Non-Malleable Extractors and Symmetric Key Cryptography from Weak Secrets

Yevgeniy Dodis* Daniel Wichs†

April 5, 2009

Abstract

We study the question of basing symmetric key cryptography on weak secrets. In this setting, Alice and Bob share an n-bit secret W, which might *not* be uniformly random, but the adversary has at least k bits of uncertainty about it (formalized using conditional min-entropy). Since standard symmetric-key primitives require uniformly random secret keys, we would like to construct an *authenticated key agreement protocol* in which Alice and Bob use W to agree on a nearly uniform key R, by communicating over a public channel controlled by an *active* adversary Eve. We study this question in the *information theoretic* setting where the attacker is computationally unbounded. We show that single-round (i.e. one message) protocols do not work when $k \leq \frac{n}{2}$, and require poor parameters even when $\frac{n}{2} < k \ll n$.

On the other hand, for arbitrary values of k, we design a communication efficient *two-round* (challenge-response) protocol extracting nearly k random bits. This dramatically improves the previous construction of Renner and Wolf [RW03], which requires $\Theta(\lambda + \log(n))$ rounds where λ is the security parameter. Our solution takes a new approach by studying and constructing *"non-malleable" seeded randomness extractors* — if an attacker sees a random seed X and comes up with an arbitrarily related seed X', then we bound the relationship between $R = \mathsf{Ext}(W; X)$ and $R' = \mathsf{Ext}(W; X')$.

We also extend our two-round key agreement protocol to the "fuzzy" setting, where Alice and Bob share "close" (but not equal) secrets W_A and W_B, and to the Bounded Retrieval Model (BRM) where the size of the secret W is huge.

*Computer Science Dept. NYU. Email: `dodis@cs.nyu.edu`.

†Computer Science Dept. NYU. Email: `wichs@cs.nyu.edu`.

1 Introduction

In this paper, we study the fundamental problem of symmetric key cryptography: Alice and Bob share a secret W and wish to communicate securely over a public channel controlled by an active adversary Eve. In particular, we want the communication to be *private* and *authentic*. Of course, this problem is well studied and can be solved using basic cryptographic primitives, either under computational assumptions, or even in the information theoretic setting. However, the standard solutions for both settings assume that the secret W is *perfectly (uniformly) random*.

In practice, many secrets, such as human-memorable passphrases and biometrics, are not uniformly random. Even keys that start out perfectly random may become compromised, for example through side-channel attacks against hardware or due to a malware infiltration of the storage device. Although all security is lost if the adversary learns the secret in its entirety, it is often reasonable to assume that the compromise is only *partial*. This assumption is natural for side-channel attacks (and was formalized in [MR04, DP08, AGV09]) where the adversary does not have full access to the device, and for malware infiltration in the *Bounded Retrieval Model* [Dzi06, CLW06], where the secret is made *intentionally huge* so that a malicious program cannot communicate it fully to an adversary. Lastly, it is conceivable that Alice and Bob, who do not share a secret initially, can use some *physical means* to agree on a key about which an eavesdropping adversary will only have *partial* information. This is, for example, the case in *Quantum Key Agreement* [BB84] and in the *wiretap channel model* [Wyn75]. In this work, we study a general setting which encompasses all of the above examples. We assume that Alice and Bob share a *weak secret*, modeled as a random variable W arbitrarily distributed over bit-strings of length n, about which an adversary Eve has some *side information*, modeled as a random variable Z correlated with W. We want to base symmetric key cryptography on *minimal assumptions about the secrecy of* W, and only require that W has at least k bits of entropy (conditioned on the side-information Z), where k is roughly proportional to the security parameter. As already mentioned, standard symmetric key primitives can be used in the case where Alice and Bob share a truly random key, and therefore we ask the following natural question.

Question 1: Can Alice and Bob use a shared *weak secret* W to securely agree on a *nearly uniform* secret key R, by communicating over a *public* and *unauthenticated* channel, controlled by an *active* attacker Eve?

One possible solution to this problem, is to use *password authenticated key exchange* (PAK) [BMP00, BPR00, KOY01, GL01, CHK$^+$05, GL06], where the secret W is used as a password. PAK allows Alice and Bob to agree on arbitrarily many random *session keys* using the secret W, and achieves strong security guarantees even when the entropy k is very low. On the other hand, all of the practical constructions of PAK either use the *random oracle model* or rely on a *trusted common reference string*. The only exception is the construction of [GL01] which, instead, requires many rounds of interaction and is not practically efficient. In addition, all of the constructions require the use of *public key cryptography*. Thus, even though we are in a symmetric key setting where Alice and Bob share a secure secret W, the use of PAK requires public key assumptions (and expensive public key operations) to take advantage of it. Also, PAK is a computational primitive, and only provides security when the attacker Eve is *computationally bounded*.

In contrast, we will study Question 1 in the *information theoretic setting*, where the adversary Eve is *computationally unbounded*. We call protocols that solve the problem of Question 1 in our setting *(information-theoretic) authenticated key agreement* (IT-AKA) protocols. Of course, IT-AKA cannot achieve all of the security guarantees of PAK. For example, IT-AKA can only be used *once* to convert a weak secret W into a uniformly random key R, and cannot be used to generate arbitrarily many session keys. Also, authenticated key agreement does not provide any security guarantees when the entropy k is very low (i.e. when the secret can be guessed with a reasonable probability). On the other hand, IT-AKA achieves information theoretic security and thus allows us to base *all of symmetric key cryptography* (information-theoretic as well computational) on weak secrets. Moreover, our constructions will be efficient (no public key operations) and do not require a common reference string or any other setup. For the rest of the paper, we will therefore assume that the adversary Eve is *computationally unbounded*.

A weaker variant of the our problem, called *privacy amplification* [BBR88, Mau92, BBCM95], requires that Alice and Bob communicate over an *authenticated channel* (alternatively, that the attacker Eve is passive). In this setting, key agreement can be solved using a (strong) *randomness extractor* [NZ96], which uses a seed X that is made public to the adversary, to extract nearly uniform randomness $R = \mathsf{Ext}(W; X)$ from a weak secret W. Privacy amplification can therefore be done in a one-round protocol, where Alice sends a seed X to Bob and both parties share the extracted key R.

The question of authenticated key agreement (when there is *no* authenticated channel and the adversary is active) was first studied by Maurer and Wolf in [MW97], who constructed an $\mathsf{IT\text{-}AKA}$ protocol for the case when W has entropy $k > \frac{2n}{3}$ (where n is the bit-length of W). This was later improved to $k > \frac{n}{2}$ in the work of [DKRS06]. Both of the above constructions are single-round, but only achieve authenticity at a price in the communication complexity (requiring at least $n - k$ bits) and the size of extracted key (which is at most $\ell < 2k - n$ bits long, and thus far below the full entropy of W). The most troubling aspect of these constructions, however, is the requirement that the entropy must exceed $k > \frac{n}{2}$, which conflicts with our goal of basing symmetric key cryptography on *minimal* secrecy assumptions. Moreover, many natural sources of secret randomness, such as biometrics, are unlikely to satisfy this requirement.

In terms of negative results, Dodis and Spencer [DS02] showed impossibility of one-round *message authentication* if the *only* randomness available to Alice and Bob comes from a weak secret W whose entropy is $k \leq \frac{n}{2}$. However, in our setting, we assume that the parties also have access to a local (non-shared) source of perfect randomness. These two settings are very different and, when no perfect randomness is available, most cryptographic primitives (including privacy amplification) are impossible *even* if $k > \frac{n}{2}$ [MP90, DOPS04, BD07]. Therefore, we feel that the result of Dodis and Spencer has often been incorrectly interpreted (for example in [RW03, DKRS06, CDF$^+$08]) as showing the impossibility of one-round authenticated key agreement protocols in *our more general setting*, where perfect (non-shared) randomness *is* available. In this paper we rectify this discrepancy by proving a (non-trivial) generalization of the [DS02] lower bound for *our* setting, thus showing that, unfortunately, single-round protocols do not exist when the entropy is $k \leq \frac{n}{2}$.

In terms of positive results, an *interactive* $\mathsf{IT\text{-}AKA}$ for *arbitrarily weak secrets* (i.e. allowing entropy $k \leq \frac{n}{2}$) was constructed by Renner and Wolf in [RW03] using a protocol which requires $\Theta(\lambda + \log(n))$ rounds of interaction, where λ is the security parameter. Several optimizations to the above protocol were proposed by Kanukurthi and Reyzin [KR09], leading to important practical efficiency gains, but without improving the (large) asymptotic round complexity of the original protocol. Thus, there is a huge gap between the lower bound (which shows that at least *two* rounds of interaction are required) and the best prior constructions. We therefore turn our attention to the following question, which will be the central question of this work.

Question 2: What is the minimal amount of interaction required to achieve authenticated key agreement ($\mathsf{IT\text{-}AKA}$) from arbitrarily weak secrets? In particular, is a two-round protocol possible?

In this paper, we answer Question 2 in the affirmative by giving an efficient construction of the first two-round $\mathsf{IT\text{-}AKA}$ protocol for arbitrarily weak secrets, and so bridge the gap between lower bound and construction. Our protocol only requires $k \geq \mathsf{poly}(\lambda, \log(n))$, where λ is the security parameter, and thus allows for entropy k which is sub-linear in the size n of the secret.[1] Hence our construction

[1] Our main (efficient) construction requires $k \geq O(\lambda^2 + \log^2(n))$ which asymptotically matches the requirements on entropy needed in [KR09]. We show that this can be improved further, by giving a (non-constructive) argument for the existence of $\mathsf{IT\text{-}AKA}$ protocols requiring only $k \geq O(\lambda + \log(n))$.

is optimal in the amount of interaction and requires (essentially) minimal assumptions on the entropy of the secret W. Our protocol is also efficient in terms of communication complexity and extracts essentially all of the entropy of W into the final shared key. Therefore, even in the setting $\frac{n}{2} < k \ll n$, where less efficient one-round protocols are possible, our two-round construction may be preferred.

Our results employ a new technique which differs significantly from the prior work. The main novelty in our construction is the design of *non-malleable extractors*, which are an interesting primitive of independent interest. For non-malleability, we consider an attacker who sees a random extractor seed X and produces an arbitrarily related seed X'. We require that the relationship between $R = \mathsf{Ext}(W;X)$ and $R' = \mathsf{Ext}(W;X')$ is "bounded" in some well-defined manner. To our knowledge, this is the first work to explore the (non-)malleability properties of extractors, a problem which is particularly difficult since we must analyze security with respect to a very large class of distributions for W and methods for modifying the seed X. Our main construction of non-malleable extractors is based on the (seemingly unrelated) concept of alternating extraction, recently introduced in [DP07]. Using non-malleable extractors, we show how Alice can authenticate a message to Bob in a simple two-round (challenge-response) protocol. Lastly, we use this message authentication protocol as a tool for our construction of two-round authenticated key agreement.

We also present two orthogonal extensions of our basic scheme. In the first extension, we consider the *fuzzy case* where Alice and Bob have two different but correlated secrets W_A, W_B. In the second extension, we consider the case where the shared secret W is huge (e.g. as in the bounded retrieval model) and hence efficient protocols require *locality* — i.e. Alice and Bob can only access a small portion of W to run their protocol.

2 Notation and Preliminaries

Notation. If W is a probability distribution or a random variable then $w \leftarrow W$ denotes that a value w is sampled randomly according to W. For a randomized algorithm or function f, we use the semicolon to make the randomness explicit i.e. $f(w;r)$ is the output of f with input w using randomness r. Otherwise, we let $f(w)$ denote a random variable for the output of f on the value w. Similarly, for a random variable W, we let $f(W)$ denote the output of f on an input sampled according to W. We use U_ℓ to denote a uniformly random distribution over ℓ bit strings.

Min-entropy and Statistical Distance. The *statistical distance* between two random variables A, B is defined by $\mathbf{SD}(A,B) = \frac{1}{2}\sum_v |\Pr[A=v] - \Pr[B=v]|$. We use $A \approx_\varepsilon B$ as shorthand for $\mathbf{SD}(A,B) \leq \varepsilon$. The *min-entropy* of a random variable W is $\mathbf{H}_\infty(W) \stackrel{\text{def}}{=} -\log(\max_w \Pr[W=w])$. This notion of entropy is useful in cryptography since it measures the *predictability* of W by an adversary. However, cryptographic secrets cannot usually be analyzed in a vacuum and we have to consider the *conditional* predictability of W when sampled according to some joint distribution (W,Z) where the adversary sees Z. Following [DORS08], the correct corresponding notion is *average conditional min entropy* defined by $\widetilde{\mathbf{H}}_\infty(W|Z) \stackrel{\text{def}}{=} -\log\left(\mathbb{E}_{z\leftarrow Z} \max_w \Pr[W=w|Z=z]\right)$. We say that a random variable W is an (n,k)-source if it is distributed over $\{0,1\}^n$ and $\mathbf{H}_\infty(W) \geq k$. We say that $(W|Z)$ is an (n,k) source if W takes values over $\{0,1\}^n$ and $\widetilde{\mathbf{H}}_\infty(W|Z) \geq k$. Several important background lemmas regarding min-entropy and statistical distance are given in Appendix B.

Extractors and MACs: We review two information theoretic primitives that we will use extensively throughout the paper: randomness extractors and (one-time) MACs. A randomness extractor uses a random seed X as a catalyst to extract nearly uniform randomness $R = \mathsf{Ext}(W;X)$ from a weak source W. A message authentication code (MAC) uses a private key R to produce a tag σ for a message μ such that an adversary who sees μ, σ cannot produce a valid tag σ' for a modified message $\mu' \neq \mu$.

Definition 1. *We say that an efficient function* $\mathsf{Ext} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^\ell$ *is an* $(n, k, d, \ell, \varepsilon)$-extractor *if for all* (n, k)*-sources* $(W|Z)$, $(Z, X, \mathsf{Ext}(W; X)) \approx_\varepsilon (Z, X, U_\ell)$ *where* X *is uniform on* $\{0,1\}^d$.

Definition 2. *We say that a family of functions* $\{\mathsf{MAC}_r : \{0,1\}^m \to \{0,1\}^s\}_{r\in\{0,1\}^n}$ *is a* δ*-secure* (one-time) message authentication code (MAC) *if for any* $\mu \neq \mu'$, σ, σ', $\Pr[\mathsf{MAC}_R(\mu) = \sigma \mid \mathsf{MAC}_R(\mu') = \sigma'] \leq \delta$ *where* R *is uniformly random on* $\{0,1\}^n$.

Some further notes about the above definitions and the parameters of known constructions for extractors and MACs are deferred to Appendix A.

3 Interactive Message Authentication

In this section we study the problem of message authentication when Alice and Bob share an arbitrarily weak secret W about which an adversary Eve has some side-information Z. Alice wants to send an authenticated message μ_A to Bob, in the presence of an active attacker Eve, who has complete control over the network and can modify protocol messages arbitrarily. Bob should either correctly receive μ_A, or detect an active attack and quit by outputting $\perp$.

Definition 3. *An* (n, k, m, δ)-message authentication protocol AUTH *is a protocol in which Alice starts with a* source message $\mu_A \in \{0,1\}^m$ *and, at the conclusion of the protocol, Bob outputs a* received message $\mu_B \in \{0,1\}^m \cup \{\perp\}$. *We require the following properties:*
Correctness. *If the adversary Eve is passive then, for any source message* $\mu_A \in \{0,1\}^m$, $\Pr[\mu_B = \mu_A] = 1$.
Security. *If* $(W|Z)$ *is an* (n, k)*-source then, for any source message* $\mu_A \in \{0,1\}^m$ *and any active adversarial strategy employed by Eve,* $\Pr[\mu_B \notin \{\mu_A, \perp\}] \leq \delta$.

For the case of perfectly random secrets W, it is well-known how to solve the above problem using *message authentication codes (MAC)*, where the authentication protocol consists of a single round in which Alice sends her message μ_A along with a *tag* $\sigma = \mathsf{MAC}_W(\mu_A)$. We show that this strategy *does not* (in general) extend to the case of weak secrets. Namely, one-round message authentication protocols are only possible if the entropy of the secret is at least $k > \frac{n}{2}$. In addition, even when this condition does hold, a single-round protocol will have a communication complexity of roughly $n-k$ bits. This lower bound often makes one-round protocols *impossible*, as in the setting of biometrics where the entropy-rate is often $k < \frac{n}{2}$, or *impractical*, as in the Bounded Retrieval Model where a communication complexity of $n - k$ bits would be huge and on the order of several gigabytes. Our lower bound applies to authentication protocols in which Alice can authenticate even a single bit. As mentioned in the introduction, this result can be thought of as a (non-trivial) extension of [DS02] to the setting where Alice and Bob have access to a local (non-shared) source of perfect randomness. The proof of the following theorem appears in Appendix C.

Theorem 4. *Any one-round* (n, k, m, δ)*-message authentication protocol with security* $\delta < \frac{1}{4}$ *must satisfy* $k > \frac{n}{2}$ *and must have a communication complexity of at least* $n - k - 2$ *bits.*

In the rest of this section, we construct an efficient two-round authentication protocol that can tolerate entropy $k \ll \frac{n}{2}$, thus showing that the above lower bound does not extend beyond a single round. Our protocol (see Figure 1) has a simple challenge-response structure; Bob initiates the conversation by sending a random challenge to Alice, who then uses the secret W to compute a response that authenticates her message. In our protocol, the challenge that Bob sends to Alice is a seed X for some randomness extractor Ext. If the adversary does not modify the seed, then Alice and Bob will use it to derive a shared random key $R = \mathsf{Ext}(W; X)$. Alice can then authenticate her message μ_A, by using R as a key for a message authentication code MAC and sending the tag $\sigma = \mathsf{MAC}_R(\mu_A)$ along with

μ_A as her response to Bob. This gives us a very natural construction of a two-round authentication protocol based on an extractor and a MAC. Unfortunately, the construction is not secure in general. The problem is that Eve can modify the extractor seed X to some arbitrarily related value X', causing Alice to derive some incorrect, but possibly related, key $R' = \mathsf{Ext}(W; X')$. Alice then uses R' to (incorrectly) compute her response $\sigma' = \mathsf{MAC}_{R'}(\mu_A)$. In general, the incorrectly computed tag σ' may allow the adversary to forge a valid tag $\tilde{\sigma} = \mathsf{MAC}_R(\mu_B)$ for a new message $\mu_B \neq \mu_A$ under the correct key R. One can think of this as a *related key attack* where Eve learns the tag computed under a related key and forges a tag for a new message under the original key. Therefore, we must somehow restrict the adversarial attacks that Eve can perform by modifying the seed X. We use a two-pronged approach to combat this problem. Firstly, we construct an extractor which has some "non-malleability" property meaning that if an attacker sees a random seed X and comes up with a related seed X' then we bound the relationship between the Bob's key $R = \mathsf{Ext}(W; X)$ and Alice's incorrect key $R' = \mathsf{Ext}(W; X')$. Secondly, we construct special MACs which are resistant to the limited types of related key attacks that our extractor allows. We then plug our special constructions of extractors and a MACs into the framework shown in Figure 1, to construct a two round authentication protocol.

Alice: W, μ_A	Eve: Z	Bob: W
		Sample X. $R = \mathsf{Ext}(W; X)$
	$X' \longleftarrow X$	
$R' = \mathsf{Ext}(W; X')$ $\sigma' \leftarrow \mathsf{MAC}_{R'}(\mu_A)$		
	$(\mu_A, \sigma') \longrightarrow (\mu_B, \tilde{\sigma})$	
		$\tilde{\sigma} \stackrel{?}{=} \mathsf{MAC}_R(\mu_B)$

Figure 1: A Framework for Message Authentication Protocols.

We present two instantiations of the above framework. As our first instantiation, we define *fully non-malleable* extractors, which essentially guarantee that randomness extracted under a modified seed is completely unrelated to that extracted under the original seed. We prove that (surprisingly) such extractors do indeed exist and can achieve very good parameters. We do so using a probabilistic method argument and therefore this approach does not help us in finding an efficient implementation. The strong non-malleability property essentially prevents Eve from performing any kind of related key attack and therefore, in the first approach, we can use standard one-time MACs for the response. In our second approach, we define a weaker non-malleability property that we call *look-ahead* and give an efficient construction of look-ahead extractors. We then construct a new message authentication code which is specifically tailored to withstand the limited types of related key attacks that look-ahead extractors allow.

3.1 Approach 1: Fully Non-Malleable Extractors (non-constructive)

In this section, we define a powerful primitive called a (fully) non-malleable extractor. This is a seeded extractor which takes a weak secret W and extracts randomness R using a seed X. For the non-malleability property, we consider the following attack game. The adversary gets the seed X and comes up with an *arbitrarily related seed* $X' \neq X$. The adversary then learns the value R' extracted from W under the seed X'. We require that the original randomness R still looks *uniformly random even when given* R', and thus the two values are *completely unrelated*.

Definition 5. *A function* $\mathsf{nmExt} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^\ell$ *is a* $(n, k, d, \ell, \varepsilon)$ non-malleable extractor

(NM-EXT) if, for any (n,k)-source $(W|Z)$ and any adversarial function $\mathcal{A}$:

$$(Z, X, \mathsf{nmExt}(W; \mathcal{A}(X,Z)), \mathsf{nmExt}(W;X)) \approx_\varepsilon (Z, X, \mathsf{nmExt}(W; \mathcal{A}(X,Z)), U_\ell)$$

where X is uniformly random over $\{0,1\}^d$ and $\mathcal{A}(X,Z) \neq X$.

Upon seeing the definition, it is not clear if non-malleable extractors can exist at all. In fact, one obvious attack would be for the adversary to choose a random seed X' unrelated to X and thus learn some ℓ bits of information about W from R'. In order for $\mathsf{nmExt}(W;X)$ to then look random, we need to make sure that W still has at least ℓ bits of residual entropy left after ℓ bits are revealed, showing that we need $\ell < \frac{k}{2}$ (i.e. we can extract at most half of the entropy) just to protect against an adversary who sees the value of the extractor at a random and *unrelated* seed X'. Of course, an adversary that can choose an *arbitrarily related* seed X' has significantly more power and there is no immediate reason to believe that we can defend against such an adversary at all. Surprisingly, using the probabilistic method, we show that non-malleable extractors do indeed exist and that the condition $\ell < \frac{k}{2}$ is essentially sufficient. The proof appears in Appendix D.1. It requires a careful analysis of the dependencies introduced by the inclusion of a related-seed attacker $\mathcal{A}$ and thus is significantly more involved than the simple probabilistic method argument for standard extractors.

Theorem 6. *There exists an (n,k,d,ℓ,ε) non-malleable extractor for any integers $n \geq k, d, \ell$ and any $\varepsilon > 0$ as long as $k > 2\ell + 3\log(1/\varepsilon) + \log(d) + 9$ and $d > \log(n-k+1) + 2\log(1/\varepsilon) + 7$.*

Plugging in a non-malleable extractor and a one-time MAC into our main construction (Figure 1) gives us a two-round authentication protocol: Bob picks an extractor seed X, computes $R = \mathsf{nmExt}(W;X)$ and sends X to Alice. Alice receives a (possibly modified) seed X' and computes $R' = \mathsf{nmExt}(W;X')$. She then uses R' as a key to a standard MAC to authenticate her message μ_A to Bob. It is fairly simple to analyze the security of the protocol. If $X' \neq X$ then, by non-malleability, the value R' is unrelated to the random key R and hence the value $\sigma' = \mathsf{MAC}_{R'}(\mu_A)$ will not help the adversary produce a valid tag $\tilde{\sigma}$ under the key R — not even to authenticate Alice's actual message μ_A! On the other hand, if $X' = X$ then $R' = R$ and hence we can rely directly on the security of the MAC to ensure that $\mu_B = \mu_A$. Therefore we get the following theorem and corollary for the existence of two-round message authentication protocols with nearly optimal parameters. See Appendix D.2 and Appendix D.3 for proofs.

Theorem 7. *Assume that nmExt is a (n,k,d,ℓ,ε) non-malleable extractor and that the collection $\{\mathsf{MAC}_r : \{0,1\}^m \to \{0,1\}^s\}$, indexed by keys $r \in \{0,1\}^\ell$, is a δ-secure one-time MAC. Then our construction outlined above gives us a $(n,k,m,2(\delta+\varepsilon))$-message authentication protocol with two rounds of interaction and a communication complexity of $d+s+m$ bits .*

Corollary 1. *There exist (n,k,m,δ)-message authentication protocols with two rounds of interaction for any integers $n \geq k, m$ and any $\delta > 0$ as long as $k > O\left(\log(\log(n)) + \log(m) + \log\left(\frac{1}{\delta}\right)\right)$. Moreover, the communication complexity of such protocols is $m + O\left(\log(n) + \log(m) + \log\left(\frac{1}{\delta}\right)\right)$.*

3.2 Approach 2: Look-Ahead Extractors (efficient construction)

In this section, we define a weaker notion of non-malleability called *look-ahead*. A look-ahead extractor uses a random seed X to extract t *blocks* of randomness $R_1, \ldots, R_t$ from a secret W. Assume that a seed X' is arbitrarily related to X and that the blocks $R'_1, \ldots, R'_t$ are extracted from W using X'. We insist that any *suffix* $R_{i+1}, \ldots, R_t$ of the original sequence looks uniformly random, even when given the *prefix* $R'_1, \ldots, R'_i$ in the related sequence. In other words, the adversary cannot modify the seed and use the (incorrectly) extracted blocks to *look ahead* into the the original sequence of blocks.

Definition 8. *Let* $\mathsf{laExt} : \{0,1\}^n \times \{0,1\}^d \to (\{0,1\}^\ell)^t$ *be a function such that* $\mathsf{laExt}(W;X)$ *outputs* blocks $R_1, \ldots, R_t$ *with* $R_i \in \{0,1\}^\ell$. *We say that* laExt *is a* $(n,k,d,\ell,t,\varepsilon)$-look-ahead extractor *if, for any* (n,k)*-source* $(W|Z)$*, any adversarial function* $\mathcal{A}$ *and any* $i \in \{0,\ldots,t-1\}$*,*

$$\big(Z, X, [R'_1, \ldots, R'_i], [R_{i+1}, \ldots, R_t]\big) \approx_\varepsilon \big(Z, X, [R'_1, \ldots, R'_i], U_{\ell(t-i)}\big) \tag{1}$$

where $[R_1, \ldots, R_t] = \mathsf{laExt}(W;X)$*,* $X' = \mathcal{A}(X,Z)$*,* $[R'_1, \ldots, R'_t] = \mathsf{laExt}(W;X')$*.*

We note that this is a significantly weaker property than full non-malleability. For example, given a random seed X, there might be a related seed X' such that $\mathsf{laExt}(W;X) = \mathsf{laExt}(W;X')$ with high probability. Nevertheless, we will show that look-ahead suffices for our needs. Our construction of a look-ahead extractor is based on the idea of alternating extraction, which was introduced by Dziembowski and Pietrzak in [DP07] as a tool for building an intrusion resilient secret sharing scheme. In the following section we review this concept using our own terminology and present an *alternating-extraction theorem* which captures the main ideas implicit in [DP07], in an abstracted and (slightly) generalized form.

Quentin: Q, S_1		Wendy W
S_1	$\xrightarrow{S_1}$	
	$\xleftarrow{R_1}$	$R_1 = \mathsf{Ext}_w(W;S_1)$
$S_2 = \mathsf{Ext}_q(Q;R_1)$	$\xrightarrow{S_2}$	
	$\xleftarrow{R_2}$	$R_2 = \mathsf{Ext}_w(W;S_2)$
	$\ldots$	
$S_t = \mathsf{Ext}_q(Q;R_{t-1})$	$\xrightarrow{S_t}$	
		$R_t = \mathsf{Ext}_w(W;S_t)$

Figure 2: Alternating Extraction

Alternating Extraction. Assume that two parties, Quentin and Wendy, have values Q, W respectively such that W is kept secret from Quentin and Q is kept secret from Wendy. Let $\mathsf{Ext}_q, \mathsf{Ext}_w$ be randomness extractors (with possibly different parameters) and assume that Quentin also has a random seed S_1 for the extractor Ext_w. The *alternating extraction protocol* (see Figure 2) is an interactive process between Quentin and Wendy, which runs in t iterations. In the first iteration, Quentin sends his seed S_1 to Wendy, Wendy computes $R_1 = \mathsf{Ext}_w(W;S_1)$, sends R_1 to Quentin, and Quentin computes $S_2 = \mathsf{Ext}_q(Q;R_1)$. In each subsequent iteration i, Quentin sends S_i to Wendy, who replies with $R_i = \mathsf{Ext}_w(W;S_i)$, and Quentin computes $S_{i+1} = \mathsf{Ext}_q(Q;R_i)$. Thus Quentin and Wendy together produce the sequence:

$$S_1, R_1 = \mathsf{Ext}_w(W;S_1), S_2 = \mathsf{Ext}_q(Q;R_1), \ldots, S_t = \mathsf{Ext}_q(Q;R_{t-1}), R_t = \mathsf{Ext}_w(W;S_t) \tag{2}$$

The *alternating-extraction theorem* says that there is no better strategy that Quentin and Wendy can use to compute the above sequence. More precisely, let us assume that, in each iteration, Quentin is limited to sending at most s_q bits to Wendy who can then reply by sending at most s_w bits to Quentin where s_q and s_w are much smaller than the entropy of Q, W (preventing Quentin from sending his entire value Q). Then, for any possible strategy cooperatively employed by Quentin and Wendy in the first i iterations, the values $R_{i+1}, R_{i+2}, \ldots, R_t$ look uniformly random to Quentin (and, symmetrically, $S_{i+1}, S_{i+2}, \ldots, S_t$ look random to Wendy). In other words, Quentin and Wendy acting together cannot speed up the process in some clever way so that Quentin would learn R_j (or even distinguish it from random) in fewer than j iterations. We prove the following theorem in Appendix E.1, essentially using the techniques of [DP07].[2]

Theorem 9 (Alternating Extraction). *Let $(W|Z)$ be an (n_w, k_w)-source and Q be an (n_q, k_q)-source independent of W, Z. For any integers s_q, s_w, t, ℓ, let $\mathsf{Ext}_w, \mathsf{Ext}_q$ be extractors with respective parameters*

$$(n_w, k_w - (s_w + \ell)t, \ell, \ell, \varepsilon_w) \quad , \quad (n_q, k_q - (s_q + \ell)t, \ell, \ell, \varepsilon_q)$$

so that the seed size and extracted key length is ℓ in both cases. Let S_1 be uniformly random on $\{0,1\}^\ell$ and define $R_1, S_2, R_2, \ldots, S_t, R_t$ as in equation (2). Let $\mathcal{A}_q(Q, S_1, Z), \mathcal{A}_w(W, Z)$ be interactive machines such that, in each iteration, $\mathcal{A}_q$ sends at most s_q bits to $\mathcal{A}_w$ which replies with at most s_w bits to $\mathcal{A}_q$. Let V_w^i, V_q^i denote the views *of the machines $\mathcal{A}_w, \mathcal{A}_q$ respectively, including their inputs and transcripts of communication, after the first i iterations. Then, for all $0 \le i \le t-1$,*

$$\left(V_q^i, R_{i+1}, R_{i+2}, \ldots, R_t\right) \approx_\varepsilon \left(V_q^i, U_{\ell(t-i)}\right) \qquad \text{and} \qquad \left(V_w^i, S_{i+1}, S_{i+2}, \ldots, S_t\right) \approx_\varepsilon \left(V_w^i, U_{\ell(t-i)}\right) \tag{3}$$

where $\varepsilon = t^2(\varepsilon_w + \varepsilon_q)$.

Construction of a Look-Ahead Extractor. At first it may seem surprising that alternating extraction (which is an interactive protocol) can help us in the construction of a non-malleable extractor (which is a non-interactive primitive). Our construction of a look-ahead extractor is relatively simple. We let $X = (Q, S_1)$ be a seed, and define

$$\mathsf{laExt}(W; (Q, S_1)) \stackrel{\text{def}}{=} R_1, \ldots, R_t. \tag{4}$$

where $R_1, \ldots, R_t$ are generated as in equation (2). Essentially, the extractor uses the seed $X = (Q, S_1)$ to run Quentin's side and the secret W to run Wendy's side in the alternating-extraction protocol for t iterations and outputs all of Wendy's blocks $R_1, \ldots, R_t$ at the conclusion. We use the alternating-extraction theorem to analyze resistance of this construction to malleability attacks. Suppose that a modified seed $X' = (Q', S_1') = \mathcal{A}((Q, S_1), Z)$ is used to extract $R_1', \ldots, R_t'$. Then that corresponds to an adversarial strategy $\mathcal{A}_q$ for Quentin where he runs $\mathcal{A}$ on his inputs, and then continues running the protocol with the values S_1', Q'. Wendy's strategy is unchanged and she sends the values $R_1', \ldots, R_t'$ to Quentin. Note that Quentin's view is therefore $V_q^i = (Z, X, R_1', \ldots, R_i')$ and hence the look-ahead property (equation (1)) follows directly from the alternating-extraction theorem (equation (3)).

Theorem 10. *Given an $(n_w, k_w - (2\ell)t, \ell, \ell, \varepsilon_w)$-extractor Ext_w and an $(n_q, n_q - (2\ell)t, \ell, \ell, \varepsilon_q)$-extractor Ext_q, our construction yields an $(n_w, k_w, n_q + \ell, \ell, t, t^2(\varepsilon_w + \varepsilon_q))$-look-ahead extractor.*

Proof. Follows from the above discussion showing how to construct a strategy $\mathcal{A}_q$ for Quentin given a malleability attacker $\mathcal{A}$. Notice that the strategy $\mathcal{A}_q$ sends $s_q = \ell$ bits in each iteration. Also, we assume that Q is chosen to be uniformly random over $\{0,1\}^{n_q}$ and therefore $k_q = n_q$. The rest of the parameters follow directly from Theorem 9. □

As shown in Appendix E.2, we can plug in the concrete efficient extractor construction of [GUV07] and get the following parameters.

Theorem 11. *For all integers $n \ge k$ and all $\varepsilon > 0$ there exist $(n, k, d, \ell, t, \varepsilon)$-look-ahead extractors as long as*

$$k \ge 2(t+2)\max(\ell, O(\log(n) + \log(t) + \log(1/\varepsilon))) \ge O(t(\ell + \log(n) + \log(t) + \log(1/\varepsilon)))$$

and $d \ge O(t(\ell + \log(n) + \log(t) + \log(1/\varepsilon)))$.

[2] One difference between us and [DP07], is that we need *all of* $R_{i+1}, \ldots, R_t$ to look random and not just R_{i+1}. The other difference is that they should look random even given the view V_q^i which includes Q.

Authentication using Look-Ahead. We will plug the look-ahead extractor into our framework (Figure 1) to construct a message authentication protocol. However, if Eve now modifies the extractor seed during the initial flow then she gets to perform some (limited) *related key attack* and, therefore, we cannot analyze the security of the construction using standard MACs. Instead, we must carefully construct and analyze a new message authentication code *with look-ahead security* – i.e. one which is secure under the types of related key attacks allowed by the look-ahead extractor.

Definition 12. *A family of functions* $\{\mathsf{MAC}_r : \{0,1\}^m \to \{0,1\}^s\}$ *indexed by keys* $r \in (\{0,1\}^\ell)^t$ *is a* $(m, s, \ell, t, \varepsilon, \delta)$-MAC with look-ahead security *if, for any random variables* $R = [R_1, \ldots, R_t], R' = [R'_1, \ldots, R'_t], V$ *which satisfy the look-ahead property:*

$$\big(V, [R'_1, \ldots, R'_i], [R_{i+1}, \ldots, R_t]\big) \approx_\varepsilon \big(V, [R'_1, \ldots, R'_i], U_{(t-i)\ell}\big) \qquad \forall i \in \{0, \ldots, t-1\} \tag{5}$$

any $\mu_A \in \{0,1\}^m$ *and any adversarial function* $\mathcal{A}$*, we have*

$$\Pr\left[\mu_B \neq \mu_A, \mathsf{MAC}_R(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} \sigma' \leftarrow \mathsf{MAC}_{R'}(\mu_A) \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}(V, \sigma') \end{array}\right] \leq \delta$$

It is simple to show that our construction (Figure 1) is a secure message authentication protocol if we plug-in a look-ahead extractor and a MAC with look-ahead security.

Theorem 13. *Plugging a* $(n, k, d, \ell, t, \varepsilon)$*-look-ahead extractor and a* $(m, s, \ell, t, \varepsilon, \delta)$*-MAC with look-ahead security into our framework (Figure 1) yields a* (n, k, m, δ)*-message authentication protocol with a communication complexity of* $d + m + s$ *bits.*

Proof. We can describe Eve through two adversarial functions $\mathcal{A}_1, \mathcal{A}_2$ where $X' = \mathcal{A}_1(X, Z)$ is the function used to modify the initial flow, and $(\mu_B, \tilde{\sigma}) = \mathcal{A}_2(X, Z, \mathsf{MAC}_{R'}(\mu_A))$ is the function used to modify the response flow. Now, for any function $\mathcal{A}_1$ (including ones which can leave the initial flow unmodified) the definition of look-ahead extractors ensures that the variables $V = (X, Z)$, $R = \mathsf{laExt}(W; X)$, $R' = \mathsf{laExt}(W; X')$ satisfy the look-ahead property ((5) in Definition 12). Therefore, Definition 12 ensures that the probability of $\mathcal{A}_2$ successfully producing $(\mu_B, \tilde{\sigma})$ such that $\mu_B \neq \mu_A$ and Bob accepts $(\mu_B, \tilde{\sigma})$ is upper-bounded by δ. □

We now proceed to construct a MAC with look-ahead security. To show the intuition behind our construction, we first (informally) analyze a simple variant for 1 bit messages. For a key $R = [R_1, R_2, R_3, R_4]$, let us define $\mathsf{MAC}_R(0) = [R_1, R_4]$ and $\mathsf{MAC}_R(1) = [R_2, R_3]$. Then, if the adversary learns $\mathsf{MAC}_{R'}(1) = [R'_2, R'_3]$, the random variable R_4 still looks random and so it is hard to predict $\mathsf{MAC}_R(0) = [R_1, R_4]$. On the other hand, if the adversary learns $\mathsf{MAC}_{R'}(0) = [R'_1, R'_4]$, the variable R'_1 is useless in helping predict $[R_2, R_3]$, and R'_4 is too short (only ℓ bits long) to reveal enough information about $[R_2, R_3]$ (which has almost 2ℓ bits of entropy). In the rest of the section, we formalize the above idea and generalize it to longer messages. All proofs appear in Appendix E.

Definition 14. *Given* $S_1, S_2 \subseteq \{1, \ldots, t\}$*, we say that the ordered pair* (S_1, S_2) *is* top-heavy *if there is some integer* j *such that,* $\left|S_1^{\geq j}\right| > \left|S_2^{\geq j}\right|$*, where* $S^{\geq j} \stackrel{def}{=} \{s \in S \mid s \geq j\}$*. Note that it is possible that* (S_1, S_2) *and* (S_2, S_1) *are both top-heavy. For a collection* Ψ *of sets* $S_i \subseteq \{1, \ldots, t\}$ *we say that* Ψ *is* pairwise top-heavy *if every ordered pair* (S_i, S_j) *of sets* $S_i, S_j \in \Psi$ *with* $i \neq j$*, is top-heavy.*

For example, if $S_1 := \{1, 4\}$, $S_2 := \{2, 3\}$, then both of the ordered pairs (S_1, S_2) and (S_2, S_1) are top heavy. Therefore the collection $\Psi = \{S_1, S_2\}$ is pairwise top-heavy. We show that any collection of pairwise top-heavy sets can be used to construct a MAC with look-ahead security.

Lemma 15. *Assume that a collection $\Psi = \{S_1, \ldots, S_{2^m}\}$ of sets $S_i \subseteq \{1, \ldots, t\}$ is* pairwise top-heavy. *Then the family of functions $\mathsf{MAC}_r(\mu) \stackrel{\text{def}}{=} [r_i \mid i \in S_\mu]$, indexed by $r \in (\{0,1\}^\ell)^t$, is a $(m, s, \ell, t, \varepsilon, \delta)$-MAC with look-ahead security where $s = \ell \max_{S_i \in \Psi}(|S_i|)$, $\delta \leq \left(2^{m-\ell} + 2^m\varepsilon\right)$. Furthermore, if there is an efficient mapping of $\mu \in \{0,1\}^m$ to S_μ, then the construction is efficient.*

Therefore, to construct efficient MACs with look-ahead security, we must construct a large collection of sets which is pairwise top-heavy. We generalize our example of $\Psi = \{\ \{1,4\}\ ,\ \{2,3\}\ \}$ to many bits, by mapping an m bit message $\mu = (b_1, \ldots, b_m) \in \{0,1\}^m$ to a subset $S \subseteq \{1, \ldots, 4m\}$ using the function

$$f(b_1, \ldots, b_m) \stackrel{\text{def}}{=} \{4i - 3 + b_i, 4i - b_i \mid i = 1, \ldots, m\} \tag{6}$$

i.e. each bit b_i decides if to include the values $\{4i-3, 4i\}$ (if $b_i = 0$) or the values $\{4i-2, 4i-1\}$ (if $b_i = 1$).

Lemma 16. *The above construction gives us a pairwise top-heavy collection Ψ of 2^m sets $S \subseteq \{1, \ldots, t\}$ where $t = 4m$. Furthermore, the function f is an efficient mapping of $\mu \in \{0,1\}^m$ to S_μ.*

Corollary 2. *We get an $(m, s, \ell, t, \varepsilon, \delta)$-MAC with look-ahead security for any m, ℓ, ε, with $t = 4m$, $s = 4m\ell$, $\delta \leq \left(2^{m-\ell} + 2^m\varepsilon\right)$.*

Plugging in our parameters for look-ahead extractors (Theorem 11) with those for MACs with look-ahead security (Corollary 2), we construct message authentication protocols with the following parameters.

Theorem 17. *We construct an efficient two-round (n, k, m, δ)-message authentication protocol for any integers $n \geq k, m$ and any $\delta > 0$ as long as $k > O(m(m + \log(n) + \log(1/\delta)))$. The protocol has communication complexity $O(m(m + \log(n) + \log(1/\delta)))$. Moreover, the size of the MAC key (and thus the entropy loss of the protocol) is bounded by $\tau = 4m(m + \log(1/\delta))$.*

The parameters of our above construction are vastly sub-optimal for all but very short messages (especially compared to our non-constructive existential results). However, we will see that we can use the above protocol efficiently as building block for authenticated key agreement by authenticating *only a very short message*. In turn, authenticated key agreement will allow us to build an authentication protocol for longer messages. Therefore, in Theorem 21, we will see that we can get efficient two-round message authentication proctors with significantly better parameters by constructing authenticated key agreement protocols first.

4 Authenticated Key Agreement

We now turn to the problem of authenticated key agreement ($\mathsf{IT\text{-}AKA}$). As before, Alice and Bob share a secret W about which Eve has some side-information Z. They would like to run a protocol, in which they agree on a *shared random key*. More concretely, Alice and Bob each have *candidate keys* r_A, r_B respectively, which are initially set to the special value $\perp$. At some point during the protocol execution, Alice can reach a `KeyDerived` state and Bob can reach a `KeyConfirmed` state. Upon reaching either of these states, a party sets its candidate key to some ℓ-bit value (not $\perp$) and does not modify it afterwards. Informally, the `KeyDerived`, `KeyConfirmed` states should be interpreted as follows:
(1) If Alice reaches the `KeyDerived` state, then she possesses a uniformly *random candidate key*, which remains private no matter how the adversary acts during the remainder of the protocol execution. However, she is not sure if her key is shared with Bob, or if Bob is even involved in the protocol execution at all.
(2) If Bob reaches the `KeyConfirmed` state and gets a candidate key r_B, then Alice must have been

involved in the protocol execution, must have reached the `KeyDerived` state, and the two parties have *shared* key $r_A = r_B$ which is private from Eve.

Definition 18. *In a* $(n, k, \ell, \varepsilon, \delta)$-(information theoretic) authenticated key agreement protocol (IT-AKA)*, Alice and Bob have candidate keys* $r_A, r_B \in \{0,1\}^\ell \cup \{\perp\}$ *respectively. For any active adversarial strategy* $\mathcal{A}$ *employed by Eve, let* R_A, R_B *be random variables which denote the values of the candidate keys* r_A, r_B *at the conclusion of the protocol execution and let* T *be a random variable which denotes the transcript of the (entire) protocol execution as seen by Eve. We require that the protocol satisfies the following three properties:*
(Correctness.) *If Eve is passive, then Alice reaches the* `KeyDerived` *state, Bob reaches the* `KeyConfirmed` *state, and* $R_A = R_B$ *(with probability 1).*
(Key Privacy.) *If* $(W|Z)$ *is an* (n, k)*-source then, for any adversarial strategy* $\mathcal{A}$ *employed by Eve, if Alice reaches the* `KeyDerived` *state during the protocol execution, then* $(Z, T, R_A) \approx_\varepsilon (Z, T, U_\ell)$*.*
(Key Authenticity.) *We say that the protocol has* pre-application authenticity *if for any* (n, k)*-source* $(W|Z)$ *and any adversarial strategy* $\mathcal{A}$ *employed by Eve, the probability that Bob reaches the* `KeyConfirmed` *state and* $R_A \neq R_B$ *is at most* δ*. We say that the protocol has* post-application authenticity *if the above holds* even *if the adversary is given* R_A *immediately after Alice reaches the* `KeyDerived` *state.*

Notes on the Definition. To understand the definition, we need to think of key agreement in a broader context where the key is used for some cryptographic task — for example to encrypt and authenticate a message. Generally, the sender (Alice) would like to be assured that her key is private (and will remain private), but she does not need the key to be shared at the time that she prepares/sends her authenticated-ciphertext. On the other hand, the recipient (Bob) would like to know that the key he uses for decryption/validation is the same *shared private* key which was used by Alice. For this reason, we make our definition asymmetric, only requiring that Alice reaches `KeyDerived` (at which point she can prepare/send her authenticated-ciphertext) and Bob alone reaches `KeyConfirmed` (at which point he can validate/decrypt). Notice, that this definition captures and generalizes prior definitions for one-round key agreement protocols ([MW03, DKRS06]) where Alice distills a key r_A on her own, goes into the `KeyDerived` state, and sends a single message to Bob. We therefore also generalize the notion of *pre/post*-application authenticity from [DKRS06], where it was noted that, if Alice wants to use her key r_A immediately after reaching `KeyDerived` (i.e. to encrypt and authenticate a message to Bob), we need to make sure that her use of the key does not *help the adversary Eve break authenticity.* Therefore, we will construct a two-round protocol meeting the stronger post-application authenticity guarantee where, even if the adversary is given (the entire) key r_A, she cannot cause Bob to derive $r_B \neq r_A$. In particular, using this protocol, Alice can encrypt and authenticate a message to Bob in two rounds of interaction.

We begin with a lower-bound showing that single-round authenticated key agreement (even with pre-application security) is essentially impossible when $k < \frac{n}{2}$ and inefficient (in communication complexity) when $\frac{n}{2} < k \ll \frac{n}{2}$.

Theorem 19. *A one-round* $(n, k, \ell, \varepsilon, \delta)$-IT-AKA *with* pre-*application authenticity having key length* $\ell \geq 4$*, and security* $\delta < \frac{1}{2}, \varepsilon < \frac{1}{16}$*, must satisfy* $k > \frac{n}{2}$ *and have a communication complexity is at least* $n - k - 2$ *bits.*

Construction. We proceed to construct an efficient, two-round, IT-AKA protocol where Bob sends a message to Alice, Alice goes into `KeyDerived` and sends a reply to Bob, and Bob goes into `KeyConfirmed`. Our construction uses the message-authentication protocols from Section 3 as building blocks. The main idea behind our construction is fairly simple; Alice uses the authentication protocol to authenticate an extractor seed X_{key} to Bob who then uses it to extract a shared key with Alice. Essentially the same idea was used in [RW03] to construct authenticated key agreement from authentication protocols (with many rounds of interaction). However, the basic idea of authenticating a seed X_{key} might not work in general, since the adversary Eve can potentially learn some information about W which *is dependant* on X_{key} during the course of the authentication protocol, thus compromising the secrecy of the final key. Indeed, to overcome this complication, [RW03] needed to add extra rounds to their construction on top of the authentication protocol. In contrast, we show that this complication does not arise when the authentication protocol follows our framework (Figure 1) and so our construction of IT-AKA as described above and shown in Figure 3 is secure and we do not need additional rounds.

Alice: W	Eve: Z	Bob: W
Sample X_{key} $R_A := \mathsf{Ext}_{\text{key}}(W; X_{\text{key}})$		Sample X_{auth} $R_{\text{auth}} := \mathsf{Ext}_{\text{auth}}(W; X_{\text{auth}})$
	$X'_{\text{auth}} \longleftarrow X_{\text{auth}}$	
`KeyDerived` $R'_{\text{auth}} := \mathsf{Ext}_{\text{auth}}(W; X'_{\text{auth}})$ $\sigma' \leftarrow \mathsf{MAC}_{R'_{\text{auth}}}(X_{\text{key}})$		
	$(X_{\text{key}}, \sigma') \longrightarrow (X'_{\text{key}}, \tilde{\sigma})$	If $\tilde{\sigma} \stackrel{?}{=} \mathsf{MAC}_{R_{\text{auth}}}(X'_{\text{key}})$ `KeyConfirmed` $R_B := \mathsf{Ext}_{\text{key}}(W; X'_{\text{key}})$

Figure 3: Authenticated Key Agreement Protocol

The security of the above construction is easy to explain on an intuitive level. By the security of the authentication protocol, if Bob reaches the `KeyConfirmed` state, then $X'_{\text{key}} = X_{\text{key}}$ and therefore $R_A = R_B$, showing authenticity (even if Eve sees R_A). For privacy, on the other hand, the only information that an active adversary might possibly get about W and which depends on X_{key}, is the tag $\sigma' = \mathsf{MAC}_{R'_{\text{auth}}}(X_{\text{key}})$. However, σ' is independent of W when conditioned on R'_{auth}. Therefore, the keys R_A, R_B are secure as long as there is enough entropy left over in W conditioned on R'_{auth} and Z. We formalize this argument in Theorem 20. We then plug in the parameters for our two authentication protocols (non-constructive and constructive) to state the final parameters achieved by our IT-AKA protocols in Corollaries 3 and 4. The proofs appear in Appendix F.

Theorem 20. *Let* AUTH *be an* (n, k, m, δ)*-message authentication protocol which instantiates our framework with the functions* $\mathsf{Ext}_{\text{auth}}, \mathsf{MAC}$ *such that key size for* MAC *is* τ *bits long. Let* $\mathsf{Ext}_{\text{key}}$ *be an* $(n, k - \tau, d = m, \ell, \varepsilon)$*-extractor. Then the our construction in Figure 3 is an* $(n, k, \ell, \varepsilon, \delta)$-IT-AKA *with* pre-*application authenticity. If we assume that* AUTH *is an* $(n, k - \ell, m, \delta)$*–message authentication protocol, then we get* post-*application authenticity.*

Corollary 3. *There* exists a (possibly inefficient) *two-round* $(n, k, \ell, \varepsilon, \delta)$*-authenticated key agreement protocol with* post-*application authenticity for any integers* $n \geq k$*, any* $\varepsilon > 0, \delta > 0$ *with key length*

$$\ell = k - O(\log(n) + \log(1/\delta) + \log(1/\varepsilon))$$

and communication complexity $O(\log(n) + \log(1/\delta) + \log(1/\varepsilon))$.

Corollary 4. *We* construct an efficient *two-round* $(n, k, \ell, \varepsilon, \delta)$*-authenticated key agreement protocol with* post-*application authenticity for any constant* $\alpha > 0$*, and any integers* $n \geq k$*, any* $\varepsilon > 0, \delta > 0$ *with key length*

$$\ell = (1 - \alpha)k - O\left(\log^2(n) + \log^2(1/\delta) + \log^2(1/\varepsilon)\right)$$

and communication complexity $O\left(\log^2(n) + \log^2(1/\delta) + \log^2(1/\varepsilon)\right)$.

As mentioned at the end of Section 3.2, we can use our construction of IT-AKA (which uses interactive message authentication as a building block) to improve the efficiency of message authentication based on the look-ahead extractor. The idea is to perform key agreement with *post*-application authenticity and let Alice use her key r_A as a key for a *standard* MAC to authenticate a long message efficiently in the second flow. We prove the following theorem in Appendix F.4.

Theorem 21. *We* construct an efficient *two-round* (n, k, m, δ)*-message authentication protocols for any integers* $n \geq k, m$ *and any* $\delta > 0$ *as long as* $k > O\left(\log^2(n) + \log^2(1/\delta) + \log(m)\right)$.

5 Extensions: The Fuzzy Case and Bounded Retrieval Model

The Fuzzy Case. We now extend our result to the "fuzzy case" where Alice and Bob have some highly-correlated, but possibly *unequal*, secrets W_A, W_B respectively. This can happen, for example, when the secret is a biometric and the variables W_A, W_B represent different (but hopefully very similar) scans of the same biometric.[3] Non-interactive (one-round) solutions for this setting in the case of *passive attackers* are called *fuzzy extractors* and were originally studied by [DRS04, DORS08]. For the case of *active attackers*, such solutions are called *robust fuzzy extractors* and were originally constructed by [DKRS06] and improved upon in [KR08]. Of course, such solutions inherit our lower-bound, and require that the entropy of the secrets is at least $k > n/2$. Interactive solutions for this setting, which allow $k \leq n/2$ appear in [RW04] and are optimized in [KR09]. Again, as in the non-fuzzy case, the prior solutions require many rounds of interaction (proportional to the security parameter). We now give a high-level outline for extending our two-round IT-AKA protocol to the fuzzy setting.

In the fuzzy setting, Alice and Bob need to perform *information reconciliation* to agree on the same shared secret. Using terminology from [DORS08], non-interactive information-reconciliation is called a *secure-sketch* and consists of two procedures $(\mathsf{SS}, \mathsf{Rec})$. Bob first computes a *sketch* $\mathsf{Skt} = \mathsf{SS}(W_B)$ of his secret value W_B, and sends this sketch to Alice. Alice then runs an efficient *recovery procedure* to compute Bob's version of the secret $W_B = \mathsf{Rec}(W_A, \mathsf{Skt})$. The sketch is secure if it does not reveal much information about W_B so that, for any Z, $\widetilde{\mathbf{H}}_\infty(W_B|Z, \mathsf{SS}(W_B)) \geq \widetilde{\mathbf{H}}_\infty(W_B|Z) - \alpha$ for some small value α called the *entropy loss*. See [DORS08] for a formal definition of secure sketches and efficient secure sketch constructions for several specific types of correlations of W_A, W_B (e.g. closeness with respect to hamming distance). Also, see the work of [RW04] for a general, but inefficient, construction of secure sketches for arbitrarily correlated variables (based on hash functions). We show how to implement (efficient) two-round authenticated key agreement in the fuzzy setting for any correlation of W_A, W_B for which there is an (efficient) secure sketch construction.

One idea of a construction for this setting is to first perform information-reconciliation (where Bob sends a sketch Skt of his secret to Alice) and then have Alice and Bob run the standard authenticated key agreement protocol using a shared secret W_B.[4] Unfortunately, this may not be secure in general since Eve gets additional attack power by being able to modify the value of the sketch $\mathsf{Skt} = \mathsf{SS}(W_B)$ sent by Bob to Alice.

We argue that the above idea *is* secure when implemented with our IT-AKA protocol based on the alternating-extraction construction of look-ahead extractors. The key realization is that the look-ahead property (see Definition 12) holds between the values $[R'_1, \ldots, R'_t] = \mathsf{laExt}(W'_B; X')$ extracted by Alice and the values $[R_1, \ldots, R_t] = \mathsf{laExt}(W_B; X)$ extracted by Bob, even if Alice uses a modified seed X' *and* a modified secret $W'_B = \mathsf{Rec}(W_A, \mathsf{Skt}')$ where X', Skt' are adversarially chosen based on Skt, X. Intuitively, any such substitution attack translates directly to an adversarial strategy for Quentin in the

[3]For example, Alice is a client who stores an initial scan W_B of her biometric with some server. Later, Alice takes a new scan W_A and would like to agree on a key with the server.

[4]For our constructions, this means that Bob sends Skt during the first round along with the random seed X, and hence preserves the two-round structure.

alternating-extraction protocol and hence cannot break the look-ahead property.

Theorem 22. *Assume that* (W_A, W_B, Z) *is some joint distribution such that* $(W_B|Z)$, $(W_A|Z)$ *are both* (n, k)*-sources and that* $(\mathsf{SS}, \mathsf{Rec})$ *is a secure sketch construction for the joint distribution* (W_A, W_B), *where the size of the sketch is bounded by* α. *Let* Ext_w *be an* $(n, k - \alpha - (2\ell t), \ell, \ell, \varepsilon_w)$*-extractor and* Ext_q *be an* $(n_q, n_q - (2\ell + \alpha)t, \ell, \ell, \varepsilon_q)$*-extractor for some* ℓ, t *and let* laExt *be the look-ahead extractor constructed from* Ext_w *and* Ext_q *as in Theorem 10, using* t *iterations of alternating-extraction. Then we get:*

$$\big(Z, \mathsf{SS}(W_B), X, [R'_1, \ldots, R'_i], [R_{i+1}, \ldots, R_t]\big) \approx_\varepsilon \big(Z, \mathsf{SS}(W_B), X, [R'_1, \ldots, R'_i], U_{\ell(t-i)}\big) \tag{7}$$

where

$$[R_1, \ldots, R_t] = \mathsf{laExt}(W_B; X), (X', \mathsf{Skt}') = \mathcal{A}(Z, \mathsf{SS}(W_B), X), W' = \mathsf{Rec}(W_A, \mathsf{Skt}'), [R'_1, \ldots, R'_t] = \mathsf{laExt}(W'; X')$$

and the achieved security is $\varepsilon \le t^2(\varepsilon_q + \varepsilon_w)$.

Proof. We use the alternating-extraction theorem where, in the honest execution, Quentin uses $X = (Q, S_1)$ and Wendy uses W_B. Let $Z' = (Z, \mathsf{SS}(W_B))$. Then an adversarial strategy in which Eve modifies $X = (Q, S_1), \mathsf{Skt} = \mathsf{SS}(W_B)$ to $X' = (Q', S'_1)$ and Skt' corresponds to a joint adversarial strategy by Quentin and Wendy where Quentin uses $X' = (Q', S'_1)$ and also sends Skt' to Wendy in the first iteration. Wendy samples from the distribution $(W_A|W_B = w_B)$ where w_B is his secret (i.e he samples from what Alice's secret would be conditioned on Bob's value). He then applies $W'_B = \mathsf{Rec}(W_A, \mathsf{Skt}')$ and follows the rest of the alternating-extraction protocol honestly. Notice that Quentin's view in this protocol is $Z', Q, R'_1, \ldots, R'_t$ whose joint distribution is identical to that in the statement of the theorem. Therefore, our theorem follows directly from alternating extraction. For parameters, notice that $\widetilde{\mathbf{H}}_\infty(W_B|Z') \ge k_w - \alpha$ and the communication from Quentin to Wendy is limited to $\ell + \alpha$ bits. □

Informally, since the look-ahead property is all we needed to prove the security of our authentication protocol and finally $\mathsf{IT\text{-}AKA}$, we see that the security of these protocols carries over to the fuzzy setting.

The Bounded Retrieval Model. The Bounded Retrieval Model was first proposed (concurrently) by [Dzi06, CLW06] and has since also been studied by [CDD+07, DP07]. The main idea is to make Alice and Bob share an intentionally *huge* secret key (e.g. 10 GB). The size of the key is crucial in protecting against intrusion attacks where the adversary gets complete control over the storage device through some malware (i.e. a virus or trojan horse) which infiltrates Alice's or Bob's storage. It is assumed that, although the malware has complete access to secret data, it *cannot communicate* too much of it to the adversary (e.g. more than 4 GB), because of limits on bandwidth or implemented security measures against excessive communication. Therefore this scenario falls into our framework where Alice and Bob share a (now huge) secret W about which the adversary has side-information Z, such that the entropy of W given Z is large (e.g. more than 6 GB). Our lower bounds show that, even if the entropy of W is $k \ge n/2$, the communication complexity of non-interactive (i.e. single-round) protocols will be at least $n - k$ bits (e.g. $4GB$), which is unrealistic. Interaction is therefore essential in this setting and, as presented, our protocols already achieve low communication complexity relative to the size of the secret W. However, the current solutions may not be efficient since they require the parties to *read* the entire secret to run the protocol. Therefore, we would like to have more efficient construction which also achieve *locality*: the parties need only read a small number of positions in W to run the protocol. We notice that, in our $\mathsf{IT\text{-}AKA}$ protocol based on the alternating-extraction construction of look-ahead extractors, the secret W is only read by the (standard) extractor $\mathsf{Ext}_{\text{key}}$ and a look-ahead extractor, which is constructed using two (standard) extractors $\mathsf{Ext}_q, \mathsf{Ext}_w$. By substituting *local* extractors (defined and constructed by Vadhan [Vad04]) for all of the above implementations, we get a construction of message authentication and $\mathsf{IT\text{-}AKA}$ protocols which also achieve locality.

Acknowledgements. We thank Leo Reyzin and his student Bhavana Kanukurthi for many helpful insightful discussions about this work. We also thank Stefan Dziebmowski for clarifying some of our confusion about alternating extraction.

References

[AGV09] Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *TCC*, pages 474–495, 2009.

[BB84] Charles H. Bennett and Gilles Brassard. Quantum cryptography: public key distribution and coin tossing. *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.

[BBCM95] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Ueli M. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41(6):1915–1923, 1995.

[BBR88] Charles H. Bennett, Gilles Brassard, and Jean-Marc Robert. Privacy amplification by public discussion. *SIAM J. Comput.*, 17(2):210–229, 1988.

[BD07] Carl Bosley and Yevgeniy Dodis. Does privacy require true randomness? In *TCC*, pages 1–20, 2007.

[BMP00] Victor Boyko, Philip D. MacKenzie, and Sarvar Patel. Provably secure password-authenticated key exchange using diffie-hellman. In *EUROCRYPT*, pages 156–171, 2000.

[BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In *EUROCRYPT*, pages 139–155, 2000.

[CDD+07] David Cash, Yan Zong Ding, Yevgeniy Dodis, Wenke Lee, Richard J. Lipton, and Shabsi Walfish. Intrusion-resilient key exchange in the bounded retrieval model. In *TCC*, pages 479–498, 2007.

[CDF+08] Ronald Cramer, Yevgeniy Dodis, Serge Fehr, Carles Padró, and Daniel Wichs. Detection of algebraic manipulation with applications to robust secret sharing and fuzzy extractors. In *EUROCRYPT*, pages 471–488, 2008.

[CHK+05] Ran Canetti, Shai Halevi, Jonathan Katz, Yehuda Lindell, and Philip D. MacKenzie. Universally composable password-based key exchange. In *EUROCRYPT*, pages 404–421, 2005.

[CLW06] Giovanni Di Crescenzo, Richard J. Lipton, and Shabsi Walfish. Perfectly secure password protocols in the bounded retrieval model. In *TCC*, pages 225–244, 2006.

[DKRS06] Yevgeniy Dodis, Jonathan Katz, Leonid Reyzin, and Adam Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. In *CRYPTO*, pages 232–250, 2006.

[DOPS04] Yevgeniy Dodis, Shien Jin Ong, Manoj Prabhakaran, and Amit Sahai. On the (im)possibility of cryptography with imperfect randomness. In *FOCS*, pages 196–205, 2004.

[DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.

[DP07] Stefan Dziembowski and Krzysztof Pietrzak. Intrusion-resilient secret sharing. In *FOCS*, pages 227–237, 2007.

[DP08] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302, 2008.

[DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *EUROCRYPT*, pages 523–540, 2004.

[DS02] Yevgeniy Dodis and Joel Spencer. On the (non)universality of the one-time pad. In *FOCS*, pages 376–, 2002.

[Dzi06] Stefan Dziembowski. Intrusion-resilience via the bounded-storage model. In *TCC*, pages 207–224, 2006.

[GL01] Oded Goldreich and Yehuda Lindell. Session-key generation using human passwords only. In *CRYPTO*, pages 408–432, 2001.

[GL06] Rosario Gennaro and Yehuda Lindell. A framework for password-based authenticated key exchange *ACM Trans. Inf. Syst. Secur.*, 9(2):181–234, 2006.

[GUV07] Venkatesan Guruswami, Christopher Umans, and Salil P. Vadhan. Unbalanced expanders and randomness extractors from parvaresh-vardy codes. In *IEEE Conference on Computational Complexity*, pages 96–108, 2007.

[KOY01] Jonathan Katz, Rafail Ostrovsky, and Moti Yung. Efficient password-authenticated key exchange using human-memorable passwords. In *EUROCRYPT*, pages 475–494, 2001.

[KR08] Bhavana Kanukurthi and Leonid Reyzin. An improved robust fuzzy extractor. In *SCN*, pages 156–171, 2008.

[KR09] Bhavana Kanukurthi and Leonid Reyzin. Key agreement from close secrets over unsecured channels. In *EUROCRYPT*, 2009. To Appear. Full version at http://eprint.iacr.org/2008/494.

[Mau92] Ueli M. Maurer. Protocols for secret key agreement by public discussion based on common information. In *CRYPTO*, pages 461–470, 1992.

[MP90] James L. McInnes and Benny Pinkas. On the impossibility of private key cryptography with weakly random keys. In *CRYPTO*, pages 421–435, 1990.

[MR04] Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In *TCC*, pages 278–296, 2004.

[MW97] Ueli M. Maurer and Stefan Wolf. Privacy amplification secure against active adversaries. In *CRYPTO*, pages 307–321, 1997.

[MW03] Ueli M. Maurer and Stefan Wolf. Secret-key agreement over unauthenticated public channels iii: Privacy amplification. *IEEE Transactions on Information Theory*, 49(4):839–851, 2003.

[NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *J. Comput. Syst. Sci.*, 52(1):43–52, 1996.

[RW03] Renato Renner and Stefan Wolf. Unconditional authenticity and privacy from an arbitrarily weak secret. In *CRYPTO*, pages 78–95, 2003.

[RW04] Renato Renner and Stefan Wolf. The exact price for unconditionally secure asymmetric cryptography. In *EUROCRYPT*, pages 109–125, 2004.

[Vad04] Salil P. Vadhan. Constructing locally computable extractors and cryptosystems in the bounded-storage model. *J. Cryptology*, 17(1):43–77, 2004.

[Wyn75] A.D. Wyner. The wire-tap channel. *Bell Systems Technical Journal*, 54(8):1355–1387, 1975.

[Zuc97] David Zuckerman. Randomness-optimal oblivious sampling. *Random Struct. Algorithms*, 11(4):345–367, 1997.

A Background on Randomness Extractors and MACs

Our definition of extractor (Definition 1) is slightly different than the standard one. Firstly, we require that the extracted randomness looks uniform, even given the extractor seed X. This is generally referred to as a *strong* randomness extractor in the literature but, since we only consider this notion in our paper, we drop the qualifier "strong". Secondly, extractors are usually defined without the side information Z and only assume that a variable W has *worst-case* min-entropy $\mathbf{H}_\infty(W) \geq k$. We refer to this notion as *worst-case extractor*. In contrast, our default notion of extractors, which we refer to as *average-case extractors* when we wish to differentiate the two notions, follows the definition of [DORS08], where it is assumed that W has *average-case* conditional min-entropy $\widetilde{\mathbf{H}}_\infty(W|Z) \geq k$ and the extracted randomness should be uniform given Z as well. As shown in [DORS08], (essentially following part (a) of Lemma 28) a worst-case extractor is also a good average-case extractor.

Lemma 23. *For any $\delta > 0$ if* Ext *is a* worst-case *$(n, k, d, \ell, \varepsilon)$-extractor then it is also an $(n, k + \log(1/\delta), d, \ell, \varepsilon + \delta)$ extractor.*

We will use the following two recent constructions of randomness extractors. Following Lemma 23 the parameters are the same for worst-case and average case extractors. The first construction, by Guruswami, Umans and Vadhan achieves the following parameters.

Lemma 24. *([GUV07]) For every constant $\alpha > 0$ all integers $n \geq k$ and all $\varepsilon \geq 0$, there is an explicit (efficient) $(n, k, d, \ell, \varepsilon)$-extractor with $\ell = (1-\alpha)k - O(\log(n) + \log(1/\varepsilon))$, and $d = O(\log(n) + \log(1/\varepsilon))$.*

The following extractor also has *locality* meaning that only a small portion of the secret W is accessed during extraction. We use this lemma which appeared in [Vad04] and is based on the the extractor of [Zuc97].

Lemma 25. *([Zuc97, Vad04]) Let $\kappa, \alpha > 0$ be arbitrary constants. Then for every $n \in \mathbb{N}$ and every $\varepsilon > \exp(-n/2^{O(\log^*(n))})$ there is an explicit (efficient) $(n, k, d, \ell, \varepsilon)$-extractor where $k = \rho n$, $d = O(\log(n) + \log(1/\varepsilon))$ and $\ell = (1-\kappa)\alpha n$. Furthermore, the extractor can achieve locality $\tau = (1+\kappa)\ell/\alpha + O(\log(1/\varepsilon))$.*

We also mention that explicit efficient constructions of message authentication codes (based on the polynomial evaluation ε-universal hash function) achieve the following parameters.

Lemma 26. *For any $m, \delta > 0$ there is an efficient δ-secure MAC family $\{\mathsf{MAC}_r : \{0,1\}^m \to \{0,1\}^s\}_{r \in \{0,1\}^n}$ with $s \leq \left(\log(m) + \log\left(\frac{1}{\delta}\right)\right), n \leq 2s$.*

B Background Lemmas for (conditional) Min-Entropy and Statistical Distance

The following two lemmas follows directly from the definition of statistical distance and conditional min entropy respectively.

Lemma 27. *Assume that A, B are random variables such that $A \approx_\varepsilon B$ and f is a (randomized) function. Then $(A, f(A)) \approx_\varepsilon (B, f(B))$.*

Lemma 28. *For any random variable W, $\mathbf{H}_\infty(W) = -\log\left(\max_{\mathcal{A}} \Pr[\mathcal{A}() = W]\right)$. For any random variables W, Z, $\widetilde{\mathbf{H}}_\infty(W|Z) = -\log\left(\max_{\mathcal{A}} \Pr[\mathcal{A}(Z) = W]\right)$. In both cases the maximum is taken over all functions $\mathcal{A}$.*

We will use the following lemma from [DORS08].

Lemma 29. *Let A, B, C be random variables.*
(a) For any $\delta > 0$, $\Pr_{b \leftarrow B}\left[\mathbf{H}_\infty(A|B = b) < \widetilde{\mathbf{H}}_\infty(A|B) - \log\left(\frac{1}{\delta}\right)\right] \leq \delta$.
(b) If B takes on values in a set of size at most 2^λ then $\widetilde{\mathbf{H}}_\infty(A|(B, C)) \geq \widetilde{\mathbf{H}}_\infty((A, B)|C) - \lambda \geq \widetilde{\mathbf{H}}_\infty(A|C) - \lambda$ and, in particular, $\widetilde{\mathbf{H}}_\infty(A|B) \geq \widetilde{\mathbf{H}}_\infty(A) - \lambda$.

We also use the following two lemmas from [DP07].

Lemma 30. *Assume that A, B, C, C' are random variables such that $A \to B \to C$ is a Markov chain and $(B, C) \approx_\varepsilon (B, C')$. Then $(A, B, C) \approx_\varepsilon (A, B, C')$.*

Lemma 31. *Assume that A, B, C, C', F are random variables and f is a function such that $(A, C, f(C, B)) \approx_\varepsilon (A, C, F)$ and $(A, C) \approx_\delta (A, C')$. Then $(A, C', f(C', B)) \approx_{\varepsilon+\delta} (A, C', F)$.*

Lastly, we use the following (slightly more complicated) lemma whose prove we provide.

Lemma 32. *Assume that (A, B, C) are random variables such that $(A, C) \approx_\varepsilon (A, U_q)$ and B is distributed over $\{0,1\}^\lambda$. Then,* $\max_{\mathcal{A}} \Pr[C = \mathcal{A}(A, B)] \leq 2^{\lambda - q} + \varepsilon$.

Proof. For any correlated random variables A, B, C we can write $(A, C, B) \approx_0 (A, C, f(A, C))$ where f is some (possibly inefficient) randomized function whose range is $\{0,1\}^\lambda$. In particular, f samples from the distribution of B conditioned on A, C. Therefore, applying Lemma 27, we get

$$(A, C, B) \approx_0 (A, C, f(A, C)) \approx_\varepsilon (A, U_q, f(A, U_q))$$

Lastly, for any $\mathcal{A}$,

$$\begin{aligned} \Pr[C = \mathcal{A}(A, B)] &\leq \Pr[C = \mathcal{A}(U_q, f(A, U_q))] + \varepsilon \\ &\leq 2^{-\widetilde{\mathbf{H}}_\infty(U_q | A, f(A, U_q))} + \varepsilon \\ &\leq 2^{\lambda - \widetilde{\mathbf{H}}_\infty(U_q | A)} + \varepsilon \\ &\leq 2^{\lambda - q} + \varepsilon. \end{aligned}$$

□

C Lower Bounds for Non-Interactive Protocols

Both of our lower bounds follow as consequences of the following lemma.

Lemma 33. *For any randomized functions* $\mathsf{Auth} : \{0,1\}^n \to \{0,1\}^s$, $\mathsf{Ver} : \{0,1\}^n \times \{0,1\}^s \to \{0,1\}$, *and any values* $0 \leq \rho \leq 1$, *one of the following three conditions holds:*
(1) There is an (n, k)-source W such that $\Pr[\mathsf{Ver}(W, \mathsf{Auth}(W)) = 1] < \rho$.
(2) There is an (n, k)-source W and a value $\sigma \in \{0,1\}^s$ such that $\Pr[\mathsf{Ver}(W, \sigma) = 1] > \rho/2$.
(3) There is an (n, k)-source W such that $\widetilde{\mathbf{H}}_\infty(W | \mathsf{Auth}(W)) \leq \max(0, 2k - n) + \log\left(\frac{1}{\rho}\right) + 2$.

Proof. Let us pick some specific functions $\mathsf{Auth}, \mathsf{Ver}$ and some value ρ. Assume that, for these choices, conditions (1) and (2) do *not* hold. We show that condition (3) must hold.

First, for any $\sigma \in \{0,1\}^s$, let us define $S(\sigma) := \{w \in \{0,1\}^n \mid \Pr[\mathsf{Ver}(w, \sigma) = 1] \geq \rho/2\}$. Essentially $S(\sigma)$ denotes the set of values w under which σ will correctly verify with high probability. Therefore, if for some σ, $|S(\sigma)| \geq 2^k$, then the random variable W which is distributed uniformly on $S(\sigma)$ satisfies condition (2) and we get a contradiction. Hence the size of $S(\sigma)$ is upper bounded by 2^k for each σ.

Assume that the function Auth uses d random coins. Then, for each $w \in \{0,1\}^n, r \in \{0,1\}^d$, we define

$$\tilde{S}(w, r) := S(\mathsf{Auth}(w; r)) = \{\tilde{w} \in \{0,1\}^n \mid \Pr[\mathsf{Ver}(\tilde{w}, \mathsf{Auth}(w; r)) = 1] \geq \rho/2\} \quad (8)$$

We define the predicate $\mathbf{Good}(w, r)$ such that

$$\mathbf{Good}(w, r) \Leftrightarrow \Pr[\mathsf{Ver}(w, \mathsf{Auth}(w; r)) = 1] \geq \rho/2 \Leftrightarrow w \in S(w, r) \quad (9)$$

On a high level, $\mathbf{Good}(w, r)$ indicates that the value $\sigma = \mathsf{Auth}(w; r)$ is likely to verify correctly and, since condition (1) does not hold, we expect that $\mathbf{Good}(w, r)$ occurs with high probability. Specifically, let W be some arbitrary (n, k) source and let R be uniformly distributed over $\{0,1\}^d$. Then, since W does not satisfy condition (1),

$$\begin{aligned} \rho &\leq \Pr[\mathsf{Ver}(W, \mathsf{Auth}(W; R)) = 1] \\ &\leq \Pr[\mathbf{Good}(W, R)] + \Pr[\mathsf{Ver}(W, \mathsf{Auth}(W; R)) = 1 \mid \neg\mathbf{Good}(W, R)] \\ &\leq \Pr[\mathbf{Good}(W, R)] + \rho/2 \\ \Longrightarrow \quad & \mathbf{Good}(W, R) \geq \rho/2 \end{aligned} \quad (10)$$

We now use the above analysis to bound $p = 2^{-\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W;R))}$ by

$$\begin{aligned}
p &= \mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W;R)} \max_w \Pr[W = w \mid \mathsf{Auth}(W;R) = \sigma] && (11)\\
&\geq \mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W;R)} \max_w \left(\Pr[W = w \mid \mathsf{Auth}(W;R) = \sigma, \mathbf{Good}(W,R)] \Pr[\mathbf{Good}(W,R)]\right) \\
&\geq (\rho/2) \mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W;R)} \max_w \Pr[W = w \mid \mathsf{Auth}(W;R) = \sigma, \mathbf{Good}(W,R)] && (12)\\
&\geq (\rho/2) \mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W)} \max_w \Pr[W = w \mid \mathsf{Auth}(W) = \sigma, W \in S(\sigma)] && (13)\\
&\geq (\rho/2) \mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W)} \max_w \Pr[W = w \mid W \in S(\sigma)] && (14)
\end{aligned}$$

where (11) is the definition of conditional min-entropy, (12) follows from the analysis of $\mathbf{Good}(W,R)$ in (10), and (13) follows from the definition of $\mathbf{Good}(w,r)$ in (9).

Now, let us further assume that W is uniformly distributed over some subset $\mathcal{W} \subset \{0,1\}^n$ of size $|\mathcal{W}| = 2^k$. Then, continuing from (14), we get

$$\begin{aligned}
p &\geq (\rho/2) \mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W)} \frac{1}{|S(\sigma)\cap\mathcal{W}|} \\
&\geq (\rho/2) \left(\mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W)} |S(\sigma)\cap\mathcal{W}|\right)^{-1} && (15)
\end{aligned}$$

Where (15) follows by Jensen's inequality. Now we'd like to say that there exists some set $\mathcal{W}$ such that the value $\mathbb{E}_{\sigma\leftarrow\mathsf{Auth}(W;R)} |S(\sigma)\cap\mathcal{W}|$ is small (recall, we define W as the uniform distribution on $\mathcal{W}$). We show that such a set exists using a probabilistic method argument. Let $\mathrm{Sets}(n,k)$ bet these of all subsets $\mathcal{W} \subset \{0,1\}^n$ of size $|\mathcal{W}| = 2^k$. Then, when $\mathcal{W}$ is chosen randomly from $\mathrm{Sets}(n,k)$, we claim that

$$\begin{aligned}
\mathop{\mathbb{E}}_{\mathcal{W}\leftarrow\mathrm{Sets}(n,k)} \left(\mathop{\mathbb{E}}_{\sigma\leftarrow\mathsf{Auth}(W;R)} |S(\sigma)\cap\mathcal{W}|\right) &\leq \mathop{\mathbb{E}}_{\mathcal{W}\leftarrow\mathrm{Sets}(n,k),\sigma\leftarrow\mathsf{Auth}(W)} |S(\sigma)\cap\mathcal{W}| && (16)\\
&\leq 1 + (2^k - 1)\frac{\max_\sigma |S(\sigma)|}{(2^n - 1)} \leq 1 + 2^{2k-n} && (17)
\end{aligned}$$

To see this, we notice that, in the experiment described in the right-hand side of (16), a random set $\mathcal{W}$ is chosen, then a random $w \in \mathcal{W}$ and $r \in \{0,1\}^d$ and we compute $|S(w,r)\cap\mathcal{W}|$. However, a syntactically different but semantically equivalent way of describing such an experiment, would be to first choose a random $w \in \{0,1\}^n$ $r \in \{0,1\}^d$ and compute $S(w,r)$; then choose the remaining $2^k - 1$ elements randomly from $\{0,1\}^n \setminus \{w\}$ to form $\mathcal{W}$. The expected value of each individual remaining element falling into $S(w,r)$ is $|S(w,r)|/(2^n-1)$ and, by the linearity of expectation, we then get the first part of (17). Recalling that $|S(\sigma)| \leq 2^k$ and $k \leq n$, the second part of (17) follows.

Therefore, it follows that there exists some *specific* set $\mathcal{W} \subseteq \{0,1\}^n$ of size 2^k, and hence a corresponding (n,k)-source W, such that (combining (15), (17)) we get

$$p = 2^{-\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W))} \geq \frac{\rho}{2(1+2^{2k-n})}$$

and hence

$$\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W)) \leq \log(1+2^{2k-n}) + \log\left(\frac{1}{\rho}\right) + 1 \leq \max(0, 2k-n) + \log\left(\frac{1}{\rho}\right) + 2.$$

□

C.1 Proof of Theorem 4.

A single-phase protocol consists of Alice sending a message σ to Bob. Let us fix Alice's source message to the bit $\mu_A = 1$, and let us define the *randomized* function Auth which maps Alice's secret w (along with some random coins) to the value σ that she will send to Bob. We also define the *randomized* function $\mathsf{Ver}(w, \sigma)$ used by Bob to verify if σ authenticates the bit 1. Following, Lemma 33, one of the following three conditions must hold: (1) the scheme does not achieve correctness and $\Pr[\mathsf{Ver}(W, \mathsf{Auth}(W))] < \rho$, (2) a message σ which authenticates 1 is easy to guess, or (3) the message σ which authenticated 1 reduces the entropy of the secret by $n-k$ bits. In the case of (2), the adversary can successfully authenticate the bit 1 to Bob without any help from Alice. In the case of (3), if Alice attempts to authenticate the bit 1 to Bob, then the adversary's uncertainty about Alice's secret w is reduced to $k-(n-k) = 2k-n$ bits and, if $k < n/2$, the adversary completely learns w. Hence, upon seing the message σ that authenticates the bit 1, the adversary can forge a message σ' which authenticates the bit 0. This intuition is formalized in the proof below. We prove a slightly more general version of Theorem 4 where we also allow *imperfect correctness* – i.e. Bob is only required to output the correct message μ_A with probability ρ.

Theorem 34. *Any single-phase (n, k, m, δ)-message authentication protocol with correctness ρ and security $\delta < \frac{\rho^2}{4}$ must satisfy $k > \frac{n}{2}$ and must have communication complexity at least $n - k - \log\left(\frac{1}{\rho}\right) - 2$ bits. In particular, when $\rho = 1$ as specified in Definition 3, then security $\delta < \frac{1}{4}$ can only be achieved if $k > \frac{n}{2}$ and with a communication complexity of at least $n-k-2$ bits.*

Proof. As in our discussion, let Auth be the (randomized) functions used by Alice to authenticate the bit 1 to Bob and let Ver be the (randomized) function used by Bob to detect if the received message authenticates 1. Since we have correctness ρ, all (n, k) sources W satisfy $\Pr[\mathsf{Ver}(W, \mathsf{Auth}(W)) = 1] \geq \rho$. By Lemma 33, one of conditions (2) or (3) must then hold.

If condition (2) holds, then there is an (n, k) source W and a value σ such that $\Pr[\mathsf{Ver}(W, \sigma) = 1] \geq \rho/2$. Hence, if the adversary sends σ to Bob, Bob will output $\mu_B = 1$ with probability at least $\rho/2$ and, therefore $\delta \geq \rho/2$. Assuming $\delta < \frac{\rho^2}{4} < \rho/2$, condition (3) must hold. So there is an (n, k)-source W such that $\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W)) \leq \max(0, 2k-n) + \log\left(\frac{1}{\rho}\right) + 2$.

First let us assume that $k < n/2$. Then $2^{-\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W))} \geq \rho/4$. By Lemma 28, there then exists an adversary $\mathcal{A}$ such that $\Pr[\mathcal{A}(\mathsf{Auth}(W)) = W] \geq \rho/4$. Assume that Alice's source message is $\mu_A = 1$. The adversary Eve waits to receive $\sigma = \mathsf{Auth}(W)$, then computes $\tilde{W} \leftarrow \mathcal{A}(\sigma)$ and $\tilde{\sigma}$ to be a randomly computed authentication of the bit 0 using the secret $\tilde{W}$. Then, $\Pr[\tilde{W} = W] \geq \rho/4$ and, by correctness, the probability that Bob outputs $\mu_B = 0$ upon receiving $\tilde{\sigma}$ conditioned on $\tilde{W} = W$ is at least ρ. Hence, Eve succeeds with probability $\delta \geq \frac{\rho^2}{4}$.

Lastly, assume that the communication complexity of the protocol is strictly less than $n - k - \log\left(\frac{1}{\rho}\right) - 2$. Then, $\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W)) > k - (n - k - \log\left(\frac{1}{\rho}\right) - 2) > 2k - n + \log\left(\frac{1}{\rho}\right) + 2$ contradicting our assumption that condition (3) holds.

□

C.2 Proof of Theorem 19.

We again prove a slightly stronger version of the theorem where we also assume imperfect correctness (i.e. the probability that, in an honest execution, Alice reaches `KeyDerived`, Bob reaches `KeyConfirmed` and the parties agree on a key is at least ρ).

Theorem 35. *Any* single-phase *(non-interactive)* $(n, k, \ell, \varepsilon, \delta)$-$\mathsf{IT\text{-}AKA}$ *with* pre-*application authenticity, correctness $\rho > \frac{9}{10}$, key length $\ell \geq 4$, and security $\delta < \frac{\rho}{2}, \varepsilon < \frac{1}{16}$, must satisfy $k > \frac{n}{2}$ and have a* *communication complexity is at least* $n - k - \log\left(\frac{1}{\rho}\right) - 2$ *bits.*

Proof. Without loss of generality, a single-phase protocol has Alice go into `KeyDerived` and send a single message to Bob who goes into the `KeyConfirmed` state. Let Auth be the functions used by Alice to prepare her message for Bob, and Ver be the function which returns 1 if Bob goes into the `KeyConfirmed` state. Then, one of the three conditions of Lemma 33 must hold. Condition (1) cannot hold by the correctness of our protocol. If condition (2) holds, then the adversary can break authenticity by sending σ to Bob without Alice's participation with probability $\delta \geq \rho/2$ and therefore it cannot hold either. Therefore, condition (3) holds and $\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W)) \leq \max(0, 2k - n) + \log\left(\frac{1}{\rho}\right) + 2$. Assuming $k \leq \frac{n}{2}$, we get $\widetilde{\mathbf{H}}_\infty(W|\mathsf{Auth}(W)) \log\left(\frac{1}{\rho}\right) + 2$ and hence, by Lemma 28, there is a function $\mathcal{A}$ such that $\Pr[\mathcal{A}(\mathsf{Auth}(W)) = W] \geq \rho/4$.

Then we construct an attacker $\mathcal{B}(\mathsf{Auth}(W))$ which predicts R_A (given $\mathsf{Auth}(W)$) as follows: run $\tilde{W} = \mathcal{A}(\mathsf{Auth}(W))$, and then follow Bob's procedure using $\tilde{W}, \mathsf{Auth}(W)$. Let $\mathcal{E}_1$ be the even that $\tilde{W} = \mathcal{A}(\mathsf{Auth}(W))$ and $\mathcal{E}_2$ be the event, after Alice sends $\mathsf{Auth}(W)$ in a passive execution, Bob recovers the key $R_B = R_A$. Then the probability that $\mathcal{B}$ succeeds is at least $\Pr[\mathcal{E}_1 \cap \mathcal{E}_2] \geq \Pr[\mathcal{E}_1] + \Pr[\mathcal{E}_2] - 1 \geq \rho/4 + \rho - 1 \geq 1/8$. However, if R_A is at least 4 bits long and ε close to uniform (conditioned on $\mathsf{Auth}(W)$) then, $\Pr[\mathcal{B}(\mathsf{Auth}(W)) = R_A] \leq 1/16 + \varepsilon$. Therefore $\varepsilon \geq 1/16$.

Lastly, we reuse the argument in the proof of Theorem 34 which show that (3) can only hold if the communication complexity is at least $n - k - \log\left(\frac{1}{\rho}\right) - 2$ bits.
□

D Proofs for Authentication Based on Fully Non-Malleable Extractors

D.1 Existence of Non-Malleable Extractors

As with regular extractors, we first define a simpler notion of a worst-case non-malleable extractor (Definition 36) and then show that it implies our standard notion of an (average case) non-malleable extractor in Definition 5.

D.1.1 Existence of Non-Malleable *Worst Case* Extractors

Definition 36. *We say that a function* $\mathsf{nmExt} : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^\ell$ *is a* $(n, k, d, \ell, \varepsilon)$ non-malleable worst-case extractor *if, for any* (n, k)*-source* W*, any adversarial function* $\mathcal{A}$*, we have:*

$$(X, \mathsf{nmExt}(W; \mathcal{A}(X)), \mathsf{nmExt}(W; X)) \approx_\varepsilon (X, \mathsf{nmExt}(W; \mathcal{A}(X, Z)), U_\ell)$$

where X *is uniformly random over* $\{0,1\}^d$ *and* $\mathcal{A}(X) \neq X$*.*

The main theorem of this section will be to show the existence of non-malleable worst-case extractors.

Theorem 37. *There exists an* $(n, k, d, \ell, \varepsilon)$ *non-malleable worst-case extractor as long as*

$$d > \log(n - k + 1) + 2\log(1/\varepsilon) + 5 \quad (18)$$
$$k > 2\ell + 2\log(1/\varepsilon) + \log(d) + 6 \quad (19)$$

We prove Theorem 37 using the probabilistic method showing that a random function R is a non-malleable (worst-case) extractor with overwhelming probability. First, a function $R : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^\ell$ is an $(n, k, d, \ell, \varepsilon)$ non-malleable worst-case extractor if for all distinguishers $\mathcal{D}$, all adversarial function $\mathcal{A}$, all (n, k)-sources W:

$$\Pr[\mathcal{D}(X, R(W, \mathcal{A}(X)), R(W, X)) = 1] - \Pr[\mathcal{D}(X, R(W, \mathcal{A}(X)), U_\ell) = 1] \leq \varepsilon \tag{20}$$

Moreover, we can only consider (n, k)-sources W which are uniformly distributed on some subset $\mathcal{W} \subseteq \{0,1\}^n$ of size $|\mathcal{W}| = 2^k$. This is because if (20) fails on some arbitrary (n, k)-source W then, the uniform distribution on the 2^k elements w in the support of W which maximize

$$\Pr[\mathcal{D}(X, R(w, \mathcal{A}(X)), R(w, X)) = 1] - \Pr[\mathcal{D}(X, R(w, \mathcal{A}(X)), U_\ell) = 1]$$

also causes (20) to fail.

Let us, for now fix some functions $\mathcal{D}, \mathcal{A}$ and a set $\mathcal{W} \subseteq \{0,1\}^n$ of size $|\mathcal{W}| = 2^k$ and let W be uniformly distributed on $\mathcal{W}$. We use the bold-face $\mathbf{R}$ to denote a random variable which is distributed uniformly on the space of all functions $R : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^\ell$.

For each $x \in \{0,1\}^d, u \in \{0,1\}^\ell$, we define

$$\mathsf{Count}(x, u) := \left|\left\{u_2 \in \{0,1\}^\ell \ : \ \mathcal{D}(x, u, u_2) = 1\right\}\right| \tag{21}$$

For each $w \in \mathcal{W}$, $x \in \{0,1\}^d$ we define the following random variables (with randomness coming from the random variable $\mathbf{R}$):

$$\begin{aligned} \mathbf{Left}(w, x) &:= \mathcal{D}(x, \mathbf{R}(w, \mathcal{A}(x)), \mathbf{R}(w, x)) & (22)\\ \mathbf{Right}(w, x) &:= \left(\frac{\mathsf{Count}(x, \mathbf{R}(w, \mathcal{A}(x)))}{2^\ell}\right) & (23)\\ \mathbf{Q}(w, x) &:= \mathbf{Left}(w, x) - \mathbf{Right}(w, x) & (24) \end{aligned}$$

and set

$$\overline{\mathbf{Q}} := \frac{\sum_{w,x} \mathbf{Q}(w, x)}{2^{k+d}} \tag{25}$$

Essentially, $\overline{\mathbf{Q}}$ is a random variable which maps each choice of the function $R \leftarrow \mathbf{R}$ to the value

$$p(R) := \Pr[\mathcal{D}(X, R(W, \mathcal{A}(X)), R(W, X)) = 1] - \Pr[\mathcal{D}(X, R(W, \mathcal{A}(X)), U_\ell) = 1] \tag{26}$$

Therefore, we want to upper bound

$$\Pr[\overline{\mathbf{Q}} > \varepsilon] = \Pr_{R \leftarrow \mathbf{R}}[p(R) > \varepsilon] \tag{27}$$

We notice that, for any w, x, we have $\mathbb{E}[\mathbf{Left}(w, x)] = \mathbb{E}[\mathbf{Right}(w, x)]$ and therefore $\mathbb{E}[\mathbf{Q}(w, x)] = 0$ and $\mathbb{E}[\overline{\mathbf{Q}}] = 0$. However, the values $\mathbf{Q}(w, x)$ are not necessarily independent from each other, preventing us from using a simple Chernoff Bound on (27). For example if $\mathcal{A}(\mathcal{A}(x)) = x$ then

$$\mathbf{Left}(w, x) = \mathcal{D}(x, \mathbf{R}(w, \mathcal{A}(x)), \mathbf{R}(w, x)) \quad \text{and} \quad \mathbf{Left}(w, \mathcal{A}(x)) = \mathcal{D}(x, \mathbf{R}(w, x), \mathbf{R}(w, \mathcal{A}(x)))$$

are not independent and hence neither are $\mathbf{Q}(w, x), \mathbf{Q}(w, \mathcal{A}(x))$. We show that all bad dependance is essentially of this form. More precisely, let us represent the function $\mathcal{A}$ as a directed graph $G = (V, E)$ on the vertex set $V = \{0,1\}^d$ and edges $E := \{(\mathcal{A}(x), x) \ : \ x \in \{0,1\}^d\}$ i.e there is an edge from x' to x iff $\mathcal{A}(x) = x'$. Since $\mathcal{A}$ is a function, the in-degree of each vertex is 1. We show that, if we limit ourselves to values of x contained in a subset of V that does not have cycles then the variables $\mathbf{Q}(x, w)$ have very limited sort of dependence.

Lemma 38. *For $V' \subseteq V$, let $G' \subseteq G$ be a restriction of G to the vertices V' and assume that the graph G' is an* acyclic *subgraph of G. Then the set $\{\mathbf{Q}(w, x)\}_{w \in \mathcal{W}, x \in V'}$ of random variables can be enumerated by $\mathbf{Q}_1, \ldots, \mathbf{Q}_m$ for $m = |V'|2^k$ such that $\mathbb{E}[\mathbf{Q}_i | \mathbf{Q}_1, \ldots, \mathbf{Q}_{i-1}] = 0$ for all $1 \leq i \leq m$.*

Proof. The graph G' is a directed acyclic graph and hence defines a partial order "$\leq$" on the vertices V' so that, if $(x', x) \in V'$ then $x' \leq x$. We use the partial order on V' to define a partial order on the set $\{\mathbf{Q}(w,x)\}_{w\in\mathcal{W},x\in V'}$. Lastly, we can extend this partial order to a total order and thus enumerate the above set as $\mathbf{Q}_1, \ldots, \mathbf{Q}_m$ such that if $x' \leq x$ and $\mathbf{Q}_i = \mathbf{Q}(w,x'), \mathbf{Q}_j = \mathbf{Q}(w,x)$ then $i \leq j$. Now we show that, for all $1 \leq i \leq m$, we have $\mathbb{E}[\mathbf{Q}_i|\mathbf{Q}_1, \ldots, \mathbf{Q}_{i-1}] = 0$. The randomness of these variables comes solely from the choice of $R \leftarrow \mathbf{R}$. We can think of a uniformly random function R as being choosing a random output for every input in the domain of R. Then conditioned on any choice of the value of R for all points *other* than (w, x) we have

$$\mathbb{E}[\mathbf{Q}_i] = \mathbb{E}[\mathbf{Q}(w,x)] = \mathbb{E}\left[\mathcal{D}(x, u', \mathbf{R}(w,x)) - \left(\frac{\mathsf{Count}(x,u')}{2^{\ell}}\right)\right] = 0 \tag{28}$$

Moreover, by the properties of our ordering, the variables $\mathbf{Q}_1, \ldots, \mathbf{Q}_{i-1}$ are independent of $\mathbf{R}(w,x)$ and hence the statement of the lemma follows. □

The good news of Lemma 38 is that restrictions of G which are acyclic do not contain bad dependance. We now show that we can partition the entire vertext set $V = \{0,1\}^d$ into two subsets V_1, V_2 of equal size such that the restriction of G to either of these sets is acyclic.

Lemma 39. *For any directed graph $G = (V,E)$ where all vertices have an ind-degree of 1 and where $|V|$ is even, there is a partition of V into V_1, V_2 such that $|V_1| = |V_2|$ and, letting G_b be the restriction of G to the set V_b, both graphs G_1, G_2 are acyclic.*

Proof. The key realization is that each vertex $v \in V$ can belong to at most one cycle. We can break apart each cycle by placing half the vertices into V_1 and the other half into V_2. We can do this for all cycles (one-by-one) keeping V_1 and V_2 balanced (during this stage, we allow $|V_1| = |V_2| + 1$ to break up cycles of odd lengths). At the conclusion, we will end up with two equally sized sets V_1, V_2 neither of which contains a cycle. □

Now, combining Lemma 38, Lemma 39 we can partition $\{\mathbf{Q}(w,x)\}$ into two (enumerated) sets $\{\mathbf{Q}^1_1, \ldots, \mathbf{Q}^1_m\}, \{\mathbf{Q}^2_1, \ldots, \mathbf{Q}^2_m\}$ where $m = 2^{d-1}$ such that, for $b \in \{1,2\}$, $1 \leq i \leq m$, $\mathbb{E}[\mathbf{Q}^b_i|\mathbf{Q}^b_1, \ldots, \mathbf{Q}^b_{i-1}] = 0$. Let us define the random variables $S^b_i = \sum_{j=1}^{i} \mathbf{Q}^b_j$ for all $b \in \{1,2\}$, $1 \leq i \leq m$. Then (for $b = 1,2$) the sequence $S^b_1, \ldots, S^b_m$ is a martingale. Now, going back to equation (27), we get

$$\begin{aligned}
\Pr[\overline{\mathbf{Q}} > \varepsilon] = \Pr\left[\frac{(S^1_m + S^2_m)}{2^{k+d}} > \varepsilon\right] &\leq \Pr[S^1_m > \varepsilon 2^{k+d-1}] + \Pr[S^2_m > \varepsilon 2^{k+d-1}] &(29)\\
&\leq 2e^{-\frac{1}{16}2^{d+k}\varepsilon^2} &(30)
\end{aligned}$$

Where (30) follows from applying Azuma's inequality to both terms on the right-hand side of (29), and noting that $|S^b_i - S^b_{i-1}| = \mathbf{Q}^b_i \leq 2$. We now use this analysis to prove Theorem 37.

Proof. *(of Theorem 37)* Thus far we have considered some fixed adversary $\mathcal{A}$, distinguisher $\mathcal{D}$ and set $\mathcal{W}$ so that (30) bounds the probability that these are *bad* (i.e. that (20) does not hold for these) for a random function R. We now make this explicit by referring to the random variable $\overline{\mathbf{Q}}$ as $\overline{\mathbf{Q}}(\mathcal{W}, \mathcal{A}, \mathcal{D})$ and will now quantify over all possible sets $\mathcal{W}$ and all functions $\mathcal{A}, \mathcal{D}$. In particular, let us define the event $\mathcal{R}$ that, for a random function $R \leftarrow \mathbf{R}$, there *exists some* set $\mathcal{W}$, adversary $\mathcal{A}$ and distinguisher $\mathcal{D}$ for which $\overline{\mathbf{Q}}(\mathcal{W}, \mathcal{A}, \mathcal{D}) \geq \varepsilon$.

We will apply the union bound over all possible values of $\mathcal{W}, \mathcal{A}, \mathcal{D}$. For ease of exposition, let $N = 2^n, K = 2^k, D = 2^d, L = 2^{\ell}$. Then, there are $\binom{N}{K}$ possible sets $\mathcal{W} \subseteq \{0,1\}^n$ of size $|\mathcal{W}| = 2^k$, there are D^D adversaries $\mathcal{A} : \{0,1\}^d \to \{0,1\}^d$ and there are 2^{DL^2} distinguishers $\mathcal{D} : \{0,1\}^d \times \{0,1\}^{\ell} \times \{0,1\}^{\ell} \to \{0,1\}$. Therefore

$$
\begin{aligned}
\Pr[\mathcal{R}] &\leq \Pr\left[\bigcup \overline{\mathbf{Q}}(\mathcal{W}, \mathcal{A}\,\mathcal{D})\right] \leq \sum \Pr[\overline{\mathbf{Q}}(\mathcal{W}, \mathcal{A}, \mathcal{D})] && (31)\\
&\leq \binom{N}{K} D^D 2^{DM^2} 2e^{-\frac{1}{16}2^{d+k}\varepsilon^2} && (32)\\
&\leq e^{K\left(1+\ln\left(\frac{N}{K}\right)\right)+D(\ln D+\ln(2)M^2)+\ln 2-\frac{1}{16}DK\varepsilon^2} && (33)
\end{aligned}
$$

Now the above is strictly less than 1 if the exponent is less than 0 and therefore it suffices to show that

$$
\left[K\left(1+\ln\left(\frac{N}{K}\right)\right) - \frac{1}{32}DK\varepsilon^2 < 0\right] \text{ and } \left[D(\ln D + \ln(2)M^2) + \ln 2 - \frac{1}{32}DK\varepsilon^2 < 0\right] \quad (34)
$$

and it is easy to check that (34) is satisfied as long as as (18), (19) hold and hence $\Pr[\mathcal{R}] < 1$. This implies that under conditions (18) and (19), there *must exist some particular function* R for which the event $\mathcal{R}$ does not occur and hence this is a non-malleable (worst-case) extractor.

It is easy to see that, with slight degradation of the parameters in (18) and (19), we can in fact ensure that $\Pr[\mathcal{R}]$ is negligible and hence a uniformly random function *is* a non-malleable (worst-case) extractor with overwhelming probability.
□

D.1.2 Worst Case Implies Average Case

Now we simply need to show that a non-malleable (worst-case) extractor is *also* a good non-malleable (average-case) extractor.

Theorem 40. *For any* $\rho > 0$*, if* nmExt *is a* $(n, k - \log\left(\frac{1}{\rho}\right), d, \ell, \varepsilon - \rho)$*-non-malleable* worst-case *extractor then it is also a* $(n, k, d, \ell, \varepsilon)$*-non-malleable* average case *extractor.*

Proof. Let $(W|Z)$ be an arbitrary average-case (n, k)-source. Let $W_z = (W \mid Z = z)$. We call a value z *"bad"* if $\widetilde{\mathbf{H}}_\infty(W_z) < k - \log\left(\frac{1}{\rho}\right)$ and *"good"* otherwise. Then by Lemma 29, $\Pr[Z \text{ is bad }] \leq \rho$. Conditioning on the Z being good,

$$
\begin{aligned}
&\mathbf{SD}((Z, X, \mathsf{nmExt}(W; \mathcal{A}(X,Z)), \mathsf{nmExt}(W;X)) \;\;,\;\; (Z, X, \mathsf{nmExt}(W; \mathcal{A}(X,Z)), U_\ell))\\
\leq\; & \sum_z \Pr[Z=z] \cdot \mathbf{SD}((X, \mathsf{nmExt}(W_z, \mathcal{A}(X,z)), \mathsf{nmExt}(W;X)) \;\;,\;\; (X, \mathsf{nmExt}(W_z, \mathcal{A}(X,z)), U_\ell))\\
\leq\; & \Pr[Z \text{ is bad}] + \sum_{\text{good } z} \mathbf{SD}((X, \mathsf{nmExt}(W; \mathcal{A}_z(X)), \mathsf{nmExt}(W;X)) \;\;,\;\; (X, \mathsf{nmExt}(W; \mathcal{A}_z(X)), U_\ell))\\
\leq\; & \rho + (\varepsilon - \rho) \leq \varepsilon
\end{aligned}
$$

□

D.1.3 Proof of Theorem 6.

Proof. By Theorem 37, we see that $(n, k - \log\left(\frac{1}{\varepsilon/2}\right), d, \ell, \varepsilon/2)$-non-malleable *worst-case* extractors exist if

$$
\begin{aligned}
d &> \log(n-k+1) + 2\log\left(\frac{1}{\varepsilon}\right) + 7\\
k &> 2\ell + 3\log\left(\frac{1}{\varepsilon}\right) + \log(d) + 9
\end{aligned}
$$

By Theorem 40, setting $\rho = \varepsilon/2$, these conditions also guarantee the existence of $(n, k, d, \ell, \varepsilon)$-non-malleable *average-case* extractors. □

D.2 Proof of Theorem 7

Proof. Let us fix a value $\mu_A \in \{0,1\}^m$ and some adversarial strategy used by Eve. Let $\mathcal{E}_1$ be the event that Eve succeeds (i.e. $\mu_B \neq \mu_A$ and $\mathsf{MAC}_R(\mu_B) = \tilde{\sigma}$) and let $\mathcal{E}_2$ be the event that Eve is active during the initial flow (i.e. $X' \neq X$). Then

$$\begin{aligned}
\Pr[\mathcal{E}_1 \cap \mathcal{E}_2] &= \Pr\left[\mathsf{MAC}_R(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} R' = \mathsf{nmExt}(W; \mathcal{A}_1(X,Z)), \sigma' \leftarrow \mathsf{MAC}_{R'}(\mu_A), \\ R = \mathsf{nmExt}(W;X) \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}_2(X,Z,\sigma) \end{array}\right] \\
&\leq \varepsilon + \Pr\left[\mathsf{MAC}_{U_\ell}(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} R' = \mathsf{nmExt}(W; \mathcal{A}_1(X,Z)), \sigma' \leftarrow \mathsf{MAC}_{R'}(\mu_A), \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}_2(X,Z,\sigma) \end{array}\right] \qquad (35) \\
&\leq \varepsilon + \delta \qquad (36)
\end{aligned}$$

where $\mathcal{A}_1$ is some function such that $\mathcal{A}_1(X,Z) \neq X$. Then (35) follows from the definition of a non-malleable extractor and (36) from that of a MAC.

Also

$$\begin{aligned}
\Pr[\mathcal{E}_1 \cap \neg\mathcal{E}_2] &= \Pr\left[\mathsf{MAC}_R(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} R = \mathsf{nmExt}(W;X), \sigma \leftarrow \mathsf{MAC}_R(\mu_A) \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}_2(X,Z,\sigma) \end{array}\right] \\
&\leq \varepsilon + \Pr\left[\mathsf{MAC}_{U_\ell}(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} \sigma \leftarrow \mathsf{MAC}_{U_\ell}(\mu_A) \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}_2(X,Z,\sigma) \end{array}\right] \qquad (37) \\
&\leq \varepsilon + \delta \qquad (38)
\end{aligned}$$

where, again, (37) follows from the definition of a non-malleable extractor and (38) from that of a MAC. Putting the two inequalities together we get $\Pr[\mathcal{E}_1] \leq 2(\varepsilon + \delta)$ as we wanted to show. □

D.3 Proof of Corollary Corollary 1

Proof. We apply Theorem 7 to the achievable parameters of non-malleable extractors from Theorem 6 and those of MACs from Lemma 26. □

E Proofs for Authentication Based on Look-Ahead Extractors

E.1 Proof of the Alternating Extraction Theorem

The main part of Theorem 9 is proved in the following slightly simpler lemma.

Lemma 41. *Let everything be as in Theorem 9, but only assume that* Ext_w *be an* $(n_w, k_w - (s_w)t, \ell, \ell, \varepsilon_w)$*-extractor and* Ext_q *be an* $(n_q, k_q - (s_q)t, \ell, \ell, \varepsilon_q)$*-extractor. Then*

$$\begin{aligned}
\left(V_w^i, S_{i+1}\right) &\approx_{\rho_w(i)} \left(V_w^i, U_\ell\right) \qquad (39) \\
\left(V_q^i, R_{i+1}\right) &\approx_{\rho_q(i)} \left(V_q^i, U_\ell\right) \qquad (40)
\end{aligned}$$

where $\rho_w(i) \stackrel{def}{=} i(\varepsilon_w + \varepsilon_q), \rho_q(i) \stackrel{def}{=} \rho_w(i) + \varepsilon_w$.

Proof. Our proof proceeds by induction. For $i = 0$, S_1 is uniform and independent of $V_w^0 = (W,Z)$ and hence $(V_w^0, S_1) = (V_w^0, U_\ell)$. On the other had, $V_q^0 = (Z, Q, S_1)$ and therefore

$$\left(V_q^0, R_1 = \mathsf{Ext}_w(W; S_1)\right) \approx_{\varepsilon_w} \left(V_q^0, U_\ell\right)$$

since $\widetilde{\mathbf{H}}_\infty(W|(Q,Z)) \geq k_w \geq k_w - (s_w)t$. Hence, the lemma holds for $i = 0$.

Assume that the lemma holds for $i-1$. We proceed in two steps. First we show that $(V_w^i, S_{i+1}) \approx (V_w^i, U_\ell)$ i.e. no matter what message $\mathcal{A}_q$ sends in the ith iteration to $\mathcal{A}_w$, the value S_{i+1} still looks random. For our analysis we introduce several new variables: let msg_q^i be the message sent by $\mathcal{A}_q$ in iteration i, and let $T_q^i = (\mathrm{msg}_q^1, \ldots, \mathrm{msg}_q^i)$. We define msg_w^i and T_w^i analogously. Then

$$
\begin{aligned}
& \left(V_q^{i-1}, R_i\right) \approx_{\rho_q(i-1)} \left(V_q^{i-1}, U_\ell\right) && (41)\\
\Rightarrow\;& \left(V_q^{i-1}, \mathrm{msg}_q^i, R_i, \mathsf{Ext}_q(Q;R_i)\right) \approx_{\rho_q(i-1)} \left(V_q^{i-1}, \mathrm{msg}_q^i, U_\ell, \mathsf{Ext}(Q;U_\ell)\right) && (42)\\
\Rightarrow\;& \left(T_q^i, R_i, \mathsf{Ext}_q(Q;R_i)\right) \approx_{\rho_q(i-1)} \left(T_q^i, U_\ell, \mathsf{Ext}_q(Q;U_\ell)\right) && (43)\\
\Rightarrow\;& \left(T_q^i, R_i, \mathsf{Ext}_q(Q;R_i)\right) \approx_{\rho_q(i-1)+\varepsilon_q} \left(T_q^i, R_i, U_\ell\right) && (44)\\
\Rightarrow\;& \left(W, T_q^i, R_i, \mathsf{Ext}_q(Q;R_i)\right) \approx_{\rho_w(i)} \left(W, T_q^i, R_i, U_\ell\right) && (45)\\
\Rightarrow\;& \left(V_w^i, S_{i+1}\right) \approx_{\rho_w(i)} \left(V_w^i, U_\ell\right) && (46)
\end{aligned}
$$

Equation (41) is given by the inductive hypothesis. Equation (42) follows by Lemma 27 where we apply the function used by $\mathcal{A}_q$ to compute the next message along with the Ext_q function. Equation (43) follows by another application of Lemma 27 where we delete Q from $V_q^{i-1}, \mathrm{msg}_q^i$ to get T_q^i. Equation (44) follows from Lemma 31 and the fact that $|T_q^i| \le (s_q)t$. Equation (45) follows from Lemma 30. Lastly, (46) follows from another application of Lemma 27.

Now, we re-use essentially the same analysis to show $(V_q^i, R_{i+1}) \approx (V_q^i, U_\ell)$

$$
\begin{aligned}
& \left(V_w^i, S_{i+1}\right) \approx_{\rho_w(i)} \left(V_w^i, U_\ell\right) && (47)\\
\Rightarrow\;& \left(V_w^i, \mathrm{msg}_w^i, S_{i+1}, \mathsf{Ext}_w(W;S_{i+1})\right) \approx_{\rho_w(i)} \left(V_w^i, \mathrm{msg}_w^i, U_\ell, \mathsf{Ext}_w(W;U_\ell)\right) && (48)\\
\Rightarrow\;& \left(T_w^i, S_{i+1}, \mathsf{Ext}_w(W;S_{i+1})\right) \approx_{\rho_w(i)} \left(T_w^i, U_\ell, \mathsf{Ext}_w(W;U_\ell)\right) && (49)\\
\Rightarrow\;& \left(T_w^i, S_{i+1}, \mathsf{Ext}_w(W;S_{i+1})\right) \approx_{\rho_w(i)+\varepsilon_w} \left(T_w^i, S_{i+1}, U_\ell\right) && (50)\\
\Rightarrow\;& \left(Q, T_w^i, S_{i+1}, \mathsf{Ext}_w(W;S_{i+1})\right) \approx_{\rho_q(i)} \left(Q, T_w^i, S_{i+1}, U_\ell)\right) && (51)\\
\Rightarrow\;& \left(V_q^i, R_{i+1}\right) \approx_{\rho_q(i)} \left(V_q^i, U_\ell)\right) && (52)
\end{aligned}
$$

Where equations (47) - (52) follow the same reasoning as (41) - (46). □

E.1.1 Proof of Theorem 9

Proof. Given $\mathcal{A}_w, \mathcal{A}_q$ which are restricted to communicating s_w, s_q bits respectively, we construct the machines $\mathcal{A}'_w, \mathcal{A}'_q$ which, on each iteration, run $\mathcal{A}_w, \mathcal{A}_q$ but also, in parallel, run the honest alternating-extraction procedure for Quentin and Wendy. Then $\mathcal{A}'_w, \mathcal{A}'_q$ have communication $s'_w = s_w + \ell$, $s'_q = s_q + \ell$. Applying Lemma 41 to $\mathcal{A}'_w, \mathcal{A}'_q$, we get

$$
\begin{aligned}
\left(V_q^i, R_{i+1}, R_{i+2}, \ldots, R_{t-1}, R_t\right) &\approx_{\rho_q(t-1)} \left(V_q^i, R_{i+1}, R_{i+2}, \ldots, R_{t-2}, R_{t-1}, U_\ell\right) && (53)\\
\left(V_q^i, R_{i+1}, R_{i+2}, \ldots, R_{t-1}, U_\ell\right) &\approx_{\rho_q(t-2)} \left(V_q^i, R_{i+1}, R_{i+2}, \ldots, R_{t-2}, U_{2\ell}\right) && (54)\\
&\cdots\\
\left(V_q^i, R_{i+1}, U_{\ell(t-i+1)}\right) &\approx_{\rho_q(i)} \left(V_q^i, U_{\ell(t-i)}\right) && (55)\\
& && (56)
\end{aligned}
$$

Therefore, by the hybrid argument,

$$
\mathbf{SD}\left(\left(V_q^i, R_{i+1}, \ldots, R_t\right)\ ,\ \left(V_q^i, U_{\ell(t-i)}\right)\right) \le t\rho_q(t-1) \le t^2(\varepsilon_w + \varepsilon_q) \quad (57)
$$

We can use the exact same argument to show that

$$
\mathbf{SD}\left(\left(V_w^i, S_{i+1}, \ldots, S_t\right)\ ,\ \left(V_q^i, U_{\ell(t-i)}\right)\right) \le t\rho_w(t-1) \le t^2(\varepsilon_w + \varepsilon_q) \quad (58)
$$

□

E.2 Proof of Theorem 11.

Proof. By Theorem 10, we need to construct an $(n, k-2\ell' t, \ell', d', \varepsilon' = \varepsilon/2t^2)$-extractor Ext_w where and a $(n', n'-2\ell' t, \ell', d', \varepsilon' \varepsilon/2t^2)$-extractor Ext_q where $\ell' = \max(\ell, d')$. By Lemma 24, such extractors Ext_w can be explicitly constructed for

$$\ell' \leq (k-2\ell' t)/2 - d' \quad \Leftarrow \quad k \geq 2(\ell' + d') + 2\ell' t \quad \Leftarrow \quad k \geq 2(t+2)\max(\ell, d')$$

where $d' = O(\log(n) + \log(1/\varepsilon')) = O(\log(n) + \log(1/\varepsilon) + \log(t))$. Setting $n' = 2(t+2)\max(\ell, d')$ we can get the same parameters for Ext_q. The last part follows since $d = d' + n'$. □

E.3 Proof of Lemma 15

Proof. Let V, R', R be random variables satisfying the look-ahead property of equation (5), and let $\mu_A \in \{0,1\}^m$ be an arbitrary message and $\mathcal{A}$ an arbitrary adversarial function. Then we need to find a bound for:

$$\Pr\left[\mu_B \neq \mu_A, \mathsf{MAC}_R(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} \sigma' \leftarrow \mathsf{MAC}_{R'}(\mu_A) \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}(V, \sigma') \end{array}\right] \tag{59}$$

Let us split $\mathcal{A}$ into two functions $\mathcal{A}_1, \mathcal{A}_2$ where $\mathcal{A}_1$ computes the first argument μ_B and $\mathcal{A}_2$ computes the second argument $\tilde{\sigma}$. Without loss of generality, we may assume that $\mathcal{A}_2$ never outputs $\mu_B = \mu_A$. We also define $\phi(\mu_B, \mu_A)$ for any $\mu_A \neq \mu_B$ to be the (first) value of $j \in \{1, \ldots, t\}$ such that $|S^{\geq j}_{\mu_B}| > |S^{\geq j}_{\mu_A}|$ (which is well defined since Ψ is pairwise top-heavy).

Then we can rewrite (59) as

$$\begin{aligned}
& \Pr\left[\mu_B \neq \mu_A, \mathsf{MAC}_R(\mu_B) = \tilde{\sigma} \;\middle|\; \begin{array}{c} \sigma' \leftarrow \mathsf{MAC}_{R'}(\mu_A) \\ (\mu_B, \tilde{\sigma}) \leftarrow \mathcal{A}(V, \sigma') \end{array}\right] \\
= \;& \Pr\left[[R_i \mid i \in S_{\mu_B}] = \mathcal{A}_2(V, [R'_i \mid i \in S_{\mu_A}]) \;\middle|\; \mu_B \leftarrow \mathcal{A}(V, \mathsf{MAC}_{R'}(\mu_A)) \right] \\
\leq \;& \Pr\left[\exists \mu_B \neq \mu_A \text{ s.t. } [R_i \mid i \in S_{\mu_B}] = \mathcal{A}_2(V, [R'_i \mid i \in S_{\mu_A}])\right] \\
\leq \;& \sum_{\mu_B} \Pr\left[[R_i \mid i \in S_{\mu_B}] = \mathcal{A}_2(V, [R'_i \mid i \in S_{\mu_A}])\right] \\
\leq \;& \sum_{\mu_B} \Pr\left[[R_i \mid i \in S^{\geq j}_{\mu_B}] = \mathcal{A}_2(V, [R'_i \mid i \in S^{\geq j}_{\mu_A}], [R'_i \mid i \in S^{<j}_{\mu_A}])\right] \text{ where } j = \phi(\mu_B, \mu_A) \\
\leq \;& \sum_{\mu_B} \left(2^{-\ell} + \varepsilon\right) \qquad\qquad (60) \\
\leq \;& \left(2^{m-\ell} + 2^m \varepsilon\right)
\end{aligned}$$

where (60) follows from Lemma 32 by setting $A = \left(V, [R'_i \mid i \in S^{<j}_{\mu_A}]\right)$, $B = [R'_i \mid i \in S^{\geq j}_{\mu_A}]$, $C = [R_i \mid i \in S^{\geq j}_{\mu_B}]$. Then $(A, C) \approx_\varepsilon (A, U_q)$ by look-ahead (for some q which depends on μ_A, μ_B) and B takes values in $\{0,1\}^\lambda$ for some λ such that $q - \lambda \geq \ell$. □

E.4 Proof of Lemma 16

Proof. Assume that $\mu_A \neq \mu_B$ and let S_A, S_B be the corresponding sets in Ψ. Let i be the first index for which the bits of μ_A and μ_B disagree: i.e. $b^A_i \neq b^B_i$ where b^A_i, b^B_i is the ith bits of μ_A, μ_B respectively. If $b^A_i = 0$ then, letting $j = 4i$, $|S^{\geq j}_A| = 1 + 2(m-i)$ and $|S^{\geq j}_B| = 2(m-i)$ so (S_A, S_B) is top-heavy. If $b^A_i = 1$ then, letting $j = 4i - 2$, $|S^{\geq j}_A| = 2 + 2(m-i)$ and $|S^{\geq j}_B| = 1 + 2(m-i)$ so again (S_A, S_B) is top-heavy. □

E.5 Proof of Theorem 17.

Proof. By Theorem 13, we need to plug in a $(n, k, d, \ell, t, \varepsilon)$-look-ahead extractor and a $(m, s, \ell, t, \varepsilon, \delta)$-MAC with look-ahead security to get an AUTH protocol. By Corollary 2, we can get such a MAC with message size m and security δ by setting

$$\ell = m + \log(1/\delta) + 1, \varepsilon = \delta/2^{m+1}, t = 4m, s = t\ell = 4m(m + \log(1/\delta) + 1) \tag{61}$$

By Theorem 11, we can construct a look-ahead extractor for any $n \geq k$ and the values ℓ, ε, t above as long as $k, d \geq O(m(m + \log(n) + \log(1/\delta)))$. □

F Proofs for Authenticated Key Agreement

F.1 Proof of Theorem 20

Proof. The correctness property is obvious. For pre-application authenticity, we are analyzing the following (equivalent) experiment. First, a value $\mu_A \leftarrow X_{\text{key}}$ is chosen by Alice (we won't carte that it is random). Then Alice and Bob run an authentication protocol where Alice uses the value μ_A and, if Bob outputs $\mu_B \neq \mu_A$ then the adversary wins. By the security of the authentication protocol this occurs with probability at most δ, proving pre-application authenticity. For post-application authenticity, we must analyze the game where Alice picks $\mu_A \leftarrow X_{\text{key}}$ and the adversary also gets $R_A = \mathsf{Ext}_{\text{key}}(W; \mu_A)$. But this just means that we need to analyze the security of the authentication protocol where the adversary has side information $Z' = (Z, R_A)$. Since $|R_A| = \ell$, we have $\widetilde{\mathbf{H}}_\infty(W|Z') \geq \widetilde{\mathbf{H}}_\infty(W|Z) - \ell \geq k - \ell$. Hence security follows if our authentication protocol is $(n, k - \ell, m, \delta)$ secure.

For privacy:

$$\begin{aligned}
&\mathbf{SD}\Big(\ (Z, X_{\text{auth}}, \mathsf{MAC}_{R'_{\text{auth}}}(X_{\text{key}}), X_{\text{key}}, R_A) \ \ , \ \ (Z, X_{\text{auth}}, \mathsf{MAC}_{R'_{\text{auth}}}(X_{\text{key}}), X_{\text{key}}, U_\ell) \ \Big) \\
&\quad \leq \ \mathbf{SD}\big(\ (Z, X_{\text{auth}}, R'_{\text{auth}}, X_{\text{key}}, R_A) \ \ , \ \ (Z, X_{\text{auth}}, R'_{\text{auth}}, X_{\text{key}}, U_\ell) \ \big) \\
&\quad \leq \ \mathbf{SD}\big(\ (Z', X_{\text{key}}, \mathsf{Ext}_{\text{key}}(W; X_{\text{key}})) \ \ , \ \ (Z', X_{\text{key}}, U_\ell) \ \big) &(62) \\
&\quad \leq \ \varepsilon &(63)
\end{aligned}$$

Where, in (62), $Z' = (Z, X_{\text{auth}}, R'_{\text{auth}})$ and so X_{key} is random and independent of Z'. Moreover

$$\widetilde{\mathbf{H}}_\infty(W|Z') \geq \widetilde{\mathbf{H}}_\infty(W|Z, X_{\text{auth}}) - \tau \geq k - \tau$$

since $|R_{\text{auth}}| = \tau$ and X_{auth} is independent from W. Therefore (63) follows since $\mathsf{Ext}_{\text{key}}$ is an $(n, k - \tau, m, \ell, \varepsilon)$ extractor. □

F.2 Proof of Corollary 3

Proof. By Theorem 20 we need to plug in an $(n, k - \tau, d, \ell, \varepsilon)$-extractor and a $(n, k - \ell, m = d, \delta)$-authentication protocol. Existentially, such extractors are known to exist as long as

$$k > \ell + \tau + O(\log(1/\varepsilon)) \tag{64}$$

and have seeds of length $d = O(\log(n) + \log(1/\varepsilon))$. Furthermore, in Corollary 1, we showed that $(n, k - \ell, d, \delta)$-authentication protocols exist where the MAC key is $\tau = O(\log(m) + \log(1/\delta)) = O(\log(\log(n)) + \log(1/\delta) + \log(1/\varepsilon))$, and require

$$k > \ell + O(\log(n) + \log(d) + \log(1/\delta)) = O(\log(n) + \log(1/\delta) + \log(1/\varepsilon)). \tag{65}$$

Therefore, our bound on ℓ satisfies both (64) and (65). □

F.3 Proof of Corollary 4

Proof. By Theorem 20 we need to plug in an $(n, k-\tau, d, \ell, \varepsilon)$-extractor and a $(n, k-\ell, m=d, \delta)$-authentication protocol. By Lemma 24, such extractors exist for any constant $\alpha > 0$ with $d = O(\log(n) + \log(1/\varepsilon))$. By Theorem 17, for $m = d$ we can get an authentication protocol with $\tau = 4d(d + \log(1/\delta))$. Therefore we can extract at most $\ell = (1-\alpha)k - \tau$ which gets us the bound for ℓ. Lastly, the authentication protocol requires $k > \ell + O(d(d + \log(1/\delta)))$ but that's already implied by our bound on ℓ. □

F.4 Proof of Theorem 21

Proof. We need to argue the security of the scenario where Alice and Bob run a $(n, k, \ell, \varepsilon, \delta_1)$-key agreement protocol for a key r_A of size ℓ to a (standard) δ_2-secure one-time MAC, and Alice then uses this key to authenticate her message (sending the tag $\sigma = \mathsf{MAC}_{r_A}(\mu_A)$ in the second phase of the key agreement protocol, immediately after reaching `KeyDerived`). Correctness is obvious. If Eve breaks security, then either she causes Bob to distill a key $r_B \neq r_A$ or else she forges a tag for the MAC under the key r_A. The first event occurs with probability at most δ_1 (even if Eve was given all of r_A and not just σ). The second event occurs with probability at most $\varepsilon + \delta_2$ by the privacy of r_A and the security of the MAC. Therefore our protocol is $\delta_1 + \delta_2 + \varepsilon$ secure. Setting $\varepsilon = \delta_1 = \delta_2 = \delta/3$ we get the desired security and parameters. □