)_{i \in [z]} \leftarrow \mathcal{A}^{\mathcal{O}^{\text{RSA}}, \mathcal{O}^{\mathcal{X}}, \mathcal{O}^{\text{exp}}}(N)$ Return 1 if and only if all the following hold: <ul style="list-style-type: none"> - $\text{cnt}_e < z$ - $e \in \mathcal{S}_{\text{exp}}$ - $\forall i \neq j \in [z], X_i \neq X_j$ - $\forall i \in [z], X_i \in \mathcal{S}_{\mathcal{X}}$ - $\forall i \in [z], Y_i^e = X_i \text{ mod } N$ 	If $e \notin \mathcal{S}_{\text{exp}}$: Return \perp $\text{cnt}_e \leftarrow \text{cnt}_e + 1$ $d \leftarrow e^{-1} \text{ mod } \phi(N)$ $Y \leftarrow X^d \text{ mod } N$ Return Y	$X \leftarrow_R \mathbb{Z}_N^*$ $\mathcal{S}_{\mathcal{X}} \leftarrow \mathcal{S}_{\mathcal{X}} \cup \{X\}$ Return X	$e \leftarrow_R \mathcal{D}_N$ $\mathcal{S}_{\text{exp}} \leftarrow \mathcal{S}_{\text{exp}} \cup \{e\}$ For $e' \in \mathcal{S}_{\text{exp}}$: $x \leftarrow e'$ For $e'' \in \mathcal{S}_{\text{exp}} \setminus \{e'\}$: $y \leftarrow \text{GCD}(x, e'')$ While $y > 1$: $x \leftarrow x/y$ $y \leftarrow \text{GCD}(x, e'')$ If $x = 1$: $\mathcal{S}_{\text{exp}} \leftarrow \mathcal{S}_{\text{exp}} \setminus \{e\}$ Go back to first step of \mathcal{O}^{exp} to re-sample e . Return e

Figure 6: Chosen-Target, **Restricted-Exponent** RSA Inversion Game with all differences with the chosen-target RSA inversion game are highlighted in blue and ignoring highlighted parts in red. If you add the portions highlighted in red as well, we obtain the Chosen-Target, **Algebraically-Restricted-Exponent** RSA Inversion Game.

RSA assumption from the prior section. As the first step, we will also create an exponent oracle \mathcal{O}^{exp} that will output challenge exponents (analogous to the oracle $\mathcal{O}^{\mathcal{X}}$ for challenge targets). Next, we generalize the decryption oracle, \mathcal{O}^{RSA} , to receive both an element X and an exponent e . However, \mathcal{O}^{RSA} makes the restriction that the input exponent e must be a challenge exponent output by \mathcal{O}^{exp} . In other words, the decryption oracle only works for the exponents output by \mathcal{O}^{exp} .

This immediately leads to the first natural definition of a one-more multi-exponent RSA assumption where an adversary makes at most cnt_e decryption oracle queries for some exponent e and must output $\text{cnt}_e + 1$ valid decryptions of the form $(X_i, Y_i)_{i \in [\text{cnt}_e + 1]}$ satisfying that $Y_i^e = X_i \text{ mod } N$ where e must be some challenge exponent output by \mathcal{O}^{exp} . We keep the same restriction as the one-more RSA assumption that the targets X_i must be valid outputs from the message oracle $\mathcal{O}^{\mathcal{X}}$ (satisfying the chosen-target property).

We formally define this game in Figure 6 that we denote as the chosen-target, restricted-exponent RSA inversion game. In this definition, we use *restricted-exponent* to denote the fact that the decryption oracle \mathcal{O}^{RSA} restricts the adversary to only pick exponents output by \mathcal{O}^{exp} .

Definition 9 ($(\epsilon, t, \ell, \ell')$ -Chosen-Target, Restricted-Exponent RSA Inversion Assumption). *Let λ be the security parameter and consider the game $G_{\text{Gen}, \mathcal{A}}^{\text{CT-RE-RSA}}(\lambda)$ in Figure 6. The $(\epsilon, t, \ell, \ell')$ -chosen-target, restricted-exponent RSA inversion assumption is true for Gen , if for any adversary \mathcal{A} that runs in time t , makes at most ℓ decryption queries and ℓ' exponent queries, the following holds*

$$\Pr[G_{\text{Gen}, \mathcal{A}}^{\text{CT-RE-RSA}}(\lambda) = 1] \leq \epsilon.$$

Algebraically Restricted Exponents. We present a variant where the exponent oracle will perform additional checks to ensure exponents do not satisfy a certain algebraic property. In particular, the oracle checks that there is no sequence of exponents e, e_1, \dots, e_z such that $e = p_1 \cdot p_2 \cdots p_z$ and $p_i \mid e_i$ for all $i \in [z]$. First, we note that the primes p_1, \dots, p_z do not necessarily need to be distinct and e can be the product of prime powers. Furthermore, e is distinct from all of e_1, \dots, e_z . However, the z exponents e_1, \dots, e_z do not need to be distinct. The exponent oracle continues to pick random exponents until this algebraic property does not hold. This ends up being important to avoid various attacks on multi-exponent one-more RSA assumptions (see Appendix F for details). We present this game in Figure 6 and the definition below.

Definition 10 ($(\epsilon, t, \ell, \ell')$ -Chosen-Target, Algebraically-Restricted-Exponent RSA Inversion Assumption). *Let λ be the security parameter and consider the game $G_{\text{Gen}, \mathcal{A}}^{\text{CT-ARE-RSA}}(\lambda)$ in Figure 6. The $(\epsilon, t, \ell, \ell')$ -chosen-target, algebraically-restricted-exponent RSA inversion assumption is true for Gen , if for any adversary \mathcal{A} that runs in time t , makes ℓ decryption queries, and makes ℓ' exponent queries, $\Pr[G_{\text{Gen}, \mathcal{A}}^{\text{CT-ARE-RSA}}(\lambda) = 1] \leq \epsilon$.*

We relate this to the chosen-target, restricted-exponent game without the algebraic restrictions (Def. 9) in Appendix F.2.

Game $G_{\text{Gen}, \mathcal{A}}^{\text{SRSA}}(\lambda)$:
$(N, (p, q)) \leftarrow_R \text{Gen}(1^\lambda)$
$X \leftarrow_R \mathbb{Z}_N^*$
$(e, Y) \leftarrow \mathcal{A}(\text{crs}, N, X)$
Return 1 if and only if $e \geq 3$ and $Y^e = X \pmod N$.

Figure 7: Strong RSA Game.

Relation to Chosen-Target RSA Inversion Assumption. One could attempt to try and prove an equivalence between the chosen-target, (algebraically-)restricted-exponent RSA inversion assumption (Definition 9) and the original chosen-target RSA inversion assumption (Definition 7) from [15]. Clearly, the chosen-target, (algebraically-)restricted-exponent variant is a no weaker assumption than the original chosen-target variant. Recall that, in the proof of Theorem 2, it was required to generate fake challenge exponents co-prime to $\phi(N)$. This can be easily done when N is a strong RSA modulus. For this reduction, we would need to additionally be able to produce decryptions with respect to the fake challenge exponents. Unfortunately, this seems quite challenging. For example, suppose an adversary tried to produce pairs (e', d') such that $e'd' = 1 \pmod{\phi(N)}$. Then, the adversary could submit e' as a challenge exponent and use d' to answer decryption queries. It is well known that producing a pair (e', d') satisfying this property is equivalent to factoring N and, thus, breaking any RSA assumption immediately. Therefore, this seems like a very challenging task for an adversary without knowledge of $\phi(N)$.

We leave it as an open problem to determine the relationship between the chosen-target, (algebraically-)restricted-exponent inversion assumption and the original chosen-target inversion assumption.

Winning Condition. In some prior works, the adversary's winning condition is outputting a pair (X, Y) for e that is different from any query to \mathcal{O}^{RSA} . In Appendix B, we show this is equivalent to our condition of outputting the tuple $(X_i, Y_i)_{i \in [z]}$ where $z > \text{cnt}_e$.

4.3 Connections to Strong RSA Assumption

Finally, we make connections between the above multi-exponent RSA assumptions and the strong RSA assumption introduced in [14, 34]. We emphasize that the assumptions in this section that will not be used for proving security of any of our protocols. We will only rely on the multi-exponent variants from Section 4.2 for our security proofs. Instead, we perform an exploration to see the connections between the assumptions from our prior section and its relation to the strong RSA assumption.

Strong RSA Assumption. In comparison with the standard RSA assumption, the strong RSA assumption provides more flexibility to the adversary to choose its own public exponent. The strong RSA problem states that the original RSA problem is intractable even when \mathcal{A} is allowed to choose the public exponent $e \geq 3$. More specifically, given a RSA modulus N , and a random target X , it is infeasible to find any pair (Y, e) such that $X = Y^e \pmod N$ and $e \geq 3$. We provide the formal definition of the strong RSA assumption below:

Definition 11 ((ϵ, t) -Strong RSA Assumption). *Let λ be the security parameter and consider the game $G_{\text{Gen}, \mathcal{A}}^{\text{SRSA}}(\lambda)$ in Figure 7. The (ϵ, t) -strong RSA assumption is true for Gen , if for any PPT adversary \mathcal{A} that runs in time t , the following holds*

$$\Pr[G_{\text{Gen}, \mathcal{A}}^{\text{SRSA}}(\lambda) = 1] \leq \epsilon.$$

Extending towards One-More Strong RSA Type Assumptions. A natural next step is to extend the strong RSA assumption using the one-more style definitions. To our knowledge, we are unaware of prior work that has formally defined this notion. We will present two natural definitions that seem to be reasonable. We show that one such definition seems plausible while the other can be broken by an adversary.

Before presenting these definitions, we quickly overview why the prior chosen-target, restricted-exponent RSA assumption (Definition 9) from Section 4.2 seems significantly weaker. Recall that, in the strong RSA assumption, the adversary is able to choose any exponent $e \geq 3$. In contrast, for the chosen-target,

Game $G_{\text{Gen}, \mathcal{A}}^{\text{CT-CE-RSA}}(\lambda)$:	Oracle $\mathcal{O}^{\text{RSA}}(X, e)$:	Oracle $\mathcal{O}^{\mathcal{X}}()$:	Oracle $\mathcal{O}^{\text{exp}}()$:
$(N, \mathcal{D}_N, (p, q)) \leftarrow_R \text{Gen}(1^\lambda)$ $\mathcal{S}_{\mathcal{X}} \leftarrow \emptyset, \mathcal{S}_{\text{exp}} \leftarrow \emptyset$ $\phi(N) \leftarrow (p-1)(q-1)$ $\text{cnt}_e \leftarrow 0, \forall e \in \mathbb{Z}_{\phi(N)}^*$ $e, (X_i, Y_i)_{i \in [x]} \leftarrow \mathcal{A}^{\mathcal{O}^{\text{RSA}}, \mathcal{O}^{\mathcal{X}}, \mathcal{O}^{\text{exp}}}(N)$ Return 1 if and only if all the following hold: <ul style="list-style-type: none"> - $\text{cnt}_e < x$ - $e \in \mathcal{S}_{\text{exp}}$ - $\forall i \neq j \in [x], X_i \neq X_j$ - $\forall i \in [x], X_i \in \mathcal{S}_{\mathcal{X}}$ - $\forall i \in [x], Y_i^e = X_i \pmod N$ 	If $e \notin \mathbb{Z}_{\phi(N)}^*$: Return \perp $\text{cnt}_e \leftarrow \text{cnt}_e + 1$ $d \leftarrow e^{-1} \pmod{\phi(N)}$ $Y \leftarrow X^d \pmod N$ Return Y	$X \leftarrow_R \mathbb{Z}_N^*$ $\mathcal{S}_{\mathcal{X}} \leftarrow \mathcal{S}_{\mathcal{X}} \cup \{X\}$ Return X	$e \leftarrow_R \mathcal{D}_N$ $\mathcal{S}_{\text{exp}} \leftarrow \mathcal{S}_{\text{exp}} \cup \{e\}$ Return e

Figure 8: Chosen-Target, **Chosen-Exponent** RSA Inversion Game. All differences with the chosen-target, restricted-exponent RSA inversion game, $G_{\text{Gen}, \mathcal{A}}^{\text{CT-RE-RSA}}(\lambda)$, are highlighted in blue.

restricted-exponent game, $G_{\text{Gen}, \mathcal{A}, \ell}^{\text{CT-RE-RSA}}(\lambda)$, the adversary is significantly restricted in the exponents of choice. The chosen-target, restricted-exponent definition places two requirements on the adversaries that seem more stringent than the strong RSA assumption. First, the output exponent must be one of the outputs of the challenge oracle \mathcal{O}^{exp} . Secondly, the adversary is not able to submit any decryption requests for exponents except for those output by \mathcal{O}^{exp} . Therefore, in our opinion, the chosen-target, restricted-exponent assumption is weaker and not a truly natural extension of the strong RSA assumption. Intuitively, the chosen-target, restricted-exponent assumption seems significantly weaker than any true one-more strong RSA assumption.

Next, we will consider assumptions that are (seemingly) stronger than the chosen-target, restricted-exponent game that get closer to the spirit of an one-more strong RSA assumption.

One-More RSA with Chosen Exponents. As our first attempt to consider stronger assumptions, we will consider a new game where the decryption oracle, \mathcal{O}^{RSA} , allows the adversary to submit arbitrary exponents for decryption. In particular, we no longer restrict queries to the decryption oracle to consist of exponents output by the challenge oracle \mathcal{O}^{exp} . As a technical detail, we do require queried exponents to be co-prime to $\phi(N)$ to enable decryption. In the case that the input exponent is not co-prime to $\phi(N)$, \mathcal{O}^{RSA} will output \perp . Note, we will still require the adversary to output exponents from \mathcal{O}^{exp} to win the game though. Clearly, this is a stronger assumption compared to the chosen-target, restricted-exponent assumption (Definition 9) as the adversary has significantly more freedom when choosing decryption oracle queries.

We formally present this game in Figure 8 that we denote as the chosen-target, chosen-exponent RSA inversion game. We use the notion *chosen-exponent* to denote that the adversary is now free to pick any exponents to submit to the decryption oracle \mathcal{O}^{RSA} .

Definition 12 ($(\epsilon, t, \ell, \ell')$ -Chosen-Target, Chosen-Exponent RSA Inversion Assumption). *Let λ be the security parameter and consider the game $G_{\text{Gen}, \mathcal{A}}^{\text{CT-CE-RSA}}(\lambda)$ in Figure 8. The $(\epsilon, t, \ell, \ell')$ -chosen-target, chosen-exponent RSA inversion assumption is true for Gen if, for any adversary \mathcal{A} that runs in time t and makes at most ℓ decryption queries and ℓ' exponent queries, the following holds*

$$\Pr[G_{\text{Gen}, \mathcal{A}}^{\text{CT-CE-RSA}}(\lambda) = 1] \leq \epsilon.$$

This definition still limits the adversary's abilities to pick exponents e . In particular, the adversary is forced to use one of the outputs of the challenge oracle \mathcal{O}^{exp} in order to win the game. While this seems like a reasonable assumption, we can show that this already provides far too much power to an adversary.

Theorem 3. *There exists a polynomial time adversary \mathcal{A} that wins the chosen-target, chosen-exponent RSA inversion game.*

Proof. We present a simple adversary \mathcal{A} as follows. First, \mathcal{A} gets a target $X \leftarrow \mathcal{O}^{\mathcal{X}}()$. Next, the \mathcal{A} gets $e \leftarrow \mathcal{O}^{\text{exp}}()$ and picks a random number e' that are both co-prime to $\phi(N)$. Without loss of generality, let $e' > e$.

The adversary calls the decryption oracle to receive $Y \leftarrow \mathcal{O}^{\text{RSA}}(X^{e'}, e \cdot e')$. Then, $Y = (X^{e'})^{1/(e \cdot e')} = X^{1/e}$. Finally, \mathcal{A} outputs $e, (X, Y)$ and wins the game as the adversary never made a single decryption oracle call for e . \square

Interestingly, the ability for the adversary to choose arbitrary exponents to send to the decryption oracle enables far too much adversarial power. The same reasoning was used in prior RSA assumptions where if the adversary is able to pick arbitrary messages to decrypt, then the adversary can easily break prior RSA assumptions. While this does not happen for exponents in the standard assumptions without decryption oracles, it ends up being the case once you provide decryption oracles to the adversary. Therefore, the challenge exponent oracle \mathcal{O}^{exp} seems necessary for any of these one-more RSA assumptions.

One-More RSA with Arbitrary Exponents. In our second attempt, we no longer provide a challenge exponent oracle, \mathcal{O}^{exp} , to the adversary. Instead, the adversary is free to choose any exponent $e \geq 3$ and produce a decryption with respect to e . We denote this game as the chosen-target, arbitrary-exponent RSA inversion game. Clearly the adversary can win this game as well, as it won the previous game where it had less adversarial power.

Inability to Extend Attack. One could try to extend these attacks to our other assumptions such as the chosen-target, (algebraically-)restricted-exponent assumption (Definition 9). In the attack, the adversary is required to find two exponents e and e' such that $e \mid e'$. In the chosen-exponent game, it is unlikely the adversary will ever receive such a pair of exponents from the oracle \mathcal{O}^{exp} . For example, if N is a strong RSA modulus such that $N = pq$ where $p = 2p' + 1$ and $q = 2q' + 1$ are safe primes, then $\phi(N) = 4p'q'$. Assuming that p' and q' are sufficiently large, then it is very unlikely for an adversary to receive random elements e, e' from $\mathbb{Z}_{\phi(N)}^*$ (via the exponent oracle, \mathcal{O}^{exp}) satisfying $e \mid e'$. Therefore, we were unable to extend this attack to any of the prior games due to the requirement that the adversary must use exponents output by oracle \mathcal{O}^{exp} .

5 RSA Signatures with Public Metadata

In this section, we start by presenting a protocol for RSA (non-blind) signatures with public metadata. The goal is to show that our protocol augmentation for public metadata does not degrade concrete security in any meaningful way compared to the original protocol without public metadata.

5.1 Preliminaries

Before we present our protocol, we start by presenting various preliminaries and building blocks that we will utilize throughout our protocols. We will use these primitives in our RSA blind signature protocol in Section 6 as well.

Unique Concatenation. Throughout our work, we will utilize concatenation of multiple values where each unique tuple should result in a unique concatenation. Formally, we may consider tuples (A, B) such that the concatenation $A \parallel B$ must be unique. That is, for any pair of tuples, (A, B) and (A', B') , it must be that $A \parallel B$ is identical to $A' \parallel B'$ if and only if $A = A'$ and $B = B'$. This could also extend to tuples of size larger than two. A straightforward way to do this is to encode tuples (A, B) as $|A| \parallel A \parallel B$ where $|A|$ is the length of A stored in some fixed-length integer (such as 64 bits). This may be extended to longer tuples naturally as (X_1, X_2, \dots, X_n) may be concatenated as $|X_1| \parallel |X_2| \parallel \dots \parallel |X_{n-1}| \parallel X_1 \parallel X_2 \parallel \dots \parallel X_n$. For the rest of the work, we will assume that when we concatenate any tuple of values $X_1 \parallel X_2 \parallel \dots \parallel X_n$, we are using the unique version with length prepending.

Message Encoding. In our protocols, we will utilize standard encoding algorithms for messages used in RSA signatures. Formally, we will assume that the message encoding consists of two algorithms: $H_{\mathcal{M}}$ and $\text{Verify}_{\mathcal{M}}$. The hash function $H_{\mathcal{M}}(X) : \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$ maps strings to elements in the multiplicative group modulo the RSA modulus N . The verification algorithm $\text{Verify}_{\mathcal{M}}(X, Y) : \{0, 1\}^* \times \mathbb{Z}_N^* \rightarrow \{0, 1\}$ checks and returns 1 if and only if Y corresponds to the output of $H_{\mathcal{M}}(X)$.

In our work, we will consider two different encodings: full domain hash (FDH) and probabilistic signature scheme (PSS). At a high level, FDH is a deterministic encoding algorithm that maps each message to a random element. In contrast, PSS is a randomized encoding algorithm that will map the same message to different outputs when evoked multiple times. Bellare and Rogaway [16] showed that the concrete security of PSS encodings was stronger than FDH encodings when applied to standard RSA signatures (without public metadata). In our implementations, we use PSS encodings due to its superior concrete security and randomized properties. Furthermore, we wish to align our implementation with current IRTF specifications for RSA blind signatures [33] that do not consider public metadata. The version of PSS encodings that we rely upon are those from prior specifications [44] that essentially use a variant of PSS encodings with message recovery presented in [16].

For our proofs, we will prove security of our schemes assuming the usage of the FDH encoding. We chose this approach as to not over-complicate our proof with PSS encoding techniques that are directly borrowed from prior work [16]. Instead, we wish to focus our proof on the new techniques required to prove security with respect to the additional public metadata functionality. Although, one can directly apply the proof techniques for PSS encodings from [16] to our security proofs. As we will use FDH, we present it formally.

Definition 13 (Full Domain Hash). *The full domain hash (FDH) message encoding consists of two deterministic algorithms $H_{\mathcal{M}} : \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$ and $\text{Verify}_{\mathcal{M}} : \{0, 1\}^* \times \mathbb{Z}_N^* \rightarrow \{0, 1\}$ as follows:*

- For each string X , $H_{\mathcal{M}}(X)$ is a random element from \mathbb{Z}_N^* .
- For a string X and $Y \in \mathbb{Z}_N^*$, $\text{Verify}_{\mathcal{M}}(X, Y) = 1$ if and only if $Y = H_{\mathcal{M}}(X)$.

Instantiations of Hash Functions. Throughout the rest of this section and our proofs, we will assume that all hash functions are modeled as random oracles unless otherwise specified. In practice, we will instantiate these hash functions using a variety of different approaches depending on prior standards or the specific requirements. We defer to prior specifications including [33, 12] for recommendations on specific instantiations of each hash function.

5.2 Protocol for RSA Signatures with Public Metadata

In this section, we formally present the protocol for RSA signatures with public metadata. Our goal of starting with the non-blind version is to show that our public metadata augmentation does not degrade the concrete security compared to the standard RSA signature scheme without public metadata.

Setup. For the reader's convenience, we quickly summarize our setup. We will denote the strong RSA modulus by N of bit length 2κ such that the two safe primes p, q (factors of N) are both of bit length κ . The public key will be $\text{pk} \leftarrow N$ which is the strong RSA modulus. The private key will be $\text{sk} \leftarrow \phi(N)$. As $p = 2p' + 1$ and $q = 2q' + 1$ are safe primes, we know that both p' and q' are also primes. Therefore, $\phi(N) = (p - 1) \cdot (q - 1) = 4p'q'$.

Public Metadata Hash Function. Next, we define our hash function that is used to map public metadata to values in the group of RSA exponents $\mathbb{Z}_{\phi(N)}^*$. Recall that these are all elements that are invertible modulo $\phi(N)$. This is a straightforward task if $\phi(N)$ was publicly known by all parties. However, it is critical that $\phi(N)$ is hidden to all parties except the signer. Therefore, we must come up with a way to perform this mapping for users that do not know the value of $\phi(N)$. We present a simple way to do this below that relies on the structure of $\mathbb{Z}_{\phi(N)}^*$ and the fact that N is a strong RSA modulus.

$H_{\text{MD}}(D)$: Let G be a hash function.

1. Compute $G(D)$ of length γ bits.
2. Return $1 \parallel G(D)$. In other words, we are returning $2 \cdot G(D) + 1$.

In words, the above function simply returns a random odd number of length $\gamma + 1$ bits. Picking γ correctly, we can guarantee that the output of H_{MD} will always be smaller than the prime values p' and q' such that $\phi(N) = 4p'q'$. This turns out to be sufficient to guarantee that outputs of H_{MD} will always be elements from $\mathbb{Z}_{\phi(N)}^*$.

Lemma 1. *Suppose N is a strong RSA modulus such that $N = pq$ where $p = 2p' + 1$ and $q = 2q' + 1$ are safe primes each of length κ and $\gamma \leq \kappa - 3$. Then, for all $D \in \{0, 1\}^*$, $H_{\text{MD}}(D)$ is co-prime to $\phi(N)$.*

Proof. Fix any $x \in \{0, 1\}^*$. First, we note that the output of $H_{\text{MD}}(x)$ is always odd. Therefore, we know that if $H_{\text{MD}}(x)$ is not co-prime to $\phi(N)$, then it must be that either $p' \mid H_{\text{MD}}(x)$ or $q' \mid H_{\text{MD}}(x)$. We know that the bit length of p' and q' is at least $\kappa - 1$ given that $p = 2p' + 1$ and $q = 2q' + 1$ and both p and q are κ bit long primes. Also note that the bit length of $H_{\text{MD}}(x)$ is $\gamma + 1 \leq \kappa - 2$. Thus, we know that $H_{\text{MD}}(x) < p'$ and $H_{\text{MD}}(x) < q'$. Therefore, $H_{\text{MD}}(x)$ is always co-prime to $\phi(N)$. \square

RSA Signatures with Public Metadata. Using the public metadata hash function H_{MD} , we are ready to present a description of our RSA signatures with public metadata. We will assume that $H_{\mathcal{M}}$ and $\text{Verify}_{\mathcal{M}}$ correspond to some message encoding algorithm (such as FDH in Definition 13). We note that the Setup algorithm is identical to the standard RSA signatures without public metadata protocols except that in our Setup algorithm we only accept N if it is a strong RSA modulus. The signing algorithm is modified to use $e_{\text{MD}} \leftarrow H_{\text{MD}}(D)$ where D is the public metadata. As the signer knows $\phi(N)$, it can compute $d_{\text{MD}} \leftarrow (e_{\text{MD}})^{-1} \bmod \phi(N)$ efficiently to produce a signature. The verification algorithm is identical when using e_{MD} as the public exponent. As we are considering non-blind signatures (i.e., non-anonymous tokens), the Blind and Finalize algorithms are trivial and, thus, omitted (see Section 2 for definitions). We present the formal algorithms for all of the necessary algorithms below.

$\text{RSA}_{\text{MD}}.\text{Setup}(1^\lambda)$:

1. Randomly generate $(N, p, q) \leftarrow_R \text{Gen}(1^\lambda)$.
2. Compute $\phi(N) \leftarrow (p - 1)(q - 1)$.
3. Return $\text{pk} \leftarrow N, \text{sk} \leftarrow \phi(N)$.

$\text{RSA}_{\text{MD}}.\text{Sign}(M, D, \text{pk} \leftarrow N, \text{sk} \leftarrow \phi(N))$:

1. Compute $e_{\text{MD}} \leftarrow H_{\text{MD}}(D)$.
2. Compute $d_{\text{MD}} \leftarrow (e_{\text{MD}})^{-1} \bmod \phi(N)$.
3. Compute $M_{\text{MD}} \leftarrow H_{\mathcal{M}}(M \parallel D)$.
4. Return signature $S \leftarrow (M_{\text{MD}})^{d_{\text{MD}}} \bmod N$.

$\text{RSA}_{\text{MD}}.\text{Verify}(S, M, D, \text{pk} \leftarrow N)$:

1. Compute $e_{\text{MD}} \leftarrow H_{\text{MD}}(D)$.
2. Compute $X \leftarrow S^{e_{\text{MD}}} \bmod N$.
3. Return $\text{Verify}_{\mathcal{M}}(M \parallel D, X)$.

Correctness. First, we show our above protocols are well-defined. In particular, we need to show that $d_{\text{MD}} \leftarrow (e_{\text{MD}})^{-1} \bmod \phi(N)$ is actually computable. We will rely on Lemma 1 stating that all outputs of H_{MD} are co-prime with $\phi(N)$. Therefore, e_{MD} has a valid inverse modulo $\phi(N)$ and the signing algorithm is always well-defined.

Next, we will show that the verification succeeds for well-formed signatures. Let S be a signature that is produced by following the **Setup** and **Sign** algorithms properly for input message M and public metadata D . Then, we know that the signature satisfies:

$$S = (H_{\mathcal{M}}(M \parallel D))^{1/H_{\text{MD}}(D)} \bmod N.$$

Then, the verification algorithm will return the following result

$$\text{Verify}_{\mathcal{M}}(M \parallel D, S^{H_{\text{MD}}(D)}) = \text{Verify}_{\mathcal{M}}(M \parallel D, H_{\mathcal{M}}(M \parallel D)) = 1$$

as we know that $S^{H_{\text{MD}}(D)} = H_{\mathcal{M}}(M \parallel D) \bmod N$ and $\text{Verify}_{\mathcal{M}}(X, H_{\mathcal{M}}(X)) = 1$. Therefore, a well-formed signature will always be verified correctly.

5.3 Unforgeability

In this section, we prove the concrete security of our RSA signature with public metadata from the RSA assumption. In particular, we aim to answer the question of whether our new public metadata variant may have weaker security compared to standard RSA signatures. We essentially show that this is not the case as we show that the concrete security bounds are similar to those proven by Bellare and Rogaway [16] with FDH message encodings except for a multiplicative factor loss in number of hashing oracle calls to H_{MD} .

At a high level, our security proof will construct an adversary \mathcal{A} for the multi-exponent RSA game given exponents from $[3, 2^{\kappa-2}]$. Afterwards, we can apply Theorem 2 to obtain security based on the standard RSA assumption. We start with the first step reducing security to the multi-exponent RSA assumption:

Lemma 2. *Suppose that $(H_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ correspond to full domain hash (FDH) message encoding. Assuming the $(\epsilon, t, q_{\text{MD}})$ -multi-exponent RSA assumption (Definition 8) with exponents in $[3, 2^{\kappa-2}]$ and the random oracle model, then RSA_{MD} is $(\epsilon_F, t_F, \ell_F)$ -strong one-more unforgeable (Definition 2) where*

- $t_F = t - O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$
- $\epsilon_F = q_{\mathcal{M}} \cdot \epsilon$

and the adversary \mathcal{A} runs in time at most t_F and makes at most ℓ_F , $q_{\mathcal{M}}$ and q_{MD} queries to the signing oracle and hash functions $H_{\mathcal{M}}$ and H_{MD} respectively.

Proof. Let \mathcal{A}_F be an adversary that can break the $(\epsilon_F, t_F, \ell_F)$ -SOMUF of RSA_{MD} with at most ℓ_F , $q_{\mathcal{M}}$ and q_{MD} queries to the signing oracle, $H_{\mathcal{M}}$ and H_{MD} . That is, \mathcal{A}_F can forge a signature with success probability of ϵ_F with running time of t_F . To complete the proof, we show that one can use \mathcal{A}_F to construct an adversary \mathcal{A} that can break the multi-exponent RSA game with exponents from $[3, 2^{\kappa-2}]$ using $O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$ additional running time and wins the game with probability at least $\epsilon_F/q_{\mathcal{M}}$.

For any adversary \mathcal{A} trying to solve the multi-exponent RSA problem, \mathcal{A} receives as input the challenge RSA exponents $(e_1, \dots, e_{q_{\text{MD}}})$, strong RSA modulus N and a challenge target X that is chosen uniformly at random \mathbb{Z}_N^* . Recall that the goal is to pick any $i \in [q_{\text{MD}}]$ and output the e_i -th root of X modulo N . In other words, compute Y such that $Y^{e_i} = X \bmod N$.

\mathcal{A} will execute the adversary \mathcal{A}_F that outputs a forgery of the form $D, (S_i, M_i)_{i \in [x]}$. Without loss of generality, we will assume the following. First, $x \leq \ell_F + 1$. As \mathcal{A}_F makes at most ℓ_F signing oracle queries, \mathcal{A} will only need to produce at most $\ell_F + 1$ valid signatures to win the strong one-more unforgeability game. Secondly, we will assume that \mathcal{A}_F will have queried $H_{\mathcal{M}}(M_i)$ on all output messages. Again, this is without loss of generality as it will only increase the running time of \mathcal{A}_F by at most $\ell_F + 1$ hash function queries. Furthermore, we will suppose that \mathcal{A}_F will have queried both $H_{\mathcal{M}}(M)$ and $H_{\text{MD}}(D)$ before submitting a signing query for M and D . Again, this is without loss of generality, as it only increases the running time of \mathcal{A}_F by at most $2\ell_F$ hash queries.

During the execution of \mathcal{A}_F , \mathcal{A} must successfully simulate queries for signing as well as the message encoding hash function $H_{\mathcal{M}}$ and the public metadata hash function H_{MD} . Before doing so, \mathcal{A} keeps track of

a count of the number of unique inputs to the message oracle $H_{\mathcal{M}}$. We know that at most $q_{\mathcal{M}}$ queries are made to $H_{\mathcal{M}}$. \mathcal{A} will pick a random integer $z \leftarrow_R [q_{\mathcal{M}}]$ and will embed the challenge X as the output of the z -th query to $H_{\mathcal{M}}$. To keep track, \mathcal{A} keeps counter $\text{cnt}_{\mathcal{M}}$ to count the number of unique inputs to $H_{\mathcal{M}}$ where $\text{cnt}_{\mathcal{M}}$ is initialized to zero. Similarly, \mathcal{A} uses cnt_{MD} to count the number of unique inputs to H_{MD} with cnt_{MD} also initialized to zero. Finally, \mathcal{A} keeps a map of inputs to important information for the hash function H_{MD} and $H_{\mathcal{M}}$ denoted by \mathbf{M}_{MD} and $\mathbf{M}_{\mathcal{M}}$. \mathcal{A} will simulate each of these queries as follows.

For any query to $H_{\text{MD}}(D)$:

1. If $\mathbf{M}_{\text{MD}}[D]$ is set, return $\mathbf{M}_{\text{MD}}[D]$.
2. Increment $\text{cnt}_{\text{MD}} \leftarrow \text{cnt}_{\text{MD}} + 1$.
3. Set $\mathbf{M}_{\text{MD}}[D] \leftarrow e_{\text{cnt}_{\text{MD}}}$.
4. Return $\mathbf{M}_{\text{MD}}[D]$.

For any query to $H_{\mathcal{M}}(M, D)$:

1. Compute $e_{\text{MD}} \leftarrow H_{\text{MD}}(D)$.
2. If $\mathbf{M}_{\mathcal{M}}[M, D]$ is set, return first value of tuple $\mathbf{M}_{\mathcal{M}}[M, D]$.
3. Increment $\text{cnt}_{\mathcal{M}} \leftarrow \text{cnt}_{\mathcal{M}} + 1$.
4. If $\text{cnt}_{\mathcal{M}} = z$:
 - (a) Set $\mathbf{M}_{\mathcal{M}}[M, D] \leftarrow (X, \perp)$.
5. Otherwise when $\text{cnt}_{\mathcal{M}} \neq z$:
 - (a) Generate random $A \leftarrow_R \mathbb{Z}_N^*$.
 - (b) Compute $B = A^{e_{\text{MD}}} \bmod N$.
 - (c) Set $\mathbf{M}_{\mathcal{M}}[M, D] \leftarrow (B, A)$.
6. Return first value of tuple $\mathbf{M}_{\mathcal{M}}[M, D]$.

For any query to $\mathcal{O}^{\text{Sign}}(M, D)$:

1. Retrieve $(B, A) \leftarrow \mathbf{M}_{\mathcal{M}}[M, D]$.
2. If $A = \perp$, abort.
3. Return A .

Finally, we will suppose that \mathcal{A} does not abort and \mathcal{A}_F outputs some forgery $D, (S_i, M_i)_{i \in [x]}$ while making strictly less than x signing queries for metadata D . If $H_{\mathcal{M}}(M_i || D) = X$ for any $i \in [x]$, then we know the following holds:

$$H_{\mathcal{M}}(M_i || D) = X = S_i^{H_{\text{MD}}(D)}$$

assuming that \mathcal{A}_F is successful at forging. Re-writing the above, we get that

$$S_i = X^{1/H_{\text{MD}}(D)} = X^{1/e_j}$$

for some $j \in [q_{\text{MD}}]$ and e_j is the j -th challenge exponent. In other words, this is a successfully decryption of the ciphertext X that would win the multi-exponent RSA game. Therefore, \mathcal{A} will simply return (j, S_i) that would win the multi-exponent RSA game.

Adversarial Advantage. First, we show that the simulated view and real view of \mathcal{A}_F are identical. In the real game, \mathcal{A}_F receives random exponents from $[3, 2^{\kappa-2}]$. In the simulated game, \mathcal{A}_F receives random challenge exponents $e_1, \dots, e_{q_{\text{MD}}}$ that are also random exponents from $[3, 2^{\kappa-2}]$. The view from $H_{\mathcal{M}}$ and the signing oracle are also identical as long as \mathcal{A} does not abort.

To upper bound the aborting probability, we first note that \mathcal{A}_F must return at least one message M_i such that (M_i, D) was never sent to the signing oracle. Assuming that \mathcal{A} successfully executed \mathcal{A}_F without aborting, the probability that public metadata D and message M_i are output by \mathcal{A}_F where $H_{\mathcal{M}}(M_i || D) = X$ such that (M_i, D) was not a signing oracle query by \mathcal{A}_F is at least $1/q_{\mathcal{M}}$. This is because \mathcal{A}_F makes at most $q_{\mathcal{M}}$ queries to $H_{\mathcal{M}}$ and must output at least one message. In this case, it is clear that \mathcal{A} will not abort when executing \mathcal{A}_F . Finally, we note that \mathcal{A} produces a forgery with probability at most ϵ_F meaning that \mathcal{A} wins the RSA game with probability at least $\epsilon_F/q_{\mathcal{M}}$.

Adversarial Running Time. Suppose that \mathcal{A}_F runs in time t_F . By our assumptions, we note that we increase the running time of \mathcal{A}_F by at most $O(\ell_F)$ hash queries. For each signing query, \mathcal{A} must perform a single exponentiation. For each query to $H_{\mathcal{M}}$ and H_{MD} , \mathcal{A} requires $O(1)$ time to simulate each answer correctly. Therefore, \mathcal{A} requires $t_F + O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$ time. \square

Using the above lemma, we can apply Theorem 2 to reduce security to standard RSA.

Theorem 4. *Suppose that $(H_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ correspond to full domain hash (FDH) message encoding. Assuming the (ϵ, t) -RSA assumption (Definition 6) with exponents in $[3, 2^{\kappa-2}]$ and the random oracle model, then RSA_{MD} is $(\epsilon_F, t_F, \ell_F)$ -strong one-more unforgeable (Definition 2) where*

- $t_F = t - O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$
- $\epsilon_F = q_{\text{MD}} \cdot q_{\mathcal{M}} \cdot \epsilon$

and the adversary \mathcal{A} runs in time at most t_F and makes at most ℓ_F , $q_{\mathcal{M}}$ and q_{MD} queries to the signing oracle and hash functions $H_{\mathcal{M}}$ and H_{MD} respectively.

Proof. First, we apply Lemma 2 to reduce security to the $(\epsilon_F/q_{\mathcal{M}}, t_F + O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F), q_{\text{MD}})$ -multi-exponent RSA assumption with exponents from $[3, 2^{\kappa-2}]$. By applying Theorem 2, we obtain that this is equivalent to the $(\epsilon_F/(q_{\text{MD}} \cdot q_{\mathcal{M}}), t_F + O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F))$ -RSA assumption with exponents from $[3, 2^{\kappa-2}]$. \square

Comparison with Standard RSA Signatures. We note that the concrete security of our variant with public metadata has almost identical bounds as those proven in [16] for RSA signatures. Re-formulating their analysis into our terminology, they showed that RSA signatures with FDH encoding has $\epsilon_F = q_{\mathcal{M}} \cdot \epsilon$ assuming the (ϵ, t) -RSA assumption with at most $q_{\mathcal{M}}$ queries to $H_{\mathcal{M}}$. For our version with public metadata, we lose multiplicative factors in the number of metadata hashing queries, q_{MD} . In other words, this provides evidence that our public metadata variant of RSA signatures provides similar unforgeability guarantees as standard RSA signatures.

Modified Protocol with Improved Concrete Security. In Appendix C, we present a slight modification to the protocol presented in this section. For the modified protocol, we show that concrete security bounds are identical to standard RSA signatures proven in [16]. In particular, this modified protocol can be proven to obtain $\epsilon_F = q_{\mathcal{M}} \cdot \epsilon$ assuming the (ϵ, t) -RSA assumption with at most $q_{\mathcal{M}}$ queries to $H_{\mathcal{M}}$ (identical to [16]).

We chose to present the simpler protocol with worse concrete security in the main body because it more aligns with our main RSA blind signatures with public metadata protocol. There is also a straightforward way to extend our modified protocol to become a blind signature. To our knowledge, this version does not provide better security bounds but adds unnecessary complexity. Therefore, we choose to extend the simpler protocol for our blind signature protocol (see Section 6). Nevertheless, the existence of a slightly modified version with identical concrete security bounds provides some evidence that our public metadata protocols should not degrade security significantly compared to the standard (no public metadata) protocols.

Discussion on Multi-Exponent Attacks. In Section 7, we outline multi-exponent attacks that assume the availability of an oracle for RSA decryption of arbitrary elements. Those attacks are not applicable

in this setting since the signer receives a message M and metadata D and signs (i.e., RSA decryption) on $H_{\mathcal{M}}(M \parallel D)$. If an attacker wishes to RSA decrypt an element X with respect to exponent $e = H_{\text{MD}}(D)$, it must find an input message M such that $H_{\mathcal{M}}(M \parallel D) = X$. This makes the attacks intractable for our non-blind RSA signatures. See Appendix D for more details.

Encoding of Message and Public Metadata. We note that our proof critically uses the fact that we pass both the message and public metadata into the FDH encoding algorithm. When \mathcal{A} programs the random oracle for $H_{\mathcal{M}}$, \mathcal{A} exponentiates according to $H_{\text{MD}}(D)$. This is only possible because D is also passed as an input to $H_{\mathcal{M}}$. If, instead, we only passed the message M into $H_{\mathcal{M}}$, then the best that \mathcal{A} can do is simply aim to guess the public metadata that will be later sent with M to the signing oracle. As there is no requirement that \mathcal{A}_F has even picked D yet, it is impossible for \mathcal{A} to guess this correctly. Therefore, it is integral that our algorithm passes the public metadata D into $H_{\mathcal{M}}$ for security.

Extension to PSS Encoding. Finally, we note that our proof can be modified when using PSS message encoding. In particular, the random oracle $H_{\mathcal{M}}$ is currently programmed in a trivial way to output random elements that is indistinguishable from FDH encodings. Instead, we can follow the identical proof techniques of Bellare and Rogaway [16] to program $H_{\mathcal{M}}$ to follow the PSS encoding. The main difference is the following. In the FDH proof, the adversary had to pick one of the query to $H_{\mathcal{M}}$ to choose to embed the random input target. By using the structure of PSS encoding, the adversary can instead embed the random input target into every query to $H_{\mathcal{M}}$. This significantly improves the ability of the adversary to win the game. As a result, we can obtain identical concrete security when using PSS encoding as those proved in [16] except for a similar additive factor exponentially small in κ when simulating the public metadata hash function H_{MD} .

6 RSA Blind Signatures with Public Metadata

In this section, we present our main protocol for RSA blind signatures with public metadata $\text{BlindRSA}_{\text{MD}}$. We also prove the unlinkability and unforgeability of our scheme. We note our scheme is a modified variant of the protocol presented in [9].

6.1 Protocol

We will present our RSA blind signature with public metadata formally. We will use the same public metadata hash function H_{MD} from Section 5.2. Additionally, we will assume some message encoding protocol defined by $H_{\mathcal{M}}$ and $\text{Verify}_{\mathcal{M}}$. We point readers back to Section 5.1 for more details on various options such as the FDH and PSS algorithms. As a note, we will utilize a modification presented by Lysyanskaya [42] that enables providing unlinkability even against malicious signers. In particular, it was shown that appending a random string to the message of bit length η is sufficient to provide anonymity against even adversarially chosen RSA modulus. We incorporate this technique into our protocol as well.

Choosing Public Metadata Set. Additionally, we also need checks in the setup portion to ensure that exponents corresponding the public metadata set MD do not satisfy certain algebraic properties. In Section 7, we will discuss why choosing public metadata set ahead of time is a good idea to guarantee unforgeability property of the blind variant of our protocol. However, we also note that we do not need to choose this set ahead of time, if we allow for sufficiently large outputs for public metadata hash function H_{MD} and a reasonable upper bound on the set size (see Section 7 for more details). However, for the rest of this section we assume that we will always need to choose the public metadata set ahead of time.

As far as the key requirements for choosing the public metadata set MD are concerned, it will be critical that MD is not too large for both anonymity. Recall that we ensure anonymity (unlinkability) amongst all users with the same public metadata. In the extreme case when MD is too large, each user could be assigned their own public metadata completely eliminating anonymity. Additionally, the set MD must be chosen ahead of time during the setup phase. Furthermore, the issuer should not sign any blinded messages for public metadata D outside of the permitted set $D \notin \text{MD}$. These will avoid potential forgeability attacks (as we discuss later). The restriction of choosing the public metadata set during setup does not impede most

applications. In our two example applications, the public metadata set can be fixed to be all countries and all valid expiration timestamps between scheduled key rotations.

$\text{BlindRSA}_{\text{MD}}.\text{Setup}(1^\lambda, \mathbf{MD})$:

1. Execute $(N, \phi(N)) \leftarrow \text{RSA}_{\text{MD}}.\text{Setup}(1^\lambda)$.
2. Generate random string $\text{salt} \in \{0, 1\}^\lambda$.
3. For $D \in \mathbf{MD}$:
 - (a) Compute $e \leftarrow H_{\text{MD}}(\text{salt} \parallel D)$.
 - (b) For $D' \in \mathbf{MD} \setminus \{D\}$:
 - i. Compute $e' \leftarrow H_{\text{MD}}(\text{salt} \parallel D')$.
 - ii. Compute $g \leftarrow \text{GCD}(e, e')$ as integers.
 - iii. While $g > 1$:
 - A. Set $e \leftarrow e/g$ as integers.
 - B. Compute $g \leftarrow \text{GCD}(e, e')$ as integers.
 - iv. If $e = 1$, go back and repeat from Step 2.
4. Return $\text{pk} \leftarrow (N, \text{salt}), \text{sk} \leftarrow \phi(N)$.

$\text{BlindRSA}_{\text{MD}}.\text{Blind}(M, D, \text{pk} \leftarrow (N, \text{salt}))$:

1. If $D \notin \mathbf{MD}$, return \perp .
2. Pick a random string rand of length η .
3. Set $M' \leftarrow M \parallel \text{rand}$.
4. Compute $e_{\text{MD}} \leftarrow H_{\text{MD}}(\text{salt} \parallel D)$.
5. Pick R uniformly at random from \mathbb{Z}_N^* .
6. Set state $\text{st} \leftarrow (\text{rand}, R)$.
7. Compute *blinded message* $B_M \leftarrow R^{e_{\text{MD}}} \cdot H_{\mathcal{M}}(M' \parallel D) \bmod N$.
8. Return (st, B_M) .

$\text{BlindRSA}_{\text{MD}}.\text{Sign}(B_M, D, \text{pk} \leftarrow (N, \text{salt}), \text{sk} \leftarrow \phi(N))$:

1. If $D \notin \mathbf{MD}$, return \perp .
2. Compute $e_{\text{MD}} \leftarrow H_{\text{MD}}(\text{salt} \parallel D)$.
3. Compute $d_{\text{MD}} \leftarrow (e_{\text{MD}})^{-1} \bmod \phi(N)$.
4. Return $S' \leftarrow (B_M)^{d_{\text{MD}}} \bmod N$.

$\text{BlindRSA}_{\text{MD}}.\text{Finalize}(\text{st} \leftarrow (\text{rand}, R), S', M, D, \text{pk} \leftarrow (N, \text{salt}))$:

1. Compute $S \leftarrow S' \cdot R^{-1} \bmod N$.
2. If $\text{BlindRSA}_{\text{MD}}.\text{Verify}((S, \text{rand}), M, D, \text{pk}) = 0$, return \perp .⁴

⁴This is optional for correctness but should be done to check that the server signed with the correct key.

3. Return *signature* (S, rand).

$\text{BlindRSA}_{\text{MD}}.\text{Verify}((S, \text{rand}), M, D, \text{pk} \leftarrow (N, \text{salt}))$:

1. Compute $M' \leftarrow M \parallel \text{rand}$.
2. Compute $e_{\text{MD}} \leftarrow H_{\text{MD}}(\text{salt} \parallel D)$.
3. Compute $X \leftarrow S^{e_{\text{MD}}} \bmod N$.
4. Return $\text{Verify}_{\mathcal{M}}(M' \parallel D, X)$.

Correctness. By Lemma 1, we know that the algorithm above is well-defined as the signer can always compute an inverse of $e_{\text{MD}} = H_{\text{MD}}(\text{salt} \parallel D)$. To show that the above protocol correctly verifies well-formed signatures, we can consider an execution of the blind signing protocol. The blinded message output by Blind will be

$$R^{e_{\text{MD}}} \cdot H_{\mathcal{M}}(M' \parallel D).$$

The output of Sign will be

$$(R^{e_{\text{MD}}} \cdot H_{\mathcal{M}}(M' \parallel D))^{(e_{\text{MD}})^{-1}} = R \cdot H_{\mathcal{M}}(M' \parallel D)^{(e_{\text{MD}})^{-1}}.$$

Finally, the output of Finalize is

$$H_{\mathcal{M}}(M' \parallel D)^{(e_{\text{MD}})^{-1}}.$$

Consider a final execution of Verify that will output

$$\text{Verify}_{\mathcal{M}}(M' \parallel D, H_{\mathcal{M}}(M' \parallel D)^{(e_{\text{MD}})^{-1} \cdot e_{\text{MD}}}) = \text{Verify}_{\mathcal{M}}(M' \parallel D, H_{\mathcal{M}}(M' \parallel D)) = 1.$$

Therefore, a well-formed signature will always be verified correctly.

Efficiency. We note that the efficiency is similar to with only a couple differences. First, we perform RSA operations with a larger exponent than standard RSA signatures. There are additional calls to H_{MD} and the signer must perform an inversion. However, we show that this only incurs minimal overhead (see Section 8). The major difference is setup that requires generating strong RSA modulus and checking exponents requiring $O(|\text{MD}|^2)$ time, but this is done once offline outside of the issuance and verification phases. We show in Section 7.1 that the number of salt resamples is very small for most choices of public metadata set sizes MD .

Our construction requires larger exponents than typically meaning more computation. In Appendix E, we discuss why restricting MD prevents potential DoS attacks leveraging large exponents.

6.2 Unforgeability

Next, we prove the unforgeability of $\text{BlindRSA}_{\text{MD}}$. To do this, we will show that any adversary that can forge signatures in $\text{BlindRSA}_{\text{MD}}$ is able to break the chosen-target, algebraically-restricted-exponent RSA inversion assumption. At a high level, we show that one can simulate signing oracle queries using the RSA decryption oracle in the chosen-target, algebraically-“restricted-exponent RSA inversion game.

Theorem 5. *Suppose that $(H_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ correspond to full domain hash (FDH) message encoding. Assuming the $(\epsilon, t, \ell, |\text{MD}|)$ -chosen-target, algebraically-restricted-exponent RSA inversion assumption (Definition 10) with exponents in $[3, e_{\text{max}}]$ such that $e_{\text{max}} = 2^{\kappa-2}$ and the random oracle model, then $\text{BlindRSA}_{\text{MD}}$ is $(\epsilon_F, t_F, \ell_F)$ -strong one-more unforgeable (Definition 2) where*

- $\epsilon_F = \epsilon + 2^{-\kappa+2 \log q_{\mathcal{M}}}$
- $t_F = t - O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$

- $\ell_F = \ell$

and the adversary \mathcal{A} runs in time at most t_F and makes at most ℓ_F , $q_{\mathcal{M}}$ and $q_{\mathbf{MD}}$ queries to the signing oracle and hash functions $H_{\mathcal{M}}$ and $H_{\mathbf{MD}}$ respectively.

Proof. To prove this, we will show that if there exists some adversary \mathcal{A}_F that successfully forges signatures for $\text{BlindRSA}_{\mathbf{MD}}$ to win the strong one-more unforgeability game, then we can use \mathcal{A}_F to construct an adversary \mathcal{A} to break the chosen-target, algebraically-restricted-exponent RSA inversion assumption. The forger \mathcal{A}_F runs in time t_F , wins the game with probability ϵ_F and uses at most ℓ_F signing queries. Our reduction will only increase the running time of the adversary \mathcal{A} by a factor of $O(q_{\mathcal{M}} + q_{\mathbf{MD}} + \ell_F)$, reduces the winning probability by at most $2^{-\kappa+2\log q_{\mathcal{M}}}$ and maintains the same number of oracle queries between signing and RSA decryption and makes at most $|\mathbf{MD}|$ queries to the exponent oracle. Furthermore the exponent distributions are identical as our setup protocol performs similar checks.

Our adversary \mathcal{A} will simulate \mathcal{A}_F to break the chosen-target, algebraically-restricted-exponent RSA game. Note that \mathcal{A} needs to be able to successfully simulate all hash and signing oracle queries performed by \mathcal{A}_F . \mathcal{A} simulates these queries as follows:

For any query to $H_{\mathbf{MD}}(D)$:

1. If $\mathbf{M}_{\mathbf{MD}}[D]$ is set, return $\mathbf{M}_{\mathbf{MD}}[D]$.
2. Compute $\mathbf{M}_{\mathbf{MD}}[D] \leftarrow \mathcal{O}^{\text{exp}}()$.
3. Return $\mathbf{M}_{\mathbf{MD}}[D]$.

For any query to $H_{\mathcal{M}}(M, D)$:

1. If $\mathbf{M}_{\mathcal{M}}[M, D]$ is set, return $\mathbf{M}_{\mathcal{M}}[M, D]$.
2. Compute $\mathbf{M}_{\mathcal{M}}[M, D] \leftarrow \mathcal{O}^{\mathcal{X}}()$.
3. Return $\mathbf{M}_{\mathcal{M}}[M, D]$.

For any query to $\mathcal{O}^{\text{Sign}}(X, D)$:

1. Compute $e_{\mathbf{MD}} \leftarrow H_{\mathbf{MD}}(\text{salt} \parallel D)$.
2. Return $\mathcal{O}^{\text{RSA}}(X, e_{\mathbf{MD}})$.

Adversarial Advantage. First, we note that the simulated outputs of \mathcal{O}^{exp} and the real outputs of $H_{\mathbf{MD}}$ are identical as $e_{\text{max}} = 2^{\kappa-2}$. Similarly, $\mathcal{O}^{\mathcal{X}}$ returns random elements identical to the FDH encoding. Finally, we note that signing always returns a valid signature as \mathcal{O}^{RSA} only receives exponents that are output by \mathcal{O}^{exp} .

Finally, \mathcal{A}_F will output the values $D, (S_i, M_i)_{i \in [x]}$ where $x \geq \text{cnt}_D + 1$ such that cnt_D is the maximum number of signing oracles performed with public metadata D . If \mathcal{A}_F successfully forges, then we know that $S_i^{H_{\mathbf{MD}}(\text{salt} \parallel D)} = H_{\mathcal{M}}(M_i \parallel D)$ for all $i \in [x]$. Then, \mathcal{A} will output $H_{\mathbf{MD}}(\text{salt} \parallel D), (S_i, H_{\mathcal{M}}(M_i \parallel D))_{i \in [x]}$ as its output to the chosen-target, algebraically-restricted-exponent RSA inversion game. Note, \mathcal{A} wins the game as long as they are all valid decryptions and there are no collisions in the output of $H_{\mathcal{M}}(M_i \parallel D)$. As there are at most $q_{\mathcal{M}}$ queries to $H_{\mathcal{M}}$, we know the collision probability is at most $q_{\mathcal{M}}^2/2^{-\kappa} = 2^{-\kappa+2\log q_{\mathcal{M}}}$ since $|\mathbb{Z}_N^*| \geq 2^{-\kappa}$.

Therefore, we know that \mathcal{A} wins the game if \mathcal{A}_F wins the unforgeability game and there are no collisions in the outputs of $H_{\mathcal{M}}$. So, the probability that \mathcal{A} wins is $\epsilon_F - 2^{-\kappa+2\log q_{\mathcal{M}}}$.

Running Time. Next, we analyze the running time of \mathcal{A} . We know that \mathcal{A}_F runs time at most t_F . For each hash and signing oracle query, \mathcal{A} performs $O(1)$ additional operations to correctly simulate responses. As there are most $q_{\mathcal{M}} + q_{\mathbf{MD}} + \ell_F$ queries overall, we see that \mathcal{A} requires $O(q_{\mathcal{M}} + q_{\mathbf{MD}} + \ell_F)$ additional time. Furthermore, we note that each signing query results in exactly one decryption oracle query. Therefore, we know that at most ℓ_F decryption oracle queries are made. Finally, each unique metadata $D \in \mathbf{MD}$ results in at most one exponent query meaning at most $|\mathbf{MD}|$ exponent queries. \square

Message Encoding Discussion. In Section 5.3, we mentioned that using the PSS message encoding as opposed to FDH message encoding would improve concrete security (see Section 5.1 for details on message encoding). To our knowledge, there is no benefit to concrete security when plugging in PSS for the security proof (Theorem 5) of our RSA blind signatures with public metadata. It is possible to program $H_{\mathcal{M}}$ such that its output will simulate PSS encodings (using the techniques from [15]), but the resulting reduction would have identical parameters. This phenomenon also exists in standard RSA blind signatures as the original work [15] proved security using only FDH as PSS did not provide any security improvements.

Discussion about Exponent Setup Check. In the setup, we sample random salts until the generated exponents do not satisfy a certain algebraic property. In particular, we check that there does not exist some metadata $D \in \mathbf{MD}$ where $e = H_{\mathbf{MD}}(\text{salt} \parallel D) = p_1 \cdot p_2 \cdots p_z$ is the product of z primes that are not necessarily distinct (that is, e can be the product of prime powers). Additionally, there must exist (not necessarily distinct) sequence of metadata $D_1, \dots, D_z \in \mathbf{MD} \setminus \{D\}$ such that $p_i \mid H_{\mathbf{MD}}(\text{salt} \parallel D_i)$. We perform this check as there are forgery attacks [26, 20] if an adversary can find such a sequence D, D_1, \dots, D_z of public metadata (see Appendix F for full attack details). Our requirement that the issuer only operates over valid metadata $D \in \mathbf{MD}$ ensures that the adversary is restricted to signing oracle (RSA inversion oracle) queries in the set $\{H_{\mathbf{MD}}(\text{salt} \parallel D) : D \in \mathbf{MD}\}$. We show these checks fail with small probability in Appendix F meaning salts are not re-sampled often.

Omitting Exponent Checks. It is possible to omit the exponent checks during setup if one is willing to withstand a small probability of a forgery. In Section 7.1, we analyze the probability that a random salt satisfies the algebraic properties necessary for a forging attack. If the public metadata set \mathbf{MD} is not too large, we show this probability is very small (see Figure 9). In our country example where there $|\mathbf{MD}| \leq 200$ countries, the probability that the forging attack in [20] succeeds is at most 2^{-52} and 2^{-72} for RSA modulus of length 2048 and 3072 respectively. Omitting exponent checks means that unforgeability probabilistically reduces to Def. 10 (see Section 7.1 for probabilities). We note that omitting exponent checks is identical to directly reducing to the chosen-target, restricted-exponent game (which does not have the checks in the exponent oracle) in Def. 9. We formalize this relationship in Appendix F.2.

6.3 Unlinkability

To prove that $\text{BlindRSA}_{\mathbf{MD}}$ provides unlinkability, we will rely on prior results in [42] that show that appending random strings to the message is sufficient to provide protections against maliciously generated parameters. At a high level, the work in [42] shows that for any maliciously chosen RSA modulus N and public exponent e , unlinkability for standard RSA blind signatures with appended random strings of bit length η holds with probability except $t \cdot 2^{-\eta}$ for adversaries running in time t .

For our case, the adversary is able to choose the public metadata D . However, this ends up being a more restrictive way for the adversary to choose as the final exponent $e_{\mathbf{MD}} = H_{\mathbf{MD}}(\text{salt} \parallel D)$. Clearly, it is an easier setting for the adversary to choose $e_{\mathbf{MD}}$ directly as opposed to D . As a result, we can rely on the results from [42] almost immediately as follows.

Theorem 6. *Assuming the random oracle model, $\text{BlindRSA}_{\mathbf{MD}}$ satisfies $(t \cdot 2^{-\eta}, t)$ -unlinkability (Definition 3).*

Proof. Consider the unlinkability games for $\text{BlindRSA}_{\mathbf{MD}}$ with public metadata and the standard RSA blind signatures in [42]. We note there is only one difference between the two settings. In the former, the adversary must choose D . The final RSA public exponent becomes $e_{\mathbf{MD}} = H_{\mathbf{MD}}(\text{salt} \parallel D)$. In the latter, the adversary is able to directly choose the final public exponent. Note, we can do this because $\text{BlindRSA}_{\mathbf{MD}}$ is nearly identical to standard RSA blind signatures after augmenting the public exponent to be $H_{\mathbf{MD}}(\text{salt} \parallel D)$. Clearly, the scenarios studied in [42] are more favorable for the adversary.

Suppose there exists an adversary \mathcal{A} that can compromise the unlinkability of $\text{BlindRSA}_{\mathbf{MD}}$ with probability ϵ and running time t . Note, \mathcal{A} outputs values public key N , messages M_0 and M_1 and public metadata D in the first step. We construct \mathcal{A}' that instead outputs $e = H_{\mathbf{MD}}(\text{salt} \parallel D)$ as the public exponent while keeping the RSA modulus N and messages M_0, M_1 identical. As the remainder of the games are identical, \mathcal{A}' also wins with probability ϵ and running time t . \square

7 Analyzing Multi-Exponent Attacks

First, we note our public metadata hash function $H_{\mathbf{MD}}$ is similar to division intractable hash functions used in [35] and analyzed in [26]. At a high level, these attacks work by trying to obtain a chain of exponents e, e_1, \dots, e_z with the following properties. Suppose $e = p_1 \cdot p_2 \cdots p_z$ is the product of z primes that are not necessarily distinct (in other words, e can be the product of prime powers). First, e is distinct from all of e_1, \dots, e_z . However, we note that the remaining z exponents, e_1, \dots, e_z , do not need to be distinct and can repeat the same exponent. Secondly, p_i divides e_i evenly as an integer, $p_i \mid e_i$. Borisov [20] showed that a forging attack is possible in this case by leveraging similar algebraic properties (this could also be translated to our chosen-target, restricted-exponent RSA inversion assumption in Definition 9 as well). We present this attack below. We also note that it is similar to the adversaries we presented for our chosen-exponent RSA assumptions in Section 4.3.

The attack by Borisov [20] aims to find a fixed public metadata D and a sequence of (not necessarily distinct) public metadata D_1, \dots, D_z with the following properties. Let $e = H_{\mathbf{MD}}(\text{salt} \parallel D) = p_1 \cdot p_2 \cdots p_z$ that is the product of z (not necessarily distinct) primes. Then, p_i should divide $e_i = H_{\mathbf{MD}}(\text{salt} \parallel D_i)$. That is, $p_i \mid e_i$. Then, the forging adversary does the following:

1. Set $x_0 = H(M \parallel D)$.
2. For each $i \in [z]$:
 - (a) Submit $x_{i-1}^{e_i/p_i}$ for blind signing with D_i and obtain x_i .
3. Return forgery x_z for message M and public metadata D .

As each intermediate blind signing output for every $i \geq 1$ satisfies:

$$x_i = H(M \parallel D)^{(e_1/p_1) \cdots (e_i/p_i)(1/e_1) \cdots (1/e_i)} = H(M \parallel D)^{1/(p_1 \cdots p_i)}$$

Then, $x_z = H(M \parallel D)^{1/(p_1 \cdots p_z)} = H(M \parallel D)^{1/e}$ is the forgery.

7.1 Analyzing Exponent Check

In our setup phase, we run exponent checks to guarantee that this algebraic property does not exist amongst valid public metadata in the set \mathbf{MD} . In this section, we analyze the existence of such bad sequences for a random choice of salt assuming $H_{\mathbf{MD}}$ uses a random oracle. This analysis will be useful in two scenarios. First, it bounds the number of salt re-samples during the setup phase. Secondly, in the case that the exponent check is omitted, this upper bounds the probability that the forging attack in [20] has the necessary algebraic property as a function of κ (prime bit-length) and $|\mathbf{MD}|$.

Theorem 7. *Consider security parameter λ and public metadata set \mathbf{MD} . Suppose that $H_{\mathbf{MD}}$ outputs γ -bit odd integers where $\gamma \leq \kappa - 2$. Let E be the event that there exists $D, D_1, \dots, D_z \in \mathbf{MD}$ such that:*

1. $D \neq D_i$ for all $i \geq 1$, but the remaining z metadata, D_1, \dots, D_z , are not necessarily distinct;
2. $e = H_{\mathbf{MD}}(\text{salt} \parallel D) = p_1 \cdot p_2 \cdots p_z$ where each p_i is a prime number that are not necessarily distinct;
3. $e \mid e_i = H_{\mathbf{MD}}(\text{salt} \parallel D_i)$ for all $i \geq 1$.

Then, for any randomly chosen salt and for every integer $x \geq 2$, $\Pr[E] \leq \frac{|\mathbf{MD}|^2}{x} + 2|\mathbf{MD}| \cdot \rho(\gamma/\log x)$ where ρ is Dickman's function.

Proof. In our analysis, we will rely on Dickman's function ρ for the analysis of smooth numbers. In particular, the probability that a random odd integer from $[1, y]$ has only prime factors no larger than x is $2 \cdot \rho(\log y/\log x)$. Let E' be the event that there exists any $D \in \mathbf{MD}$ such that $H_{\mathbf{MD}}(\text{salt} \parallel D)$ has no prime factors larger than x . Then, we know that $\Pr[E'] \leq 2|\mathbf{MD}| \cdot \rho(\gamma/\log x)$ since $H_{\mathbf{MD}}$ outputs γ -bit

Public Metadata Set Size ($ \mathbf{MD} $)	$\kappa = 1024$	$\kappa = 1536$	$\kappa = 2048$
100	2^{-54}	2^{-72}	2^{-87}
200	2^{-52}	2^{-70}	2^{-85}
500	2^{-50}	2^{-68}	2^{-83}
1,000	2^{-49}	2^{-66}	2^{-82}
5,000	2^{-45}	2^{-63}	2^{-78}
10,000	2^{-43}	2^{-61}	2^{-76}
50,000	2^{-40}	2^{-58}	2^{-73}
100,000	2^{-38}	2^{-56}	2^{-71}
1,000,000	2^{-33}	2^{-51}	2^{-66}

Figure 9: Table with upper bounds of $\Pr[E]$ from Theorem 7 with $\gamma = \kappa - 3$ and varying sizes of $|\mathbf{MD}|$. Recall that the RSA modulus bit-length is 2κ .

odd integers. Suppose that event E' does not occur. Fix any $D \in \mathbf{MD}$ and we want to bound the probability that there exists some sequence starting with D satisfying event E . Pick any prime factor p of $H_{\mathbf{MD}}(\text{salt} \parallel D)$ where $p > x$. If D is the start of the sequence satisfying the event E , then it better be that there exists some metadata $D' \in \mathbf{MD}$ such that p divides $H_{\mathbf{MD}}(\text{salt} \parallel D')$. The probability that a prime p divides a random odd integer from $[3, 2^{\kappa-2}]$ (output of $H_{\mathbf{MD}}$ for some $D' \in \mathbf{MD}$) is $1/p < 1/x$. Then, the probability that there exists any $D' \in \mathbf{MD}$ such that p divides $H_{\mathbf{MD}}(\text{salt} \parallel D')$ is at most $|\mathbf{MD}|/x$. In other words, the probability D is the start of some sequence satisfying event E is at most $|\mathbf{MD}|/x$. Using a Union bound over D , we get that $\Pr[E \mid \neg E'] \leq |\mathbf{MD}|^2/x$. Putting it altogether, we get the following: $\Pr[E] \leq \Pr[E \mid \neg E'] + \Pr[E'] \leq |\mathbf{MD}|^2/x + 2|\mathbf{MD}| \cdot \rho(\gamma/\log x)$. \square

To use the above, one should attempt to find an integer x that minimizes the probability $\Pr[E]$. We do this for different κ with $\gamma = \kappa - 3$ and $|\mathbf{MD}|$ in Figure 9. In general, the probabilities are very small meaning that the expected number of times that `salt` needs to be re-sampled is very small (salts are almost never re-sampled). We also prove an asymptotic bound (see Appendix F.1) on the above probability, but they are much looser than the values in Figure 9.

7.2 Analyzing Running Time of Attack

The analysis in the prior section only analyzes the probability that the chosen set of metadata and hash functions would allow such a bad chain of exponents exists. Any metadata forgery attack would still be required to compute and find such a bad chain of exponents to perform metadata forgery. We note that the best known attacks for metadata forgery by Borisov [20] still remain relatively inefficient. At a high level, this attack proceeds by attempting to find outputs of the metadata hash function $H_{\mathbf{MD}}$ that are smooth. This works by trying random metadata and then factoring them to check if they are smooth. Suppose, we found an integer that is x -smooth with no prime factors larger than x . Afterwards, the attack tries to find metadata whose output from $H_{\mathbf{MD}}$ have prime factors that divide the found integer. Note that this second step already seems to require x computations of $H_{\mathbf{MD}}$. From estimates of these attacks by Borisov [20], the runtime against our scheme with $\kappa = 1024$ and $\gamma = \kappa - 3 = 1021$ would require approximately 2^{113} time to find such a bad chain of exponents. Furthermore, the likelihood that such a bad exponent chain even exists is already quite low probability as we have shown above.

From the above, we can conclude that it seems sufficient to omit the exponent setup checks if the metadata set \mathbf{MD} cannot be fixed ahead of time. As long as one picks large enough output lengths for $H_{\mathbf{MD}}$ as well as a reasonable upper bound on the total size of the metadata set \mathbf{MD} , then it is already low probability that a bad chain of exponents exists making it susceptible to metadata forgery. As a note, the second condition of small \mathbf{MD} is critically important to maintain sufficient unlinkability anyways. Furthermore, even if there is a bad chain, the best attacks [20] for finding such bad exponent chains remain intractable to our knowledge.

Modulus Bit Length	Algorithm	Public Metadata	No Public Metadata
2048	Blind	1.4 ± 0.01	0.31 ± 0.01
2048	Sign	4.3 ± 0.01	1.6 ± 0.01
2048	Finalize	1.1 ± 0.01	0.028 ± 0.001
2048	Verify	1.1 ± 0.01	0.025 ± 0.001
3072	Blind	4.1 ± 0.11	0.61 ± 0.01
3072	Sign	12 ± 0.16	4.1 ± 0.01
3072	Finalize	3.4 ± 0.01	0.053 ± 0.001
3072	Verify	3.4 ± 0.01	0.053 ± 0.001
4096	Blind	8.9 ± 0.12	1.2 ± 0.01
4096	Sign	26 ± 0.16	8.1 ± 0.01
4096	Finalize	7.9 ± 0.01	0.088 ± 0.001
4096	Verify	7.9 ± 0.01	0.088 ± 0.001

Figure 10: Comparison of RSA blind signature with and without public metadata. Computational costs are presented in milliseconds. Each experiment was repeated 100 times.

	[57]	[62]	Ours (2048)	Ours (3072)
Blind	1.6 ± 0.08	1.1 ± 0.13	1.4 ± 0.01	4.1 ± 0.11
Sign	1.0 ± 0.03	5.6 ± 0.38	4.3 ± 0.01	12 ± 0.16
Finalize	3.3 ± 0.21	6.6 ± 0.47	1.1 ± 0.01	3.4 ± 0.01
Verify	3.9 ± 0.28	1.3 ± 0.05	1.1 ± 0.01	3.4 ± 0.01
Signature Size	48	145	256	384
Security Level	128	192	112	128

Figure 11: Comparison of $\text{BlindRSA}_{\text{MD}}$ with 2048-bit and 3072-bit modulus and solutions using pairings [57] and POPRFs with only secret key verification [62]. Computational costs are presented in milliseconds and signature sizes in bytes. Security levels are presented for the underlying curve or RSA group. Each experiment was repeated 100 times.

8 Experimental Evaluation

In this section, we present an experimental evaluation of our RSA blind signatures with public metadata.

Our Implementation. Our open-source code can be found at [1] containing an implementation of RSA blind signatures with and without public metadata. Both implementations utilize PSS message encodings and follow their respective specifications [12, 44] for other details such as instantiating hash functions and ensuring domain separation. Our implementations rely on BoringSSL [2] for many underlying cryptographic primitives such as basic RSA functionality. For our protocol, the random string for unlinkability will be $\eta = 384$ bits following [42, 44] along with SHA384 and 384-bit salts for our hash function.

Experimental Setup. We conducted our experiments using Ubuntu PCs with 4 cores, AMD EPYC 7B12 2.2 GHz and 32 GB of RAM. Our experiments are all executed using a single thread.

Key Generation. For a standard RSA key generation protocol that does not support public metadata, key generation times range from 0.09 to 1 second for RSA modulus sizes from 2048 to 4096 bits. For our protocol the key needs to have a strong RSA modulus, RSA key generation takes around 2 seconds for 2048, 45 seconds for 3072 and 90 seconds for 4096 bit strong moduli. We present the computational cost of our exponent checks in Figure 12. As setup is an infrequent operation, the increase in setup time is inconsequential.

Comparison with RSA Blind Signatures. We report all our results of the computational costs in Figure 10. Note that Verify is also a sub-routine for the Finalize routine. In general, we note that the protocol with public metadata is slower. This is expected as Sign requires an additional inverse operation




Figure 12: Setup time for different sizes of moduli against different Public Metadata set sizes. Every experiment was repeated 10 times.

modulo $\phi(N)$ to compute the private exponent. Moreover, the output of H_{MD} and hence the public exponent can be quite large that is needed for security. Therefore, it is not surprising that the public metadata variant is slower. Nevertheless, the RSA blind signatures with public metadata are still more than sufficiently fast enough for real-world applications. Both protocols have identical signature sizes depending on modulus length.

Comparison with Pairing Variant. We compare with the public metadata solution in [57] using pairings in Figure 11. We implemented the protocol in C++ using the curve BLS12-381 with the AMCL library [60] following the IRTF draft [18]. We compare with our 2048-bit modulus solution here following the draft [12] and with 3072-bit modulus solution in Figure 11. Our scheme was similar during Blind, 4x slower during Sign but 3x faster during Finalize. Our scheme is also 3.5x faster for Verify. The signature size for this protocol is 5x smaller than ours. We note again that pairings are not readily available in production cryptographic libraries.

Comparison with OPRF Variant. We compare with OPRF-based protocols with public metadata [62, 38] that only support secret key verification in Figure 11. We used the implementation available in Rust [50] with the curve P384 following the RFC [30] and compare with our 2048-bit modulus protocol here following the draft [12] and with 3072-bit modulus solution in Figure 11. Our scheme was slightly slower during Blind, slightly faster for Sign, Verify and 6x faster for Finalize. This protocol’s signature size is smaller than our schemes.

9 Deployment Telemetry

We report anonymous telemetry over 30 days from April 25, 2024 until May 23, 2024 on our deployment of RSA blind signatures with public metadata with 2048-bit modulus for GoogleOne VPN.

Query Volume. We recorded query volume over each minute period in Figure 13. Our systems served 5,953 queries on average every minute. At peak, there were 7,276 queries within one minute. This works out to 99.2 queries per second on average and 121.3 queries per second at peak. On average, we received more than 85 million queries per day and more than 250 million total queries over the entire month. Queries were




Figure 13: Queries per second measured every minute period.

split 50.6% for issuance and the remainder for redemption. Issuance is naturally higher as some issued tokens may never be redeemed. The diurnal query pattern reflects the concentration of users in North America and Europe.

Latency. Overall, the median latency is stable for both operations at 33 ms for issuance and 91 ms for redemption. We also report that the 90-th (99-th) percentile latencies were 48 ms (57 ms) and 240 ms (352 ms) for issuance and redemption respectively. The higher latency (and variance) of redemption come from database reads and write for tracking and preventing double spending of tokens.

Cost Modeling. For our deployment, the majority of our costs are tied to computational costs and long-term storage for double spending checks. The cost of storing all spent tokens over the month cost less than \$400 (or less than \$0.16 every 100,000 redeemed tokens). Computational costs were less than \$2300 (or less than \$0.92 every 100,000 queries). All costs estimated using standard Google cloud computing and database pricing found here [36].

10 Conclusion and Open Problems

In this paper, we present a protocol for RSA blind signatures that support public metadata. We prove our schemes secure and provide strong evidence that concrete security of our scheme is nearly identical to standard RSA blind signatures. Furthermore, we show that our public metadata protocol is concretely efficient with comparable overhead as standard RSA blind signatures and other solutions that use less available cryptography or without public key verification. We also report on anonymous deployment telemetry to showcase the scalability of our protocol in real-world settings.

We leave the following open questions:

- Can one prove security of our variant of RSA blind signatures with public metadata from the chosen-target RSA inversion assumption (Definition 7) instead of the chosen-target, (algebraically-)restricted-exponent RSA inversion assumptions (Definitions 9 and 10) that we introduced? This seems possible since we are able to prove the non-blind version of RSA signatures with public metadata directly from the RSA assumption.

- What is the relationship between the chosen-target (Definition 7) and chosen-target, (algebraically) restricted-exponent (Definitions 9 and 10) variants of the RSA inversion assumption?
- What is the relationship between the chosen-target, restricted-exponent (Definition 9) and chosen-target, chosen-exponent (Definition 12) variants of the RSA inversion assumption?

Acknowledgements. The authors would like to thank Scott Hendrickson for proposing and helping with the problem, Cathie Yun for help implementing standard RSA blind signatures, and David Benjamin, Steven Valdez, Jim Laskey, Robert Liu, and Chris Wood for helping review the protocol and our implementations. The authors would also like to thank Sarvar Patel and Giuseppe Persiano for very insightful discussions on the new RSA assumptions and feedback on earlier versions of this paper as well as Jonathan Katz for advice and discussion on estimating security parameters. Finally, the authors would like to thank Nikita Borisov for pointing out and sharing their metadata forgery attacks.

References

- [1] Anonymous tokens (AT). <https://github.com/google/anonymous-tokens>.
- [2] BoringSSL. <https://github.com/google/boringssl>.
- [3] Getting started with trust tokens. <https://web.dev/trust-tokens/>.
- [4] iCloud private relay overview. https://www.apple.com/icloud/docs/iCloud_Private_Relay_Overview_Dec2021.pdf.
- [5] Introducing private click measurement, PCM. <https://webkit.org/blog/11529/introducing-private-click-measurement-pcm/>.
- [6] VPN by Google One, explained. <https://one.google.com/about/vpn/howitworks>.
- [7] Michel Abdalla, Chanathip Namprempre, and Gregory Neven. On the (im)possibility of blind message authentication codes. In David Pointcheval, editor, *CT-RSA 2006*, volume 3860 of *LNCS*, pages 262–279. Springer, Heidelberg, February 2006.
- [8] Masayuki Abe and Jan Camenisch. Partially blind signature schemes. In *1997 Symposium on Cryptography and Information Security*, 1997.
- [9] Masayuki Abe and Eiichiro Fujisaki. How to date blind signatures. In Kwangjo Kim and Tsutomu Matsumoto, editors, *ASIACRYPT'96*, volume 1163 of *LNCS*, pages 244–251. Springer, Heidelberg, November 1996.
- [10] Masayuki Abe and Tatsuaki Okamoto. Provably secure partially blind signatures. In *Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings*, volume 1880 of *Lecture Notes in Computer Science*, pages 271–286. Springer, 2000.
- [11] Martin R Albrecht, Alex Davidson, Amit Deo, and Daniel Gardham. Crypto dark matter on the torus: Oblivious prfs from shallow prfs and tthe. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 447–476. Springer, 2024.
- [12] Ghous Amjad, Scott Hendrickson, Christopher Wood, and Kevin Yeo. Partially blind RSA signatures. <https://datatracker.ietf.org/doc/draft-amjad-cfrg-partially-blind-rsa/>, 2023.
- [13] <https://github.com/arkworks-rs>, no date.
- [14] Niko Bari and Birgit Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 480–494. Springer, Heidelberg, May 1997.
- [15] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The power of RSA inversion oracles and the security of Chaum's RSA-based blind signature scheme. In Paul F. Syverson, editor, *FC 2001*, volume 2339 of *LNCS*, pages 319–338. Springer, Heidelberg, February 2002.
- [16] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 399–416. Springer, Heidelberg, May 1996.

- [17] Fabrice Benhamouda, Tancrede Lepoint, Michele Orrù, and Mariana Raykova. Publicly verifiable anonymous tokens with private metadata bit. Cryptology ePrint Archive, Paper 2022/004, 2022. <https://eprint.iacr.org/2022/004>.
- [18] Dan Boneh, Sergey Gorbunov, Riad Wahby, Hoeteck Wee, Christopher Woor, and Zhenfei Zhang. Bls signatures. <https://datatracker.ietf.org/doc/draft-irtf-cfrg-bls-signature/>, 2023.
- [19] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Heidelberg, December 2001.
- [20] Nikita Borisov. Metadata forgery in partially-blind rsa signatures. Personal Communication, no date.
- [21] Carlos Cardenas. Cloud security: The challenges with key management in the cloud (and everywhere else). <https://www.tritondatacenter.com/blog/cloud-security-the-challenges-with-key-management-in-the-cloud-and-everywhere-else>, no date.
- [22] Melissa Chase, F. Betül Durak, and Serge Vaudenay. Anonymous tokens with stronger metadata bit hiding from algebraic macs. In *Advances in Cryptology – CRYPTO 2023*, pages 418–449. Springer Nature Switzerland, 2023.
- [23] David Chaum. Blind signatures for untraceable payments. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *CRYPTO’82*, pages 199–203. Plenum Press, New York, USA, 1982.
- [24] David Chaum. Blind signature system. In David Chaum, editor, *CRYPTO’83*, page 153. Plenum Press, New York, USA, 1983.
- [25] <https://github.com/cloudflare/circl>, no date.
- [26] Jean-Sébastien Coron and David Naccache. Security analysis of the gennaro-halevi-rabin signature scheme. In *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*, pages 91–101. Springer, 2000.
- [27] Geoffroy Couteau, Thomas Peters, and David Pointcheval. Removing the strong RSA assumption from arguments over the integers. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 321–350. Springer, Heidelberg, April / May 2017.
- [28] <https://csrc.nist.gov/projects/cryptographic-module-validation-program/modules-in-process/modules-in-process-list>, no date. Accessed on 08.08.2024.
- [29] Ivan Damgård and Eiichiro Fujisaki. A statistically-hiding integer commitment scheme based on groups with hidden order. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 125–142. Springer, Heidelberg, December 2002.
- [30] Alex Davidson, Armando Faz-Hernandez, Nick Sullivan, and Christopher Wood. Oblivious pseudorandom functions (oprfs) using prime-order groups. <https://datatracker.ietf.org/doc/rfc9497/>, 2024.
- [31] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *PoPETs*, 2018(3):164–180, July 2018.
- [32] Randall Degges. The hardest thing about data encryption. <https://developer.okta.com/blog/2019/07/25/the-hardest-thing-about-data-encryption>, no date.
- [33] Frank Denis, Frederic Jacobs, and Christopher Wood. RSA blind signatures. <https://datatracker.ietf.org/doc/draft-irtf-cfrg-rsa-blind-signatures/>, 2023.
- [34] Eiichiro Fujisaki and Tatsuaki Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In Burton S. Kaliski Jr., editor, *CRYPTO’97*, volume 1294 of *LNCS*, pages 16–30. Springer, Heidelberg, August 1997.
- [35] Rosario Gennaro, Shai Halevi, and Tal Rabin. Secure hash-and-sign signatures without the random oracle. In Jacques Stern, editor, *Advances in Cryptology — EUROCRYPT ’99*, pages 123–139, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.
- [36] <https://cloud.google.com/pricing>, no date.
- [37] Scott Hendrickson and Christopher Wood. Privacy pass: Public metadata issuance. <https://datatracker.ietf.org/doc/draft-ietf-privacypass-public-metadata-issuance/00/>, 2023.
- [38] Scott Hendrickson and Christopher A. Wood. Privacy Pass Issuance Protocols with Public Metadata. Internet-Draft draft-ietf-privacypass-public-metadata-issuance-00, Internet Engineering Task Force, April 2024.

- [39] Sharon Huang, Subodh Iyengar, Sundar Jeyaraman, Shiv Kushwah, Chen-Kuei Lee, Zutian Luo, Payman Mohassel, Ananth Raghunathan, Shaahid Shaikh, Yen-Chieh Sung, et al. Dit: De-identified authenticated telemetry at scale, 2021.
- [40] Julia Kastner, Julian Loss, and Jiayu Xu. On pairing-free blind signature schemes in the algebraic group model. In *IACR International Conference on Public-Key Cryptography*, pages 468–497. Springer, 2022.
- [41] Ben Kreuter, Tancrede Lepoint, Michele Orrù, and Mariana Raykova. Anonymous tokens with private metadata bit. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 308–336. Springer, Heidelberg, August 2020.
- [42] Anna Lysyanskaya. Security analysis of RSA-BSSA. In *IACR International Conference on Public-Key Cryptography*, pages 251–280. Springer, 2023.
- [43] Pieter Moree. Integers without large prime factors: from ramanujan to de bruijn. *Integers*, 14, 2014.
- [44] K. Moriarty, B. Kaliski, J. Jonsson, and A. Rusch. PKCS #1: RSA cryptography specifications version 2.2. <https://datatracker.ietf.org/doc/html/rfc8017>, 2016.
- [45] <https://github.com/nss-dev/nss>, no date.
- [46] <https://github.com/openssl/openssl>, no date.
- [47] Michele Orrù, Stefano Tessaro, Greg Zaverucha, and Chenzhi Zhu. Oblivious issuance of proofs. *IACR Crypto 2024*, Springer-Verlag, 2024.
- [48] <https://github.com/nccgroup/pairing-bls12381>, no date.
- [49] David Pointcheval and Jacques Stern. Provably secure blind signature schemes. In Kwangjo Kim and Tsutomu Matsumoto, editors, *ASIACRYPT'96*, volume 1163 of *LNCS*, pages 252–265. Springer, Heidelberg, November 1996.
- [50] <https://github.com/facebook/voprf>, no date.
- [51] Luke Probasco. Key management: The hardest part of encryption. <https://info.townsendsecurity.com/bid/74336/key-management-the-hardest-part-of-encryption>, no date.
- [52] <https://github.com/relic-toolkit/relic>, no date.
- [53] Dominique Schröder and Dominique Unruh. Security of blind signatures revisited. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 662–679. Springer, Heidelberg, May 2012.
- [54] Atle Selberg. An elementary proof of the prime-number theorem. *Annals of Mathematics*, pages 305–313, 1949.
- [55] Victor Shoup. Practical threshold signatures. In *Advances in Cryptology—EUROCRYPT 2000: International Conference on the Theory and Application of Cryptographic Techniques Bruges, Belgium, May 14–18, 2000 Proceedings 19*, pages 207–220. Springer, 2000.
- [56] Victor Shoup. *A computational introduction to number theory and algebra*. Cambridge university press, 2009.
- [57] Tjerand Silde and Martin Strand. Anonymous tokens with public metadata and applications to private contact tracing. In *Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, May 2–6, 2022, Revised Selected Papers*, pages 179–199. Springer, 2022.
- [58] Stefano Tessaro and Chenzhi Zhu. Short pairing-free blind signatures with exponential security. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 782–811. Springer, Heidelberg, May / June 2022.
- [59] <https://github.com/bcgit/bc-java>, no date.
- [60] <https://github.com/miracl/core>, no date.
- [61] <https://www.nccgroup.com/us/research-blog/public-report-security-review-of-rsa-blind-signatures-with-public-metadata/>, 2023.
- [62] Nirvan Tyagi, Sofia Celi, Thomas Ristenpart, Nick Sullivan, Stefano Tessaro, and Christopher A. Wood. A fast and simple partially oblivious prf, with applications. In *Advances in Cryptology – EUROCRYPT 2022*, pages 674–705. Springer International Publishing, 2022.
- [63] <https://github.com/wolfSSL/wolfssl>, no date.
- [64] Fangguo Zhang, Reihaneh Safavi-Naini, and Willy Susilo. Efficient verifiably encrypted signature and partially blind signature from bilinear pairings. In Thomas Johansson and Subhamoy Maitra, editors, *INDOCRYPT 2003*, volume 2904 of *LNCS*, pages 191–204. Springer, Heidelberg, December 2003.

A Discussion about Strong RSA Modulus

In practice today, the `Gen` algorithm is used to generate only the RSA modulus N . The distribution \mathcal{D}_N to generate the public exponent is typically a singleton element. To generate the RSA modulus N , the standard approach is to pick two random prime numbers that are sufficiently large. For more details, we point to the RSA modulus generation algorithm within `BoringSSL` [2].

Generating Strong RSA Modulus. For our new protocol, we need to modify the `Gen` algorithm to always output strong RSA modulus $N = pq$ where p and q are safe primes. The straightforward approach is to use the same `Gen` algorithm of generating random primes and simply check whether the chosen primes are safe. For any prime p , one can simply use primality testing algorithms to check whether $(p - 1)/2$ is prime to see if p is a safe prime. As safe primes are more scarce than primes, this algorithm will be slower. However, we can estimate the asymptotic additional time needed to generate safe primes compared to normal primes. By the prime number theorem [54], we know that the density of prime numbers is $O(1/\kappa)$ for numbers at most 2^κ . For safe primes, we can rely on one of several conjectures such as Dickson’s conjecture or the twin prime conjecture (see [56]) that state the density of safe primes is $O(1/\kappa^2)$. Therefore, we could expect that the `Gen` algorithm finds a safe prime after generating $O(\kappa)$ random primes.

Comparing Assumptions. In Section 3.2, we present two RSA assumptions: the standard one and one restricted to only strong RSA modulus. We now prove the relation of these two assumptions that was formally presented in Theorem 1. At a high level, we will use the same ideas as the prior section comparing the densities of primes and safe primes.

Proof of Theorem 1. First, we will assume there is an adversary \mathcal{A} that breaks the strong modulus version of the assumption with probability ϵ and running time t . We show that \mathcal{A} will also break the standard RSA game with probability $O(\epsilon/\kappa^2)$ and running time t to prove the statement. Recall that κ is the length of the prime factors of the RSA modulus.

To prove this, we simply analyze the probability that a randomly generated RSA modulus is a strong RSA modulus. Assuming Dickson’s conjecture [56], the probability that a randomly generated prime number of length κ is also a safe prime is $\Theta(1/\kappa)$. For a RSA modulus to be strong, both prime factors must be safe. Therefore, this occurs with probability $\Theta(1/\kappa^2)$. Conditioned on the RSA modulus being strong, the adversary \mathcal{A} wins the standard RSA game with probability ϵ . Therefore, \mathcal{A} wins the game with probability at least $O(\epsilon/\kappa^2)$. \square

To our knowledge, there is no straightforward relation in the opposite direction. In particular, it is possible that the RSA assumption for strong modulus is true while the standard RSA assumption is false. This could be possible if there exists an attack that targets specifically non-strong RSA modulus. We leave it as an open question to determine the relation in this direction.

B Equivalence of Unforgeability Definitions

In this section, we consider the alternative definition of unforgeability with public metadata that was introduced in [57]. At a high level, the difference is that this alternative definition does not limit the total number of queries explicitly. Instead, the adversary is limited to at most ℓ oracle signing queries for each choice of public metadata D . Moreover, to win the game, adversary needs to output $\ell + 1$ signature message pairs under a fixed choice of metadata. We denote this definition as Alternative Strong One-More Unforgeability with Public Metadata, $\mathsf{G}_{\text{Tok}, \mathcal{A}, \ell}^{\text{ASOMUF}}(\lambda)$.

Definition 14 ((ϵ, t, ℓ) -Alternative Strong One-More Unforgeability with Public Metadata). *Let λ be the security parameter and consider the game $\mathsf{G}_{\text{Tok}, \mathcal{A}, \ell}^{\text{ASOMUF}}$ in Figure 14. A token scheme Tok is (ϵ, t, ℓ) -alternative strong one-more unforgeable if, for any adversary \mathcal{A} that runs in probabilistic time t and for any $\ell \geq 0$,*

$$\Pr[\mathsf{G}_{\text{Tok}, \mathcal{A}, \ell}^{\text{ASOMUF}}(\lambda) = 1] \leq \epsilon.$$

Game $G_{\text{Tok}, \mathcal{A}, \ell}^{\text{ASOMUF}}(\lambda)$:	Oracle $\mathcal{O}^{\text{Sign}}(\text{msg}, D)$:
$(\text{pk}, \text{sk}) \leftarrow \text{Tok.Setup}(1^\lambda)$ Initialize $\text{cnt}_D \leftarrow 0$ for all choices of public metadata D . $D, (S_i, M_i)_{i \in [\ell+1]} \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Sign}}}(\text{crs}, \text{pk})$ Return 1 if and only if all the following hold: - $\forall i \neq j \in [\ell+1], S_i \neq S_j \vee M_i \neq M_j$ - $\forall i \in [\ell+1], \text{Tok.Verify}(S_i, M_i, D, \text{pk}) = 1$	If $\text{cnt}_D \geq \ell$: Return \perp . $\text{cnt}_D \leftarrow \text{cnt}_D + 1$ Return $\text{Tok.Sign}(\text{msg}, D, \text{pk}, \text{sk})$.

Figure 14: Alternative Strong One-More Unforgeability (ASOMUF) Game with Public Metadata.

We show that definitions 14 and 2 are essentially equivalent up to some loss in the number of oracle queries due to being unable to bound the total number of queries executed by the adversary in the alternative definition.

Theorem 8. *A token scheme Tok satisfies alternative strong one-more unforgeability with public metadata if and only if the token scheme Tok also satisfies Def. 2. The reductions only lose up to $O(t + \ell)$ additive factors in the adversary’s running time and number of oracle signing queries.*

Proof. Assume Tok does not satisfy (ϵ, t, ℓ) -alternative strong one-more unforgeability and there exists an adversary \mathcal{A} that wins the game G^{ASOMUF} . The same adversary also wins the game G^{SOMUF} meaning Tok is not (ϵ, t, t) -strong one-more unforgeable as the maximum number of oracle queries sent by the adversary is at most its running time.

For the other direction, assume that Tok is not (ϵ, t, ℓ) -strong one-more unforgeable. Suppose adversary \mathcal{A} wins the game G^{SOMUF} . If the output of \mathcal{A} consists of $\ell + 1$ tuples, then it also wins G^{ASOMUF} . Otherwise, we can afford to use additional queries to the signing oracle in G^{ASOMUF} . Suppose that \mathcal{A} outputs $\text{cnt}_D + 1$ tuples using at most cnt_D signing oracle queries for public metadata D . Then, we can win G^{SOMUF} by the adversary making an additional $\ell - \text{cnt}_D - 1$ signing oracle queries for different messages. Therefore, Tok is not $(\epsilon, t + O(\ell), 2\ell)$ -alternative strong one-more unforgeable as we bound the additional running time by $O(\ell)$ and the additional number of signing oracle queries by at most ℓ . \square

Winning Conditions for RSA Inversion Games. We note the same equivalence can be shown to exist for the RSA assumption games. In our games, the adversary’s condition for winning is outputting the tuple $(X_i, Y_i)_{i \in [z]}$ where $z > \text{cnt}_e$ where cnt_e oracle queries are done. For some prior works, the adversary’s winning condition is outputting a pair (X, Y) for e that is different from any query to \mathcal{O}^{RSA} . We show the two are essentially equivalent. Since there are only cnt_e queries to \mathcal{O}^{RSA} for e , the tuple $(X_i, Y_i)_{i \in [z]}$ must contain at least one pair (X, Y) that was never queried. If an adversary is able to find a single pair (X, Y) that was never queried to \mathcal{O}^{RSA} , the adversary can pad the results of the cnt_e oracle queries for e to obtain a tuple of $z > \text{cnt}_e$ pairs.

C Modified RSA Signatures with Public Metadata with Better Concrete Security

In this section, we show that there is a slightly modified variant of the RSA signatures with public metadata protocol from Section 5 with better concrete security reductions. We chose to present this variant in the appendix as the main goal of our paper is the blind protocol. To our knowledge, this modified variant when extended to the blind version provides no additional security benefit, but adds unnecessary complexity. Therefore, we chose to present the simpler version with worse security bounds in the main body (Section 5) and the modified version here.

The resulting security bounds of this modified protocol identically match those obtained by Bellare and Rogaway [16] for RSA signatures (without public metadata) using FDH encodings.

C.1 Modified Protocol

Our modified protocol makes two minor changes to the protocol presented in Section 5. First, the public key is augmented to contain both the strong RSA modulus N as well as a standard RSA public exponent e . Henceforth, we assume the public key is of the form $\mathbf{pk} = (N, e)$. Additionally, we augment each public exponent by multiplying with e . We present the modified protocol with all changes highlighted in blue (the hash function H_{MD} remains unchanged).

$\text{RSA}_{\text{MD}}.\text{Setup}(1^\lambda)$:

1. Randomly generate $(N, p, q, \mathcal{D}_N) \leftarrow_R \text{Gen}(1^\lambda)$.
2. Draw $e \leftarrow \mathcal{D}_N$.
3. Compute $\phi(N) \leftarrow (p-1)(q-1)$.
4. Return $\mathbf{pk} \leftarrow (N, e)$, $\mathbf{sk} \leftarrow \phi(N)$.

$\text{RSA}_{\text{MD}}.\text{Sign}(M, D, \mathbf{pk} \leftarrow (N, e), \mathbf{sk} \leftarrow \phi(N))$:

1. Compute $e_{\text{MD}} \leftarrow e \cdot H_{\text{MD}}(D)$.
2. Compute $d_{\text{MD}} \leftarrow (e_{\text{MD}})^{-1} \bmod \phi(N)$.
3. Compute $M_{\text{MD}} \leftarrow H_{\mathcal{M}}(M \parallel D)$.
4. Return signature $S \leftarrow (M_{\text{MD}})^{d_{\text{MD}}} \bmod N$.

$\text{RSA}_{\text{MD}}.\text{Verify}(S, M, D, \mathbf{pk} \leftarrow (N, e))$:

1. Compute $e_{\text{MD}} \leftarrow e \cdot H_{\text{MD}}(D)$.
2. Compute $X \leftarrow S^{e_{\text{MD}}} \bmod N$.
3. Return $\text{Verify}_{\mathcal{M}}(M \parallel D, X)$.

C.2 New Reduction from RSA to Multi-Exponent RSA Assumption

First, we present a modified reduction from the RSA assumption to the multi-exponent RSA assumption from Theorem 2 under different exponent distributions. For the RSA assumption, we will assume that a fixed public exponent e is used as done in practice. For the multi-exponent RSA assumption, we will use the exponent distribution $\mathcal{D}'_{N,e}$ defined as follows. An uniformly random odd number e' is chosen from $[3, 2^{\kappa-2}]$. The final output is $e \cdot e'$.

We make the following two observations about $\mathcal{D}'_{N,e}$. First, we note that the output can be much larger than $\kappa - 2$ bits. In fact, the output of $\mathcal{D}'_{N,e}$ could be as large as $2(\kappa - 2) = 2\kappa - 4$ bits. Nevertheless, we note that this is not an issue because any output of $\mathcal{D}'_{N,e}$ will be invertible modulo $\phi(N)$ as it is the product $e \cdot e'$. We know that e is invertible modulo $\phi(N)$ since it is a valid RSA public exponent with respect to modulus N . As e' is an odd number from the range $[3, 2^{\kappa-2}]$, we know that e' is also invertible modulo $\phi(N)$. As a result, $e \cdot e'$ is also invertible modulo $\phi(N)$.

With these two modifications, we present a reduction that has no loss. Later, we show that $\mathcal{D}'_{N,e}$ is identical to the public exponents that are generated in the modified protocol from the prior section. We will use this fact to prove better security bounds for the modified protocol.

Theorem 9. *If the (ϵ, t) -RSA assumption for strong modulus (Definition 6) is true with a fixed exponent e , then the $(\epsilon, t - O(\ell), \ell)$ -multi-exponent RSA assumption (Definition 8) is true with exponents drawn from $\mathcal{D}'_{N,e}$ defined above.*

Proof. Towards a contradiction, suppose that there exists an adversary \mathcal{A}' that wins the multi-exponent RSA game with probability ϵ' with running time t' when given ℓ challenge exponents. We build an adversary \mathcal{A} for the RSA game that wins with probability ϵ' with the running time $t = t' + O(\ell)$.

Recall that \mathcal{A} is given an exponent e and random target X and must produce Y satisfying $Y^e = X \pmod N$. To do this, \mathcal{A} will pick ℓ random odd numbers from the set $[3, 2^{\kappa-2}]$ that we denote as z_1, \dots, z_ℓ . We choose the ℓ challenge exponents as $e_i = e \cdot z_i$ for all $i \in [\ell]$. Note, each e_i are essentially drawn from the distribution $\mathcal{D}'_{N,e}$. Finally, \mathcal{A} passes the strong RSA modulus N , the ℓ exponents as well as the random target X to $\mathcal{A}'(N, e_1, \dots, e_\ell, X)$. We know that \mathcal{A}' returns i and a valid decryption Y of X under e_i with probability ϵ' . Additionally, we know that $Y = X^{1/e_i} = X^{1/(e \cdot z_i)}$ in this case. Therefore, \mathcal{A} can output $Y^{z_i} = X^{1/e}$ to win the RSA game with probability ϵ' .

For the running time, \mathcal{A} uses an additional $O(\ell)$ time to generate random exponents. If \mathcal{A}' runs in time t' , then \mathcal{A} runs in time $t' + O(\ell)$. \square

C.3 Unforgeability

With the new reduction, we can essentially re-do the proof of Theorem 4 using the above multi-exponent RSA assumption with the modified exponent distribution. The security bounds are identical to the ones proven in [16] for standard non-public metadata RSA signatures.

Lemma 3. *Suppose that $(H_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ correspond to full domain hash (FDH) message encoding. Assuming the $(\epsilon, t, q_{\text{MD}})$ -multi-exponent RSA assumption (Definition 8) with exponents from $\mathcal{D}'_{N,e}$ and the random oracle model, then RSA_{MD} is $(\epsilon_F, t_F, \ell_F)$ -strong one-more unforgeable (Definition 2) where*

- $t_F = t - O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$
- $\epsilon_F = q_{\mathcal{M}} \cdot \epsilon$

and the adversary \mathcal{A} runs in time at most t_F and makes at most ℓ_F , $q_{\mathcal{M}}$ and q_{MD} queries to the signing oracle and hash functions $H_{\mathcal{M}}$ and H_{MD} respectively.

Proof. The proof follows identically as the proof of Lemma 2. The only difference required is the observation that the public exponents of the modified algorithm are generated identically to $\mathcal{D}'_{N,e}$ defined in Appendix C.2. \square

Theorem 10. *Suppose that $(H_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ correspond to full domain hash (FDH) message encoding. Assuming the (ϵ, t) -RSA assumption (Definition 6) with fixed public exponent e and the random oracle model, then RSA_{MD} is $(\epsilon_F, t_F, \ell_F)$ -strong one-more unforgeable (Definition 2) where*

- $t_F = t - O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$
- $\epsilon_F = q_{\mathcal{M}} \cdot \epsilon$

and the adversary \mathcal{A} runs in time at most t_F and makes at most ℓ_F , $q_{\mathcal{M}}$ and q_{MD} queries to the signing oracle and hash functions $H_{\mathcal{M}}$ and H_{MD} respectively.

Proof. First, we apply Lemma 3 to reduce security to the $(\epsilon_F/q_{\mathcal{M}}, t_F + O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F), q_{\text{MD}})$ -multi-exponent RSA assumption with exponents $\mathcal{D}'_{N,e}$ for some public exponent e . By applying Theorem 9, we obtain that this is equivalent to the $(\epsilon_F/q_{\mathcal{M}}, t_F + O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F))$ -RSA assumption with fixed exponent e . \square

D Non-Applicability of Multi-Exponent Attacks for RSA Signatures With Public Metadata

In this section, we explain in detail why the multi-exponent attacks in Section 7 that leverage certain algebraic properties that appear in exponents are not applicable for our standard (non-blind) RSA signatures with

public metadata. At a high level, the reasoning lies in one of the core difference between standard and blind RSA signatures. In standard (non-blind) signature schemes, the signer receives the input message M along with the public metadata D in plaintext. In contrast, for blind RSA signatures, the signer receives an element X in the RSA group that is a blinded version of the input message M (the public metadata D is still received in plaintext). This critical difference is leveraged by the multi-exponents attacks in Section 7 as we will show.

Taking a closer look at the attack, it requires find a sequence D, D_1, \dots, D_z of public metadata with the following properties. Let $e = H_{\mathbf{MD}}(D) = p_1 \cdots p_z$ be the product of (not necessarily distinct) primes such that $p_i \mid e_i$ where $e_i = H_{\mathbf{MD}}(D_i)$. Next, we can consider the first iteration of the loop of the attack. In particular, we set $x_0 = H_{\mathcal{M}}(M \parallel D)$ for some input message M . In the next step, we need to obtain a valid signature of $x_0^{e_1/p_1}$ for the exponent $e_1 = H_{\mathbf{MD}}(D_1)$ associated with public metadata D_1 . The output will essentially be x_1 . The step is repeated to obtain x_i that is output of the blind signing protocol for element $x_{i-1}^{e_i/p_i}$ for all $i \in [z]$. The final output is x_z that is a forgery for message M and public metadata D . For the blind RSA signatures with public metadata protocol, we note that the signer receives a blinded version of the input message M that is equivalent to some element in the underlying RSA group. Therefore, it is perfectly valid to send $x_{i-1}^{e_i/p_i}$ as the input element for the blind signing protocol.

We can now attempt to do the same attack for the standard (non-blind) RSA signatures with public metadata. The key point is that the adversary can no longer submit the element $x_{i-1}^{e_i/p_i}$ directly as input to the signing protocol. If the adversary wished to perform blind signing for this element, it must find some input message M_i such that $H_{\mathcal{M}}(M_i \parallel D_i) = x_{i-1}^{e_i/p_i}$. As we assume that $H_{\mathcal{M}}$ is a random oracle (that is, an ideal one-way function), it is computationally intractable for the adversary to find such a message M_i such that $H_{\mathcal{M}}(M_i \parallel D_i) = x_{i-1}^{e_i/p_i}$. Therefore, this attack does not apply directly to our standard RSA signatures with public metadata.

As a result, we note that several features required for the blind RSA signatures with public metadata are not needed for our non-blind protocol. First, we do not require the exponent checks that re-sample salts for $H_{\mathbf{MD}}$ when the set of valid exponents satisfy the algebraic properties required for the attack above. Secondly, our underlying RSA assumptions can simply use random odd exponents (instead of only focusing on exponent sets without the algebraic property). Therefore, it is not surprising that we prove our standard (non-blind) protocol directly from the RSA assumption (Definition 6) while we require the algebraically-restricted exponent assumption (Definition 10) for our blind signatures.

E Potential Denial-of-Service Risks

The protocols in this paper require performing RSA operations with larger exponents than standard RSA. The usage of large exponents may incur additional computation for the signer in the protocols but note that a user cannot change the size of the exponent (which is fixed at setup) arbitrarily. In particular, one could consider denial-of-service (DoS) risks where users wish to pick bad exponents (i.e., metadata) that may incur additional computation by the signer. In general, these risks are mitigated by the fact that the signer only performs computation on valid metadata $D \in \mathbf{MD}$. If an attacker picks bad metadata $D \notin \mathbf{MD}$, the signer rejects and avoids additional computation. Therefore, an attacker is restricted to trying to pick the worse metadata (and corresponding exponent) that appears in the valid metadata set \mathbf{MD} . If \mathbf{MD} is small, this severely restricts the freedom of the attacker to find bad inputs to cause larger signer computation. Furthermore, if this is a concern, one can also re-sample the salt corresponding to $H_{\mathbf{MD}}$ and check that no exponents corresponding to valid metadata incurs large computation.

F Supplementary Material for Analyzing Multi-Exponent Attacks

F.1 Asymptotic Analysis of Exponent Check

Theorem 11. Consider security parameter λ and public metadata set \mathbf{MD} of size $|\mathbf{MD}| = \text{poly}(\lambda)$. Suppose that $H_{\mathbf{MD}}$ outputs γ -bit odd integers where $\gamma = \kappa - 2$ and $\gamma \geq 32$. Let E be the event defined in Theorem 7. Then, for any randomly chosen salt, $\Pr[E] \leq O(2^{-\sqrt{\kappa}})$.

Proof. We use the asymptotic upper bound from [43] stating that $\rho(u) = u^{-u+o(u)}$. We set $x = 2^{-\sqrt{\kappa}}$ in Theorem 7:

$$\begin{aligned} \Pr[E] &\leq \frac{|\mathbf{MD}|^2}{2^{\sqrt{\kappa}}} + 2|\mathbf{MD}| \cdot \rho(\gamma/\sqrt{\kappa}) \\ &\leq \frac{\text{poly}(\lambda)}{2^{\sqrt{\kappa}}} + \frac{\text{poly}(\lambda)}{(\gamma/\sqrt{\kappa})^{\gamma/\sqrt{\kappa}}} \end{aligned}$$

since $|\mathbf{MD}| = \text{poly}(\lambda)$. Next, we plug in $\gamma = \kappa - 2$ and use the fact that $\gamma \geq 32$ to get

$$\begin{aligned} \Pr[E] &\leq \frac{\text{poly}(\lambda)}{2^{\sqrt{\kappa}}} + \frac{\text{poly}(\lambda)}{(\sqrt{\kappa}/2)^{\sqrt{\kappa}/2}} \\ &\leq \frac{\text{poly}(\lambda)}{2^{\sqrt{\kappa}}} + \frac{\text{poly}(\lambda)}{2^{\sqrt{\kappa}/2 \cdot \log(\sqrt{\kappa}/2)}} \\ &\leq \frac{\text{poly}(\lambda)}{2^{\sqrt{\kappa}}} + \frac{\text{poly}(\lambda)}{2^{\sqrt{\kappa}}} \\ &= O\left(2^{-\sqrt{\kappa}}\right) \end{aligned}$$

to complete the proof. □

Roughly speaking, the above lemma states that if one wishes to have this probability be $2^{-\lambda}$, then the security parameter should be chosen as approximately $\kappa = O(\lambda^2)$. However, this guidance is worse than those in Figure 9 using Dickman's rho function directly.

F.2 Relationship between Restricted-Exponent and Algebraically-Restricted-Exponent Assumptions

We utilize the above result to formalize the relationship between the chosen-target, restricted-exponent assumption (Definition 9) and the chosen-target, algebraically-restricted-exponent assumption (Definition 10). Recall the only difference is that the exponent oracle performs additional checks in the latter game. We show that the two games are equivalent where the adversarial advantage difference is at most the probability upper bound that we proved in Theorem 7.

Theorem 12. If the $(\epsilon, t, \ell, \ell')$ -chosen-target, restricted-exponent RSA inversion assumption (Def. 9) is true with \mathcal{D}_N choosing uniformly random odd exponents from $[3, e_{\max}]$ where $e_{\max} \leq 2^{\kappa-2}$, then the $(\epsilon', t, \ell, \ell')$ -chosen-target, algebraically-restricted-exponent RSA inversion assumption (Def. 10) is true with the same \mathcal{D}_N where

$$\epsilon' = \epsilon + \frac{\ell'^2}{x} + (2 \cdot \ell' \cdot \rho(\kappa - 2, \log x))$$

for some integer $x \geq 2$ and ρ is Dickman's function.

Proof. Towards a contradiction, suppose there exists an adversary \mathcal{A}' running in time t' with ℓ and ℓ' queries to the decryption and exponent oracles respectively and advantage ϵ' for the chosen-target, algebraically-restricted-exponent game in Definition 10. We construct an adversary \mathcal{A} for the chosen-target, restricted-exponent game in Definition 9. Essentially, \mathcal{A} executes \mathcal{A}' identically. We note that the advantage of \mathcal{A} is

identical to \mathcal{A}' whenever the exponents output by the exponent oracle satisfy the checks. As the output of \mathcal{D}_N is identical to H_{MD} , we can immediately apply Theorem 7 to obtain that the advantage of \mathcal{A} satisfies

$$\epsilon \leq \epsilon' - \frac{\ell'^2}{x} - (2 \cdot \ell' \cdot \rho(\kappa - 2, \log x))$$

where we replace $|\text{MD}|$ with ℓ' to limit the number of exponents exposed to the adversary. Note, we are free to pick any choice of integer $x \geq 2$ to minimize the right hand side of the equation. \square

We can now use this relation to prove unforgeability of our blind RSA signatures with public metadata protocol directly from the chosen-target, restricted-exponent RSA inversion assumption.

Theorem 13. *Suppose that $(H_{\mathcal{M}}, \text{Verify}_{\mathcal{M}})$ correspond to full domain hash (FDH) message encoding. Assuming the $(\epsilon, t, \ell, |\text{MD}|)$ -chosen-target, restricted-exponent RSA inversion assumption (Definition 9) with exponents in $[3, e_{\text{max}}]$ such that $e_{\text{max}} = 2^{\kappa-2}$ and the random oracle model, then $\text{BlindRSA}_{\text{MD}}$ is $(\epsilon_F, t_F, \ell_F)$ -strong one-more unforgeable (Definition 2) where*

- $\epsilon_F = \epsilon + 2^{-\kappa+2 \log q_{\mathcal{M}}} + \frac{|\text{MD}|^2}{x} + (2 \cdot |\text{MD}| \cdot \rho(\kappa - 2, \log x))$ for any integer $x \geq 2$
- $t_F = t - O(q_{\mathcal{M}} + q_{\text{MD}} + \ell_F)$
- $\ell_F = \ell$

and the adversary \mathcal{A} runs in time at most t_F and makes at most $\ell_F, q_{\mathcal{M}}$ and q_{MD} queries to the signing oracle and hash functions $H_{\mathcal{M}}$ and H_{MD} respectively.

Proof. First, we reduce the unforgeability game to the chosen-target, algebraically-restricted-exponent game in Theorem 5. Afterwards, we reduce to the final chosen-target, restricted-exponent game using Theorem 13. \square

We note that one could also plug Theorem 11 into the above theorem to obtain the bound on adversarial advantage as

$$\epsilon_F = \epsilon + 2^{-\kappa+2 \log q_{\mathcal{M}}} + O\left(2^{-\sqrt{\kappa}}\right).$$

G Survey of Cryptography Libraries for Production Usage

In this section, we survey the availability of various cryptographic operations across different libraries that are production ready (including those that are implemented to be specifically resistant to side-channel attacks). In particular, we will avoid any experimental and/or research libraries with explicit disclaimers that their implementations of cryptographic operations is not ready for production usage.

For our analysis, we will focus on the availability of the necessary cryptographic operations for our protocol (presented in this paper) as well as those based on pairings [57]. With respect to our protocol, we will focus on RSA operations with respect to large exponents. For pairings, we will simply analyze whether the cryptographic library exposes the necessary pairing-friendly curves for implementing [57]. As a caveat, we note our protocol also requires the additional cryptographic operations of generation safe primes during key generation that may not be exposed by some cryptographic libraries. We ignore safe prime generation because it is done exclusively by the signer in an offline manner (thus, the efficiency of safe prime generation is less important and the operation is less subject to side-channel attacks). Furthermore, it is not complex to implement safe prime generation using standard prime generation along with standard primality tests (both of which are available in every cryptographic library that we surveyed).

To choose production ready libraries, we consider several libraries that have active FIPS 140-2 certifications with current reviews for FIPS 140-3 validations (see [28] for comprehensive list). We caveat that not all cryptography libraries that are suitable for production usage have FIPS certifications. For example, we will include the ACML library [60] that supports pairings that was built for production usage without

Library	RSA with Large Exponents	Pairing-Friendly Curves
BoringSSL [2]	✓	×
OpenSSL [46]	✓	×
wolfCrypt [63]	✓	×
BouncyCastle [59]	×	×
NSS [45]	✓	×
ACML [60]*	✓	✓

Figure 15: Comparison of the availability of RSA with large exponent operations and pairing-friendly curves in various cryptographic libraries suitable for production usage. Amongst this group, we note ACML is the only library without FIPS certifications.

FIPS certifications. Nevertheless, we believe this is a reasonable methodology for determining libraries whose cryptographic implementations are aiming to be usable in production systems. For our analysis, we will consider the ACML [60], BoringSSL [2], OpenSSL [46], wolfCrypt [63], BouncyCastle [59] and NSS [45] libraries. We present our findings in Figure 15 about the availability of RSA operations with large exponents and pairing-friendly curves.

RSA Operations with Large Exponents. We note that both BoringSSL and OpenSSL have the necessary public exposed functions to enable our protocol (our implementation [1] is built using BoringSSL). The required RSA operations are available for wolfCrypt that is also a derivative of OpenSSL (and we expect this to be true for most OpenSSL derivatives). Similarly, both the NSS and ACML libraries also offer the necessary RSA operations. We note that the necessary public functions are not available in BouncyCastle, but the required implementations exist (and would only need to be exposed to enable this implementation). This is not true for pairings where no implementation exists in BouncyCastle.

Pairing-Friendly Curves. In general, we can see that many of the cryptographic libraries ready for production usage do not support pairing-friendly curves (except ACML [60] that does not have FIPS certifications). This highlights the currently limited availability of pairing-based cryptography for production systems. We do caveat that pairings are available in many experimental and/or research-oriented library including Arkworks [13], CIRCL [25], NCC Group’s implementation of BLS12-381 [48] and RELIC [52]. However, all of these have explicit disclaimers that the implementations should not be used in production. Additionally, we note there has been recent interest in pairing-based cryptography in IRTF [18] that may lead to wider availability of pairings in the future.