

Collusion Resistant Trace-and-Revoke for Arbitrary Identities from Standard Assumptions

Sam Kim
Stanford University
skim13@cs.stanford.edu

David J. Wu
University of Virginia
dwu4@virginia.edu

Abstract

A traitor tracing scheme is a multi-user public-key encryption scheme where each user in the system holds a decryption key that is associated with the user’s identity. Using the public key, a content distributor can encrypt a message to all of the users in the system. At the same time, if a malicious group of users combine their respective decryption keys to build a “pirate decoder,” there is an efficient tracing algorithm that the content distributor can use to identify at least one of the keys used to construct the decoder. A *trace-and-revoke* scheme is an extension of a standard traitor tracing scheme where there is an additional key-revocation mechanism that the content distributor can use to disable the decryption capabilities of compromised keys. Namely, during encryption, the content distributor can encrypt a message with respect to a list of revoked users such that only non-revoked users can decrypt the resulting ciphertext.

Trace-and-revoke schemes are challenging to construct. Existing constructions from standard assumptions can only tolerate bounded collusions (i.e., there is an *a priori* bound on the number of keys an adversary obtains), have system parameters that scale *exponentially* in the bit-length of the identities, or satisfy weaker notions of traceability that are vulnerable to certain types of “pirate evolution” attacks. In this work, we provide the first construction of a trace-and-revoke scheme that is fully collusion resistant and capable of supporting arbitrary identities (i.e., the identities can be drawn from an exponential-size space). Our scheme supports public encryption and secret tracing, and can be based on the sub-exponential hardness of the LWE problem (with a super-polynomial modulus-to-noise ratio). The ciphertext size in our construction scales logarithmically in the size of the identity space and linearly in the size of the revocation list. Our scheme leverages techniques from both combinatorial and algebraic constructions for traitor tracing.

1 Introduction

Traitor tracing schemes [CFN94] provide content distributors a way to identify malicious receivers and pirates. Specifically, a traitor tracing scheme is a public-key encryption scheme that is defined over a set of global public parameters \mathbf{pp} and many secret decryption keys $\{\mathbf{sk}_{\text{id}}\}$. Each of the decryption keys \mathbf{sk}_{id} is associated with an identifier id (e.g., a user’s name or profile picture). Anyone is able to encrypt a message using the public parameters \mathbf{pp} and any user who holds a valid decryption key \mathbf{sk}_{id} can decrypt the resulting ciphertext. The main security property is *traceability*, which says that if a coalition of users combine their respective decryption keys to create a new decryption algorithm (i.e., a “pirate decoder”), there is an efficient tracing algorithm that, given (black-box) access to the decoder, will successfully identify at least one of the secret keys that was used to

construct the pirate decoder. As such, traitor tracing schemes provide an effective way for content distributors to combat piracy.

In practice, simply identifying the keys that went into a pirate decoder is not enough; we also require a way for the content distributor to disable the decryption capabilities of a compromised key. Traitor tracing schemes that support efficient key-revocation mechanisms are called *trace-and-revoke* schemes [NP00]. In a trace-and-revoke scheme, the encryption algorithm additionally takes in a list of revoked users \mathcal{L} . A ciphertext that is generated with respect to a revocation list \mathcal{L} can only be decrypted by keys for identities $\text{id} \notin \mathcal{L}$. Furthermore, the revocation mechanism should remain compatible with tracing: namely, if an adversary builds a pirate decoder that can still decrypt ciphertexts encrypted with respect to a revocation list \mathcal{L} , the tracing algorithm should successfully identify at least one of the non-revoked decryption keys (i.e., some $\text{id} \notin \mathcal{L}$) that went into the construction of the pirate decoder. We give the formal definition in Section 4.

Properties of trace-and-revoke schemes. There are a number of possible properties that a trace-and-revoke scheme could provide. We enumerate several important ones below:

- **Collusion resistance:** A trace-and-revoke scheme is t -collusion resistant if tracing works as long as the pirate obtains fewer than t decryption keys, and the scheme parameters are allowed to depend on t . When t can be an arbitrary polynomial, the scheme is *fully collusion resistant*.
- ***A priori* unbounded revocation:** Some trace-and-revoke schemes support bounded revocation where at setup time, there is an *a priori* bound r on the maximum number of revoked users the scheme supports. A scheme supports *a priori unbounded revocation* if the number of revoked users can be an arbitrary polynomial. We note here that while we can require an even stronger property that supports revoking a super-polynomial number of users, the scheme we develop in this work does not support this stronger property (except in certain restricted settings; see Section 1.1).
- **Black box tracing:** A trace-and-revoke scheme supports black box tracing if the tracing algorithm only requires oracle access to the pirate decoder. This means we do not need to impose any restrictions on the structure of the adversary’s decoder. Tracing must work on *any* decoder that is able to decrypt (or even better, *distinguish*) ciphertexts.
- **Identity-based:** A trace-and-revoke scheme is “identity-based” or supports *arbitrary* identities if the set of possible identities \mathcal{ID} the scheme supports can be exponential in size [NWZ16]. In most trace-and-revoke schemes, the set of possible identities is assumed to have polynomial size (i.e., identities are represented by an element of the set $[N] = \{1, \dots, N\}$). This means that there is an a priori bound on the maximum number of users supported by the system, and moreover, in practical scenarios, the tracing authority needs to separately maintain a database mapping from a numeric index $\text{id} \in [N]$ to a user’s actual identifier (which may not fit into a string of length $\log N$). In addition, as noted in [NWZ16], an added benefit of trace-and-revoke schemes that support arbitrary identities is *anonymity*: namely, a user can obtain a decryption key for their identity without needing to reveal their identity to the key issuer.

Our results. In this work, we focus on constructing trace-and-revoke schemes that provide each of the above guarantees. Namely, we seek schemes that are flexible (e.g., can support arbitrary identities of polynomial length and an arbitrary polynomial number of revocations) while providing strong security (i.e., full collusion resistance and security against arbitrary adversarial strategies). We achieve these properties assuming sub-exponential hardness of the learning with errors (LWE) assumption [Reg05]. Specifically, we show the following:

Theorem 1.1 (informal). *Let λ be a security parameter and $\mathcal{ID} = \{0, 1\}^n$ be the set of possible identities. Assuming sub-exponential hardness of LWE, there exists a fully collusion resistant trace-and-revoke scheme where the secret key for an identity $\text{id} \in \{0, 1\}^n$ has size $n \cdot \text{poly}(\lambda, \log n)$ and a ciphertext encrypting a message m with respect to a revocation list $\mathcal{L} \subseteq \{0, 1\}^n$ has size $|m| + |\mathcal{L}| \cdot \text{poly}(\lambda, \log n)$. Encryption in our scheme is a public operation while tracing requires knowledge of a secret key.*

Previous trace-and-revoke constructions were either not collusion resistant [NWZ16, ABP⁺17], could only support a polynomial-size identity space [BW06, GKSW10, GQWW19], achieved weaker models of tracing [NNL01, DF02], or relied on strong assumptions such as indistinguishability obfuscation [NWZ16] or (positional) witness encryption [GVW19]. We refer to Section 1.2 for a more detailed comparison of our construction with existing ones.

Open questions. Before giving an overview of our construction, we highlight several interesting directions to further improve upon our trace-and-revoke scheme:

- *Public tracing:* Our tracing algorithm requires a secret key. It is an interesting open problem to obtain fully collusion resistant trace-and-revoke for arbitrary identities with public tracing from standard assumptions. In fact, even obtaining a collusion resistant traitor tracing scheme with succinct keys and public tracing from standard assumptions is currently open.
- *Succinct broadcast:* The length of the ciphertexts in our construction scales *linearly* in the size of the revocation list, and as such, our scheme only supports revocation for a polynomial number of users. It is an open question is to develop an scheme that supports arbitrary identities and where the ciphertext size scales *sublinearly* in the number of revoked users (and more generally, where the ciphertext size scales with the *description length* of the revocation list rather than its size). Schemes with these properties are often called “broadcast, trace, and revoke” schemes [BW06] as they combine both the succinctness of a “broadcast encryption” [FN93] with the tracing capability of a traitor tracing scheme. Existing broadcast, trace, and revoke constructions [BW06, GKSW10, GQWW19] from standard assumptions can only handle a polynomial number of users. We provide a more thorough comparison in Section 1.2.
- *Polynomial hardness:* Security of our tracing construction relies on the *sub-exponential* hardness of LWE. Our reliance on sub-exponential hardness assumptions is due to our use of complexity leveraging [BB04] to instantiate *adaptively-secure* variants of the underlying cryptographic primitives we require in our construction. An important open problem is to base security on polynomial hardness. The work of Goyal et al. [GKW19] show how to obtain traitor tracing for an exponential-size identity space from a polynomial hardness assumption, but their scheme does not support revocation.

1.1 Construction Overview

In this section, we provide a high-level overview of our construction. Our approach combines an identity-based traitor tracing scheme based on the techniques developed in [NWZ16, GKW18] with the combinatorial revocation scheme from [NNL01]. We describe each of these components below.

Traitor tracing from private linear broadcast. Boneh et al. [BSW06] showed how to construct a collusion resistant traitor tracing scheme from a private linear broadcast encryption (PLBE) scheme. A PLBE scheme is an encryption scheme where decryption keys are associated with an index $i \in [N]$, and ciphertexts are associated with a secret index $j \in [N]$ and a message m . The correctness property guarantees that a decryption key sk_i for index i can decrypt all ciphertexts encrypted to indices j where $i \leq j$. There are two ways to generate a ciphertext. The *public* encryption algorithm allows anyone to encrypt to the index N , which can be decrypted by secret keys sk_i for all $i \in [N]$. The *secret* encryption algorithm allows the tracing authority who holds a tracing key to encrypt to indices $j \leq N$. The “index-hiding” requirement guarantees that an adversary who does not have a key for index j cannot distinguish an encryption to index j from an encryption to index $j + 1$. Finally, the “message-hiding” requirement says that ciphertexts encrypted to index 0 are semantically secure (given any subset of decryption keys for indices $1 \leq j \leq N$). These properties form the basis of the tracing algorithm described in [BSW06]. Boneh et al. showed how to construct PLBE from pairing-based assumptions where the ciphertexts have size $O(\sqrt{N})$. Hence their scheme only supports a polynomial-size identity space.

Recently, Goyal et al. [GKW18] gave a new construction of a PLBE scheme from the LWE assumption by combining a new cryptographic notion called mixed functional encryption (mixed FE) with an attribute-based encryption (ABE) scheme [SW05, GPSW06]. Their construction has the appealing property that the size of all of the system parameters (e.g., the public parameters, decryption keys, and ciphertexts) scale with $\text{poly}(\lambda, \log N)$. Thus, the construction of Goyal et al. [GKW18] can in principle support arbitrary set of identities. However, the tracing algorithm in the PLBE framework runs in time that scales *linearly* with the size of the identity space. As a result, the [GKW18] construction does not support tracing over an exponential space of identities.

Identity-based traitor-tracing from functional encryption. In [NWZ16], Nishimaki et al. introduced a more general tracing algorithm for PLBE that supports an exponential identity space (by abstracting the tracing problem as an “oracle jump-finding” problem). Their construction relies on a PLBE scheme that satisfies a more general notion of index-hiding security. Namely a ciphertext encrypted to index j_1 should be indistinguishable from a ciphertext encrypted to index j_2 as long as the adversary does not have any keys in the interval (j_1, j_2) .¹ A limitation of this construction is that the ciphertexts scale linearly in the bit-length of the identities. Nishimaki et al. then show how to construct a traitor tracing scheme with *short* ciphertexts (i.e., one where the ciphertext size scales with $\text{poly}(\log \log N)$) from a private broadcast encryption scheme that support slightly more general broadcast sets. Finally, they note that private broadcast is just a special case of general-purpose functional encryption which can be instantiated using indistinguishability obfuscation [GGH⁺13], or, in the bounded-collusion setting, from LWE [GKP⁺13] or even just

¹This property follows from the usual index-hiding security game by a standard hybrid argument when the indices are drawn from a polynomial-size space, but not when the indices are drawn from an exponentially-large one.

public-key encryption [SS10, GVW12].

A more general view of [GKW18]. In this work, we take a more general view of the PLBE construction in [GKW18] and show that the construction in fact gives a *secret-key predicate encryption scheme with a broadcast functionality*. In turn, PLBE can be viewed as a specific instantiation of the predicate encryption scheme for the particular class of threshold predicates. This view will enable our generalization to identity-based traitor tracing with short ciphertexts (by following the approach of [NWZ16]) as well as enable an efficient mechanism for key revocation. Note that the “broadcast functionality” considered here refers to a method to *publicly* encrypt a message that can be decrypted by *all* secret keys in the system (i.e., broadcasting a message to all users in the system). We are not requiring the ability to succinctly broadcast messages to subsets of users (as in the setting of broadcast encryption [FN93]).

Specifically, in a secret-key (ciphertext-policy) predicate encryption scheme, ciphertexts are associated with a predicate f and a message m , while decryption keys are associated with an attribute x . Decrypting a ciphertext $\text{ct}_{f,m}$ associated with a predicate f and a message m with a function key for an attribute x yields m if $f(x) = 1$ and \perp otherwise. Moreover, the policy f associated with a ciphertext is hidden irrespective of whether decryption succeeds or not—this property is the analog of the “strong” attribute-hiding property considered in the study of key-policy predicate encryption [BW07, KSW08, SBC⁺07]. Finally, while the predicate encryption scheme is secret-key, there exists a *public* encryption algorithm that allows anyone to encrypt a message with respect to the “always-accept” policy (i.e., $f(x) = 1$ for all inputs x). In Section 3.1, we show how to combine mixed FE (for general circuits) and attribute-based encryption (for general circuits) to obtain a secret-key ciphertext-policy predicate encryption scheme with broadcast. This construction is a direct analog of the [GKW18] construction of PLBE from the same set of underlying primitives. Next, we note that this type of predicate encryption directly implies a fully collusion resistant traitor tracing scheme with short ciphertexts via [NWZ16]. The one difference, however, is that since the predicate encryption scheme is in the secret-key setting, only the tracing authority who holds the master secret key is able to run the tracing algorithm. Thus in contrast to [NWZ16], our scheme only supports *secret tracing*. We note that working in the secret-key setting introduces some new challenges in the security analysis of the [NWZ16] construction. These can be handled using similar techniques as those developed in [GKW18], and we discuss this in greater detail in Section 4.1.

Trace-and-revoke via revocable predicate encryption. Thus far, we have shown how to combine ideas from [GKW18] and [NWZ16] to obtain a collusion resistant traitor tracing scheme for arbitrary identities. The next step is to develop a mechanism for key revocation. Previously, Nishimaki et al. showed how to use a revocable functional encryption scheme to construct a trace-and-revoke scheme. In this work, we show that a revocable variant of our secret-key predicate encryption scheme with broadcast also suffices for this general transformation. Namely, in a revocable predicate encryption scheme, each decryption key is additionally tagged with an identity id , and at encryption time (both secret and public), the encrypter provides both the decryption policy f and the revocation list \mathcal{L} . The resulting ciphertext can then be decrypted by all keys $\text{sk}_{\text{id},x}$ associated with an identity id and an attribute x such that $f(x) = 1$ and $\text{id} \notin \mathcal{L}$.

A natural approach to support revocation is to include the revocation list \mathcal{L} as part of the ciphertext policy in the predicate encryption scheme. We would then embed the identity id as part of the decryption key, and the final decryption policy would first check that $\text{id} \notin \mathcal{L}$ and then check

that $f(x) = 1$. While this basic approach seems straightforward, it unfortunately does not apply in our setting. As noted above, the predicate encryption scheme we construct is a *secret-key* scheme, and the only public operation it supports is the broadcast functionality.² Obtaining a public-key analog of collusion resistant, strong attribute-hiding predicate encryption seems quite challenging (and in fact, implies public-key functional encryption). But as we note in Remark 3.3, even in the bounded-collusion setting (where we can construct public-key predicate encryption from standard assumptions), this basic approach seems to run into a barrier, and any such instantiation from standard assumptions would likely have to assume a bound on the maximum number of revoked users. In this work, we seek solutions from standard assumptions that are collusion resistant and support unbounded revocation.

Revocable predicate encryption via subset cover set systems. As we described above, constructing a collusion resistant trace-and-revoke scheme for arbitrary identities reduces to constructing a secret-key revocable predicate encryption scheme with a broadcast functionality. To build the necessary revocable predicate encryption scheme, we leverage ideas from combinatorial constructions of traitor tracing. We note that while we rely on combinatorial ideas in our construction, we do not provide a generic transformation of any predicate encryption scheme into a revocable analog. Rather, our construction relies on a careful integration of the algebraic approach from [GKW18] with the combinatorial approach from [NNL01].

The core combinatorial ingredient that we use for our construction is a subset-cover set system, a notion that has featured in several traitor tracing constructions [NNL01, DF02, HS02]. Let $[N]$ be the identity space. A subset-cover set system for $[N]$ is a set of indices $[K]$ with the following two properties. Each identity $\text{id} \in [N]$ is associated with a small number of indices $\mathcal{I}_{\text{id}} \subseteq [K]$. Moreover, given a revocation list $\mathcal{L} \subseteq [N]$, there is an efficient algorithm to compute a “covering” set of indices $\mathcal{J}_{\mathcal{L}} \subseteq [K]$ with the property that $\text{id} \in \mathcal{L}$ if and only if $\mathcal{I}_{\text{id}} \cap \mathcal{J}_{\mathcal{L}} = \emptyset$. If we instantiate using the subset-cover set system from [NNL01], then $K = O(N)$, $|\mathcal{I}_{\text{id}}| = O(\log N)$, and $|\mathcal{J}_{\mathcal{L}}| = O(|\mathcal{L}| \log(N/|\mathcal{L}|))$.

Given a subset-cover set system, a first attempt to construct a revocable predicate encryption scheme is as follows. We associate a set of public parameters pp_i and master secret key msk_i with each index $i \in [K]$. A key for an identity $\text{id} \in [N]$ and an attribute x would consist of predicate encryption keys $\text{sk}_{\text{id},x} \leftarrow \text{KeyGen}(\text{msk}_i, x)$ for all the predicate encryption schemes $i \in \mathcal{I}_{\text{id}}$ associated with id . Finally, an encryption of a message m with respect to the revocation list $\mathcal{L} \subseteq [N]$ would consist of a collection of ciphertexts $\{\text{ct}_i\}_{i \in \mathcal{J}_{\mathcal{L}}}$ where each ct_i is an encryption of m with respect to pp_i for $i \in \mathcal{J}_{\mathcal{L}}$. By the property described above, if $\text{id} \notin \mathcal{L}$, then $\mathcal{I}_{\text{id}} \cap \mathcal{J}_{\mathcal{L}} \neq \emptyset$. This means that all non-revoked users $\text{id} \notin \mathcal{L}$ will possess a key $\text{sk}_{i,x}$ for some $i \in \mathcal{J}_{\mathcal{L}}$, and therefore, will be able to decrypt (provided that $f(x) = 1$). For a revoked user, it will be the case that $i \notin \mathcal{J}_{\mathcal{L}}$ for all $i \in \mathcal{I}_{\text{id}}$, and they will be unable to decrypt. The problem though is that the size of the public parameters now scale *linearly* with K (which is as large as N). As such, this scheme only supports a polynomial number of identities. Thus, we need a different approach. We describe two candidate ideas below:

- If the underlying predicate encryption scheme has the property where the master secret key msk can be sampled *after* the public parameters pp , then in principle, the construction

²The recent work of Goyal et al. [GQWW19] introduces a notion of *broadcast mixed FE* that supports a *succinct* public broadcast to a restricted set of identities (of polynomial size). The notion we develop in this work supports an exponential-sized identity space, but in a *non-succinct* manner (i.e., the ciphertext size scales linearly with the size of the revocation list).

above would suffice. Namely, we would use a single set of public parameters for all of the predicate encryption schemes, and derive the master secret key msk_i for each $i \in [K]$ from a pseudorandom function (PRF). Unfortunately, such a predicate encryption scheme cannot be secure since the adversary can always generate for itself a master secret key and use it to decrypt.

- If the scheme supports a *public* encryption algorithm, then we can support revocation by including the index $i \in [K]$ as part of the policy associated with the ciphertext as well as the attribute in the decryption key. Then, the decryption policy would additionally check that the index associated with the key matched the index associated with the ciphertext. Essentially, we ensure that a decryption key for i can only be used to decrypt ciphertexts encrypted to index i . However, this revocation approach also does not seem to apply in our setting because our predicate encryption scheme is in the secret-key setting, and it is not clear how to generalize to a public-key encryption algorithm that can support more general policies (while retaining the same security properties).³

While neither of these approaches directly apply in our setting, we can combine *both* ideas in our construction to obtain a revocable predicate encryption scheme. As noted above, our basic secret-key predicate encryption scheme with broadcast combines a mixed FE scheme with an ABE scheme. Without getting into too many details, the construction has the following properties. Each ciphertext in the scheme consists of a mixed FE ciphertext and an ABE ciphertext, and analogously, each decryption key consists of a mixed FE decryption key and an ABE decryption key. The mixed FE scheme is a secret-key scheme that supports a broadcast mechanism while the ABE scheme is a standard public-key scheme. The key observation is that if *both* the underlying mixed FE scheme and the ABE scheme support revocation, then the resulting predicate encryption scheme also supports revocation. For our construction it is critical that both schemes support revocation as we rely on the mixed FE scheme to hide the ciphertext policy and the ABE scheme to hide the message. If only one of the underlying schemes supports revocation, then one or both of these security properties become incompatible with revocation. We now describe how we implement revocation for the underlying mixed FE and ABE schemes:

- The mixed FE scheme is a secret-key scheme that supports public broadcast. Unlike standard predicate encryption, the security properties of mixed FE can be satisfied by schemes where the master secret key is sampled *after* the public parameters, and this property is satisfied by existing constructions [GKW18, CVW⁺18a]. This means that we can associate a different mixed FE scheme with each index $i \in [K]$ where the master secret key associated with each instance is derived from a PRF. All of the mixed FE schemes share a common set of public parameters. We can now use the first revocation idea described above to implement revocation for the mixed FE scheme.
- Next, the ABE scheme is a public-key encryption scheme, and thus, we can use the second type of revocation described above. Namely, we require a single set of ABE parameters and

³While the notion of attribute-based mixed FE from [CVW⁺18a] seems like it would also provide this functionality, this revocation approach only preserves the message hiding property and not the mixed FE attribute hiding property of the underlying attribute-based mixed FE scheme. For our trace-and-revoke scheme, we require both message hiding and attribute hiding (which we refer to as “function hiding”). Obtaining the latter property seemingly requires a way to revoke mixed FE decryption keys.

simply include the index $i \in [K]$ in both the decryption key and the ciphertext to identify which index is being targeted.

By combining these two approaches for revocation, we show in Section 3.1 how to construct a secret-key revocable predicate encryption with broadcast scheme from the sub-exponential hardness of LWE. Notably, our final revocation mechanism relies critically on both the combinatoric properties of the subset-cover set system as well as the specific algebraic nature of the predicate encryption construction. Together, this yields the first collusion resistant trace-and-revoke scheme for arbitrary identities from the same underlying assumptions (Theorem 1.1).

A simple extension: more general revocation policies. While the basic scheme we described above supports revoking any polynomial number of identities, it naturally extends to support any revocation policy supported by the underlying subset-cover set system. Specifically, if we use the prefix-based subset-cover set system by Naor et al. [NNL01], our scheme supports revoking any number of identities that can be specified by a polynomial number of *prefix-based patterns*. For instance, we can revoke all users whose identity starts with a fixed prefix—which may consist of an *exponential* number of identities. In a concrete application, if the first few bits of a user’s identity specifies a region, then we can use prefix-based policies to efficiently revoke all of the users from one or more regions. We provide more discussion in Remark 3.10.

1.2 Related Work

In this section, we survey some of the related work on traitor tracing and trace-and-revoke schemes and compare our results to existing notions.

Traitor tracing and trace-and-revoke. Numerous works have studied constructions of both traitor tracing and trace-and-revoke schemes from a wide range of assumptions and settings. Very broadly, most existing constructions can be categorized into two main categories: *combinatorial* approaches [CFN94, NP98, SSW01, CFNP00, NNL01, HS02, DF02, SSW01, BN08] and *algebraic* approaches [KD98, NP00, BSW06, BW06, GKSW10, LPSS14, KT15, NWZ16, ABP⁺17, GKW18, CVW⁺18a, GVW19, GQWW19]. We refer to these works and the references therein for a survey of the field.

Many existing traitor-tracing and trace-and-revoke schemes (from standard assumptions) are only secure against bounded collusions [CFN94, KD98, NP00, SSW01, LPSS14, KT15, NWZ16, ABP⁺17]. Other schemes are fully collusion resistant, but can only handle a polynomial-size identity space [BSW06, BW06, GKSW10, GKW18, CVW⁺18a, GQWW19]. In this work, we focus on schemes that are fully collusion resistant and support arbitrary identity spaces. While there are schemes that are both collusion resistant and support a super-polynomial identity space [NWZ16, GVW19], these construction require strong assumptions such as indistinguishability obfuscation [BGI⁺12] or positional witness encryption and cannot currently be based on standard intractability assumptions.

Several of the aforementioned schemes from standard assumptions [BW06, GKSW10, GQWW19] additionally provide a *succinct* broadcast mechanism where anyone can encrypt a message to any subset of the users with a ciphertext whose size scales with $N^{1/2}$ [BW06, GKSW10] or with N^ϵ [GQWW19] for any constant $\epsilon > 0$, where N is the total number of users in the system. Such schemes are commonly referred to as “broadcast, trace, and revoke” schemes. Notably, the ciphertext size in these constructions is *independent* of the number of revoked users and only depends on the

total number of users. In our trace-and-revoke construction (Theorem 1.1), the ciphertext size scales *linearly* with the number of revoked users (which can be $\Omega(N)$ in the worst case). Thus, in the setting where we have a polynomial-size identity space and when the number of revoked users is a sufficiently-large fraction of the total number of users, existing broadcast, trace, and revoke constructions will have shorter ciphertexts. In the setting where there is an exponential identity space, the ciphertexts in these existing constructions are also exponential, and they do not provide a compelling solution.

Several works [NP98, CFNP00, BN08] consider a threshold notion of traitor tracing where the tracing algorithm is only guaranteed to work for decoders that succeed with probability at least $\delta = 1/\text{poly}(\lambda)$ (and the scheme parameters are allowed to depend on the parameter δ). In this work, we focus on schemes that work for any decoder that succeeds with non-negligible probability.

Some combinatorial constructions [NNL01, HS02, DF02] are fully collusion resistant, but they only satisfy a weaker notion of traceability where the tracing algorithm either succeeds in extracting a pirated key *or* identifies an encryption strategy that disables the pirate decoder (this latter strategy increases the ciphertext size). This weaker traceability notion has led to pirate evolution [KP07] and Pirate 2.0 attacks [BP09] on schemes satisfying this weaker security notion. In this work, we focus on the strong notion of traceability where the tracing algorithm always succeeds in extracting at least one pirate key from any functional decoder. This notion is not affected by the pirate evolution attacks.

Cryptographic watermarking. A closely-related notion to traitor tracing is cryptographic watermarking [BGI⁺12, CHN⁺16]. Very briefly, a cryptographic watermarking scheme allows an authority to embed arbitrary data into the secret key of a cryptographic function such that the marked program preserves the original functionality, and moreover, it is difficult to remove the watermark from the program without destroying its functionality. A collusion resistant watermarking scheme for a public-key encryption scheme would imply a collusion resistant traitor tracing scheme. Existing constructions [KW17, QWZ18, KW19] of watermarking from standard assumptions are not collusion resistant and they are also limited to watermarking PRFs, which are not sufficient for traitor tracing. The recent construction of watermarking for public-key primitives [GKM⁺19] does imply a traitor tracing scheme for general identities (with public tracing), but only provides bounded collusion resistance (in fact, in this setting, their construction precisely coincides with the bounded collusion resistant traitor tracing construction from [NWZ16]). Moreover, it is not clear that existing constructions of watermarking can be extended to support key revocation.

Concurrent work. In a recent and concurrent work, Goyal et al. [GKW19] also study the problem of identity-based traitor tracing for arbitrary identities (i.e., which they call “traitor tracing with embedded identities”). Their focus is on traitor tracing (without revocation) and achieving security based on *polynomial* hardness assumptions. In contrast, our focus is on supporting both tracing *and* revocation while still supporting arbitrary identities. Security of our construction, however, does rely on making a stronger sub-exponential hardness assumption.

2 Preliminaries

We begin by introducing some notation. We use λ (often implicitly) to denote the security parameter. We write $\text{poly}(\lambda)$ to denote a quantity that is bounded by a fixed polynomial in λ and $\text{negl}(\lambda)$ to

denote a function that is $o(1/\lambda^c)$ for all $c \in \mathbb{N}$. We say that an event occurs with overwhelming probability if its complement occurs with negligible probability. We say an algorithm is efficient if it runs in probabilistic polynomial time in the length of its input. For two families of distributions $\mathcal{D}_1 = \{\mathcal{D}_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ and $\mathcal{D}_2 = \{\mathcal{D}_{2,\lambda}\}_{\lambda \in \mathbb{N}}$, we write $\mathcal{D}_1 \stackrel{c}{\approx} \mathcal{D}_2$ if the two distributions are computationally indistinguishable (i.e., no efficient algorithm can distinguish \mathcal{D}_1 from \mathcal{D}_2 except with negligible probability).

For an integer $n \geq 1$, we write $[n]$ to denote the set of integers $\{1, \dots, n\}$. For integers $1 \leq m \leq n$, we write $[m, n]$ to denote the set of integers $\{m, m+1, \dots, n\}$, and $[m, n]_{\mathbb{R}}$ to denote the closed interval between m and n (inclusive) over the real numbers. For a distribution \mathcal{D} , we write $x \leftarrow \mathcal{D}$ to denote that x is drawn from \mathcal{D} . For a finite set S , we write $x \stackrel{R}{\leftarrow} S$ to denote that x is drawn uniformly at random from S .

Cryptographic primitives. We now recall the standard definition of pseudorandom functions and collision-resistant hash functions.

Definition 2.1 (Pseudorandom Function [GGM84]). A pseudorandom function (PRF) with key-space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in \mathbb{N}}$, domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$, and range $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ is an efficiently-computable function $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that for all efficient adversaries \mathcal{A} ,

$$\Pr[k \stackrel{R}{\leftarrow} \mathcal{K} : \mathcal{A}^{F(k,\cdot)}(1^\lambda) = 1] - \Pr[f \stackrel{R}{\leftarrow} \text{Funs}[\mathcal{X}, \mathcal{Y}] : \mathcal{A}^{f(\cdot)}(1^\lambda) = 1] = \text{negl}(\lambda).$$

Definition 2.2 (Keyed Collision-Resistant Hash Function). A keyed collision-resistant hash function with key-space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in \mathbb{N}}$, domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$, and range $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ is an efficiently-computable function $H: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that for all efficient adversaries \mathcal{A} and sampling $k \stackrel{R}{\leftarrow} \mathcal{K}$,

$$\Pr[(x_0, x_1) \leftarrow \mathcal{A}(1^\lambda, k) : x_0 \neq x_1 \text{ and } H(k, x_0) = H(k, x_1)] = \text{negl}(\lambda).$$

Subset-cover set systems. As discussed in Section 1.1, the subset-cover framework introduced by Naor et al. [NNL01] is the basis for many *combinatorial* trace-and-revoke schemes. We provide the formal definition below:

Definition 2.3 (Subset-Cover Set System [NNL01]). Let N be a positive integer. A subset-cover set system for $[N]$ is a set of indices $[K]$ where $K = \text{poly}(N)$ together with a pair of algorithms (Encode, ComputeCover) satisfying the following properties:

- **Encode**(x) $\rightarrow \mathcal{I}_x$: On input an element $x \in [N]$, the encoding algorithm outputs a set of indices $\mathcal{I}_x \subseteq [K]$.
- **ComputeCover**(\mathcal{L}) $\rightarrow \mathcal{J}_\mathcal{L}$: On input a revocation list $\mathcal{L} \subseteq [N]$, the cover-computation algorithm outputs a collection of indices $\mathcal{J}_\mathcal{L} \subseteq [K]$.

We require the following efficiency and security requirements for a subset-cover set system.

- **Efficiency**: Take any element $x \in [N]$ and any revocation list $\mathcal{L} \subseteq [N]$. Then, **Encode**(x) runs in time $\text{poly}(\log N)$ and **ComputeCover**(\mathcal{L}) runs in time $\text{poly}(|\mathcal{L}|, \log N)$.
- **Correctness**: Take any element $x \in [N]$ and revocation list $\mathcal{L} \subseteq [N]$, and let $\mathcal{I}_x \leftarrow \text{Encode}(x)$, $\mathcal{J}_\mathcal{L} \leftarrow \text{ComputeCover}(\mathcal{L})$. Then, $x \in \mathcal{L}$ if and only if $\mathcal{I}_x \cap \mathcal{J}_\mathcal{L} = \emptyset$.

In this work, we will use the “complete subtree” system from [NNL01, §3.1]. The details of this construction are not essential to our construction, so we omit them and just summarize the main properties below:

Fact 2.4 (Subset-Cover Set System [NNL01, §3.1]). Let N be a positive integer. Then there exists a subset-cover set system $[K]$ for $[N]$ where $K = 2N - 1$, and where the algorithms (Encode, ComputeCover) satisfy the following properties:

- For all elements $x \in [N]$, if $\mathcal{I}_x \leftarrow \text{Encode}(x)$, then $|\mathcal{I}_x| = \log N + 1$.
- For all revocation lists $\mathcal{L} \subseteq [N]$, if $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, then $|\mathcal{J}_{\mathcal{L}}| = O(|\mathcal{L}| \log(N/|\mathcal{L}|))$.

The generalized jump-finding problem. Next, we recall the generalized jump-finding problem introduced by Nishimaki et al. [NWZ16, §3.1] for constructing identity-based traitor tracing schemes with succinct ciphertexts. We note that [NWZ16] also introduced a simpler variant of the jump-finding problem that essentially abstracts out the algorithmic core of the traitor tracing construction from private linear broadcast. Here, we consider the generalized version because it enables shorter ciphertexts (where the ciphertext size scales logarithmically with the bit-length of the identities)

Definition 2.5 (Generalized Jump-Finding Problem [NWZ16, Definition 3.9]). For positive integers $N, r, q \in \mathbb{N}$ and $\delta, \varepsilon > 0$, the $(N, r, q, \delta, \varepsilon)$ generalized jump-finding problem is defined as follows. An adversary begins by choosing a set C of up to q tuples $(s, b_1, \dots, b_r) \in [N] \times \{0, 1\}^r$ where all of the s are distinct. Each tuple (s, b_1, \dots, b_r) describes a curve between grid points from the top to bottom of the grid $[1, r] \times [0, 2N]$, which oscillates about the column at position $2s - 1$, with $b = (b_1, \dots, b_r)$ specifying which side of the column the curve is on in each row. The curves divide the grid into $|C| + 1$ contiguous regions. For each pair $(i, x) \in [1, r] \times [0, 2N]$, the adversary chooses a probability $p_{i,x} \in [0, 1]_{\mathbb{R}}$ with the following properties:

- For any two pairs $(i, 2x), (j, 2x) \in [1, r] \times [0, 2N]$, it holds that $|p_{i,2x} - p_{j,2x}| < \delta$.
- Let $C_i = \{(s, b_1, \dots, b_r) \in C : 2s - b_i\}$ be the set of values $2s - b_i$ for tuples in C . For any two pairs $(i, x), (i, y) \in [1, r] \times [0, 2N]$ such that $(x, y) \cap C_i = \emptyset$, then $|p_{i,x} - p_{i,y}| < \delta$.
- For all $i, j \in [r]$, it holds that $p_{i,0} = p_{j,0}$ and $p_{i,2N} = p_{j,2N}$. Define $p_0 = p_{i,0}$ and $p_{2N} = p_{i,2N}$.
- Finally, $|p_{2N} - p_0| > \varepsilon$.

Next, define the oracle $Q: [1, r] \times [0, 2N] \rightarrow \{0, 1\}$ to be a *randomized* oracle that on input (i, x) outputs 1 with probability $p_{i,x}$. Repeated calls to Q on the same input (i, x) will yield a fresh and independently-sampled bit. The $(N, r, q, \delta, \varepsilon)$ generalized jump-finding problem is to output some element in C given oracle access to Q .

Theorem 2.6 (Generalized Jump-Finding Algorithm [NWZ16, Theorem 3.10]). *There is an efficient algorithm $\text{QTrace}^Q(\lambda, N, r, q, \delta, \varepsilon)$ that runs in time $t = \text{poly}(\lambda, \log N, r, q, 1/\delta)$ and makes at most t queries to Q that solves the $(N, r, q, \delta, \varepsilon)$ generalized jump-finding problem with probability $1 - \text{negl}(\lambda)$ whenever $\varepsilon \geq \delta(9 + 4(\lceil \log N \rceil - 1)q)$. Moreover, any element $(s, b_1, \dots, b_r) \in [N] \times \{0, 1\}^r$ output by QTrace^Q satisfies the following property (with overwhelming probability):*

- For all $i \in [r]$, $|P(i, 2s - b_i) - P(i, 2s - 1 - b_i)| \geq \delta$, where $P(i, x) := \Pr[Q(i, x) = 1]$.

Remark 2.7 (Cheating Oracles [NWZ16, Remark 3.8]). The algorithm QTrace^Q from Theorem 2.6 succeeds in solving the $(N, r, q, \delta, \varepsilon)$ generalized jump-finding problem even if the oracle Q does not satisfy all of the requirements in Definition 2.5. As long as the first two properties hold for all pairs (i, x) and (j, y) queried by QTrace^Q , the algorithm succeeds in outputting an element in C .

2.1 Functional Encryption

In this section, we recall the notions of attribute-based encryption (ABE) and mixed functional encryption (mixed FE) that we use in this work.

Mixed FE. A mixed FE scheme [GKW18] is a secret-key FE scheme (i.e., a secret key is needed to encrypt) where ciphertexts are associated with binary-valued functions $f: \mathcal{X} \rightarrow \{0, 1\}$ and decryption keys are associated with inputs $x \in \mathcal{X}$. When a secret key sk_x associated with an input x is used to decrypt a ciphertext encrypting a message f , the decryption algorithm outputs $f(x)$. The special property in a mixed FE scheme is that there additionally exists a *public-key* encryption algorithm that can be used to encrypt to the “always-accept” function (i.e., the function f where $f(x) = 1$ for all $x \in \mathcal{X}$). Moreover, ciphertexts encrypted using the public key are computationally indistinguishable from ciphertexts produced by using the secret key to encrypt the “always-accept” function. Finally, for our constructions, we require an additional property where the master public key and the master secret key for the mixed FE scheme can be generated *independently*. This means that we can have a family of mixed FE schemes sharing a common set of public parameters. As we discuss in Remark 2.10, all existing mixed FE schemes satisfy this requirement.

Definition 2.8 (Mixed Functional Encryption [GKW18]). A mixed functional encryption scheme Π_{MFE} with domain \mathcal{X} and function family $\mathcal{F} = \{f: \mathcal{X} \rightarrow \{0, 1\}\}$ is a tuple of algorithms $\Pi_{\text{MFE}} = (\text{PrmsGen}, \text{MSKGen}, \text{KeyGen}, \text{PKEnc}, \text{SKEnc}, \text{Dec})$ with the following properties:

- $\text{PrmsGen}(1^\lambda) \rightarrow \text{pp}$: On input the security parameter λ , the parameter generation algorithm outputs the public parameters pp .
- $\text{MSKGen}(\text{pp}) \rightarrow \text{msk}$: On input the public parameters pp , the master secret key generation algorithm outputs a master secret key msk .
- $\text{KeyGen}(\text{msk}, x) \rightarrow \text{sk}_x$: On input the master secret key msk and an input $x \in \mathcal{X}$, the key-generation algorithm outputs a secret key sk_x .
- $\text{PKEnc}(\text{pp}) \rightarrow \text{ct}$: On input the public parameters pp , the public encryption algorithm outputs a ciphertext ct .
- $\text{SKEnc}(\text{msk}, f) \rightarrow \text{ct}_f$: On input the master secret key msk and a function $f \in \mathcal{F}$, the secret encryption algorithm outputs a ciphertext ct_f .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow b$: On input a secret key sk and a ciphertext ct , the decryption algorithm outputs a bit $b \in \{0, 1\}$.

A mixed FE scheme should satisfy the following properties:

- **Correctness:** For all functions $f \in \mathcal{F}$ and all inputs $x \in \mathcal{X}$, and setting $\text{pp} \leftarrow \text{PrmsGen}(1^\lambda)$, $\text{msk} \leftarrow \text{MSKGen}(\text{pp})$, $\text{sk}_x \leftarrow \text{KeyGen}(\text{msk}, x)$, $\text{ct} \leftarrow \text{PKEnc}(\text{pp})$, $\text{ct}_f \leftarrow \text{SEnc}(\text{msk}, f)$, it follows that

$$\Pr[\text{Dec}(\text{sk}_x, \text{ct}) = 1] = 1 - \text{negl}(\lambda) \quad \text{and} \quad \Pr[\text{Dec}(\text{sk}_x, \text{ct}_f) = f(x)] = 1 - \text{negl}(\lambda).$$

- **Semantic security:** For a bit $b \in \{0, 1\}$, we define the security experiment $\text{ExptMFE}_{\text{SS}}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} . The challenger begins by sampling $\text{pp} \leftarrow \text{PrmsGen}(1^\lambda)$, $\text{msk} \leftarrow \text{MSKGen}(\text{pp})$, and gives pp to \mathcal{A} . The adversary is then given access to the following oracles:
 - **Key-generation oracle:** On input $x \in \mathcal{X}$, the challenger replies with $\text{sk}_x \leftarrow \text{KeyGen}(\text{msk}, x)$.
 - **Encryption oracle:** On input $f \in \mathcal{F}$, the challenger replies with $\text{ct}_f \leftarrow \text{SEnc}(\text{msk}, f)$.
 - **Challenge oracle:** On input two functions $f_0, f_1 \in \mathcal{F}$, the challenger replies with $\text{ct} \leftarrow \text{SEnc}(\text{msk}, f_b)$.

At the end of the game, the adversary outputs a bit $b' \in \{0, 1\}$, which is also the output of the experiment. An adversary \mathcal{A} is admissible for the mixed FE semantic security game if it makes one challenge query (f_0, f_1) , and for all inputs $x \in \mathcal{X}$ the adversary submits to the key-generation oracle, $f_0(x) = f_1(x)$. The mixed FE scheme satisfies (adaptive) semantic security if for all efficient and admissible adversaries \mathcal{A} ,

$$|\Pr[\text{ExptMFE}_{\text{SS}}[\lambda, \mathcal{A}, 0] = 1] - \Pr[\text{ExptMFE}_{\text{SS}}[\lambda, \mathcal{A}, 1] = 1]| = \text{negl}(\lambda).$$

- **Public/secret key indistinguishability:** For a bit $b \in \{0, 1\}$, we define the security experiment $\text{ExptMFE}_{\text{PK/SK}}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} . The challenger begins by sampling $\text{pp} \leftarrow \text{PrmsGen}(1^\lambda)$, $\text{msk} \leftarrow \text{MSKGen}(\text{pp})$, and gives pp to \mathcal{A} . The adversary is then given access to the following oracles:
 - **Key-generation oracle:** On input $x \in \mathcal{X}$, the challenger replies with $\text{sk}_x \leftarrow \text{KeyGen}(\text{msk}, x)$.
 - **Encryption oracle:** On input $f \in \mathcal{F}$, the challenger replies with $\text{ct}_f \leftarrow \text{SEnc}(\text{msk}, f)$.
 - **Challenge oracle:** On input a function $f \in \mathcal{F}$, the challenger computes $\text{ct}_0 \leftarrow \text{PKEnc}(\text{pp})$ and $\text{ct}_1 \leftarrow \text{SEnc}(\text{msk}, f)$ and gives sk_b to the adversary.

At the end of the game, the adversary outputs a bit $b' \in \{0, 1\}$, which is also the output of the experiment. An adversary \mathcal{A} is admissible for the public/secret key indistinguishability game if it makes a single challenge query $f \in \mathcal{F}$ and for all inputs $x \in \mathcal{X}$ the adversary submits to the key-generation oracle, $f(x) = 1$. The mixed FE scheme satisfies (adaptive) public/secret key indistinguishability if for all efficient and admissible adversaries \mathcal{A} , it holds that

$$\left| \Pr[\text{ExptMFE}_{\text{PK/SK}}[\lambda, \mathcal{A}, 0] = 1] - \Pr[\text{ExptMFE}_{\text{PK/SK}}[\lambda, \mathcal{A}, 1] = 1] \right| = \text{negl}(\lambda).$$

Remark 2.9 (Non-Adaptive q -Query Security). For each of the security notions in Definition 2.8 (semantic security and public/secret key indistinguishability), we define a notion of non-adaptive q -query security where the corresponding security notion only holds against all adversaries that make at most $q \in \mathbb{N}$ queries to the encryption oracle, and moreover, all of the non-encryption queries occur *before* the encryption queries.

Remark 2.10 (Public Parameters). In existing mixed FE definitions [GKW18, CVW⁺18a], there is a single Setup algorithm that generates both the public parameters pp and the master secret key msk . In our setting, we require a stronger property where many mixed FE schemes can share a *common* set of public parameters (output by PrmsGen). For this reason, we have separate algorithms PrmsGen and MSKGen for generating pp and msk , respectively. All existing constructions of mixed FE in fact satisfy this stronger requirement. This includes the construction of [GKW18] as well as the constructions of [CVW⁺18a] instantiated with lockable or compute-and-compare obfuscation [GKW17, WZ17], or with a suitable family of private constrained PRFs [BLW17] (i.e., the constructions of [CC17, BTVW17, CVW18b]).

Remark 2.11 (Sub-Exponential Hardness). To argue adaptive security of some of our constructions, we require the stronger notion of *sub-exponential* hardness. We say that a primitive is sub-exponentially secure if for all adversaries \mathcal{A} running in time $\text{poly}(\lambda)$, the adversary’s distinguishing advantage is bounded by $2^{-\Omega(\lambda^\varepsilon)}$ for some constant $\varepsilon > 0$. Sub-exponential hardness assumptions are commonly used to argue adaptive security of cryptographic constructions via complexity leveraging [BB04].

Fact 2.12 (Sub-Exponential Mixed FE from Sub-Exponential LWE). Assuming the sub-exponential hardness of LWE (with a super-polynomial modulus-to-noise ratio), the [GKW18] construction gives a sub-exponentially secure mixed FE scheme that supports the class of NC^1 functions and satisfies non-adaptive q -query security (Remark 2.9) for any a priori bounded $q = \text{poly}(\lambda)$. Similarly, also assuming sub-exponential hardness of LWE (with super-polynomial modulus-to-noise ratio), the [CVW⁺18a] construction (in conjunction with lockable or compute-and-compare obfuscation [GKW17, WZ17] or private constrained PRFs [CC17, BTVW17, CVW18b]) yields a sub-exponentially secure mixed FE scheme that supports all circuits of a priori bounded polynomial depth $d = d(\lambda)$ and satisfies non-adaptive q -query security for a priori bounded $q = q(\lambda)$. We describe one specific instantiation below (specialized to the setting where $q = 1$). Suppose $\mathcal{X} = \{0, 1\}^\ell$ and \mathcal{F} is a function class where each function $f \in \mathcal{F}$ can be described by a string of length $z = z(\lambda)$ and which can be computed by a Boolean circuit of depth $d = d(\lambda)$.⁴ Then the non-adaptive 1-query mixed FE scheme from [CVW⁺18a] instantiated with the key-homomorphic private constrained PRFs from [CC17, BTVW17, CVW18b] satisfies the following properties:

- **Public parameter size:** $|\text{pp}| = \ell \cdot \text{poly}(\lambda, d, z)$.
- **Secret key size:** A secret key sk_x for an input $x \in \mathcal{X}$ has size $|\text{sk}_x| = \ell + \text{poly}(\lambda, z)$.
- **Ciphertext size:** A ciphertext ct (output by either SKEnc or PKEnc) has size $\text{poly}(\lambda, d, z)$.
- **Decryption complexity:** The decryption function Dec can be computed by a Boolean circuit with depth $\text{poly}(\lambda, d, z)$.

Attribute-based encryption. We now recall the definition of (key-policy) attribute-based encryption (ABE) [SW05, GPSW06]. As we note in Fact 2.14, ABE schemes for general circuit policies can be constructed from standard lattice assumptions.

⁴While the description length z of the function f can always be upper-bounded by the size of the Boolean circuit computing f , for some function classes, the description length of f can be much smaller than the size of the Boolean circuit computing f . This will be true for the circuit classes we use to construct trace-and revoke schemes.

Definition 2.13 (Attribute-Based Encryption [SW05, GPSW06]). An attribute-based encryption (ABE) scheme over a message space \mathcal{M} , an attribute space \mathcal{X} , and a function family $\mathcal{F} = \{f: \mathcal{X} \rightarrow \{0, 1\}\}$ is a tuple of algorithms $\Pi_{\text{ABE}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ with the following properties:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{msk})$: On input the security parameter λ , the setup algorithm outputs the public parameters pp and the master secret key msk .
- $\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$: On input the master secret key msk and a function $f \in \mathcal{F}$, the key-generation algorithm outputs a decryption key sk_f .
- $\text{Enc}(\text{pp}, x, m) \rightarrow \text{ct}_{x,m}$: On input the public parameters pp , an attribute $x \in \mathcal{X}$, and a message $m \in \mathcal{M}$, the encryption algorithm outputs a ciphertext $\text{ct}_{x,m}$.
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m/\perp$: On input a decryption key sk , and a ciphertext ct , the decryption algorithm either outputs a message $m \in \mathcal{M}$ or a special symbol \perp .

An attribute-based encryption scheme should satisfy the following properties:

- **Correctness:** For all functions $f \in \mathcal{F}$, all attributes $x \in \mathcal{X}$ where $f(x) = 1$, and all messages $m \in \mathcal{M}$, if we set $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f)$, $\text{ct}_{x,m} \leftarrow \text{Enc}(\text{pp}, x, m)$, it holds that

$$\Pr[\text{Dec}(\text{sk}_f, \text{ct}_{x,m}) = m] = 1 - \text{negl}(\lambda).$$

- **Semantic security:** For a bit $b \in \{0, 1\}$, we define the experiment $\text{ExptABE}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} . The challenger begins by sampling $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives pp to \mathcal{A} . The adversary is then given access to the following oracles:
 - **Key-generation oracle:** On input a function $f \in \mathcal{F}$, the challenger replies with $\text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f)$.
 - **Challenge oracle:** On input an attribute $x \in \mathcal{X}$ and messages $m_0, m_1 \in \mathcal{M}$, the challenger replies with $\text{ct}_b \leftarrow \text{Enc}(\text{pp}, x, m_b)$.

At the end of the game, the adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment. An adversary \mathcal{A} is admissible for the security game if it makes a single challenge query (x, m_0, m_1) and for all functions f it submitted to the key-generation oracle, it holds that $f(x) = 0$. We say that an ABE scheme is semantically secure if for all efficient and admissible adversaries \mathcal{A} ,

$$|\Pr[\text{ExptABE}[\lambda, \mathcal{A}, 0]] - \Pr[\text{ExptABE}[\lambda, \mathcal{A}, 1]]| = \text{negl}(\lambda).$$

Fact 2.14 (Attribute-Based Encryption from Sub-Exponential LWE). Assuming sub-exponential hardness⁵ of LWE (with a super-polynomial modulus-to-noise ratio), there exist attribute-based encryption schemes [GVW13, BGG⁺14] that supports all function families that can be computed by Boolean circuits of a priori polynomially-bounded depth $d = d(\lambda)$. We describe one specific instantiation here. Suppose $\mathcal{X} = \{0, 1\}^\ell$, $\mathcal{M} = \{0, 1\}^t$, and \mathcal{F} is a function class where every function $f \in \mathcal{F}$ can be computed by a Boolean circuit of depth at most $d = d(\lambda)$. Then, the [BGG⁺14] ABE construction has the following properties:

⁵The ABE schemes from [GVW13, BGG⁺14] prove security in a *selective* model of security, while for our construction, we require adaptive security. Thus, we rely on *sub-exponential* hardness (Remark 2.11) and complexity leveraging [BB04].

- **Public parameter size:** For these parameters, $|\text{pp}| = \text{poly}(\lambda, d, \ell)$.
- **Secret key size:** A secret key sk_f for a function f has size $|\text{sk}_f| = |f| + \text{poly}(\lambda, d, \ell)$. Here, $|f|$ denotes the description length of the function f .
- **Ciphertext size:** A ciphertext encrypting an attribute $x \in \{0, 1\}^\ell$ and message $m \in \{0, 1\}^t$ has size $t \cdot \text{poly}(\lambda, d, \ell)$. We can always use hybrid encryption where we use the ABE scheme to encrypt a symmetric key k and then encrypt the message with k .⁶ In this case, the ciphertext size is $t + \ell \cdot \text{poly}(\lambda, d, \ell)$.

3 Revocable Predicate Encryption

In this section, we introduce our notion of a secret-key revocable predicate encryption scheme that supports a public broadcast functionality (i.e., a public-key encryption algorithm that outputs ciphertexts that can be decrypted by all secret keys in the system). This will be the primary primitive we use to construct our identity-based trace-and-revoke scheme (described in Section 4). Our definitions can be viewed as a special case of the more general notion of (public-key) revocable functional encryption from [NWZ16]. The advantage of considering this relaxed notion is that it enables constructions from standard assumptions (whereas we only know how to construct fully secure revocable functional encryption from indistinguishability obfuscation). We introduce our notion below and then show how to construct it by combining mixed FE, ABE, and a subset-cover set system in Section 3.1.

Definition 3.1 (Secret-Key Revocable Predicate Encryption with Broadcast). A *secret-key revocable predicate encryption scheme (RPE) scheme with broadcast* for an identity space \mathcal{ID} , an attribute space \mathcal{X} , a function family $\mathcal{F} = \{f: \mathcal{X} \rightarrow \{0, 1\}\}$, and a message space \mathcal{M} is a tuple of algorithms $\Pi_{\text{RPE}} = (\text{Setup}, \text{KeyGen}, \text{Broadcast}, \text{Enc}, \text{Dec})$ defined as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{msk})$: On input the security parameter λ , the setup algorithm outputs the public parameters pp and the master secret key msk .
- $\text{KeyGen}(\text{msk}, \text{id}, x) \rightarrow \text{sk}_{\text{id}, x}$: On input the master secret key msk , an identity $\text{id} \in \mathcal{ID}$, and an attribute $x \in \mathcal{X}$, the key-generation algorithm outputs a decryption key $\text{sk}_{\text{id}, x}$.
- $\text{Broadcast}(\text{pp}, m, \mathcal{L}) \rightarrow \text{ct}_{m, \mathcal{L}}$: On input the public key, a message m , and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the broadcast algorithm outputs a ciphertext $\text{ct}_{m, \mathcal{L}}$.
- $\text{Enc}(\text{msk}, f, m, \mathcal{L}) \rightarrow \text{ct}_{f, m, \mathcal{L}}$: On input the master secret key msk , a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the encryption algorithm outputs a ciphertext $\text{ct}_{f, m, \mathcal{L}}$.
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m/\perp$: On input a decryption key sk and a ciphertext ct , the decryption algorithm either outputs a message $m \in \mathcal{M}$ or a special symbol \perp .

A secret-key RPE scheme with broadcast should satisfy the following properties:

- **Correctness:** For all functions $f \in \mathcal{F}$, all identities $\text{id} \in \mathcal{ID}$, all attributes $x \in \mathcal{X}$ where $f(x) = 1$, all messages $m \in \mathcal{M}$, and all revocation lists $\mathcal{L} \subseteq \mathcal{ID}$ where $\text{id} \notin \mathcal{L}$, if we set $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{sk}_{\text{id}, x} \leftarrow \text{KeyGen}(\text{msk}, \text{id}, x)$, the following holds:

⁶We note that a symmetric encryption scheme can be constructed from any one-way function and therefore, can be based on the hardness of LWE.

– **Broadcast correctness:** If $\text{ct}_{m,\mathcal{L}} \leftarrow \text{Broadcast}(\text{pp}, m, \mathcal{L})$, then

$$\Pr[\text{Dec}(\text{sk}_{\text{id},x}, \text{ct}_{m,\mathcal{L}}) = m] = 1 - \text{negl}(\lambda).$$

– **Encryption correctness:** If $\text{ct}_{f,m,\mathcal{L}} \leftarrow \text{Enc}(\text{msk}, f, m, \mathcal{L})$, then

$$\Pr[\text{Dec}(\text{sk}_{\text{id},x}, \text{ct}_{f,m,\mathcal{L}}) = m] = 1 - \text{negl}(\lambda).$$

- **Message hiding:** For a bit $b \in \{0, 1\}$, we define the experiment $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} . The challenger begins by sampling $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives pp to \mathcal{A} . The adversary is then given access to the following oracles:

- **Key-generation oracle:** On input an identity $\text{id} \in \mathcal{ID}$ and an attribute $x \in \mathcal{X}$, the challenger replies with $\text{sk}_{\text{id},x} \leftarrow \text{KeyGen}(\text{msk}, \text{id}, x)$.
- **Encryption oracle:** On input a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger replies with $\text{ct}_{f,m,\mathcal{L}} \leftarrow \text{Enc}(\text{msk}, f, m, \mathcal{L})$.
- **Challenge oracle:** On input a function $f \in \mathcal{F}$, two messages $m_0, m_1 \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger computes $\text{ct}_b \leftarrow \text{Enc}(\text{msk}, f, m_b, \mathcal{L})$ and gives ct_b to the adversary.

At the end of the game, the adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment. An adversary \mathcal{A} is admissible for the message hiding game if it makes a single challenge query $(f, m_0, m_1, \mathcal{L})$ such that for all pairs (id, x) the adversary submitted to the key-generation oracle, it holds that $f(x) = 0$ or $\text{id} \in \mathcal{L}$. We say that Π_{RPE} satisfies (adaptive) *message hiding* if for all efficient and admissible adversaries \mathcal{A} ,

$$|\Pr[\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 0] = 1] - \Pr[\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 1] = 1]| = \text{negl}(\lambda).$$

- **Function hiding:** For a bit $b \in \{0, 1\}$, we define the experiment $\text{ExptRPE}_{\text{FH}}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} exactly as $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, b]$, except the challenge oracle is replaced with the following:

- **Challenge oracle:** On input two functions $f_0, f_1 \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger computes $\text{ct}_b \leftarrow \text{Enc}(\text{msk}, f_b, m, \mathcal{L})$ and gives ct_b to the adversary.

We say an adversary \mathcal{A} is admissible for the function-hiding game if it makes a single challenge query $(f_0, f_1, m, \mathcal{L})$ such that for all pairs (id, x) the adversary submitted to the key-generation oracle, either $f_0(x) = f_1(x)$ or $\text{id} \in \mathcal{L}$. We say that Π_{RPE} satisfies (adaptive) *function hiding* if for all efficient and admissible adversaries \mathcal{A} ,

$$|\Pr[\text{ExptRPE}_{\text{FH}}[\lambda, \mathcal{A}, 0] = 1] - \Pr[\text{ExptRPE}_{\text{FH}}[\lambda, \mathcal{A}, 1] = 1]| = \text{negl}(\lambda).$$

- **Broadcast security:** For a bit $b \in \{0, 1\}$, we define the security experiment $\text{ExptRPE}_{\text{BC}}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} exactly as $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, b]$, except the challenge oracle is replaced with the following:

- **Challenge oracle:** On input a message $m \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger computes $\text{ct}_0 \leftarrow \text{Broadcast}(\text{pp}, m, \mathcal{L})$ and $\text{ct}_1 \leftarrow \text{Enc}(\text{msk}, f, m, \mathcal{L})$ where f_{accept} is the “always-accept” function (i.e., $f_{\text{accept}}(x) = 1$ for all $x \in \mathcal{X}$). It gives ct_b to the adversary.

At the end of the game, the adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment. We say that Π_{RPE} satisfies (adaptive) *broadcast security* if for all efficient adversaries \mathcal{A} that make at most one challenge query,

$$|\Pr[\text{ExptRPE}_{\text{BC}}[\lambda, \mathcal{A}, b] = 1] - \Pr[\text{ExptRPE}_{\text{BC}}[\lambda, \mathcal{A}, 1]]| = \text{negl}(\lambda).$$

Remark 3.2 (Non-Adaptive q -Query Security). Analogously to Remark 2.9, for each of the security notions in Definition 3.1 (message hiding, function hiding, and broadcast security), we define a notion of *non-adaptive q -query security* where the corresponding security notion only holds against all adversaries that make at most $q \in \mathbb{N}$ queries to the encryption oracle, and moreover, all of the non-encryption queries occur *before* the encryption queries. Achieving this notion is easier and suffices for our main construction (*adaptively-secure trace-and-revoke*).

Remark 3.3 (Embedding the Revocation List in the Attribute). A natural approach for constructing a revocable predicate encryption scheme from any vanilla predicate encryption scheme is to include the revocation list \mathcal{L} as part of the function in the predicate encryption scheme. A decryption key for an identity id would then check that id is not contained in the revocation list \mathcal{L} associated with the ciphertext. This is the approach suggested in [NWZ16, Remark 6.2] in the context of constructing a revocable functional encryption scheme. While this approach may seem straightforward, it has a significant drawback in most settings. In existing predicate encryption schemes from standard assumptions, the decryption functionality is represented as a *circuit*, which takes *fixed-size* inputs. Thus, if the revocation list is embedded as part of the ciphertext, then a predicate encryption scheme for circuit-based predicates would only be able to support an *a priori* bounded number of revocations. In contrast, our construction allows for revoking an *arbitrary* polynomial number of users (Section 3.1). Of course, if we can construct predicate or functional encryption for Turing machine or RAM computations, then this natural revocation approach would suffice. Existing constructions of functional encryption for Turing machine computations all rely on indistinguishability obfuscation [KLW15, AJS17, AS16, GS18].

3.1 Constructing Secret-Key Revocable Predicate Encryption with Broadcast

In this section, we describe our construction of a secret-key revocable predicate encryption with broadcast scheme for general predicates by combining a mixed FE scheme, an ABE scheme, and a subset-cover set system. As discussed in 1.1, our core construction (without revocation) can be viewed as a direct generalization of the construction of private linear broadcast encryption from mixed FE and ABE from [GKW18]. We next augment our construction with a subset cover set system to support revocation. Our techniques allow revoking an arbitrary number of users (in contrast to previous trace-and-revoke schemes from standard assumptions that could only handle bounded revocations [NWZ16, ABP⁺17]). We give our full construction and its analysis below:

Construction 3.4 (Secret-Key Revocable Predicate Encryption with Broadcast). Fix an identity space $\mathcal{ID} = \{0, 1\}^n$, attribute space \mathcal{X} , function family $\mathcal{F} = \{f: \mathcal{X} \rightarrow \{0, 1\}\}$ and message space \mathcal{M} , where $n = n(\lambda)$.

- Let $[K]$ be the subset-cover set system for the set $\mathcal{ID} = \{0, 1\}^n$. Let $\Pi_{\text{SC}} = (\text{Encode}, \text{ComputeCover})$ be the algorithms associated with the set system.
- Let $\Pi_{\text{MFE}} = (\text{MFE.PrmsGen}, \text{MFE.MSKGen}, \text{MFE.KeyGen}, \text{MFE.PKEnc}, \text{MFE.SKEnc}, \text{MFE.Dec})$ be a mixed FE scheme with domain \mathcal{X} and function family \mathcal{F} . Let $\rho = \rho(\lambda)$ be the randomness complexity of the master secret key generation algorithm MFE.MSKGen , let \mathcal{CT} denote the ciphertext space of Π_{MFE} (i.e., the range of MFE.PKEnc and MFE.SKEnc), and let \mathcal{SK} denote the secret key space of Π_{MFE} (i.e., the range of MFE.KeyGen). We will require that Π_{MFE} be sub-exponentially secure (Remark 2.11), so let $\varepsilon > 0$ be a constant such that $2^{-\Omega(\lambda^\varepsilon)}$ bounds the advantage of any efficient adversary \mathcal{A} for the security of Π_{MFE} .
- For a secret key $\text{mfe.sk} \in \mathcal{SK}$ and an index $i^* \in [K]$, define the function $g_{\text{mfe.sk}, i^*} : \mathcal{CT} \times [K] \rightarrow \{0, 1\}$ to be the function

$$g_{\text{mfe.sk}, i^*}(\text{ct}, i) = \begin{cases} 1 & \text{MFE.Dec}(\text{mfe.sk}, \text{ct}) = 1 \text{ and } i = i^* \\ 0 & \text{otherwise.} \end{cases}$$

- Let $\Pi_{\text{ABE}} = (\text{ABE.Setup}, \text{ABE.KeyGen}, \text{ABE.Enc}, \text{ABE.Dec})$ be an attribute-based encryption scheme over message space \mathcal{M} , attribute space $\mathcal{X}' = \mathcal{CT} \times [K]$ and function family $\mathcal{F}' = \{\text{mfe.sk} \in \mathcal{SK}, i^* \in [K] : g_{\text{mfe.sk}, i^*}\}$.
- Let $F : \mathcal{K} \times [K] \rightarrow \{0, 1\}^\rho$ be a pseudorandom function.

We construct a secret-key revocable predicate encryption scheme as follows:

- **Setup**(1^λ): On input the security parameter λ , the setup algorithm sets $\lambda' = \max(\lambda, (\log K)^{2/\varepsilon})$. It then generates mixed FE public parameters $\text{mfe.pp} \leftarrow \text{MFE.PrmsGen}(1^{\lambda'})$. It also instantiates an attribute-based encryption scheme $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda)$, samples a PRF key $k \xleftarrow{R} \mathcal{K}$, and outputs

$$\text{pp} = (\text{mfe.pp}, \text{abe.pp}) \quad \text{and} \quad \text{msk} = (\text{pp}, \text{abe.msk}, k).$$

- **KeyGen**(msk, id, x): On input a master secret key msk , an identity $\text{id} \in \mathcal{ID}$, and an attribute $x \in \mathcal{X}$, the key-generation algorithm does the following:
 1. Compute a subset-cover encoding of the identity $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$.
 2. For each index $i \in \mathcal{I}_{\text{id}}$, the algorithm samples randomness $r_i \leftarrow F(k, i)$. It then generates a mixed FE master secret key $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$ and a mixed FE decryption key $\text{mfe.sk}_{i,x} \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}_i, x)$.
 3. Finally, for each $i \in \mathcal{I}_{\text{id}}$, it constructs an ABE decryption key with respect to the function $g_{\text{mfe.msk}_i, x, i}$ as follows: $\text{abe.sk}_{i,x} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, g_{\text{mfe.msk}_i, x, i})$.
 4. It outputs the collection of keys $\text{sk}_{\text{id}, x} = \{(i, \text{abe.sk}_{i,x})\}_{i \in \mathcal{I}_{\text{id}}}$.
- **Broadcast**($\text{pp}, m, \mathcal{L}$): On input the public parameters $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$, a message m , and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the broadcast algorithm does the following:
 1. Obtain a cover for $\mathcal{ID} \setminus \mathcal{L}$ by computing $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$.

2. For each $i \in \mathcal{J}_{\mathcal{L}}$, it generates a mixed FE ciphertext $\text{mfe.ct}_i \leftarrow \text{MFE.PKEnc}(\text{mfe.pp})$ and an ABE ciphertext $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$.
 3. It outputs the ciphertext $\text{ct}_{m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$.
- $\text{Enc}(\text{msk}, f, m, \mathcal{L})$: On input the master secret key $\text{msk} = (\text{pp}, \text{abe.msk}, k)$, a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, where $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$, the encryption algorithm does the following:
 1. Obtain a cover for $\mathcal{ID} \setminus \mathcal{L}$ by computing $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$.
 2. Then, for each $i \in \mathcal{J}_{\mathcal{L}}$, it computes $r_i \leftarrow F(k, i)$ and derives the corresponding mixed FE master secret key $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$. It then encrypts $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f)$.
 3. For each $i \in \mathcal{J}_{\mathcal{L}}$, it computes $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$, and outputs the ciphertext $\text{ct}_{f,m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$.
 - $\text{Dec}(\text{sk}, \text{ct})$: On input a key $\text{sk} = \{(i, \text{abe.sk}_i)\}_{i \in \mathcal{I}}$ and a ciphertext $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}}$, the decryption algorithm first checks if $\mathcal{I} \cap \mathcal{J} = \emptyset$. If so, it outputs \perp . Otherwise, it chooses an arbitrary index $i \in \mathcal{I} \cap \mathcal{J}$ and outputs $m \leftarrow \text{ABE.Dec}(\text{abe.sk}_i, \text{abe.ct}_i)$.

Correctness and security analysis. We state our main theorems on the properties of Construction 3.4 below, but defer their analysis to Appendix A.

Theorem 3.5 (Correctness). *Suppose that Π_{MFE} , Π_{ABE} , and Π_{SC} are correct. Then, the predicate encryption scheme Π_{RPE} from Construction 3.4 is correct.*

Theorem 3.6 (Message Hiding). *Suppose that Π_{MFE} and Π_{SC} are correct, and Π_{ABE} satisfies semantic security. Then, the predicate encryption scheme Π_{RPE} from Construction 3.4 satisfies message hiding.*

Theorem 3.7 (Function Hiding). *Suppose that Π_{MFE} satisfies sub-exponential non-adaptive q -query (resp., adaptive) semantic security. Specifically, suppose that the advantage of any adversary running in time $\text{poly}(\lambda)$ in the semantic security game is bounded by $2^{-\Omega(\lambda^\epsilon)}$. In addition, suppose that Π_{ABE} is secure, F is a secure PRF, and Π_{SC} is correct. Then, the predicate encryption scheme in Construction 3.4 satisfies non-adaptive q -query (resp., adaptive) function hiding security.*

Theorem 3.8 (Broadcast Security). *Suppose that Π_{MFE} satisfies sub-exponential non-adaptive q -query (resp., adaptive) public/secret key indistinguishability. Specifically, suppose that the advantage of any adversary running in time $\text{poly}(\lambda)$ in the public/secret key indistinguishability game is bounded by $2^{-\Omega(\lambda^\epsilon)}$. In addition, suppose that F is a secure PRF. Then the predicate encryption scheme Π_{RPE} in Construction 3.4 satisfies non-adaptive q -query (resp., adaptive) broadcast security.*

3.2 Instantiating Secret-Key Revocable Predicate Encryption with Broadcast

In this section, we describe one possible instantiation of secret-key revocable predicate encryption with broadcast from Construction 3.4. In particular, combining Construction 3.4 with Theorems 3.5 through 3.8 yields the following corollary:

Corollary 3.9 (Secret-Key Revocable Predicate Encryption from LWE). *Take an identity-space $\mathcal{ID} = \{0, 1\}^n$, attribute space $\mathcal{X} = \{0, 1\}^\ell$, and message space $\mathcal{M} = \{0, 1\}^t$ where $n = n(\lambda)$, $\ell = \ell(\lambda)$, and $t = t(\lambda)$. Let $\mathcal{F} = \{f: \mathcal{X} \rightarrow \{0, 1\}\}$ be a function family where every function $f \in \mathcal{F}$ can be specified by a string of length $z = z(\lambda)$ and computed by a Boolean circuit of depth $d = d(\lambda)$. Then, assuming sub-exponential hardness of LWE (with a super-polynomial modulus-to-noise ratio), there exists a non-adaptive 1-key secure secret-key revocable predicate encryption scheme with broadcast Π_{RPE} over the identity space \mathcal{ID} , attribute space \mathcal{X} , and function family \mathcal{F} . Moreover, Π_{RPE} satisfies the following properties:*

- **Public parameter size:** $|\text{pp}| = \ell \cdot \text{poly}(\lambda, d, n, z)$.
- **Secret key size:** The secret key $\text{sk}_{\text{id},x}$ for an identity $\text{id} \in \{0, 1\}^n$ and an attribute $x \in \{0, 1\}^\ell$ has size $|\text{sk}_{\text{id},x}| = \ell + \text{poly}(\lambda, d, n, z)$.
- **Ciphertext size:** An encryption $\text{ct}_{m,\mathcal{L}}$ of a message $m \in \{0, 1\}^t$ with revocation list \mathcal{L} has size $|\text{ct}_{m,\mathcal{L}}| = t + |\mathcal{L}| \cdot \text{poly}(\lambda, d, n, z)$.

Proof. We instantiate Construction 3.4 using the subset-cover set system from Fact 2.4, the mixed FE scheme using the construction from Fact 2.12, the ABE scheme using the construction from Fact 2.14, and the PRF from any one-way function [GGM84]. The mixed FE scheme is instantiated with domain $\mathcal{X} = \{0, 1\}^\ell$ and function family \mathcal{F} , while the ABE scheme is instantiated with message space \mathcal{M} , attribute space $\mathcal{X}' = \mathcal{CT} \times [K]$ and function family $\mathcal{F}' = \{\text{mfe.sk} \in \mathcal{SK}, i^* \in [K] : g_{\text{mfe.sk}, i^*}\}$. We will use the following bounds in our analysis:

- From Fact 2.4, we have that $K = O(N)$, and correspondingly, $\log K = O(\log N) = O(n)$.
- By Fact 2.12, we have that the length of a mixed FE ciphertext $\text{mfe.ct} \in \mathcal{CT}$ is bounded by $|\text{mfe.ct}| = \text{poly}(\lambda, d, z)$. Correspondingly, this means that the length ℓ_{ABE} of an ABE attribute is bounded by $\ell_{\text{ABE}} = \text{poly}(\lambda, d, z) + \log K = \text{poly}(\lambda, d, n, z)$.
- By Fact 2.12, each function $g_{\text{mfe.sk}, i^*}$ can be implemented by a circuit with depth at most $\text{poly}(\lambda, d) + \log \log K = \text{poly}(\lambda, d, \log n)$. Specifically, the mixed FE decryption circuit can be evaluated by a circuit of depth $\text{poly}(\lambda', d) = \text{poly}(\lambda, d, n, z)$ and the equality-check circuit can be evaluated by a circuit of depth $\log \log K$ (since each input to the equality-check circuit is a $(\log K)$ -bit value). Thus, the functions in \mathcal{F}' can be computed by Boolean circuits with depth at most $d_{\text{ABE}} \leq \text{poly}(\lambda, d, n, z)$. The description length of functions in \mathcal{F}' is $|\text{mfe.sk}| + \log K = \ell + \text{poly}(\lambda, n, z)$.

Putting all the pieces together, we now have the following:

- **Public parameter size:** The public parameters pp consist of the ABE public parameters abe.pp and the mixed FE public parameters mfe.pp . By Fact 2.14,

$$|\text{abe.pp}| = \text{poly}(\lambda, d_{\text{ABE}}, \ell_{\text{ABE}}) = \text{poly}(\lambda, d, n, z),$$

and correspondingly, by Fact 2.12,

$$|\text{mfe.pp}| = \ell \cdot \text{poly}(\lambda', d, z) = \ell \cdot \text{poly}(\lambda, d, n, z),$$

since $\lambda' = \text{poly}(\lambda, \log K) = \text{poly}(\lambda, n)$. Thus, $|\text{pp}| = \ell \cdot \text{poly}(\lambda, d, n, z)$.

- **Secret key size:** The secret key $\text{sk}_{\text{id},x} = \{(i, \text{abe.sk}_{i,x})\}_{i \in \mathcal{I}_{\text{id}}}$ for an identity id and attribute x consists of $|\mathcal{I}_{\text{id}}|$ ABE secret keys, where $|\mathcal{I}_{\text{id}}| \leftarrow \text{Encode}(\text{id})$. By Fact 2.4, $|\mathcal{I}_{\text{id}}| = \log N + 1 = \text{poly}(n)$. Finally, by Facts 2.12 and 2.14

$$|\text{abe.sk}_{i,x}| = |g_{\text{mfe.sk}_{i,x},i}| + \text{poly}(\lambda, d_{\text{ABE}}, \ell_{\text{ABE}}) = \ell + \text{poly}(\lambda, d, n, z).$$

Thus, $|\text{sk}_{\text{id},x}| = |\mathcal{I}_{\text{id}}| \cdot |\text{abe.sk}_{i,x}| = \ell + \text{poly}(\lambda, d, n, z)$.

- **Ciphertext size:** Without loss of generality, we can always use hybrid encryption for the ciphertexts. Namely, the encryption algorithm samples a symmetric key k to encrypt the message and then encrypts k using the secret-key revocable predicate encryption scheme. The final ciphertext $\text{ct}_{m,\mathcal{L}}$ then consists of a symmetric encryption of the message m (which has size $|m| + \text{poly}(\lambda)$) and a revocable predicate encryption ciphertext $\widehat{\text{ct}}$ of the key k . In this case, $|k| = \text{poly}(\lambda)$, and the overall ciphertext size is $|\text{ct}| = |m| + \text{poly}(\lambda) + |\widehat{\text{ct}}|$, where $\widehat{\text{ct}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ is an encryption of k using Π_{RPE} . By construction, $\widehat{\text{ct}}$ consists of $|\mathcal{J}_{\mathcal{L}}|$ ABE ciphertexts, where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. By Fact 2.4, $|\mathcal{L}| = O(|\mathcal{L}| \log(N/|\mathcal{L}|)) = |\mathcal{L}| \cdot \text{poly}(n)$. By Fact 2.14, $|\text{abe.ct}_i| = |k| + \ell_{\text{ABE}} \cdot \text{poly}(\lambda, d_{\text{ABE}}, \ell_{\text{ABE}}) = \text{poly}(\lambda, d, n, z)$, and so

$$|\text{ct}_{m,\mathcal{L}}| = |m| + \text{poly}(\lambda) + |\widehat{\text{ct}}| = t + |\mathcal{L}| \cdot \text{poly}(\lambda, d, n, z). \quad \square$$

Remark 3.10 (Handling More General Revocation Policies). Construction 3.4 naturally supports any revocation policy that can be described by a polynomial-size cover in the underlying subset-cover set system. In particular, the prefix-based subset-cover set system by Naor et al. [NNL01] from Fact 2.4 can compute a cover that excludes any polynomial number of *prefixes* (in addition to full identities). For instance, we can use the set system to revoke all users whose identities start with “000” or “01” (i.e., revoke all identities of the form 000*** and 01****). This way, the number of revoked users in the set \mathcal{L} can be *exponential*, as long as they can be described by a polynomial-number of prefix-based clusters. Correspondingly, the traitor tracing scheme we construct in Section 4 will also support these types of revocation policies.

4 Identity-Based Trace-and-Revoke

In this section, we describe how to construct an identity-based trace-and-revoke scheme using a secret-key revocable predicate encryption scheme with broadcast (Definition 3.1). We begin by recalling the formal definition of a trace-and-revoke scheme. Our definitions are adapted from the corresponding ones in [BW06, NWZ16]. As we discuss in greater detail in Remark 4.2, our definition combines aspects of both definitions and is strictly stronger than both of the previous notions.

Definition 4.1 (Trace-and-Revoke [NWZ16, adapted]). A trace-and-revoke scheme for a set of identities \mathcal{ID} and a message space \mathcal{M} is a tuple of algorithms $\Pi_{\text{TR}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ defined as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{msk})$: On input the security parameter λ , the setup algorithm outputs the public parameters pp and the master secret key msk .
- $\text{KeyGen}(\text{msk}, \text{id}) \rightarrow \text{sk}_{\text{id}}$: On input the master secret key msk and an identity $\text{id} \in \mathcal{ID}$, the key-generation algorithm outputs a secret key sk_{id} .

- $\text{Enc}(\text{pp}, m, \mathcal{L}) \rightarrow \text{ct}_{m,\mathcal{L}}$: On input the public parameters pp , a message $m \in \mathcal{M}$, and a list of revoked users $\mathcal{L} \subseteq \mathcal{ID}$, the encryption algorithm outputs a ciphertext $\text{ct}_{m,\mathcal{L}}$.
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m/\perp$: On input a decryption key sk and a ciphertext ct , the decryption algorithm either outputs a message $m \in \mathcal{M}$ or a special symbol \perp .
- $\text{Trace}^{\mathcal{D}}(\text{msk}, m_0, m_1, \mathcal{L}, \varepsilon) \rightarrow \text{id}/\perp$: On input the master secret key msk , two messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, a decoder-success parameter $\varepsilon > 0$, and assuming oracle access to a decoder algorithm \mathcal{D} , the tracing algorithm either outputs an identity $\text{id} \in \mathcal{ID}$ or \perp .

Moreover, a trace-and-revoke scheme should satisfy the following properties:

- **Correctness:** For all messages $m \in \mathcal{M}$, all identities $\text{id} \in \mathcal{ID}$, and all revocation lists $\mathcal{L} \subseteq \mathcal{ID}$ where $\text{id} \notin \mathcal{L}$, if we set $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id})$, and $\text{ct}_{m,\mathcal{L}} \leftarrow \text{Enc}(\text{pp}, m, \mathcal{L})$, then

$$\Pr[\text{Dec}(\text{sk}_{\text{id}}, \text{ct}_{m,\mathcal{L}}) = m] = 1 - \text{negl}(\lambda).$$

- **Semantic Security:** For a bit $b \in \{0, 1\}$, we define the security experiment $\text{ExptTR}_{\text{SS}}[\lambda, \mathcal{A}, b]$ between a challenger and an adversary \mathcal{A} . The challenger begins by sampling $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives pp to \mathcal{A} . The adversary is then given access to the following oracles:
 - **Key-generation oracle.** On input an identity $\text{id} \in \mathcal{ID}$, the challenger replies with $\text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id})$.
 - **Challenge oracle.** On input two messages $m_0, m_1 \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger replies with $\text{ct}_b \leftarrow \text{Enc}(\text{pp}, m_b, \mathcal{L})$.

At the end of the game, the adversary outputs a bit $b' \in \{0, 1\}$, which is the output of the experiment. An adversary \mathcal{A} is admissible for the semantic security game if it makes a single challenge query (m_0, m_1, \mathcal{L}) , and moreover, for all key-generation queries id the adversary makes, $\text{id} \in \mathcal{L}$. We say that Π_{TR} is semantically secure if for all efficient and admissible adversaries \mathcal{A} ,

$$|\Pr[\text{ExptTR}_{\text{SS}}[\lambda, \mathcal{A}, 0] = 1] - \Pr[\text{ExptTR}_{\text{SS}}[\lambda, \mathcal{A}, 1] = 1]| = \text{negl}(\lambda).$$

- **Traceability:** We define the experiment $\text{ExptTR}_{\text{TR}}[\lambda, \mathcal{A}]$ between a challenger and an adversary \mathcal{A} . The challenger begins by sampling $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and gives pp to \mathcal{A} . The adversary is then given access to the key-generation oracle:
 - **Key-generation oracle.** On input an identity $\text{id} \in \mathcal{ID}$, the challenger replies with $\text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id})$.

At the end of the game, the adversary outputs a decoder algorithm \mathcal{D} , two messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, and a non-negligible decoder-success probability $\varepsilon > 0$. Let $\mathcal{R} \subseteq \mathcal{ID}$ be the set of identities the adversary submitted to the key-generation oracle and let $\text{id}^* \leftarrow \text{Trace}^{\mathcal{D}}(\text{msk}, m_0, m_1, \mathcal{L}, \varepsilon)$. Then the output of the experiment is 1 if $\text{id}^* \notin \mathcal{R} \setminus \mathcal{L}$ and 0 otherwise. We say that an adversary \mathcal{A} is admissible for the traceability game if the decoder algorithm output by \mathcal{A} satisfies

$$\Pr[b \stackrel{\mathcal{R}}{\leftarrow} \{0, 1\} : \mathcal{D}(\text{Enc}(\text{pp}, m_b, \mathcal{L})) = b] \geq 1/2 + \varepsilon.$$

Finally, we say that Π_{TR} satisfies traceability security if for all efficient and admissible adversaries \mathcal{A} ,

$$\Pr[\text{ExptTR}_{\text{TR}}[\lambda, \mathcal{A}] = 1] = \text{negl}(\lambda).$$

Remark 4.2 (Comparison to Previous Traceability Notions). Our notion of traceability in Definition 4.1 combines aspects of the notions considered in [BW06] and [NWZ16] and is stronger than both of these previous definitions. First, similar to [NWZ16], we only require that the decoder \mathcal{D} output by \mathcal{A} to be able to *distinguish* the encryptions of two adversarially-chosen messages. The previous notion in [BW06] made the more stringent requirement that the adversary’s decoder must correctly decrypt a noticeable fraction of ciphertexts. Thus, our definitions enable tracing for much weaker decoders. Next, and similar to [BW06], our tracing definition naturally incorporates revocation. Namely, if an adversary constructs a decoder that is able to distinguish encryptions of two messages with respect to a revocation list \mathcal{L} , then the tracing algorithm must identify a compromised key that is outside \mathcal{L} . In contrast, the definition in [NWZ16] only considered tracing in a standalone setting: namely, while the scheme supports revocation, the tracing definition only considered decoders that can decrypt ciphertexts encrypted to an *empty* revocation list. Overall, our definition is stronger than the previous definitions and we believe provides a more realistic modeling of the security demands in applications of trace-and-revoke systems.

Remark 4.3 (Adaptive Security). We note that all of the security requirements in Definition 4.1 are adaptive: namely, the adversary chooses its challenge messages and revocation list after seeing the public parameters and (adaptively-chosen) secret decryption keys. Our final construction is fully adaptive (Construction 4.4, Corollary 4.8), but we do rely on complexity leveraging and sub-exponential hardness assumptions. We remark here that a selective notion of security where the adversary commits to its revocation list ahead of time does not seem to directly imply adaptive security by the usual complexity leveraging technique [BB04] unless we additionally impose an a priori bound on the size of the revocation list (which we do not require in our analysis). It is an interesting problem to construct a fully collusion resistant trace-and-revoke scheme for arbitrary identities from standard polynomial hardness assumptions.

4.1 Constructing an Identity-Based Trace-and-Revoke Scheme

Our construction follows the general high-level schema as that by Nishimaki et al. [NWZ16], except our construction is secretly-traceable (but will provide *full* collusion resistance). Very briefly, we use a secret-key revocable predicate encryption scheme to embed an instance of the generalized jump-finding problem (Definition 2.5) where the position of the “jumps” correspond to non-revoked keys. The tracing algorithm relies on the generalized jump-finding algorithm (Theorem 2.6) to identify the compromised keys. We give our construction below.

Construction 4.4 (Identity-Based Trace-and-Revoke). Let $\mathcal{ID} = \{0, 1\}^n$ be the identity space and let \mathcal{M} be a message space. We additionally rely on the following primitives:

- Let $H: \mathcal{K} \times \mathcal{ID} \rightarrow [2^\ell]$ be a keyed collision-resistant hash function.
- Let $\mathcal{ID}_0 = [2^{\ell+1}]$. For a pair $(i, u) \in [n] \times [0, 2^{\ell+1}]$, define the function $f_{i,u}: \mathcal{ID}_0^n \rightarrow \{0, 1\}$ to be the function that takes as input $v = (v_1, \dots, v_n)$, where each $v_i \in \mathcal{ID}_0$, and outputs 1 if $v_i \leq u$ and 0 otherwise. When $u = 0$, $f_{i,u}(v) = 0$ for all $i \in [n]$ and $v \in \mathcal{ID}_0^n$. Similarly, when

$u = 2^{\ell+1}$, $f_{i,u}(v) = 1$ for all $i \in [n]$ and $v \in \mathcal{ID}_0^n$. We will use a canonical “all-zeroes” function to represent $f_{i,0}$ and a canonical “all-ones” function to represent $f_{i,2^{\ell+1}}$ for all $i \in [n]$.

- Let $\Pi_{\text{RPE}} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Broadcast}, \text{RPE.Enc}, \text{RPE.Dec})$ be a secret-key revocable predicate encryption scheme with broadcast with attribute space \mathcal{ID}_0^n , label space $[2^\ell]$, message space \mathcal{M} , and function space $\mathcal{F} = \{i \in [n], u \in [0, 2^{\ell+1}] : f_{i,u}\}$.

We construct a trace-and-revoke scheme $\Pi_{\text{TR}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ with identity space \mathcal{ID} and message space \mathcal{M} as follows:

- $\text{Setup}(1^\lambda)$: On input the security parameter λ , the setup algorithm samples a key $\text{hk} \xleftarrow{\text{R}} \mathcal{K}$, parameters $(\text{rpe.pp}, \text{rpe.msk}) \leftarrow \text{RPE.Setup}(1^\lambda)$, and outputs

$$\text{pp} = (\text{hk}, \text{rpe.pp}) \quad \text{and} \quad \text{msk} = (\text{hk}, \text{rpe.msk}).$$

- $\text{KeyGen}(\text{msk}, \text{id})$: On input the master secret key $\text{msk} = (\text{hk}, \text{rpe.msk})$ and an identity $\text{id} = (\text{id}_1, \dots, \text{id}_n) \in \mathcal{ID}$, the key-generation algorithm computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and defines the vector $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n) \in \mathcal{ID}_0^n$. It outputs $\text{sk}_{\text{id}} \leftarrow \text{RPE.KeyGen}(\text{rpe.msk}, s_{\text{id}}, v_{\text{id}})$.
- $\text{Enc}(\text{pp}, m, \mathcal{L})$: On input the public parameters $\text{pp} = (\text{hk}, \text{rpe.pp})$, a message m , and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the encryption algorithm first constructs a new list $\mathcal{L}' \subseteq \{0, 1\}^\ell$ where $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. Then, it outputs $\text{ct}_{m, \mathcal{L}} \leftarrow \text{RPE.Broadcast}(\text{rpe.pp}, m, \mathcal{L}')$.
- $\text{Dec}(\text{sk}, \text{ct})$: On input a secret key sk and a ciphertext ct , the decryption algorithm outputs $m \leftarrow \text{RPE.Dec}(\text{sk}, \text{ct})$.
- $\text{Trace}^{\mathcal{D}}(\text{msk}, m_0, m_1, \mathcal{L}, \varepsilon)$: On input the decryption oracle \mathcal{D} , the master secret key $\text{msk} = (\text{hk}, \text{rpe.msk})$, messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, and a success probability ε , the tracing algorithm begins by constructing the set $\mathcal{L}' \subseteq \{0, 1\}^\ell$ where $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. It then defines the following *randomized* oracle Q :

On input a pair $(i, u) \in [n] \times [0, 2^{\ell+1}]$:

1. Sample a random bit $b \xleftarrow{\text{R}} \{0, 1\}$, and construct the ciphertext $\text{ct}_b \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{i,u}, m_b, \mathcal{L}')$.
2. Run the decoder algorithm \mathcal{D} on the ciphertext ct_b to obtain a bit $b' \leftarrow \mathcal{D}(\text{ct}_b)$.
3. Output 1 if $b = b'$ and 0 otherwise.

Figure 1: The randomized oracle Q used for tracing.

Let $q = 1$, set $\delta_q = \varepsilon / (9 + 4(\ell - 1)q)$, and compute $\mathcal{T}_q \leftarrow \text{QTrace}^Q(\lambda, 2^\ell, n, q, \delta_q, \varepsilon)$. If \mathcal{T}_q is non-empty, take any element $(s_{\text{id}}, \text{id}_1, \dots, \text{id}_n) \in \mathcal{T}_q$, and output $\text{id} = (\text{id}_1, \dots, \text{id}_n) \in \mathcal{ID}$. Otherwise, update $q \leftarrow 2q$ and repeat this procedure.⁷

⁷We will argue in the proof of Theorem 4.7 that this algorithm will terminate with overwhelming probability. Alternatively, we can set an upper bound on the maximum number of iterations q_{max} . In this case, the tracing algorithm succeeds as long as the total number of keys issued is bounded by $2^{q_{\text{max}}}$. Note that this is not an *a priori* bound on the number of keys that can be issued, just a bound on the number of iterations on which to run the tracing algorithm, which can be a flexible parameter (independent of other scheme parameters).

Correctness and security analysis. We now show that Π_{TR} from Construction 4.4 satisfies correctness, semantic security, and traceability. We state the main theorems below, but defer their formal proofs to Appendix B. The analysis proceeds similarly to the corresponding analysis from [NWZ16], except we operate in the secret-traceability setting. The main challenge in the secret-key setting is that when the adversary in the traceability game outputs a pirate decoder, the reduction algorithm cannot easily tell whether the decoder is “useful” or not (where a “useful” decoder is one that can be leveraged to break the security of the underlying secret-key revocable predicate encryption scheme). The analysis in [NWZ16] solves this problem by having the reduction algorithm sample ciphertexts of its own and observe the decoder’s behavior on those ciphertexts. In this way, the reduction is able to estimate the decoder’s distinguishing advantage and identify whether the adversary produced a good decoder or not. In the secret-key setting, the reduction *cannot* sample ciphertexts of its own and as such, it cannot estimate the decoder’s success probability. To solve this problem, we adopt the approach taken in [GKW18] and allow the reduction algorithm to make a *single* encryption query to the secret-key predicate encryption scheme. Using the same type of analysis as in [GKW18], we then show that with just a single encryption query, the reduction can leverage the decoder output by the traceability adversary to break security of the underlying predicate encryption scheme. The full analysis is provided in Appendix B.3.

Theorem 4.5 (Correctness). *If H is collision-resistant and Π_{RPE} is correct, then Π_{TR} from Construction 4.4 is correct.*

Theorem 4.6 (Semantic Security). *If Π_{RPE} satisfies broadcast security and message hiding (without encryption queries), then Π_{TR} from Construction 4.4 is semantically secure.*

Theorem 4.7 (Traceability). *If H is collision-resistant and Π_{RPE} satisfies non-adaptive 1-query message hiding security, non-adaptive 1-query function hiding, and non-adaptive 1-query broadcast security, then Π_{TR} is traceable. In particular, the tracing algorithm Trace is efficient.*

4.2 Instantiating the Trace-and-Revoke Scheme

In this section, we describe our instantiation of our resulting trace-and-revoke scheme using the secret-key revocable predicate encryption scheme from Section 3.1 (Construction 3.4, Corollary 3.9). In particular, combining Construction 4.4 with Theorems 4.5 through 4.7 yields the following corollary:

Corollary 4.8 (Identity-Based Trace-and-Revoke from LWE). *Assuming sub-exponential hardness of LWE (with a super-polynomial modulus-to-noise ratio), there exists a fully secure identity-based trace-and-revoke scheme with identity space $\mathcal{ID} = \{0, 1\}^n$ and message space $\mathcal{M} = \{0, 1\}^t$ with the following properties:*

- **Public parameter size:** $|\text{pp}| = n \cdot \text{poly}(\lambda, \log n)$.
- **Secret key size:** The secret key sk_{id} for an identity $\text{id} \in \{0, 1\}^n$ has size $|\text{sk}_{\text{id}}| = n \cdot \text{poly}(\lambda, \log n)$.
- **Ciphertext size:** An encryption $\text{ct}_{m, \mathcal{L}}$ of a message $m \in \{0, 1\}^t$ with respect to a revocation list \mathcal{L} has size $|\text{ct}_{m, \mathcal{L}}| = t + |\mathcal{L}| \cdot \text{poly}(\lambda, \log n)$.

Proof. The claim follows by instantiating Construction 4.4 with the following primitives:

- We can instantiate the collision-resistant hash function H with the standard SIS-based collision-resistant hash function [Ajt96, GGH96]. In this case, the hash key hk has size $|\text{hk}| = \text{poly}(\lambda)$ and the output length of the hash function is also $\ell = \text{poly}(\lambda)$.
- We instantiate the secret-key revocable predicate encryption scheme with broadcast Π_{RPE} with the construction from Corollary 3.9. For $i \in [n]$ and $u \in [0, 2^{\ell+1}]$, the description length z of the functions $f_{i,u} \in \mathcal{F}$ satisfies

$$z = |i| + |u| \leq \log n + \ell + 3 = \text{poly}(\lambda, \log n).$$

Moreover, each function $f_{i,u}$ is computing a comparison on ℓ -bit values and selecting one out of the n components of the vector. This can be computed by a Boolean circuit with depth $d = \text{poly}(\lambda, \log n) - \text{poly}(\lambda)$ for the comparison and $\text{poly}(\log n)$ to select the element to compare. Finally, the identity-space for the underlying revocable predicate encryption scheme is $\mathcal{ID}_0 = [2^{\ell+1}]$ and the attribute space is \mathcal{ID}_0^n .

We now verify the parameter sizes for the resulting construction:

- **Public parameters size:** The public parameters pp consists of the hash key hk and the public parameters rpe.pp for the revocable predicate encryption scheme. Thus,

$$|\text{pp}| = |\text{hk}| + |\text{rpe.pp}| = \text{poly}(\lambda) + n\ell \cdot \text{poly}(\lambda, d, \ell, z) = n \cdot \text{poly}(\lambda, \log n).$$

- **Secret key size:** The secret key sk_{id} for an identity $\text{id} \in \{0, 1\}^n$ consists of a secret key for the underlying revocable predicate encryption scheme. By Corollary 3.9, we have that $|\text{sk}_{\text{id}}| = n\ell + \text{poly}(\lambda, d, \ell, z) = n \cdot \text{poly}(\lambda, \log n)$.
- **Ciphertext size:** The ciphertext $\text{ct}_{m,\mathcal{L}}$ for a message $m \in \{0, 1\}^t$ with respect to a revocation list \mathcal{L} consists of a ciphertext for the underlying revocable predicate encryption scheme. By Corollary 3.9,

$$|\text{ct}_{m,\mathcal{L}}| = t + |\mathcal{L}| \cdot \text{poly}(\lambda, d, \ell, z) = t + |\mathcal{L}| \cdot \text{poly}(\lambda, \log n). \quad \square$$

Acknowledgments

We thank Ahmadreza Rahimi for helpful discussions on this work and the anonymous reviewers for useful suggestions on improving the exposition. S. Kim is supported by NSF, DARPA, a grant from ONR, and the Simons Foundation. D. J. Wu is supported by NSF CNS-1917414 and a University of Virginia SEAS Research Innovation Award. Opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA.

References

- [ABP⁺17] Shweta Agrawal, Sanjay Bhattacharjee, Duong Hieu Phan, Damien Stehlé, and Shota Yamada. Efficient public trace and revoke from standard assumptions: Extended abstract. In *ACM CCS*, pages 2277–2293, 2017.

- [AJS17] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. Indistinguishability obfuscation for turing machines: Constant overhead and amortization. In *CRYPTO*, pages 252–279, 2017.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
- [AS16] Prabhanjan Vijendra Ananth and Amit Sahai. Functional encryption for turing machines. In *TCC*, pages 125–153, 2016.
- [BB04] Dan Boneh and Xavier Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *EUROCRYPT*, pages 223–238, 2004.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, pages 533–556, 2014.
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6:1–6:48, 2012.
- [BLW17] Dan Boneh, Kevin Lewi, and David J. Wu. Constraining pseudorandom functions privately. In *PKC*, pages 494–524, 2017.
- [BN08] Dan Boneh and Moni Naor. Traitor tracing with constant size ciphertext. In *ACM CCS*, pages 501–510, 2008.
- [BP09] Olivier Billet and Duong Hieu Phan. Traitors collaborating in public: Pirates 2.0. In *EUROCRYPT*, pages 189–205, 2009.
- [BSW06] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In *EUROCRYPT*, pages 573–592, 2006.
- [BTVW17] Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained PRFs (and more) from LWE. In *TCC*, pages 264–302, 2017.
- [BW06] Dan Boneh and Brent Waters. A fully collusion resistant broadcast, trace, and revoke system. In *ACM CCS*, pages 211–220, 2006.
- [BW07] Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In *TCC*, pages 535–554, 2007.
- [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained PRFs for nc^1 from LWE. In *EUROCRYPT*, pages 446–476, 2017.
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In *CRYPTO*, pages 257–270, 1994.
- [CFNP00] Benny Chor, Amos Fiat, Moni Naor, and Benny Pinkas. Tracing traitors. *IEEE Trans. Information Theory*, 46(3):893–910, 2000.

- [CHN⁺16] Aloni Cohen, Justin Holmgren, Ryo Nishimaki, Vinod Vaikuntanathan, and Daniel Wichs. Watermarking cryptographic capabilities. In *STOC*, pages 1115–1127, 2016.
- [CVW⁺18a] Yilei Chen, Vinod Vaikuntanathan, Brent Waters, Hoeteck Wee, and Daniel Wichs. Traitor-tracing from LWE made simple and attribute-based. In *TCC*, pages 341–369, 2018.
- [CVW18b] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In *CRYPTO*, pages 577–607, 2018.
- [DF02] Yevgeniy Dodis and Nelly Fazio. Public key broadcast encryption for stateless receivers. In *Security and Privacy in Digital Rights Management*, pages 61–80, 2002.
- [FN93] Amos Fiat and Moni Naor. Broadcast encryption. In *CRYPTO*, pages 480–491, 1993.
- [GGH96] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Collision-free hashing from lattice problems. *IACR Cryptology ePrint Archive*, 1996:9, 1996.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, pages 40–49, 2013.
- [GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions (extended abstract). In *FOCS*, pages 464–479, 1984.
- [GKM⁺19] Rishab Goyal, Sam Kim, Nathan Manohar, Brent Waters, and David J. Wu. Watermarking public-key cryptographic primitives. In *CRYPTO*, 2019.
- [GKP⁺13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *STOC*, pages 555–564, 2013.
- [GKSW10] Sanjam Garg, Abishek Kumarasubramanian, Amit Sahai, and Brent Waters. Building efficient fully collusion-resilient traitor tracing and revocation schemes. In *ACM CCS*, pages 121–130, 2010.
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *FOCS*, pages 612–621, 2017.
- [GKW18] Rishab Goyal, Venkata Koppula, and Brent Waters. Collusion resistant traitor tracing from learning with errors. In *STOC*, pages 660–670, 2018.
- [GKW19] Rishab Goyal, Venkata Koppula, and Brent Waters. New approaches to traitor tracing with embedded identities. In *TCC*, 2019.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM CCS*, pages 89–98, 2006.

- [GQWW19] Rishab Goyal, Willy Quach, Brent Waters, and Daniel Wichs. Broadcast and trace with n^ϵ ciphertext size from standard assumptions. In *CRYPTO*, pages 826–855, 2019.
- [GS18] Sanjam Garg and Akshayaram Srinivasan. A simple construction of io for turing machines. In *TCC*, pages 425–454, 2018.
- [GVW12] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In *CRYPTO*, pages 162–179, 2012.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554, 2013.
- [GVW19] Rishab Goyal, Satyanarayana Vusirikala, and Brent Waters. Collusion resistant broadcast and trace from positional witness encryption. In *PKC*, pages 3–33, 2019.
- [HS02] Dani Halevy and Adi Shamir. The LSD broadcast encryption scheme. In *CRYPTO*, pages 47–60, 2002.
- [KD98] Kaoru Kurosawa and Yvo Desmedt. Optimum traitor tracing and asymmetric schemes. In *EUROCRYPT*, pages 145–157, 1998.
- [KLW15] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In *STOC*, pages 419–428, 2015.
- [KP07] Aggelos Kiayias and Serdar Pehlivanoglu. Pirate evolution: How to make the most of your traitor keys. In *CRYPTO*, pages 448–465, 2007.
- [KSW08] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.
- [KT15] Aggelos Kiayias and Qiang Tang. Traitor deterring schemes: Using bitcoin as collateral for digital content. In *ACM CCS*, pages 231–242, 2015.
- [KW17] Sam Kim and David J. Wu. Watermarking cryptographic functionalities from standard lattice assumptions. In *CRYPTO*, pages 503–536, 2017.
- [KW19] Sam Kim and David J. Wu. Watermarking PRFs from lattices: Stronger security via extractable PRFs. In *CRYPTO*, 2019.
- [LPSS14] San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-lwe and applications in traitor tracing. In *CRYPTO*, pages 315–334, 2014.
- [NNL01] Dalit Naor, Moni Naor, and Jeffery Lotspiech. Revocation and tracing schemes for stateless receivers. In *CRYPTO*, pages 41–62, 2001.
- [NP98] Moni Naor and Benny Pinkas. Threshold traitor tracing. In *CRYPTO*, pages 502–517, 1998.

- [NP00] Moni Naor and Benny Pinkas. Efficient trace and revoke schemes. In *Financial Cryptography*, pages 1–20, 2000.
- [NWZ16] Ryo Nishimaki, Daniel Wichs, and Mark Zhandry. Anonymous traitor tracing: How to embed arbitrary information in a key. In *EUROCRYPT*, pages 388–419, 2016.
- [QWZ18] Willy Quach, Daniel Wichs, and Giorgos Zirdelis. Watermarking PRFs under standard assumptions: Public marking and security with extraction queries. In *TCC*, pages 669–698, 2018.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.
- [SBC⁺07] Elaine Shi, John Bethencourt, Hubert T.-H. Chan, Dawn Xiaodong Song, and Adrian Perrig. Multi-dimensional range query over encrypted data. In *IEEE (S&P)*, pages 350–364, 2007.
- [SS10] Amit Sahai and Hakan Seyalioglu. Worry-free encryption: functional encryption with public keys. In *ACM CCS*, pages 463–472, 2010.
- [SSW01] Jessica Staddon, Douglas R. Stinson, and Ruizhong Wei. Combinatorial properties of frameproof and traceability codes. *IEEE Trans. Information Theory*, 47(3):1042–1049, 2001.
- [SW05] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In *FOCS*, pages 600–611, 2017.

A Analysis of Construction 3.4

In this section, we provide the formal analysis of our secret-key revocable predicate encryption with broadcast scheme from Section 3.1.

A.1 Proof of Theorem 3.5 (Correctness)

Take a function $f \in \mathcal{F}$, an attribute $x \in \mathcal{X}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$ where $x \notin \mathcal{L}$ and $f(x) = 1$. Let $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$, $\text{msk} = (\text{mfe.msk}, \text{abe.msk})$ be the public parameters and the master secret key that are output by $\text{Setup}(1^\lambda)$, and take $\text{sk}_{\text{id},x} \leftarrow \text{KeyGen}(\text{msk}, \text{id}, x)$. By construction, we have $\text{sk}_{\text{id},x} = \{(i, \text{abe.sk}_{i,x})\}_{i \in \mathcal{I}_{\text{id}}}$ where $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$, $\text{abe.sk}_{i,x} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, g_{\text{mfe.msk}_{i,x},i})$, $\text{mfe.sk}_{i,x} \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}_i, x)$, $r_i \leftarrow F(k, i)$, and $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$. We now show each of the requirements individually:

- **Broadcast correctness:** Let $\text{ct}_{m,\mathcal{L}} \leftarrow \text{Broadcast}(\text{pp}, m, \mathcal{L})$. This means that $\text{ct}_{m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.PKEnc}(\text{mfe.pp})$. Since $x \notin \mathcal{L}$, by correctness of Π_{SC} , there exists an index $i \in \mathcal{I}_{\text{id}} \cap \mathcal{J}_{\mathcal{L}}$. By correctness of Π_{MFE} , with overwhelming probability, $\text{MFE.Dec}(\text{mfe.sk}_{i,x}, \text{mfe.ct}_i) = 1$. This means that $g_{\text{mfe.msk}_{i,x},i}(\text{mfe.ct}_i, i) = 1$ and hence, by correctness of Π_{ABE} , with overwhelming probability, $\text{ABE.Dec}(\text{abe.sk}_{i,x}, \text{abe.ct}_i) = m$.

- **Encryption correctness:** Let $\text{ct}_{f,m,\mathcal{L}} \leftarrow \text{Enc}(\text{msk}, f, m, \mathcal{L})$. This means that $\text{ct}_{f,m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$, where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f)$. Since $x \notin \mathcal{L}$, there exists $i \in \mathcal{I}_{\text{id}} \cap \mathcal{J}_{\mathcal{L}}$. By correctness of Π_{MFE} , $\text{MFE.Dec}(\text{mfe.sk}_{i,x}, \text{mfe.ct}_i) = f(x) = 1$ with overwhelming probability, and $g_{\text{mfe.msk}_{i,x}, i}(\text{mfe.ct}_i, i) = 1$. Then, by correctness of Π_{ABE} , $\text{ABE.Dec}(\text{abe.sk}_{i,x}, \text{abe.ct}_i) = m$ with overwhelming probability. \square

A.2 Proof of Theorem 3.6 (Message Hiding)

First, the only difference between $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 0]$ and $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 1]$ is in how the challenge ciphertext is generated. Let $(f, m_0, m_1, \mathcal{L})$ be a challenge query that the adversary \mathcal{A} makes in the message hiding game. In both experiments, the challenge ciphertext $\text{ct}_b = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ consists of $|\mathcal{J}_{\mathcal{L}}|$ ABE ciphertexts abe.ct_i , where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. In $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 0]$, each ciphertext abe.ct_i is an encryption of m_0 , while in $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 1]$, abe.ct_i is an encryption of m_1 . To argue that the ABE ciphertexts in the two experiments are computationally indistinguishable, we proceed via a sequence of intermediate hybrid experiments where in each hybrid, we switch a single ciphertext abe.ct_i from an ABE encryption of m_0 to an ABE encryption of m_1 .

First, let $Q = Q(\lambda)$ be a bound on the size of the cover $\mathcal{J}_{\mathcal{L}} \subseteq [K]$. Such a bound exists because $|\mathcal{L}| = \text{poly}(\lambda)$ since \mathcal{A} has to output \mathcal{L} , and so we can always bound $|\mathcal{L}|$ by the running time of \mathcal{A} , which is assumed to be $\text{poly}(\lambda)$. Next, since Π_{SC} is efficient, $|\mathcal{J}_{\mathcal{L}}| \leq \text{poly}(|\mathcal{L}|, n) = \text{poly}(\lambda)$. Now, for $j \in [0, Q]$, we define a sequence of hybrid experiments as follows:

- **Hyb_j:** This is the real message-hiding experiment $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 0]$ except in how the challenger responds to the challenge query. Specifically, when the adversary \mathcal{A} makes its challenge query $(f, m_0, m_1, \mathcal{L})$, the challenger proceeds as follows:
 1. Let $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L}) \subseteq [K]$. Write $\mathcal{J}_{\mathcal{L}} = \{i_1, \dots, i_t\}$ where $i_1 < i_2 < \dots < i_t$ and $t \leq Q$.
 2. Generate the mixed FE ciphertexts mfe.ct_i for all $i \in \mathcal{J}_{\mathcal{L}}$ exactly as in $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 0]$.
 3. For $i \leq i_j$, the challenger sets $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m_1)$. For $i > i_j$, the challenger sets $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m_0)$.
 4. It gives $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ to the adversary.

We write $\text{Hyb}_j(\mathcal{A})$ to denote the output of Hyb_j with \mathcal{A} . By definition, $\text{Hyb}_0 \equiv \text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 0]$ and $\text{Hyb}_Q \equiv \text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, 1]$. We now show that by semantic security of Π_{ABE} , each consecutive pair of hybrids Hyb_{j-1} and Hyb_j for $j \in [Q]$ are computationally indistinguishable.

Lemma A.1. *If Π_{ABE} is semantically secure, then for all efficient adversaries \mathcal{A} and $j \in [Q]$,*

$$|\Pr[\text{Hyb}_{j-1}(\mathcal{A}) = 1] - \Pr[\text{Hyb}_j(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Let \mathcal{A} be an adversary that distinguishes Hyb_{j-1} and Hyb_j . We use \mathcal{A} to construct an algorithm \mathcal{B} that breaks semantic security of Π_{ABE} . Algorithm \mathcal{B} works as follows:

- **Setup phase:** First, \mathcal{B} receives the public parameters abe.pp from the ABE challenger. It generates $\text{mfe.pp} \leftarrow \text{MFE.PrmsGen}(1^{\lambda'})$, $k \xleftarrow{\text{R}} \mathcal{K}$, where $\lambda' = \max(\lambda, (\log K)^{2/\varepsilon})$, and gives $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$ to \mathcal{A} .

- **Query phase:** Algorithm \mathcal{B} responds to each of \mathcal{A} 's oracle queries as follows:
 - **Key-generation oracle:** On input an identity $\text{id} \in \mathcal{ID}$ and an attribute $x \in \mathcal{X}$, algorithm \mathcal{B} first computes $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$. Then, for each $i \in \mathcal{I}_{\text{id}}$, it computes $r_i \leftarrow F(k, i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.sk}_{i,x} \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}_i, x)$. It then submits a key-generation query to the ABE challenger on the function $g_{\text{mfe.msk}_{i,x}, i}$ to obtain an ABE secret key $\text{abe.sk}_{i,x}$. Finally, it replies to \mathcal{A} with the keys $\text{sk}_{\text{id},x} = \{(i, \text{abe.sk}_{i,x})\}_{i \in \mathcal{I}_{\text{id}}}$.
 - **Encryption oracle:** On input a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, algorithm \mathcal{B} computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. For each $i \in \mathcal{J}_{\mathcal{L}}$, it computes $r_i \leftarrow F(k, i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f)$, and $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$. It responds with the ciphertext $\text{ct}_{f,m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$.
 - **Challenge oracle:** On input a function $f \in \mathcal{F}$, two messages $m_0, m_1 \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, algorithm \mathcal{B} computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L}) \subseteq [K]$. Write $\mathcal{J}_{\mathcal{L}} = \{i_1, \dots, i_t\}$ where $i_1 < i_2 < \dots < i_t$. For each $i \in \mathcal{J}_{\mathcal{L}}$, it generates $r_i \leftarrow F(k, i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f)$. It then generates the ABE ciphertexts as follows:
 - * For $i < i_j$, algorithm \mathcal{B} sets $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m_1)$.
 - * For $i = i_j$, algorithm \mathcal{B} makes a challenge query with attribute $(\text{mfe.ct}_i, i)$ and messages m_0, m_1 to receive a ciphertext abe.ct_i .
 - * For $i > i_j$, algorithm \mathcal{B} sets $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m_0)$.
It replies to \mathcal{A} with the ciphertext $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$.

- **Output phase:** At the end of the experiment, algorithm \mathcal{B} outputs whatever \mathcal{A} outputs.

We now show that \mathcal{B} perfectly simulates either Hyb_{j-1} or Hyb_j depending on whether it is interacting in $\text{ExptABE}[\lambda, \mathcal{B}, 0]$ or $\text{ExptABE}[\lambda, \mathcal{B}, 1]$. Furthermore, we show that as long as \mathcal{A} is admissible for $\text{ExptRPE}_{\text{MH}}[\lambda, \mathcal{A}, b]$, algorithm \mathcal{B} is admissible for $\text{ExptABE}[\lambda, \mathcal{B}, b]$. These two conditions show that algorithm \mathcal{B} breaks semantic security of Π_{ABE} with the same advantage of \mathcal{A} for distinguishing Hyb_{j-1} and Hyb_j .

Admissibility condition. Let $(f, m_0, m_1, \mathcal{L})$ be the challenge query made by \mathcal{A} . In response, algorithm \mathcal{B} submits $((\text{mfe.ct}_{i_j}, i_j), m_0, m_1)$ as its challenge query to the ABE challenger, where i_j is the j^{th} smallest value in $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$ and $\text{mfe.ct}_{i_j} \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_{i_j}, f)$. To show that \mathcal{B} is admissible, we must show that for each key-generation query $g_{\text{mfe.sk}_{i,x}, i}$ that \mathcal{B} makes, $g_{\text{mfe.sk}_{i,x}, i}(\text{mfe.ct}_{i_j}, i_j) = 0$. First, we note that \mathcal{B} only makes key-generation queries when \mathcal{A} makes a key-generation query. Suppose \mathcal{A} makes a key-generation query on a pair (id, x) . Then, algorithm \mathcal{B} will issue key-generation queries on functions $g_{\text{mfe.msk}_{i,x}, i}$ for all $i \in \mathcal{I}_{\text{id}}$ where $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$. By admissibility of \mathcal{A} , either $\text{id} \in \mathcal{L}$ or $f(x) = 0$. We consider these two cases:

- Suppose $\text{id} \in \mathcal{L}$. By correctness of Π_{SC} , this means that $\mathcal{J}_{\mathcal{L}} \cap \mathcal{I}_{\text{id}} = \emptyset$. Since $i_j \in \mathcal{J}_{\mathcal{L}}$, this means, that $i_j \notin \mathcal{I}_{\text{id}}$, and correspondingly, $i \neq i_j$ and so $g_{\text{mfe.msk}_{i,x}, i}(\text{mfe.ct}_{i_j}, i_j) = 0$.
- Suppose $f(x) = 0$. For all $i \neq i_j$, we have that $g_{\text{mfe.msk}_{i,x}, i}(\text{mfe.ct}_{i_j}, i_j) = 0$, so it suffices to only consider the case where $i = i_j$. In this case, by correctness of Π_{MFE} , we have that

$\text{MFE.Dec}(\text{mfe.msk}_{i_j,x}, \text{mfe.ct}_{i_j}) = f(x) = 0$ with overwhelming probability, and once again, admissibility holds.

We conclude that if \mathcal{A} is admissible, then with overwhelming probability, all of \mathcal{B} 's key-generation queries are also admissible.

Correctness of the simulation. The public parameters pp that \mathcal{B} provides to \mathcal{A} during the setup phase of the experiment is distributed exactly as in Hyb_{j-1} and Hyb_j . We now consider how \mathcal{B} simulates the responses to \mathcal{A} 's oracle queries:

- **Key-generation oracle:** Let (id, x) be a key-generation query and let $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$. For each $i \in \mathcal{I}_{\text{id}}$ algorithm \mathcal{B} generates the mixed FE decryption keys $\text{mfe.sk}_{i,x}$ for all $i \in \mathcal{I}_{\text{id}}$ exactly as described in Hyb_{j-1} and Hyb_j . To construct the ABE decryption keys, the challenger submits the function $g_{\text{mfe.sk}_{i,x},i}$ for each $i \in \mathcal{I}_{\text{id}}$ to the ABE challenger to receive keys $\text{abe.sk}_{i,x} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, g_{\text{mfe.msk}_{i,x},i})$, where abe.msk is the ABE secret key sampled by the ABE challenger (and unknown to \mathcal{B}). This is precisely the distribution of secret keys that would be output in Hyb_{j-1} and Hyb_j .
- **Encryption oracle:** Since the ABE scheme is a public-key encryption scheme, \mathcal{B} perfectly simulates the encryption queries as in Hyb_{j-1} and Hyb_j .
- **Challenge oracle:** Let $(f, m_0, m_1, \mathcal{L})$ be the challenge query made by \mathcal{A} . Let $\mathcal{J}_{\mathcal{L}} = \{i_1, \dots, i_t\}$ be the cover output by $\text{ComputeCover}(\mathcal{L})$. By construction, algorithm \mathcal{B} generates the ciphertexts abe.ct_i for $i \neq i_j$ exactly as described in Hyb_{j-1} and Hyb_j . For abe.ct_{i_j} , algorithm \mathcal{B} submits $((\text{mfe.ct}_{i_j}, i_j), m_0, m_1)$ to the ABE challenger to obtain a ciphertext abe.ct_{i_j} . If the ABE challenger replies with $\text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_{i_j}, i_j), m_0)$, then \mathcal{B} perfectly simulates Hyb_{j-1} while if the challenger replies with $\text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_{i_j}, i_j), m_1)$, then \mathcal{B} perfectly simulates the distribution in Hyb_j .

We conclude that depending on the challenge bit b for the ABE security game, algorithm \mathcal{B} either perfectly simulates Hyb_{j-1} or Hyb_j for \mathcal{A} . The lemma follows. \square

Theorem 3.6 now follows by a standard hybrid argument. \square

A.3 Proof of Theorem 3.7 (Function Hiding)

Proof overview. By construction, an encryption of a function f , a message m , and a revocation list L consist of $|\mathcal{J}_{\mathcal{L}}|$ ABE ciphertexts $\{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. Each of these ABE ciphertexts is an encryption of the message m with attribute mfe.ct_i where mfe.ct_i is a mixed FE encryption of the function f (under the mixed FE master secret key associated with index $i \in [K]$). Therefore, to argue that a ciphertext for function f_0 is indistinguishable from a ciphertext for function f_1 , a natural idea is to consider a sequence of $|\mathcal{J}_{\mathcal{L}}|$ hybrid arguments, where in each hybrid, we replace one of the mixed FE ciphertexts mfe.ct_i encrypting f_0 to one encrypting f_1 . The problem is that the adversary can choose the challenge revocation list \mathcal{L} adaptively *after* it has made a number of key generation and encryption queries. Thus, the reduction does not know in advance the set $\mathcal{J}_{\mathcal{L}} \subseteq [K]$, and correspondingly, it does not know which of the K mixed FE instances will appear in the challenge ciphertext. One approach is to have the reduction algorithm guess the index $i \in [K]$ that needs to be switched in each hybrid. This will incur a loss of $K^{|\mathcal{J}_{\mathcal{L}}|} = K^{\text{poly}(\lambda)}$ in

the security reduction. Alternatively, we can consider a sequence of K hybrid experiments where in hybrid Hyb_j , we set mfe.ct_j to be an encryption to f_0 whenever $i \leq j$ and an encryption to f_1 whenever $i > j$. This incurs a loss of K in the security reduction; when K is super-polynomial in the security parameter, we require sub-exponential security for the underlying mixed FE scheme.

Note that while we can consider a selective notion of security where the adversary must first commit to its revocation list, selective security does *not* imply adaptive security in this setting, since the size of the revocation list is not a priori bounded. This means that the reduction cannot guess the revocation list at the beginning. For this reason, we directly argue adaptive security via a sub-exponential security reduction.

Proof. We now describe the sequence of hybrid arguments we use in our analysis.

- **Hyb₀:** This is the real function-hiding experiment $\text{ExptRPE}_{\text{FH}}[\lambda, \mathcal{A}, 0]$. Namely, the challenger begins by sampling $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, where $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$ and $\text{msk} = (\text{pp}, \text{abe.msk}, k)$ and gives pp to \mathcal{A} . The challenger then responds to oracle queries as follows:
 - **Key-generation oracle:** On input an identity $\text{id} \in \mathcal{ID}$ and attribute $x \in \mathcal{X}$, the challenger computes $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$. For each $i \in \mathcal{I}_{\text{id}}$, the challenger constructs $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$ where $r_i \leftarrow F(k, i)$ and $\text{mfe.sk}_{i,x} \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}_i, x)$. It computes $\text{abe.sk}_{i,x} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, g_{\text{mfe.msk}_i, x, i})$ and replies with $\text{sk}_{\text{id}, x} = \{(i, \text{abe.sk}_{i,x})\}_{i \in \mathcal{I}_{\text{id}}}$.
 - **Encryption oracle:** On input a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger replies with $\text{ct}_{f,m,\mathcal{L}} \leftarrow \text{Enc}(\text{msk}, f, m, \mathcal{L})$. Namely, $\text{ct}_{f,m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$, where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$, $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f)$, and $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; F(k, i))$.
 - **Challenge oracle:** On input two functions $f_0, f_1 \in \mathcal{F}$, a message $m \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger replies with $\text{ct}^* \leftarrow \text{Enc}(\text{msk}, f_0, m, \mathcal{L})$, where $\text{ct}^* = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$, $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_0)$, and $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; F(k, i))$.

At the end, the adversary outputs a bit $b' \in \{0, 1\}$, which is also the output of the experiment.

- **Hyb₁:** Same as Hyb_0 except the challenger samples a truly random function $f \xleftarrow{\text{R}} \text{Funs}[[K], \{0, 1\}^\rho]$ at the beginning of the experiments and evaluates $f(\cdot)$ in place of $F(k, \cdot)$.
- **Hyb₂:** Same as Hyb_1 , except the challenger constructs the challenge ciphertext as $\text{ct}^* \leftarrow \text{Enc}(\text{msk}, f_1, m, \mathcal{L})$. This procedure is the same as in Hyb_1 , except the challenge now uses $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_1)$.
- **Hyb₃:** This is the function-hiding experiment $\text{ExptRPE}_{\text{FH}}[\lambda, \mathcal{A}, 1]$.

For an adversary \mathcal{A} , we write $\text{Hyb}_j(\mathcal{A})$ to denote the output of Hyb_j with \mathcal{A} . We now show that the output distributions of each consecutive pair of hybrid experiments are computationally indistinguishable.

Lemma A.2. *If F is a secure PRF, then for all efficient adversaries \mathcal{A} ,*

$$|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. The lemma follows immediately from the definition of PRF security. \square

Lemma A.3. *Suppose that Π_{MFE} satisfies sub-exponential non-adaptive q -query (resp., adaptive) semantic security. Specifically, suppose that the advantage of any adversary running in time $\text{poly}(\lambda)$ in the semantic security game is bounded by $2^{-\Omega(\lambda^\epsilon)}$. Suppose also that Π_{SC} is correct. Then for all efficient non-adaptive q -query (resp., adaptive) adversaries \mathcal{A} ,*

$$|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. As discussed above, we will define a collection of $K + 1$ hybrid experiments where in the j^{th} hybrid experiment, we set mfe.ct_i in the challenge ciphertext to be an encryption of f_0 if $i \leq j$ and an encryption of f_1 if $i > j$. We rely on semantic security of Π_{RPE} to argue indistinguishability of each adjacent pair of hybrids. We begin by defining our sequence of intermediate experiments:

- **Hyb $_{1,j}$:** For each $j \in [0, K]$, the experiment $\text{Hyb}_{1,j}$ is identical to Hyb_1 except in the way the challenger generates the challenge ciphertext. Specifically, when the adversary \mathcal{A} makes its challenge query $(f_0, f_1, m, \mathcal{L})$, the challenger proceeds as follows:
 1. Let $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L}) \subseteq [K]$.
 2. For each $i \in \mathcal{J}_{\mathcal{L}}$, the challenger computes $r_i \leftarrow f(i)$, and $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$. It sets mfe.ct_i as follows:
 - If $i \leq j$, set $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_1)$,
 - If $i > j$, set $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_0)$.
 3. For each $i \in \mathcal{J}_{\mathcal{L}}$, the challenger sets $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$.
 4. It gives $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ to the adversary.

By construction, $\text{Hyb}_{1,0} \equiv \text{Hyb}_1$ and $\text{Hyb}_{1,K} \equiv \text{Hyb}_2$. We now show that the output distributions of each adjacent pair of hybrid experiments $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$ are computationally indistinguishable. Let \mathcal{A} be an adversary that distinguishes $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. We use \mathcal{A} to construct an adversary \mathcal{B} that breaks semantic security of Π_{MFE} :

- **Setup phase:** At the beginning of the game, the mixed FE challenger samples $\text{mfe.pp} \leftarrow \text{MFE.PrmsGen}(1^{\lambda'})$ where $\lambda' = \max(\lambda, (\log K)^{2/\epsilon})$ and gives mfe.pp to \mathcal{B} . The challenger will also sample a key $\text{mfe.msk} \leftarrow \text{MFE.MSKGen}(\text{pp})$, which it keeps to itself. Algorithm \mathcal{B} then generates $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda)$ and gives $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$ to \mathcal{A} . In the following description, algorithm \mathcal{B} will *lazily sample* the values of $f: [K] \rightarrow \{0, 1\}^\rho$. Namely, the first time \mathcal{B} needs to compute $f(i)$ on some $i \in [K]$, algorithm \mathcal{B} samples a random value $r_i \xleftarrow{\mathcal{R}} \{0, 1\}^\rho$ as the value for $f(i)$ and will use this value for $f(i)$ thereafter.
- **Query phase:** Algorithm \mathcal{B} responds to each of \mathcal{A} 's oracle queries as follows:
 - **Key-generation oracle:** On input an identity $\text{id} \in \mathcal{ID}$ and an attribute $x \in \mathcal{X}$, algorithm \mathcal{B} computes $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$. Then, for each $i \in \mathcal{I}_{\text{id}}$, it generates $\text{mfe.sk}_{i,x}$ as follows:
 - * If $i \neq j$, algorithm \mathcal{B} sets $r_i \leftarrow f(i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.sk}_{i,x} \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}_i, x)$.
 - * If $i = j$, algorithm \mathcal{B} makes a key-generation query on x to obtain $\text{mfe.sk}_{i,x}$.

Then, for each $i \in \mathcal{I}_{\text{id}}$, algorithm \mathcal{B} computes $\text{abe.sk}_{i,x} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, g_{\text{msk}_{i,x},i})$. It replies to \mathcal{A} with $\text{sk}_{\text{id},x} = \{(i, \text{abe.sk}_{i,x})\}_{i \in \mathcal{I}_{\text{id}}}$.

- **Encryption oracle:** On input a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, algorithm \mathcal{B} first computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. Then, for each $i \in \mathcal{J}_{\mathcal{L}}$, it generates the mixed FE ciphertext mfe.ct_i as follows:

- * If $i \neq j$, algorithm \mathcal{B} computes $r_i \leftarrow f(i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}_i, x)$.
- * If $i = j$, algorithm \mathcal{B} makes an encryption query on f to obtain mfe.ct_i .

Then, for each $i \in \mathcal{J}_{\mathcal{L}}$, algorithm \mathcal{B} computes $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$. It replies to the adversary with $\text{ct}_{f,m,\mathcal{L}} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$.

- **Challenge oracle:** On input two functions $f_0, f_1 \in \mathcal{F}$, a message $m \in \mathcal{M}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, algorithm \mathcal{B} computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L}) \subseteq [K]$. For each $i \in \mathcal{J}_{\mathcal{L}}$, it generates a mixed FE ciphertext mfe.ct_i as follows:

- * If $i < j$, algorithm \mathcal{B} sets $r_i \leftarrow f(i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_1)$.
- * If $i = j$, algorithm \mathcal{B} makes a challenge query on the pair (f_0, f_1) to receive mfe.ct_i .
- * If $i > j$, algorithm \mathcal{B} sets $r_i \leftarrow f(i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_0)$.

It then computes $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$ for all $i \in \mathcal{J}_{\mathcal{L}}$, and gives $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ to \mathcal{A} .

- **Output phase:** At the end of the experiment, algorithm \mathcal{B} outputs whatever \mathcal{A} outputs.

We show that \mathcal{B} perfectly simulates either $\text{Hyb}_{1,j-1}$ or $\text{Hyb}_{1,j}$ depending on whether it is interacting according to $\text{ExptMFE}_{\text{SS}}[\lambda, \mathcal{B}, 0]$ or $\text{ExptMFE}_{\text{SS}}[\lambda, \mathcal{B}, 1]$. Furthermore, we show that as long as \mathcal{A} is admissible for $\text{ExptRPE}_{\text{FH}}[\lambda, \mathcal{A}, b]$, algorithm \mathcal{B} is admissible for $\text{ExptMFE}_{\text{SS}}[\lambda, \mathcal{B}, b]$. These two conditions show that algorithm \mathcal{B} breaks semantic security of Π_{MFE} with the same advantage of \mathcal{A} for distinguishing $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$.

Admissibility condition. Since \mathcal{A} makes at most one challenge query, the same is true for \mathcal{B} . Let $(f_0, f_1, m, \mathcal{L})$ be the challenge query made by \mathcal{A} . In response, algorithm \mathcal{B} computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, and if $j \in \mathcal{J}_{\mathcal{L}}$, then it makes a challenge query (f_0, f_1) to the mixed FE challenger. If $j \notin \mathcal{J}_{\mathcal{L}}$, then \mathcal{B} does not make any challenge query and thus, is trivially admissible.

Consider the case where $j \in \mathcal{J}_{\mathcal{L}}$. To show that \mathcal{B} is admissible, we must show that for each key-generation query $x \in \mathcal{X}$ that \mathcal{B} makes, $f_0(x) = f_1(x)$. First we note that \mathcal{B} makes a key-generation query only when \mathcal{A} makes a key-generation query. Suppose that \mathcal{A} makes a key-generation query on a pair (id, x) . Then, algorithm \mathcal{B} will compute $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$ and if $j \in \mathcal{I}_{\text{id}}$, it will make a key-generation query $x \in \mathcal{X}$ to the mixed FE challenger. Since \mathcal{A} is admissible for the semantic security game, either $\text{id} \in \mathcal{L}$ or $f_0(x) = f_1(x)$. We consider these two cases:

- Suppose that $\text{id} \in \mathcal{L}$. Then, by correctness of Π_{SC} , we have $\mathcal{J}_{\mathcal{L}} \cap \mathcal{I}_{\text{id}} = \emptyset$. Since $j \in \mathcal{J}_{\mathcal{L}}$ by assumption, we have $j \notin \mathcal{I}_{\text{id}}$, and therefore, \mathcal{B} does not make a key-generation query on x to the mixed FE challenger.
- Suppose that $f_0(x) = f_1(x)$. In this case, submitting $x \in \mathcal{X}$ to the mixed FE challenger does not violate admissibility.

We conclude that if \mathcal{A} is admissible, then \mathcal{B} is also admissible. In addition, algorithm \mathcal{B} only makes encryption queries when \mathcal{A} makes an encryption query, and it makes at most 1 encryption query in response to each of \mathcal{A} 's encryption queries. Thus, if \mathcal{A} is a non-adaptive q -query adversary, \mathcal{B} is also a non-adaptive q -query adversary.

Correctness of the simulation. The public parameters pp that \mathcal{B} provides to \mathcal{A} during the setup phase of the experiment is distributed exactly as in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. We now consider how \mathcal{B} simulates the responses to \mathcal{A} 's oracle queries:

- **Key-generation oracle:** Let (id, x) be a key-generation query and let $\mathcal{I}_{\text{id}} \leftarrow \text{Encode}(\text{id})$. For each $i \in \mathcal{I}_{\text{id}} \setminus \{j\}$, algorithm \mathcal{B} generates mixed FE decryption keys $\text{mfe.sk}_{i,x}$ exactly as described in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. For $\text{mfe.sk}_{j,x}$, algorithm \mathcal{B} submits x to the mixed FE challenger to receive $\text{mfe.sk}_{j,x} \leftarrow \text{MFE.KeyGen}(\text{mfe.msk}, x)$, where mfe.msk is the mixed FE secret key sampled by the mixed FE challenge (and unknown to \mathcal{B}). In this case, the master secret key mfe.msk sampled by the challenger plays the role of mfe.msk_j (and is correctly distributed since in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$, the key mfe.msk_j is sampled by running $\text{MFE.MSKGen}(\text{mfe.pp})$ with uniform and independent randomness $f(j)$, which is precisely the same distribution the challenger uses to sample mfe.msk). Finally, the ABE decryption keys are constructed exactly as in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$.
- **Encryption oracle:** Let (f, m, \mathcal{L}) be an encryption query. Algorithm \mathcal{B} constructs $\mathcal{J}_{\mathcal{L}}$ exactly as in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. Similarly, all of the ciphertext components mfe.ct_i for $i \in \mathcal{J}_{\mathcal{L}} \setminus \{j\}$ are correctly constructed. To simulate mfe.ct_j (if $j \in \mathcal{J}_{\mathcal{L}}$), algorithm \mathcal{B} submits a key-generation query on function f to receive from the challenger $\text{mfe.ct}_j \leftarrow \text{MFE.SKEnc}(\text{mfe.sk}, f)$, where the challenger's key mfe.msk plays the role of mfe.msk_j in the simulation. Finally, the ABE ciphertexts are computed exactly as in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$, so we conclude that \mathcal{B} correctly simulates the encryption queries.
- **Challenge oracle:** Let $(f_0, f_1, m, \mathcal{L})$ be the challenge query made by \mathcal{A} , and let $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. For each $i \in \mathcal{I}_{\text{id}} \setminus \{j\}$, algorithm \mathcal{B} generates mixed FE ciphertexts mfe.ct_i exactly as described in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. For mfe.ct_j , algorithm \mathcal{B} submits (f_0, f_1) to the mixed FE challenger to receive mfe.ct_j . If the mixed FE challenger replies with $\text{MFE.SKEnc}(\text{mfe.msk}, f_0)$, then \mathcal{B} perfectly simulates mfe.ct_j in $\text{Hyb}_{1,j-1}$ while if the challenger replies with $\text{MFE.SKEnc}(\text{mfe.msk}, f_1)$, then \mathcal{B} perfectly simulates mfe.ct_j in $\text{Hyb}_{1,j}$. Then, as \mathcal{B} generates the ABE ciphertexts exactly as in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$, it perfectly simulates either $\text{Hyb}_{1,j-1}$ or $\text{Hyb}_{1,j}$.

We conclude that \mathcal{B} breaks semantic security of Π_{MFE} with the same advantage \mathcal{A} has for distinguishing between hybrids $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. Thus, by a hybrid argument, if there exists an efficient adversary \mathcal{A} that can distinguish between Hyb_1 and Hyb_2 with non-negligible advantage $\varepsilon_0 = 1/\text{poly}(\lambda)$, then there exists an efficient adversary \mathcal{B} that runs in time $\text{poly}(\lambda) = \text{poly}(\lambda')$ that breaks semantic security of Π_{MFE} with advantage at least ε_0/K . By construction, $K \leq 2^{(\lambda')^{\varepsilon/2}}$, and so \mathcal{B} succeeds with advantage at least

$$\varepsilon_0/2^{(\lambda')^{\varepsilon/2}} = 1/\text{poly}(\lambda') \cdot 2^{-(\lambda')^{\varepsilon/2}} = 2^{-O((\lambda')^{\varepsilon/2})},$$

which contradicts the sub-exponential hardness of Π_{MFE} . \square

Lemma A.4. *If F is a secure PRF, then for all efficient adversaries \mathcal{A} ,*

$$|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. The only difference between hybrids Hyb_2 and Hyb_3 is that the adversary uses $f(\cdot)$ in place of $F(k, \cdot)$ when responding to the adversary's queries. The lemma follows immediately from the definition of PRF security. \square

Theorem 3.7 now follows by a standard hybrid argument. \square

A.4 Proof of Theorem 3.8

For similar reasons as those described in the proof of Theorem 3.7, we rely on sub-exponential hardness of mixed FE in our analysis. The analysis here proceeds very similarly as the proof of Theorem 3.7 (Appendix A.3), and many steps are in fact identical. We present the proof in full for completeness. We start by defining a sequence of hybrid experiments:

- **Hyb₀:** This is the broadcast security experiment $\text{ExptRPE}_{\text{BC}}[\lambda, \mathcal{A}, 0]$. Namely, the challenger begins by sampling $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, where $\text{pp} = (\text{mfe.pp}, \text{abe.pp})$ and $\text{msk} = (\text{pp}, \text{abe.msk}, k)$ and gives pp to \mathcal{A} . The challenger then responds to oracle queries as follows:
 - **Key-generation oracle:** On input an identity $\text{id} \in \mathcal{ID}$ and attribute $x \in \mathcal{X}$, the challenger replies with $\text{sk}_{\text{id}, x} \leftarrow \text{KeyGen}(\text{msk}, \text{id}, x)$.
 - **Encryption oracle:** On input a function $f \in \mathcal{F}$, a message $m \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger replies with $\text{ct}_{f, m, \mathcal{L}} \leftarrow \text{Enc}(\text{msk}, f, m, \mathcal{L})$.
 - **Challenge oracle:** On input a message $m \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, the challenger replies with $\text{ct}^* \leftarrow \text{Broadcast}(\text{pp}, m, \mathcal{L})$, where $\text{ct}^* = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ where $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.PKEnc}(\text{mfe.pp})$.

At the end, the adversary outputs a bit $b' \in \{0, 1\}$, which is also the output of the experiment.

- **Hyb₁:** Same as Hyb_0 , except the challenger samples a truly random function $f \xleftarrow{\text{R}} \text{Funs}[[K], \{0, 1\}^\rho]$ at the beginning of the experiment and evaluates $f(\cdot)$ in place of $F(k, \cdot)$.
- **Hyb₂:** Same as Hyb_1 , except the challenger responds to the challenge query with $\text{ct}^* \leftarrow \text{Enc}(\text{msk}, f_{\text{accept}}, m, \mathcal{L})$. Specifically, the challenger computes $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f)$ where $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}, r_i)$ and $r_i \leftarrow f(i)$ for each $i \in \mathcal{J}_{\mathcal{L}}$ and sets $\text{ct}^* = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ where $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$.
- **Hyb₃:** This is the broadcast security experiment $\text{ExptRPE}_{\text{BC}}[\lambda, \mathcal{A}, 1]$.

For an adversary \mathcal{A} , we write $\text{Hyb}_j(\mathcal{A})$ to denote the output of Hyb_j with \mathcal{A} . We now show that the output distributions of each consecutive pair of hybrid experiments are computationally indistinguishable.

Lemma A.5. *If F is a secure PRF, then for all efficient adversaries \mathcal{A} ,*

$$|\Pr[\text{Hyb}_0(\mathcal{A}) = 1] - \Pr[\text{Hyb}_1(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. The lemma follows immediately from the definition of PRF security. \square

Lemma A.6. *Suppose Π_{MFE} satisfies sub-exponential non-adaptive q -query (resp., adaptive) public/secret key indistinguishability. Specifically, suppose that the advantage of any adversary running in time $\text{poly}(\lambda)$ in the public/secret key indistinguishability game is bounded by $2^{-\Omega(\lambda^\epsilon)}$. Then, for all efficient non-adaptive q -query (resp., adaptive) adversaries \mathcal{A} ,*

$$|\Pr[\text{Hyb}_1(\mathcal{A}) = 1] - \Pr[\text{Hyb}_2(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. Similar to the proof of Theorem 3.7, we will define $K + 1$ experiments where in the j^{th} hybrid, we use the public encryption algorithm of Π_{MFE} to encrypt ciphertexts to indices $j' \leq j$ and we use the secret encryption algorithm of Π_{MFE} to encrypt ciphertexts to indices $j' > j$. We then rely on public/secret key indistinguishability to show that each pair of intermediate experiments are computationally indistinguishable. We begin by defining our sequence of intermediate experiments:

- **Hyb $_{1,j}$:** For each $j \in [0, K]$, the experiment $\text{Hyb}_{1,j}$ is identical to Hyb_1 except for the way the challenger generates the challenge ciphertext. Specifically, when the adversary \mathcal{A} makes its challenge query (m, \mathcal{L}) , the challenger proceeds as follows:
 1. Let $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L}) \subseteq [K]$.
 2. For each $i \in \mathcal{J}_{\mathcal{L}}$, the challenger computes $r_i \leftarrow f(i)$, and $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$. It sets mfe.ct_i as follows:
 - If $i \leq j$, set $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_{\text{accept}})$,
 - If $i > j$, set $\text{mfe.ct}_i \leftarrow \text{MFE.PKEnc}(\text{mfe.pp})$.
 3. For each $i \in \mathcal{J}_{\mathcal{L}}$, the challenger sets $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$.
 4. It gives $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ to the adversary.

By construction, $\text{Hyb}_{1,0} \equiv \text{Hyb}_1$ and $\text{Hyb}_{1,K} \equiv \text{Hyb}_2$. We now show that the output distributions of each adjacent pair of hybrid experiments $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$ are computationally indistinguishable. Let \mathcal{A} be an adversary that distinguishes $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. We construct an algorithm \mathcal{B} that breaks public/secret key indistinguishability of Π_{MFE} . Algorithm \mathcal{B} works as follows:

- **Setup phase:** Same as in the proof of Lemma A.3.
- **Query phase:** Algorithm \mathcal{B} responds to each of \mathcal{A} 's oracle queries as follows:
 - **Key-generation oracle:** Same as in the proof of Lemma A.3.
 - **Encryption oracle:** Same as in the proof of Lemma A.3.
 - **Challenge oracle:** On input a message $m \in \mathcal{M}$ and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, algorithm \mathcal{B} computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L}) \subseteq [K]$. For each $i \in \mathcal{J}_{\mathcal{L}}$, it generates the mixed FE ciphertext mfe.ct_i as follows:
 - * If $i < j$, algorithm \mathcal{B} sets $r_i \leftarrow f(i)$, $\text{mfe.msk}_i \leftarrow \text{MFE.MSKGen}(\text{mfe.pp}; r_i)$, and $\text{mfe.ct}_i \leftarrow \text{MFE.SKEnc}(\text{mfe.msk}_i, f_{\text{accept}})$.
 - * If $i = j$, algorithm \mathcal{B} makes a challenge query f_{accept} to receive mfe.ct_i .
 - * If $i > j$, algorithm \mathcal{B} sets $\text{mfe.ct}_i \leftarrow \text{MFE.PKEnc}(\text{mfe.pp})$.

It then computes $\text{abe.ct}_i \leftarrow \text{ABE.Enc}(\text{abe.pp}, (\text{mfe.ct}_i, i), m)$ for all $i \in \mathcal{J}_{\mathcal{L}}$, and gives $\text{ct} = \{(i, \text{abe.ct}_i)\}_{i \in \mathcal{J}_{\mathcal{L}}}$ to \mathcal{A} .

- **Output phase:** At the end of the experiment, algorithm \mathcal{B} outputs whatever \mathcal{A} outputs.

We show that \mathcal{B} perfectly simulates either $\text{Hyb}_{1,j-1}$ or $\text{Hyb}_{1,j}$ depending on whether it is interacting in $\text{ExptMFE}_{\text{PK/SK}}[\lambda, \mathcal{B}, 0]$ or $\text{ExptMFE}_{\text{PK/SK}}[\lambda, \mathcal{B}, 1]$. Furthermore, we show that as long as \mathcal{A} is admissible for $\text{ExptRPE}_{\text{BC}}[\lambda, \mathcal{A}, b]$, algorithm \mathcal{B} is admissible for $\text{ExptMFE}_{\text{PK/SK}}[\lambda, \mathcal{B}, b]$. These two conditions show that algorithm \mathcal{B} breaks public/secret key indistinguishability of Π_{MFE} with the same advantage of \mathcal{A} for distinguishing $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$.

Admissibility condition. The only time algorithm \mathcal{B} submits a challenge query is when responding to \mathcal{A} 's challenge query. Namely, when \mathcal{A} makes a single challenge query (m, \mathcal{L}) , algorithm \mathcal{B} computes $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$, and if $j \in \mathcal{J}_{\mathcal{L}}$, it makes a challenge query f_{accept} to the mixed FE challenger. By definition, \mathcal{B} is admissible if for all key-generation queries x that it makes to the mixed FE challenger, $f_{\text{accept}}(x) = 1$, which holds by definition of f_{accept} . Moreover, if \mathcal{A} is a non-adaptive q -query adversary for the broadcast hiding game, then it makes at most q encryption queries and all of these queries occur after the non-encryption queries. By construction, algorithm \mathcal{B} only makes encryption queries when \mathcal{A} makes an encryption query, and it makes at most 1 encryption query in response to each of \mathcal{A} 's encryption queries. Thus, if \mathcal{A} is a non-adaptive q -query adversary, \mathcal{B} is also a non-adaptive q -query adversary.

Correctness of the simulation. The public parameters pp that \mathcal{B} provides to \mathcal{A} during the setup phase of the experiment is distributed exactly as in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. We now consider how \mathcal{B} simulates the responses to \mathcal{A} 's oracle queries:

- **Key-generation oracle:** Same as in the proof of Lemma A.3.
- **Encryption oracle:** Same as in the proof of Lemma A.3.
- **Challenge oracle:** Let (m, \mathcal{L}) be the challenge query made by \mathcal{A} , and let $\mathcal{J}_{\mathcal{L}} \leftarrow \text{ComputeCover}(\mathcal{L})$. For each $i \in \mathcal{I}_{\text{id}} \setminus \{j\}$, algorithm \mathcal{B} generates mixed FE ciphertexts mfe.ct_i exactly as described in $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. For mfe.ct_j , algorithm \mathcal{B} submits f_{accept} to the mixed FE challenger to receive mfe.ct_j . If the mixed FE challenger replies with $\text{MFE.PKEnc}(\text{mfe.pp})$, then \mathcal{B} perfectly simulates mfe.ct_j in $\text{Hyb}_{1,j}$ while if the challenger replies with $\text{MFE.SKEnc}(\text{mfe.msk}, f_{\text{accept}})$, then \mathcal{B} perfectly simulates mfe.ct_j in $\text{Hyb}_{1,j-1}$.

We conclude that \mathcal{B} breaks public/secret key indistinguishability of Π_{MFE} with the same advantage that \mathcal{A} has for distinguishing between hybrids $\text{Hyb}_{1,j-1}$ and $\text{Hyb}_{1,j}$. Thus, by a hybrid argument, if there exists an efficient adversary \mathcal{A} that can distinguish between Hyb_1 and Hyb_2 with non-negligible advantage $\varepsilon_0 = 1/\text{poly}(\lambda)$, then there exists an efficient adversary \mathcal{B} that runs in time $\text{poly}(\lambda) = \text{poly}(\lambda')$ that breaks public/secret key indistinguishability of Π_{MFE} with advantage at least ε_0/K . By construction, $K \leq 2^{(\lambda')^{\varepsilon/2}}$, and so \mathcal{B} succeeds with advantage at least

$$\varepsilon_0/2^{(\lambda')^{\varepsilon/2}} = 1/\text{poly}(\lambda') \cdot 2^{-(\lambda')^{\varepsilon/2}} = 2^{-O((\lambda')^{\varepsilon/2})},$$

which contradicts the sub-exponential hardness of Π_{MFE} . □

Lemma A.7. *If F is a secure PRF, then for all efficient adversaries \mathcal{A} ,*

$$|\Pr[\text{Hyb}_2(\mathcal{A}) = 1] - \Pr[\text{Hyb}_3(\mathcal{A}) = 1]| = \text{negl}(\lambda).$$

Proof. The only difference between hybrids Hyb_2 and Hyb_3 is that the adversary uses $f(\cdot)$ in place of $F(k, \cdot)$ when responding to the adversary's queries. The lemma follows immediately from the definition of PRF security. \square

Theorem 3.8 now follows by a standard hybrid argument. \square

B Analysis of Construction 4.4

In this section, we provide the formal analysis of the trace-and-revoke scheme from Section 4.1.

B.1 Proof of Theorem 4.5 (Correctness)

Take a message $m \in \mathcal{M}$, an identity $\text{id} \in \mathcal{ID}$, and a revocation list $\mathcal{L} \subseteq \mathcal{ID}$ where $\text{id} \notin \mathcal{L}$. Take $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id})$ and $\text{ct}_{m, \mathcal{L}} \leftarrow \text{Enc}(\text{pp}, m, \mathcal{L})$. In this case, $\text{pp} = (\text{hk}, \text{rpe.pp})$, $\text{msk} = (\text{hk}, \text{rpe.msk})$, sk_{id} is output by $\text{RPE.KeyGen}(\text{rpe.msk}, H(\text{hk}, \text{id}), v_{\text{id}})$ for some vector $v_{\text{id}} \in \mathcal{ID}_0^n$, and $\text{ct}_{m, \mathcal{L}}$ is output by $\text{RPE.Broadcast}(\text{rpe.pp}, m, \mathcal{L}')$ where $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. Since H is collision-resistant and $\text{id} \notin \mathcal{L}$, we have $H(\text{hk}, \text{id}) \notin \mathcal{L}'$ with overwhelming probability. The claim now follows by broadcast correctness of Π_{RPE} . \square

B.2 Proof of Theorem 4.6 (Semantic Security)

We proceed with a hybrid argument:

- Hyb_0 : This is experiment $\text{ExptTR}_{\text{SS}}[\lambda, \mathcal{A}, 0]$. Namely, the challenger responds to the challenge query (m_0, m_1, \mathcal{L}) with the ciphertext $\text{ct}^* \leftarrow \text{Enc}(\text{pp}, m_0, \mathcal{L})$.
- Hyb_1 : Same as Hyb_0 , except the challenger constructs the challenge ciphertext as $\text{ct}^* \leftarrow \text{Enc}(\text{msk}, f_{\text{accept}}, m_0, \mathcal{L})$ where f_{accept} is the “always-accept” function.
- Hyb_2 : Same as Hyb_1 , except the challenger constructs the challenge ciphertext to be $\text{ct}^* \leftarrow \text{Enc}(\text{msk}, f_{\text{accept}}, m_1, \mathcal{L})$.
- Hyb_3 : Same as Hyb_2 , except the challenger constructs the challenge ciphertext as $\text{ct}^* \leftarrow \text{Enc}(\text{pp}, m_1, \mathcal{L})$.

We now show that each consecutive pair of hybrid experiments are computationally indistinguishable:

- Hybrids Hyb_0 and Hyb_1 are computationally indistinguishable by broadcast security of Π_{RPE} . Specifically, suppose that there exists an efficient adversary \mathcal{A} that distinguishes Hyb_0 from Hyb_1 . We use \mathcal{A} to construct an adversary \mathcal{B} for the broadcast security game:
 1. At the beginning of the broadcast security game, algorithm \mathcal{B} receives the public parameters rpe.pp from the broadcast security challenger. In addition it samples a hash key $\text{hk} \xleftarrow{\mathcal{R}} \mathcal{K}$ and gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to \mathcal{A} .

2. Whenever \mathcal{A} makes a key-generation query on an identity $\text{id} \in \mathcal{ID}$, algorithm \mathcal{B} computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and the vector $v_{\text{id}} \in \mathcal{ID}_0^n$ as in the real `KeyGen` algorithm. It then makes a key-generation query to its challenger on the pair $(s_{\text{id}}, v_{\text{id}})$ to obtain a key sk_{id} , which it forwards to \mathcal{A} .
3. Whenever \mathcal{A} makes a challenge query on input (m_0, m_1, \mathcal{L}) , algorithm \mathcal{B} first constructs $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$ and makes a challenge query on the pair (m_0, \mathcal{L}') . The challenger replies to \mathcal{B} with a ciphertext ct^* , which it forwards to \mathcal{A} .
4. At the end of the game, algorithm \mathcal{B} outputs whatever \mathcal{A} outputs.

By construction, if $\text{ct}^* \leftarrow \text{RPE.Broadcast}(\text{rpe.pp}, m_0, \mathcal{L}')$, then \mathcal{B} perfectly simulated Hyb_0 for \mathcal{A} , and if $\text{ct}^* \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{\text{accept}}, m_0, \mathcal{L}')$, where rpe.msk is the secret key sampled by the broadcast security challenger, then \mathcal{B} perfectly simulated Hyb_1 for \mathcal{A} . The claim follows.

- Hybrids Hyb_1 and Hyb_2 are computationally indistinguishable by message hiding of Π_{RPE} . Specifically, suppose there exists an efficient adversary \mathcal{A} that can distinguish Hyb_1 from Hyb_2 . We use \mathcal{A} to construct an adversary \mathcal{B} for the message hiding game:

1. At the beginning of the message hiding game, algorithm \mathcal{B} receives the public parameters rpe.pp from the broadcast security challenger. In addition, it samples a hash key $\text{hk} \xleftarrow{\mathcal{R}} \mathcal{K}$ and gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to \mathcal{A} .
2. Whenever \mathcal{A} makes a key-generation query on $\text{id} \in \mathcal{ID}$, algorithm \mathcal{B} computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and the vector $v_{\text{id}} \in \mathcal{ID}_0^n$ as in the real `KeyGen` algorithm. It makes a key-generation query to its challenger on the pair $(s_{\text{id}}, v_{\text{id}})$ to obtain a key sk_{id} , which it forwards to \mathcal{A} .
3. Whenever \mathcal{A} makes a challenge query on input (m_0, m_1, \mathcal{L}) algorithm \mathcal{B} first constructs $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$ and makes a challenge query $(f_{\text{accept}}, m_0, m_1, \mathcal{L}')$ to the message hiding challenger. The challenger replies to \mathcal{B} with a ciphertext ct^* which it forwards to \mathcal{A} .
4. At the end of the game, algorithm \mathcal{B} outputs whatever \mathcal{A} outputs.

First, we argue that \mathcal{B} is admissible for the message hiding game. Since \mathcal{A} is admissible for the semantic security game, this means that for all key-generation queries $\text{id} \in \mathcal{ID}$ that \mathcal{A} makes, it must be the case that $\text{id} \in \mathcal{L}$. By construction of \mathcal{L}' , this means that $s_{\text{id}} = H(\text{hk}, \text{id}) \in \mathcal{L}'$ for all id appearing in a key-generation query. Thus, \mathcal{B} is admissible for the message hiding game. Now, if the message hiding challenger replies with $\text{ct}^* \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{\text{accept}}, m_0, \mathcal{L}')$, where rpe.msk is the master secret key sampled by the message hiding challenger, then \mathcal{B} perfectly simulates Hyb_1 for \mathcal{A} . Conversely, if $\text{ct}^* \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{\text{accept}}, m_1, \mathcal{L}')$, then \mathcal{B} perfectly simulates Hyb_2 for \mathcal{A} . The claim follows.

- Hybrids Hyb_2 and Hyb_3 are computationally indistinguishable by broadcast security of Π_{RPE} (via the same argument used to show indistinguishability of hybrids Hyb_0 and Hyb_1). \square

B.3 Proof of Theorem 4.7 (Traceability)

Let Q be the randomized oracle from Figure 1, and let \mathcal{A} be an efficient adversary for the traceability game. At the beginning of the traceability game, the challenger samples $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and

gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to the adversary and keeps $\text{msk} = (\text{hk}, \text{rpe.msk})$ for itself. Let $\mathcal{R} \subseteq \mathcal{ID}$ be the set of identities \mathcal{A} submits to the key-generation oracle in the traceability game and let \mathcal{L} be the set of revoked users chosen by \mathcal{A} . At the end of the traceability game, the adversary \mathcal{A} outputs two messages m_0 and m_1 , a revocation list $\mathcal{L} \subseteq \mathcal{ID}$, and a success probability ε . For each $\text{id} \in \mathcal{R}$, let $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$. By collision-resistance of H , all of the s_{id} will be distinct with overwhelming probability. We start by showing that the oracle Q defines an instance of the generalized jump-finding problem for the (adversarially-chosen) set $C = \{\text{id} \in \mathcal{R} \setminus \mathcal{L} : (s_{\text{id}}, \text{id}_1, \dots, \text{id}_n)\}$. First, for a pair $(i, u) \in [1, n] \times [0, 2^{\ell+1}]$, let $p_{i,u} := \Pr[Q^{\mathcal{D}}(i, u) = 1]$. In addition, for $\alpha \in \{0, 1\}$, we define

$$p_{i,u,\alpha} := \Pr[\text{ct} \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{i,u}, m_\alpha, \mathcal{L}') : \mathcal{D}(\text{ct}) = \alpha],$$

where $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. We now show that the oracle Q satisfies each of the requirements in Definition 2.5:

Lemma B.1. *Suppose Π_{RPE} satisfies non-adaptive 1-query function hiding. Then, for any two (adversarially-chosen) pairs $(i, 2u), (j, 2u) \in [1, n] \times [0, 2^{\ell+1}]$, we have that $|p_{i,2u} - p_{j,2u}| = \text{negl}(\lambda)$.*

Proof. Suppose \mathcal{A} produces a decoder \mathcal{D} along with indices $i, j \in [1, n]$ and $u \in [0, 2^{\ell+1}]$ such that $|p_{i,2u} - p_{j,2u}| \geq \varepsilon_0 = 1/\text{poly}(\lambda)$ with non-negligible probability ε_1 . We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the 1-query function hiding property of Π_{RPE} :

1. At the beginning of the function hiding game, algorithm \mathcal{B} receives the public parameters rpe.pp from the function hiding challenger. It chooses a hash key $\text{hk} \xleftarrow{\mathcal{R}} \mathcal{K}$ and gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to \mathcal{A} .
2. When \mathcal{A} makes a key-generation query on an identity $\text{id} \in \mathcal{ID}$, the challenger computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and constructs the vector $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n)$. Algorithm \mathcal{B} makes a key-generation query on the pair $(s_{\text{id}}, v_{\text{id}})$ to obtain a key sk_{id} , which it forwards to \mathcal{A} .
3. At some point, \mathcal{A} outputs a decoder algorithm \mathcal{D} , two messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$ and a decoder success probability ε . It also outputs indices $i, j \in [1, n]$ and $u \in [0, 2^{\ell+1}]$.
4. Algorithm \mathcal{B} constructs the set $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. It chooses two random bits $\alpha, \beta \xleftarrow{\mathcal{R}} \{0, 1\}$ and makes a challenge query on $(m_\alpha, f_{i,2u}, f_{j,2u}, \mathcal{L}')$ to obtain a ciphertext ct^* . Then, if $\beta = 0$, it also makes an encryption query on $(m_\alpha, f_{i,2u}, \mathcal{L}')$ and if $\beta = 1$, it makes an encryption query on $(m_\alpha, f_{j,2u}, \mathcal{L}')$. Let ct_β be the resulting ciphertext. Finally, the challenger checks if $\mathcal{D}(\text{ct}_\beta) = \mathcal{D}(\text{ct}^*)$ and outputs β if so and $1 - \beta$ otherwise.

Next, we show that \mathcal{B} is admissible and breaks function hiding with non-negligible advantage.

- To see that \mathcal{B} is admissible, take any $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n)$ that \mathcal{B} submits to the key-generation oracle. Suppose that $f_{i,2u}(v_{\text{id}}) = 1$. This means that $v_{\text{id},i} = 2s_{\text{id}} - \text{id}_i \leq 2u$. Since $\text{id}_i, \text{id}_j \in \{0, 1\}$, this means that $v_{\text{id},j} = 2s_{\text{id}} - \text{id}_j \leq 2u$, and correspondingly $f_{j,2u}(v_{\text{id}}) = 1$. Thus, for all v_{id} that \mathcal{B} submits to the key-generation oracle, $f_{i,2u}(v_{\text{id}}) = f_{j,2u}(v_{\text{id}})$. Moreover, \mathcal{B} makes exactly 1 query to the encryption oracle after making all of its key-generation and challenge queries, so it is admissible for the non-adaptive 1-query function hiding game.

- We now compute the distinguishing advantage of algorithm \mathcal{B} . We consider two settings, depending on the challenge bit $b \in \{0, 1\}$. When $b = 0$, the challenger constructs the challenge ciphertext as $\text{ct}^* \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{i,2u}, m_\alpha, \mathcal{L}')$. We compute the probability that \mathcal{B} outputs 1 in this case. Let b' denote the output bit of \mathcal{B} .

$$\begin{aligned} \Pr[b' = 1] &= \Pr[\beta = 1 \wedge \mathcal{D}(\text{ct}_1) = \mathcal{D}(\text{ct}^*)] + \Pr[\beta = 0 \wedge \mathcal{D}(\text{ct}_0) \neq \mathcal{D}(\text{ct}^*)] \\ &= \frac{1}{2} (\Pr[\mathcal{D}(\text{ct}_1) = \mathcal{D}(\text{ct}^*)] + \Pr[\mathcal{D}(\text{ct}_0) \neq \mathcal{D}(\text{ct}^*)]), \end{aligned} \quad (\text{B.1})$$

since β is a uniform (and independently) random bit. Here (and in the remainder of the analysis), the probabilities are taken over the decoder's randomness, algorithm \mathcal{B} 's randomness, and the challenger's randomness. To simplify the notation in the following analysis, we will write \hat{p}_i and \hat{p}_j to denote $p_{i,2u,0}$ and $p_{j,2u,0}$, respectively. Similarly, we will write \hat{q}_i and \hat{q}_j to denote $p_{i,2u,1}$ and $p_{j,2u,1}$, respectively. In particular, this means that

$$\hat{p}_i := \Pr[\text{ct} \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{i,2u}, m_0, \mathcal{L}') : \mathcal{D}(\text{ct}) = 0]$$

Now, using the fact that ct_0 , ct_1 , and ct^* are all independently constructed, we have that

$$\begin{aligned} \Pr[\mathcal{D}(\text{ct}_1) = \mathcal{D}(\text{ct}^*)] &= \Pr[\mathcal{D}(\text{ct}_1) = 0 \wedge \mathcal{D}(\text{ct}^*) = 0] + \Pr[\mathcal{D}(\text{ct}_1) = 1 \wedge \mathcal{D}(\text{ct}^*) = 1] \\ &= \Pr[\mathcal{D}(\text{ct}_1) = 0] \Pr[\mathcal{D}(\text{ct}^*) = 0] + \Pr[\mathcal{D}(\text{ct}_1) = 1] \Pr[\mathcal{D}(\text{ct}^*) = 1]. \end{aligned}$$

Suppose $\alpha = 0$. Then, ct^* , ct_0 , and ct_1 are all encryptions of m_0 , and so we have

$$\begin{aligned} \Pr[\mathcal{D}(\text{ct}_1) = \mathcal{D}(\text{ct}^*)] &= \Pr[\mathcal{D}(\text{ct}_1) = 0] \Pr[\mathcal{D}(\text{ct}^*) = 0] + \Pr[\mathcal{D}(\text{ct}_1) = 1] \Pr[\mathcal{D}(\text{ct}^*) = 1] \\ &= \hat{p}_j \cdot \hat{p}_i + (1 - \hat{p}_j)(1 - \hat{p}_i). \end{aligned}$$

A similar calculation shows that

$$\begin{aligned} \Pr[\mathcal{D}(\text{ct}_0) \neq \mathcal{D}(\text{ct}^*)] &= \Pr[\mathcal{D}(\text{ct}_0) = 0] \Pr[\mathcal{D}(\text{ct}^*) = 1] + \Pr[\mathcal{D}(\text{ct}_0) = 1] \Pr[\mathcal{D}(\text{ct}^*) = 0] \\ &= \hat{p}_i \cdot (1 - \hat{p}_i) + (1 - \hat{p}_i) \cdot \hat{p}_i. \end{aligned}$$

Together with Eq. (B.1),

$$\Pr[b' = 1 \mid b = 0 \wedge \alpha = 0] = \frac{1}{2} (\hat{p}_i \cdot (1 - (\hat{p}_i - \hat{p}_j)) + (1 - \hat{p}_i) \cdot (1 - (\hat{p}_j - \hat{p}_i))). \quad (\text{B.2})$$

We can do a similar sequence of calculations when $\alpha = 1$. The only difference in this case is that ct^* , ct_0 , and ct_1 are now encryptions of m_1 instead. This means that

$$\begin{aligned} \Pr[\mathcal{D}(\text{ct}_1) = \mathcal{D}(\text{ct}^*)] &= (1 - \hat{q}_j)(1 - \hat{q}_i) + \hat{q}_j \cdot \hat{q}_i \\ \Pr[\mathcal{D}(\text{ct}_0) \neq \mathcal{D}(\text{ct}^*)] &= (1 - \hat{q}_i) \cdot \hat{q}_i + \hat{q}_i \cdot (1 - \hat{q}_i). \end{aligned}$$

Correspondingly,

$$\Pr[b' = 1 \mid b = 0 \wedge \alpha = 1] = \frac{1}{2} (\hat{q}_i \cdot (1 - (\hat{q}_i - \hat{q}_j)) + (1 - \hat{q}_i) \cdot (1 - (\hat{q}_j - \hat{q}_i))). \quad (\text{B.3})$$

Consider now the case where $b = 1$. In this case, the challenge ciphertext is given by $\text{ct}^* \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{j,2u}, m_\alpha, \mathcal{L}')$. When $b = 1$ and $\alpha = 0$, we now have

$$\begin{aligned} \Pr[\mathcal{D}(\text{ct}_1) = \mathcal{D}(\text{ct}^*)] &= \hat{p}_j \cdot \hat{p}_j + (1 - \hat{p}_j)(1 - \hat{p}_j) \\ \Pr[\mathcal{D}(\text{ct}_0) \neq \mathcal{D}(\text{ct}^*)] &= \hat{p}_i \cdot (1 - \hat{p}_j) + (1 - \hat{p}_i) \cdot \hat{p}_j, \end{aligned}$$

and correspondingly,

$$\Pr[b' = 1 \mid b = 1 \wedge \alpha = 0] = \frac{1}{2}(\hat{p}_j \cdot (1 - (\hat{p}_i - \hat{p}_j)) + (1 - \hat{p}_j) \cdot (1 - (\hat{p}_j - \hat{p}_i))). \quad (\text{B.4})$$

Lastly, when $b = 1$ and $\alpha = 1$, we have

$$\Pr[b' = 1 \mid b = 1 \wedge \alpha = 1] = \frac{1}{2}(\hat{q}_j \cdot (1 - (\hat{q}_i - \hat{q}_j)) + (1 - \hat{q}_j) \cdot (1 - (\hat{q}_j - \hat{q}_i))). \quad (\text{B.5})$$

From Eq. (B.2) and (B.4)

$$\Pr[b' = 1 \mid b = 1 \wedge \alpha = 0] - \Pr[b' = 1 \mid b = 0 \wedge \alpha = 0] = (\hat{p}_j - \hat{p}_i)^2.$$

Similarly, from Eq. (B.3) and (B.5),

$$\Pr[b' = 1 \mid b = 1 \wedge \alpha = 1] - \Pr[b' = 1 \mid b = 0 \wedge \alpha = 1] = (\hat{q}_j - \hat{q}_i)^2.$$

Finally, since the challenger samples $\alpha \stackrel{R}{\leftarrow} \{0, 1\}$, algorithm \mathcal{B} 's advantage is given by

$$|\Pr[b' = 1 \mid b = 1] - \Pr[b' = 1 \mid b = 0]| = \frac{1}{2}((\hat{p}_j - \hat{p}_i)^2 + (\hat{q}_j - \hat{q}_i)^2). \quad (\text{B.6})$$

It suffices to show this quantity is non-negligible. First, we note that \mathcal{B} perfectly simulates the traceability game for \mathcal{A} , so with probability ε_1 , the decoder \mathcal{D} that \mathcal{A} outputs will satisfy $|p_{j,2u} - p_{i,2u}| \geq \varepsilon_0$. Consider this case. First, by definition,

$$p_{i,2u} = \Pr[Q^{\mathcal{D}}(i, 2u) = 1] = \frac{1}{2}(p_{i,2u,0} + p_{i,2u,1}) = \frac{1}{2}(\hat{p}_i + \hat{q}_i).$$

Similarly, $p_{j,2u} = \frac{1}{2}(\hat{p}_j + \hat{q}_j)$. By the triangle inequality,

$$\varepsilon_0 \leq |p_{j,2u} - p_{i,2u}| \leq \frac{1}{2}(|\hat{p}_j - \hat{p}_i| + |\hat{q}_j - \hat{q}_i|).$$

Thus, at least one of $|\hat{p}_j - \hat{p}_i| \geq \varepsilon_0$ or $|\hat{q}_j - \hat{q}_i| \geq \varepsilon_0$ must hold. Appealing to Eq. (B.6), we conclude that algorithm \mathcal{B} 's distinguishing advantage for the 1-query function hiding game is at least

$$|\Pr[b' = 1 \mid b = 1] - \Pr[b' = 1 \mid b = 0]| \geq \frac{\varepsilon_1 \varepsilon_0^2}{2},$$

which is non-negligible. \square

Lemma B.2. *Let $C_i = \{(s_{\text{id}}, \text{id}_1, \dots, \text{id}_n) \in C : 2s_{\text{id}} - \text{id}_i\}$. Suppose Π_{RPE} satisfies non-adaptive 1-query function hiding security. Then, for any two (adversarially-chosen) pairs $(i, u_1), (i, u_2) \in [1, n] \times [0, 2^{\ell+1}]$ where $u_1 \leq u_2$ and $(u_1, u_2] \cap C_i = \emptyset$, $|p_{i,u_1} - p_{i,u_2}| = \text{negl}(\lambda)$.*

Proof. The proof follows by a similar argument as in the proof of Lemma B.1. Specifically, suppose \mathcal{A} is able to produce a decoder \mathcal{D} together with indices $i \in [1, n]$, $u_1, u_2 \in [0, 2^{\ell+1}]$ with non-negligible probability ε_1 such that all of the following conditions hold:

- $|p_{i,u_1} - p_{i,u_2}| \geq \varepsilon_0 = 1/\text{poly}(\lambda)$
- $u_1 \leq u_2$ and $(u_1, u_2] \cap C_i = \emptyset$, where C and C_i are the sets defined by the adversary's queries.

We use \mathcal{A} to construct an adversary \mathcal{B} that breaks the 1-query function hiding security of Π_{RPE} :

1. At the beginning of the function hiding game, algorithm \mathcal{B} receives the public parameters rpe.pp from the function hiding challenger. It chooses a hash key $\text{hk} \xleftarrow{\mathcal{R}} \mathcal{K}$ and gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to \mathcal{A} .
2. When \mathcal{A} makes a key-generation query on an identity $\text{id} \in \mathcal{ID}$, the challenger computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and constructs the vector $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n)$. Algorithm \mathcal{B} makes a key-generation query on the pair $(s_{\text{id}}, v_{\text{id}})$ to obtain a key sk_{id} , which it forwards to \mathcal{A} .
3. At the end of the game, algorithm \mathcal{A} output a decoder algorithm \mathcal{D} , two messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$ and a decoder success probability ε . It also chooses indices $i \in [1, n]$ and $u_1, u_2 \in [0, 2^{\ell+1}]$. If $u_1 > u_2$ or $(u_1, u_2] \cap C_i \neq \emptyset$, then \mathcal{B} aborts and outputs 0.
4. Algorithm \mathcal{B} constructs the set $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. It chooses two random bits $\alpha, \beta \xleftarrow{\mathcal{R}} \{0, 1\}$ and makes a challenge query on $(m_\alpha, f_{i,u_1}, f_{i,u_2}, \mathcal{L}')$ to obtain a ciphertext ct^* . Then, if $\beta = 0$, it also makes an encryption query on $(m_\alpha, f_{i,u_1}, \mathcal{L}')$ and if $\beta = 1$, it makes an encryption query on $(m_\alpha, f_{i,u_2}, \mathcal{L}')$. Let ct_β be the resulting ciphertext. Algorithm \mathcal{B} checks if $\mathcal{D}(\text{ct}_\beta) = \mathcal{D}(\text{ct}^*)$ and outputs β if so and $1 - \beta$ otherwise.

As in the proof of Lemma B.1, we argue that \mathcal{B} is admissible and that it breaks function hiding with non-negligible advantage.

- First, we show that \mathcal{B} is admissible. Let $(s_{\text{id}}, v_{\text{id}})$ be a key-generation query made by \mathcal{B} , where $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n)$. We need to show that for all non-revoked identities $\text{id} \notin \mathcal{L}$ (i.e., all $s_{\text{id}} \notin \mathcal{L}'$), $f_{i,u_1}(v_{\text{id}}) = f_{i,u_2}(v_{\text{id}})$. Since $\text{id} \notin \mathcal{L}$, this means that $\text{id} \in \mathcal{R} \setminus \mathcal{L}$, and correspondingly, $(s_{\text{id}}, \text{id}_1, \dots, \text{id}_n) \in C$. By construction, if \mathcal{B} does not abort, then $(u_1, u_2] \cap C_i = \emptyset$ so that means that $2s_{\text{id}} - \text{id}_i \notin (u_1, u_2]$. There are now two possibilities. If $f_{i,u_1}(v_{\text{id}}) = 1$, then $v_{\text{id},i} \leq u_1 \leq u_2$, and so $f_{i,u_2}(v_{\text{id}}) = 1$. Alternatively, if $f_{i,u_1}(v_{\text{id}}) = 0$, then $v_{\text{id},i} = 2s_{\text{id}} - \text{id}_i > u_1$. Since $2s_{\text{id}} - \text{id}_i \notin (u_1, u_2]$, it must be the case that $2s_{\text{id}} - \text{id}_i > u_2$, in which case, $f_{i,u_2}(v_{\text{id}}) = 0$. Finally, \mathcal{B} only makes 1 encryption query after making all of its key-generation and challenge queries, so we conclude that it is admissible for the non-adaptive 1-query function hiding game.
- Next, we consider the distinguishing advantage of \mathcal{B} . First, \mathcal{B} perfectly simulates the traceability game for \mathcal{A} so with probability ε_1 , \mathcal{B} will not abort and the decoder \mathcal{D} that \mathcal{A} satisfies $|p_{i,u_1} - p_{i,u_2}| \geq \varepsilon_0$. Using the same type of analysis as in the proof of Lemma B.1, we conclude that in this case, \mathcal{B} is able to win the 1-query function hiding game with advantage at least $\varepsilon_1 \varepsilon_0^2 / 2$, which is non-negligible. \square

Lemma B.3. *For all $i, j \in [n]$, we have that $p_{i,0} = p_{j,0}$. Set $p_0 = p_{1,0}$. Moreover, if Π_{RPE} satisfies non-adaptive 1-query message hiding, then $p_0 \leq 1/2 + \text{negl}(\lambda)$.*

Proof. By definition, $f_{i,0}$ and $f_{j,0}$ are identical functions (and have the same canonical representation) for all $i, j \in [n]$. Thus, for all $i, j \in [n]$, $p_{i,0} = p_{j,0}$. We now show that $p_0 \leq 1/2 + \text{negl}(\lambda)$. To show this, we first define the following process. Let $m_0, m_1 \in \mathcal{M}$ be the messages, \mathcal{L} be the revocation list, and \mathcal{D} be the decoder chosen by the adversary. Let $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. Now, sample a random bit $b \xleftarrow{\mathcal{R}} \{0, 1\}$ and construct the ciphertext $\text{ct} \leftarrow \text{RPE.Enc}(\text{rpe.msk}, f_{1,0}, m_0, \mathcal{L}')$. Let $\hat{p} := \Pr[\mathcal{D}(\text{ct}) = b]$. Since \mathcal{D} and ct are sampled independently of b , $\hat{p} = 1/2$. We now show that if \mathcal{A}

produces a decoder such that $|p_0 - \hat{p}| \geq \varepsilon_0$ with probability at least ε_1 , we can use \mathcal{A} to construct an adversary \mathcal{B} that breaks 1-query message hiding security of Π_{RPE} :

1. At the beginning of the message hiding game, algorithm \mathcal{B} receives the public parameters rpe.pp from the message hiding challenger. It choose a hash key $\text{hk} \xleftarrow{\text{R}} \mathcal{K}$ and gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to \mathcal{A} .
2. When \mathcal{A} makes a key-generation query on an identity $\text{id} \in \mathcal{ID}$, the challenger computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and constructs the vector $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n)$. Algorithm \mathcal{B} makes a key-generation query on the pair $(s_{\text{id}}, v_{\text{id}})$ to obtain a key sk_{id} , which it forwards to \mathcal{A} .
3. At the end of the game, \mathcal{A} outputs a decoder algorithm \mathcal{D} , two messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$ and a decoder success probability ε .
4. Algorithm \mathcal{B} constructs the set $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. It chooses two random bits $\alpha, \beta \xleftarrow{\text{R}} \{0, 1\}$ and makes a challenge query on $(f_{1,0}, m_\alpha, m_0, \mathcal{L}')$ to obtain a ciphertext ct^* . Next, if $\beta = 0$, it makes an encryption query on $(m_\alpha, f_{1,0}, \mathcal{L}')$ and if $\beta = 1$, it makes an encryption query on $(m_0, f_{1,0}, \mathcal{L}')$. Let ct_β be the resulting ciphertext. Algorithm \mathcal{B} checks if $\mathcal{D}(\text{ct}_\beta) = \mathcal{D}(\text{ct}^*)$ and outputs β if so and $1 - \beta$ otherwise.

Since $f_{1,0}$ is the all-zeroes function, and \mathcal{B} makes exactly 1 encryption query after making all of its non-encryption queries, algorithm \mathcal{B} is admissible for the non-adaptive 1-query message-hiding game. We now consider the distinguishing advantage of \mathcal{B} . Since \mathcal{B} perfectly simulates the traceability game for \mathcal{A} , with probability ε_1 , the decoder \mathcal{D} that \mathcal{A} outputs will satisfy $|p_0 - \hat{p}| > \varepsilon_0$. The claim now follows by the same type of analysis as in the proof of Lemma B.1. In particular, algorithm \mathcal{B} wins the 1-query message hiding game with advantage at least $\varepsilon_1 \varepsilon_0^2 / 2$, which is non-negligible. This means that $|p_0 - \hat{p}| = \text{negl}(\lambda)$, and the claim follows. \square

Lemma B.4. *For all $i, j \in [n]$, we have that $p_{i,2^{\ell+1}} = p_{j,2^{\ell+1}}$. Set $p_{2^{\ell+1}} = p_{1,2^{\ell+1}}$. Moreover, assuming Π_{RPE} satisfies non-adaptive 1-query broadcast security, $p_{2^{\ell+1}} \geq 1/2 + \varepsilon - \text{negl}(\lambda)$, where ε is the non-negligible decoder success probability output by \mathcal{A} .*

Proof. By definition, $f_{i,2^{\ell+1}}$ and $f_{j,2^{\ell+1}}$ are identical functions and have the same canonical representation (i.e., $f_{i,2^{\ell+1}} \equiv f_{\text{accept}} \equiv f_{j,2^{\ell+1}}$) for all $i, j \in [n]$. We now show that $p_{2^{\ell+1}} \geq 1/2 + \varepsilon - \text{negl}(\lambda)$. Similar to the proof of Lemma B.3, consider the following process. Let $m_0, m_1 \in \mathcal{M}$ be the messages, \mathcal{L} be the revocation list, and \mathcal{D} be the decoder chosen by the adversary. Let $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. Now, sample a random bit $b \xleftarrow{\text{R}} \{0, 1\}$ and construct the ciphertext $\text{ct} \leftarrow \text{RPE.Broadcast}(\text{rpe.pp}, m_b, \mathcal{L}')$. Let $\hat{p} := \Pr[\mathcal{D}(\text{ct}) = b]$. Since \mathcal{A} is admissible for the traceability game, $\hat{p} \geq 1/2 + \varepsilon$. We now show that if \mathcal{A} produces a decoder such that $|p_0 - \hat{p}| \geq \varepsilon_0$ with probability at least ε_1 , we can use \mathcal{A} to construct an adversary \mathcal{B} that breaks 1-query broadcast security of Π_{RPE} :

1. At the beginning of the broadcast security game, algorithm \mathcal{B} receives the public parameters rpe.pp from the challenger. It choose a hash key $\text{hk} \xleftarrow{\text{R}} \mathcal{K}$ and gives $\text{pp} = (\text{hk}, \text{rpe.pp})$ to \mathcal{A} .
2. When \mathcal{A} makes a key-generation query on an identity $\text{id} \in \mathcal{ID}$, the challenger computes $s_{\text{id}} \leftarrow H(\text{hk}, \text{id})$ and constructs the vector $v_{\text{id}} = (2s_{\text{id}} - \text{id}_1, \dots, 2s_{\text{id}} - \text{id}_n)$. Algorithm \mathcal{B} makes a key-generation query on the pair $(s_{\text{id}}, v_{\text{id}})$ to obtain a key sk_{id} , which it forwards to \mathcal{A} .
3. At the end of the game, \mathcal{A} outputs a decoder algorithm \mathcal{D} , two messages $m_0, m_1 \in \mathcal{M}$, a revocation list $\mathcal{L} \subseteq \mathcal{ID}$ and a decoder success probability ε .

4. Algorithm \mathcal{B} constructs the set $\mathcal{L}' = \{\text{id} \in \mathcal{L} : H(\text{hk}, \text{id})\}$. It chooses two random bits $\alpha, \beta \stackrel{\text{R}}{\leftarrow} \{0, 1\}$ and makes a challenge query on (m_α, \mathcal{L}') to obtain a ciphertext ct^* . Next if $\beta = 0$, it computes $\text{ct}_\beta \leftarrow \text{RPE.Broadcast}(\text{rpe.pp}, m_\alpha, \mathcal{L}')$, and if $\beta = 1$, it makes an encryption query on $(f_{\text{accept}}, m_\alpha, \mathcal{L}')$ to obtain a ciphertext ct_β . Algorithm \mathcal{B} checks if $\mathcal{D}(\text{ct}_\beta) = \mathcal{D}(\text{ct}^*)$ and outputs β if so and $1 - \beta$ otherwise.

By definition, $f_{1,2^{\ell+1}} \equiv f_{\text{accept}}$. Moreover, \mathcal{B} makes at most 1 encryption query after making all of its non-encryption queries, so it is admissible for the non-adaptive 1-query broadcast security game. We now compute the distinguishing advantage of \mathcal{B} . Since \mathcal{B} perfectly simulates the traceability game for \mathcal{A} , with probability ε_1 , the decoder \mathcal{D} that \mathcal{A} outputs will satisfy $|p_{2^{\ell+1}} - \hat{p}| > \varepsilon_0$. The claim now follows by the same type of analysis as in the proof of Lemma B.1. In particular, algorithm \mathcal{B} wins the 1-query broadcast security game with advantage at least $\varepsilon_1 \varepsilon_0^2 / 2$, which is non-negligible. This means that $|p_{2^{\ell+1}} - \hat{p}| = \text{negl}(\lambda)$ and the claim follows. \square

Combining Lemmas B.1 through B.4, we conclude that the oracle Q defines an instance of the $(2^\ell, n, |C|, \delta, \varepsilon)$ -generalized jump-finding game for any $\delta \leq \varepsilon / (9 + 4(\ell - 1)|C|)$, where C is the set of non-revoked identities queried by the adversary.⁸ By Theorem 2.6 and Remark 2.7, the Trace algorithm will recover an element in C with overwhelming probability when executed on some $q > \log |C|$. By construction of C , if this happens, then Trace will output an identity $\text{id} \in \mathcal{R} \setminus \mathcal{L}$. Thus, the Trace algorithm will terminate in at most $\log |C| = \text{poly}(\lambda)$ iterations, and each iteration requires time $\text{poly}(\lambda, \ell, n, \varepsilon, |C|) = \text{poly}(\lambda)$. This means that overall, the Trace algorithm terminates in polynomial time (with overwhelming probability). It suffices to argue that running QTrace^Q on values of $q < |C|$ does not cause the Trace algorithm to output an element that is not in C . This also follows from the correctness requirement in Theorem 2.6. Namely, QTrace^Q only outputs elements where there is a “jump” (of magnitude at least δ_q) in the decryption advantage, and for smaller values of q , the magnitude δ_q of the jumps is greater. Thus, any element output by QTrace^Q using a value of $q < |C|$ must also be contained in C . We conclude that with overwhelming probability, $\text{Trace}^{\mathcal{D}}(\text{msk}, m_0, m_1, \mathcal{L}, \varepsilon)$ will output some $\text{id}^* \in \mathcal{R} \setminus \mathcal{L}$, and the claim follows. \square

⁸Technically, we have only showed that the first two properties of Definition 2.5 hold for efficiently-sampleable pairs of points (i, x) and (j, y) , but this is sufficient for invoking Theorem 2.6 (see Remark 2.7). Namely, since QTrace^Q is an efficient algorithm, all of the properties in Definition 2.5 hold for the inputs queried by QTrace^Q .