

Multi-Key Authenticated Encryption with Corruptions: Reductions are Lossy^{*}

Tibor Jager¹, Martijn Stam², Ryan Stanley-Oakes², and Bogdan Warinschi²

¹ Department of Computer Science, Paderborn University,
Warburger Strasse 100,
33098 Paderborn, Germany.
tibor.jager@upb.de

² Department of Computer Science, University of Bristol,
Merchant Venturers Building, Woodland Road,
Bristol, BS8 1UB, United Kingdom.

{martijn.stam, ryan.stanley-oakes, bogdan.warinschi}@bristol.ac.uk

Abstract. We study the security of symmetric encryption schemes in settings with multiple users and realistic adversaries who can adaptively corrupt encryption keys. To avoid confinement to any particular definitional paradigm, we propose a general framework for multi-key security definitions. By appropriate settings of the parameters of the framework, we obtain multi-key variants of many of the existing single-key security notions. This framework is instrumental in establishing our main results. We show that for all single-key secure encryption schemes satisfying a minimal key uniqueness assumption and almost any instantiation of our general multi-key security notion, any reasonable reduction from the multi-key game to a standard single-key game necessarily incurs a linear loss in the number of keys. We prove this result for all three classical single-key security notions capturing confidentiality, authenticity and the combined authenticated encryption notion.

Keywords: Encryption · Multi-Key · Corruption · Reductions · Tightness

1 Introduction

In theory, most symmetric and public key cryptosystems are considered by default in a single-key setting, yet in reality cryptographic ecosystems provide an abundance of keys—and hence targets—for an adversary to attack. Often one can construct a reduction that shows that single-key security implies multi-key security, but typically such a reduction is lossy: an adversary’s multi-key advantage is roughly bounded by the single-key advantage times the number of keys n in the ecosystem. The ramifications of such a loss can be debated [17], but undeniably in a concrete setting with perhaps 2^{30} to 2^{40} keys in circulation, an *actual* loss of 30 to 40 bits of security would be considerable. Therefore the natural question arises to what extent this loss in the reduction is inevitable.

This inevitable loss of reductions from multi-key to single-key has previously been addressed by Bellare et al. [6] when introducing multi-key security for public key schemes. Specifically, they provided a counterexample: namely a pathological encryption scheme that has a small chance (about $\frac{1}{n}$, where n is a parameter) of leaking the key when used in a single-key environment. In a multi-key scenario, where there are n key pairs, insecurity of the scheme is amplified to the point where it becomes a constant. It follows that any *generic* reduction, i.e. a reduction that works for any scheme, from the multi-key to single-key security must lose a factor of about n . A similar example can be concocted for symmetric schemes to conclude that there cannot be a tight generic reduction from a multi-key game to a single-key game for symmetric encryption, i.e. a reduction that works for *all* encryption schemes, since the reduction will not be tight when instantiated by the pathological scheme. However, this does not rule out all reductions, since a tighter reduction could exist that exploits specific features of a certain class of (non-pathological) schemes.

Consider a setting with a security notion G for primitives (e.g. pseudorandomness for blockciphers), a security notion H for constructions (e.g. ciphertext integrity for authenticated encryption), and suppose we are given a specific construction $C[\mathcal{E}]$ building on any instantiation \mathcal{E} of the primitive. A reduction

^{*} An extended abstract of this paper will appear at TCC 2017; this is the full version.

\mathcal{R} would take adversary \mathcal{A} against the H property of the construction and turn it into one against the G property of the primitive. To be *black-box*, the reduction \mathcal{R} should not depend on \mathcal{A} , but instead only use \mathcal{A} 's input/output behaviour. However, when considering black-box reductions, it turns out there are many shades of black. Baecher et al. [4] presented a taxonomy of black-box reductions; the shades of black emerge when considering whether \mathcal{R} may depend on the construction C and/or the primitive \mathcal{E} or not. A *fully* black-box (BBB) reduction works for all C and \mathcal{E} , while partially black-box (NBN) reductions can depend on the specific choice of C and \mathcal{E} .

The pathological encryption schemes used as counterexamples are by nature rather contrived and the one used by Bellare et al. is of dubious security even in the single-key setting [6]. The counterexamples suffice to rule out tight BBB reductions, but they do *not* rule out the existence of potentially large classes of encryption schemes—perhaps practical ones, or even all *secure* ones—for which a tight NBN reduction *does* exist. Clearly, such an NBN reduction could not be generic, but instead would have to exploit some feature of the specific primitive or construction under consideration. Even when the primitive is assumed ‘ideal’ as is common in symmetric cryptology, the relevant reductions typically still depend on the details of the construction at hand, and are therefore not fully (BBB) black-box. Concluding, for secure schemes the relation between single-key and multi-key security is still largely unsettled.

Our Contribution. Focusing on authenticated encryption (AE) schemes, we make two main contributions: a general multi-key security definition including corruptions and lower-bounds on the tightness of black-box (NBN) reductions from the multi-key security to the single-key security of AE schemes.

General Security Definition. The first complication we face is the choice of security notions. As we recall in more detail in Sect. 2.1, there are many different ways of defining single-key security for AE. For instance, confidentiality can be expressed in several (not necessarily equivalent) ways, including left-or-right indistinguishability (LRIND) and ciphertexts being indistinguishable from random strings (IND). Moreover there are different ways of treating nonces; each defines a slightly different security notion.

When moving to a multi-key setting, the water becomes even more muddied, especially when considering adaptive corruptions as we do. Adaptive corruptions allow an adversary to learn some of the keys during the course of the multi-key game; it models the real-life circumstance that not all keys will remain secret and some *will* leak. In this setting, security can be formulated in (at least) two ways: firstly using a hidden bit b_i for each key K_i , with the adversary having to guess the bit b_i for a key K_i that has not been corrupted; and secondly, using a single hidden bit b determining the ‘challenge’ oracles for all n keys (e.g. left or right, real or random) with the adversary having to guess this bit b , under the restriction that no single key gets both corrupted and challenged.

As we explain in Appendix B, these two approaches do not appear to be tightly equivalent to each other. Furthermore, notions that used to be equivalent in the single-key setting suddenly start drifting apart, something previously observed in the multi-instance setting [9]. Again, this creates a bit of a conundrum as to what is the ‘right’ multi-key security notion, where we want to avoid a situation where we show that a reduction loss targeting one security notion is inevitable, while leaving the door open for tight reductions targeting another.

To avoid having to make a choice, we instead provide a general definition for multi-key security game (Definition 7) that allows us to plug in the ‘flavour’ of AE security we desire, and of which the two approaches for dealing with corruptions in a multi-key setting are special cases.

Lower Bounds on the Loss for Simple Reductions. Roughly speaking, we show that for any member G^n of a large class of n -key security games that allow for adaptive corruptions and for most AE schemes $C[\mathcal{E}]$ built on a single-key secure AE scheme \mathcal{E} (including $C[\mathcal{E}] = \mathcal{E}$), any black-box reduction from G^n for $C[\mathcal{E}]$ to a standard single-key security game H^1 for \mathcal{E} incurs a loss that is close to n . By ‘black-box’, we mean *at least* NBN: the reduction must be black-box with respect to the adversary against G^n but can depend on C and \mathcal{E} .




Fig. 1. A roadmap of our results, showing that some reductions between the security notions for authenticated encryption are necessarily lossy. A green arrow $G \rightarrow G'$ indicates that there is a non-lossy reduction from G' to G (so security in the sense of G implies security in the sense of G'). A red arrow $G \rightarrow G'$ indicates that all reductions from G' to G have a loss that is linear in n .

Fig. 1 shows both the logic of our approach and the overall results. The main idea is to first consider a *very weak* n -key security game, K^n , and show that reductions from K^n to H^1 are lossy. Then, for any n -key game G^n that tightly implies K^n , the loss from G^n to H^1 will have to match that from K^n to H^1 (or a contradiction would appear when composing the reduction from K^n to G^n with that from G^n to H^1). Our weak security notion K^n is a 1-out-of- n key recovery game where the adversary first sees encryptions of fixed messages under all n keys, then corrupts all but one key and must try to guess the uncorrupted key. The choice for the three H^1 notions AE-PAS, IND-PAS, and CTI-CPA is inspired by their ubiquity in current AE literature (the naming convention is clarified in Sect. 2.1).

To show for each choice of H^1 that reductions from K^n for $C[\mathcal{E}]$ to H^1 for \mathcal{E} are lossy, we use three *meta-reductions*. Unlike using pathological schemes as counterexamples, meta-reductions can easily deal with NBN reductions that depend on the construction C and scheme \mathcal{E} : a meta-reduction \mathcal{M} *simulates* an ideal adversary \mathcal{A} against $C[\mathcal{E}]$ for a reduction \mathcal{R} and then uses \mathcal{R} to break \mathcal{E} [3, 15, 20]. Then one finds the inevitable loss factor of \mathcal{R} by bounding the advantage of \mathcal{M} (in its interaction with \mathcal{R}) by the advantage of the best possible adversary against \mathcal{E} . We remark that this technique is vacuous for insecure schemes \mathcal{E} as the resulting bound on the advantage of \mathcal{M} is not meaningful.

More precisely, we show that for the three choices of H^1 , any black-box reduction running in time at most t from K^n for $C[\mathcal{E}]$ to H^1 for \mathcal{E} must lose $(\frac{1}{n} + \epsilon)^{-1}$, where ϵ is essentially the maximum advantage in H^1 of an adversary running in time $n \cdot t$. These results hold provided that $C[\mathcal{E}]$ is *key-unique*: given sufficient plaintext-ciphertext pairs the key is *always* uniquely determined. For almost all variants G^n of our general n -key security game, there is a tight reduction from K^n to G^n (Lemma 12); combining this tight reduction with the unavoidable loss from K^n to H^1 shows that any black-box reduction from G^n to H^1 is lossy.

In summary, we show that for almost any variant G^n of the general n -key security game and for $H^1 \in \{\text{AE-PAS}_{\mathcal{E}}^{X,1}, \text{IND-PAS}_{\mathcal{E}}^{X,1}, \text{CTI-CPA}_{\mathcal{E}}^{X,1}\}$, if \mathcal{E} is “secure” in the sense of H^1 and $C[\mathcal{E}]$ is key-unique, then any black-box reduction from G^n to H^1 with a “reasonable” runtime loses approximately n .

Related Work. The idea of using a weak auxiliary security game to prove that reductions are lossy for more meaningful games was pioneered by Bader et al. for public key primitives [3]. Bader et al. considered as their H^1 notion a non-interactive assumption, whereas our H^1 games are highly interactive. The main obstacle here is that our meta-reduction needs to simulate an appropriate environment towards n copies of the reduction, while having access only to a single set of oracles for the considered single-key game. Thus we are forced to devise an additional mechanism that allows the meta-reduction to simulate responses to the oracle queries made by \mathcal{R} and prove that \mathcal{R} cannot distinguish this simulation from the real oracles in its game.

Multi-key security was first considered in the public key setting [6], extending the LRIND-CCA notion to a single-bit multi-key setting without corruptions. A simple hybrid argument shows the loss

of security is at most linear in the number of keys; furthermore this loss is inevitable as demonstrated by a *counterexample*. Relatedly, for many schemes a generic key recovery attack exists whose success probability is linear in both time and the number of keys n [11, 12, 23]. For schemes where this generic key recovery attack is actually the *best* attack (in both the single-key and n -key games), this shows that security in the n -key setting is indeed n times less than in the single-key setting. However, even for very secure schemes it is unlikely that key recovery is the optimum strategy for e.g. distinguishing genuine ciphertexts from random strings.

The danger of ignoring the loss in reductions between security notions is by now widely understood [16, 17] and has served as motivation for work on improved security analysis that avoid the loss of generic reductions. Recent results include multi-user security for Even–Mansour [34], AES-GCM with nonce randomisation [10], double encryption [27], and block ciphers [44].

Tightness is better understood in the public key setting than in the symmetric setting. There are, for instance, many constructions of (identity-based) public-key encryption [6, 14, 18, 25, 29], digital signatures [1, 13, 28, 32, 33, 42], key exchange protocols [2], as well as several different types of lower bounds and impossibility results [19, 22, 24, 30, 36]. We emphasise that, for signature schemes and public key encryption schemes, ‘tightly secure’ means that the reduction from the scheme to some complexity assumption does not incur a multiplicative loss equal to *the number of signing or encryption queries*.

There exist several other previous works describing meta-reductions from interactive problems, such as the *one-more discrete logarithm* (OMDL) problem [20, 24, 36, 43]. However, all these works have in common that they consider a significantly simpler setting, where the reduction is rewound a much smaller number of times (typically only once), and with only a single oracle (the discrete logarithm oracle).

2 Preliminaries

Notation. For any integer $n \geq 1$ we use $[n]$ to denote the set $\{1, \dots, n\}$ and for any $i \in [n]$ we use $[n \setminus i]$ to denote the set $[n] \setminus \{i\}$. For any finite set S we write $x \leftarrow_s S$ to indicate that x is drawn uniformly at random from S . In any security experiment, if an adversary \mathcal{A} has worst-case runtime t , then we say \mathcal{A} is a t -adversary. When \mathcal{A} is clear from the context, we write $t_{\mathcal{A}}$ for its worst case runtime. Since our security notions are *concrete*, rather than asymptotic (as is standard for symmetric cryptography), we loosely use the term “secure” to mean that, for all reasonable values of t , the advantage of any t -adversary in the relevant security game is close to 0. Of course, what constitutes a “reasonable” runtime depends on the model of computation and is beyond the scope of this work.

2.1 Authenticated Encryption

Syntax. Both the syntax and security definitions for symmetric and then authenticated encryption have evolved over the years. We will use the modern perspective where encryption is deterministic and takes in not just a key and a message, but also a nonce, which could be used to provide an explicit form of randomization. Our syntax is summarised in Definition 1 and is a simplification of that used for subtle authenticated encryption [5]. For simplicity, we omit any associated data, though our later results could be extended to that setting; moreover we are not interested in the ‘subtle’ aspect, where decryption might ‘leak’, e.g. unverified plaintext or multiple error symbols.

Definition 1 (Authenticated Encryption). An authenticated encryption scheme is a pair of deterministic algorithms $(\mathcal{E}, \mathcal{D})$ satisfying

$$\begin{aligned} \mathcal{E} &: \mathsf{K} \times \mathsf{N} \times \mathsf{M} \rightarrow \mathsf{C} \\ \mathcal{D} &: \mathsf{K} \times \mathsf{N} \times \mathsf{C} \rightarrow \mathsf{M} \cup \{\perp\} \end{aligned}$$

where K , M , N and C are subsets of $\{0, 1\}^*$ whose elements are called keys, messages, nonces and ciphertexts respectively. The unique failure symbol \perp indicates that C was not a valid encryption under the key K with nonce N .

As is customary, we abbreviate $\mathcal{E}(K, N, M)$ by $\mathcal{E}_K^N(M)$ and $\mathcal{D}(K, N, C)$ by $\mathcal{D}_K^N(C)$ and assume throughout that all authenticated encryption schemes satisfy, for all $K \in \mathcal{K}, N \in \mathcal{N}, M \in \mathcal{M}$ and all $C \in \mathcal{C}$, the following three properties:

1. (*correctness*) $\mathcal{D}_K^N(\mathcal{E}_K^N(M)) = M$,
2. (*tidiness*) $\mathcal{D}_K^N(C) \neq \perp \Rightarrow \mathcal{E}_K^N(\mathcal{D}_K^N(C)) = C$,
3. (*length-regularity*) $|\mathcal{E}_K^N(M)| = \text{enclen}(|M|)$ for some fixed function enclen .

Correctness and tidiness together imply that \mathcal{D} is uniquely determined by \mathcal{E} , allowing us to refer to the pair $(\mathcal{E}, \mathcal{D})$ simply by \mathcal{E} [35].

Single-Key Security Notions. An authenticated encryption scheme should provide both *confidentiality* and *authenticity*. When defining an adversary’s advantage, we separate these orthogonal properties by looking at the IND–PAS and CTI–CPA security games, while also considering their combination AE–PAS in a single game [39]. Below we discuss these notions in more detail, however we defer formal definitions of the relevant games and advantages to the next section, where they will be viewed as a special case of the *multi-key* games given in Definition 7 (cf. Remark 9).

The notions IND–PAS, CTI–CPA and AE–PAS are commonly called IND–CPA, for indistinguishability under chosen plaintext attack; INT–CTXT, for integrity of ciphertexts; and AE, for authenticated encryption (respectively). However, we adhere to the GOAL–POWER naming scheme [5]. It makes explicit that, in the first case, the adversary’s goal is to distinguish between real ciphertexts and random strings (IND, for indistinguishability) without access to any additional oracles (PAS, for passive); in the second case, the adversary’s goal is to forge a well-formed ciphertext (CTI, for ciphertext integrity) and has access to an ‘always-real’ encryption oracle (CPA, for chosen plaintext attack); and in the third case, the adversary tries to either distinguish real ciphertexts from random strings or forge a well-formed ciphertext (AE, for authenticated encryption), without having access to any additional oracles (PAS). For the notions above, we opted for minimal adversarial powers: it is often possible to trade queries to additional oracles (such as a true encryption oracle) for queries to the challenge oracle. We refer to Barwell et al. [5] for an overview of known relations between various notions.

Nonce Usage Convention. All three of the games above have variants according to how nonces may be used by the adversary in the game:

1. In the IV-based setting, denoted IV, the adversary is required to choose nonces uniformly at random for each encryption query.
2. In the nonce-respecting setting, denoted NR, the adversary chooses nonces adaptively for each encryption query, but may never use the same nonce in more than one encryption query.
3. In the misuse-resistant setting, denoted MR, the adversary chooses nonces adaptively for each encryption query and may use the same nonce in more than one encryption query.

Remark 2. The customary definition for IV-based security lets the game select the IVs [35]. We prefer the recent alternative [5] that provides the same interface across the various notions by restricting the class of valid adversaries in the IV-based setting to those who always provide uniformly random nonces in encryption queries. (Note that there is no need to check the distribution of nonces.) This gives a subtly stronger notion, as a reduction will no longer be able to ‘program’ the IV, which it would be allowed to do in the classical definition (cf. [21, 31]).

The results in this paper hold with the alternative, customary formulation of IV-based encryption, with only cosmetic changes to the proof (to take into account the changed interface).

Different Confidentiality Goals. Above we captured the confidentiality goal IND as distinguishing between real ciphertexts and random strings of the appropriate length. However, there are several competing notions to capture confidentiality, all captured by considering a different challenge encryption oracle:

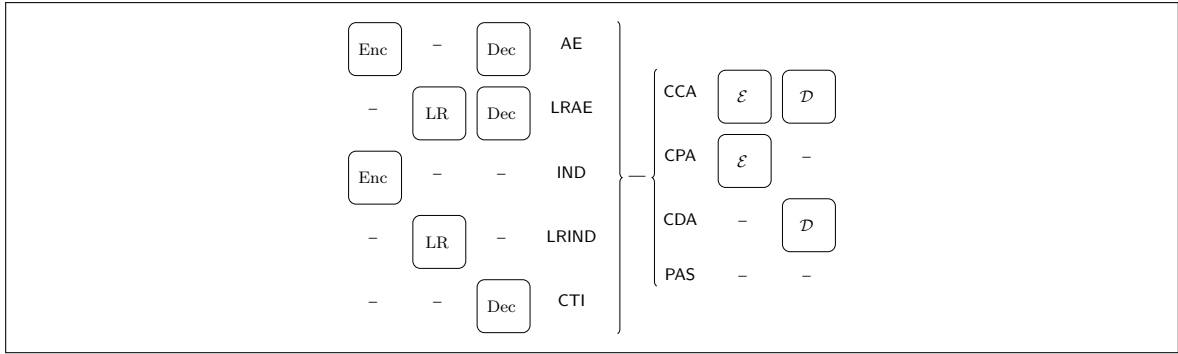


Fig. 2. The oracles available to the adversary for each GOAL-POWER security notion. Formal definitions of each oracle are given in Fig. 4. (Many thanks to Guy Barwell for providing this diagram.)

- In left-or-right indistinguishability (LRIND) the challenge oracle is LR; on input (M_0, M_1, N) , this oracle returns $\mathcal{E}_K^N(M_b)$ (here b is the hidden bit that the adversary must try to learn).
- In real-or-random indistinguishability the challenge oracle, on input (M, N) , returns either $\mathcal{E}_K^N(M)$ or $\mathcal{E}_K^N(\$)$, where $\$$ is a random string of the same length as M .
- In pseudorandom-injection indistinguishability the challenge oracle, on input (M, N) , returns either $\mathcal{E}_K^N(M)$ or $\rho^N(M)$, where ρ is a suitably sampled family of random injections [26, 39].

In the single-key setting, these four notions can be partitioned into two groups of two each, namely left-or-right and real-or-random on the one hand and IND and pseudorandom-injection indistinguishability on the other. Within each group, the two notions can be considered equivalent, as an adversary against one can be turned into an adversary against the other with the same resources and a closely related advantage. Furthermore, security in the IND setting trivially implies security in the LRIND setting, but not vice versa.

Summary. Thus, for each authenticated encryption scheme \mathcal{E} , we potentially obtain $5 \times 4 = 20$ security games (see Fig. 2) and for each we need to consider three classes of adversary depending on nonce usage behaviour. However, for single-key security, we will concentrate on *nine* notions only, namely $G_{\mathcal{E}}^{X,1}$, where

$$G \in \{\text{AE-PAS}, \text{IND-PAS}, \text{CTI-CPA}\},$$

$$X \in \{\text{IV}, \text{NR}, \text{MR}\}$$

and where the 1 in the superscript indicates that these are *single-key* security games.

Remark 3. In this paper we use meta-reductions to analyse reductions from multi-key games to single-key games for authenticated encryption. We show that, for any AE scheme that is secure in a single-key sense, any reduction from the multi-key game to the single-key game is lossy. We do not need to consider equivalent single-key notions separately, as any scheme that is secure according to one notion will be secure according to the other, and one can convert between the single-key games without (significant) additional loss. From this perspective, we can leverage known equivalences as mentioned above. However, the set $\{\text{AE-PAS}, \text{IND-PAS}, \text{CTI-CPA}\}$ does not provide a comprehensive set of meta-reduction results; for that we would have to consider for example LRIND-PAS and IND-CCA as well (the full set would contain eight games). Nevertheless, our results capture the single-key notions that are most commonly used.

2.2 Black-Box Reductions

Informally, a *reduction* \mathcal{R} is an algorithm that transforms an adversary \mathcal{A} in some security game G into an adversary $\mathcal{R}(\mathcal{A})$ in a different security game G' . One hopes that, if the advantage $\text{Adv}^G(\mathcal{A})$ of \mathcal{A}

in G is high, then the advantage $\text{Adv}^{G'}(\mathcal{R}(\mathcal{A}))$ is also high. Here \mathcal{R} breaks some scheme \mathcal{E} , given an adversary \mathcal{A} that breaks a construction $C[\mathcal{E}]$ that uses \mathcal{E} . The construction C is typically fixed, so the reduction \mathcal{R} may depend on it (though to unclutter notation we leave this dependency implicit). On the contrary, when discussing the reduction \mathcal{R} , \mathcal{E} is crucially quantified over some class of schemes \mathcal{C} .

Three properties of a reduction \mathcal{R} are usually of interest: how the resources, specifically run-time, of the resulting adversary $\mathcal{R}(\mathcal{A})$ relate to those of \mathcal{A} ; how the reduction translates the success of \mathcal{A} to that of $\mathcal{R}(\mathcal{A})$; and how ‘lossy’ this translation is, i.e. how $\text{Adv}^{G'}(\mathcal{R}(\mathcal{A}))$ compares to $\text{Adv}^G(\mathcal{A})$. The overall picture for a reduction, especially its loss, strongly depends on the class \mathcal{C} of schemes considered.

Formally, we take into account both the translation \mathbb{S} and the relation \mathbb{T} in runtime into account by considering the quotient of \mathcal{A} and $\mathcal{R}(\mathcal{A})$ ’s *work factors*, themselves defined as the quotient of time over success probability (cf. [3]).

Definition 4. We say that \mathcal{R} is a (\mathbb{S}, \mathbb{T}) reduction from G to G' if for every $t_{\mathcal{A}}$ -adversary \mathcal{A} against G , $\mathcal{R}_{\mathcal{A}}$ is an $\mathbb{T}(t_{\mathcal{A}})$ -adversary against G' and $\text{Adv}^{G'}(\mathcal{R}(\mathcal{A})) = \mathbb{S}(\text{Adv}^G(\mathcal{A}))$. Furthermore, the tightness of a reduction \mathcal{R} relative to the class of schemes \mathcal{C} is defined as

$$\sup_{\mathcal{A}, \mathcal{E}} \frac{\text{Adv}^G(\mathcal{A}) \cdot t_{\mathcal{R}(\mathcal{A})}}{\text{Adv}^{G'}(\mathcal{R}(\mathcal{A})) \cdot t_{\mathcal{A}}} = \sup_{\mathcal{A}, \mathcal{E}} \frac{\mathbb{T}(t_{\mathcal{A}}) \cdot \text{Adv}^G(\mathcal{A})}{t_{\mathcal{A}} \cdot \mathbb{S}(\text{Adv}^G(\mathcal{A}))}$$

where the supremum is taken over all schemes \mathcal{E} in \mathcal{C} and all (valid) adversaries \mathcal{A} against \mathcal{E} .

Remark 5. Our quantification over *valid* adversaries only is inspired by the AE literature’s reliance on only considering adversaries satisfying certain behaviour (e.g. to avoid trivial wins, or distinguishing between IV, NR, and MR settings). In all cases, one can recast to a security game that incorporates checks and balances to deal with arbitrary adversarial behaviour. This recasting is without loss of generality as an adversary in this more general game will be ‘aware’ that it is making a ‘bad’ query *and* this bad behaviour does not influence the state of the game (cf. [8]). Of course, when determining \mathbb{S} we do need to take into account whether the reduction \mathcal{R} preserves validity.

In this paper we are concerned with *simple, black-box reductions*: these are reductions that have only black-box access to adversary \mathcal{A} , and that run \mathcal{A} precisely once (without rewinding). For a (\mathbb{S}, \mathbb{T}) simple reduction \mathcal{R} we have that $\mathbb{T}(t_{\mathcal{A}}) = t_{\mathcal{A}} + t_{\mathcal{R}}$, where $t_{\mathcal{R}}$ is the time taken for whatever additional work \mathcal{R} does. Henceforth, we write $t_{\mathcal{R}}$ for this quantity, whenever \mathcal{R} is a simple reduction.

These reductions compose in the obvious way: if \mathcal{R}_1 is a simple $(\mathbb{S}_1, \mathbb{T}_1)$ reduction from G_1 to G_2 and \mathcal{R}_2 is a simple $(\mathbb{S}_2, \mathbb{T}_2)$ reduction from G_2 to G_3 , then we can construct a simple $(\mathbb{S}_3, \mathbb{T}_3)$ reduction \mathcal{R}_3 from G_1 to G_3 , where $\mathbb{S}_3(\epsilon) = \mathbb{S}_2(\mathbb{S}_1(\epsilon))$ and $\mathbb{T}_3(t) = \mathbb{T}_2(\mathbb{T}_1(t))$.

Bounding Tightness. Precisely evaluating the tightness of a reduction can be difficult, yet to show that for schemes in \mathcal{C} any simple reduction \mathcal{R} loses *at least* some factor L , it suffices to show that for any \mathcal{R} there exists a scheme $\mathcal{E} \in \mathcal{C}$ and a valid adversary \mathcal{A} such that

$$\frac{\text{Adv}^G(\mathcal{A})}{\text{Adv}^{G'}(\mathcal{R}(\mathcal{A}))} \geq L. \quad (1)$$

Indeed, the desired lower bound follows since, for simple reductions, $\mathbb{T}(t_{\mathcal{A}}) \geq t_{\mathcal{A}}$.

We briefly discuss two distinct techniques to establish a bound such as the one above, in which the order of quantifiers is $(\forall \mathcal{R} \exists \mathcal{E} \exists \mathcal{A})$:

- *Counterexample* ($\exists \mathcal{E} \forall \mathcal{A} \forall \mathcal{R}$). Here, one shows that there exists a scheme $\mathcal{E} \in \mathcal{C}$ such that for any adversary \mathcal{A} and any reduction \mathcal{R} , inequality 1 is satisfied. One drawback of such results is that they only imply the desired lowerbound for a class of schemes \mathcal{C} containing \mathcal{E} ; tighter reductions might be possible in the class $\mathcal{C}' := \mathcal{C} \setminus \{\mathcal{E}\}$. Moreover, if the counterexample scheme \mathcal{E} is an artificially insecure scheme (e.g. the one used by Bellare et al. [6]), then the lowerbound might not hold within the class of *secure* schemes, which are obviously of greater significance in practice.

- *Meta-reduction Lowerbound* ($\forall \mathcal{E} \exists \mathcal{A} \forall \mathcal{R}$). For any $\mathcal{E} \in \mathcal{C}$, this technique constructs an idealised adversary \mathcal{A} with advantage 1 and then shows, via a meta-reduction simulating \mathcal{A} , that any simple reduction interacting with \mathcal{A} must have advantage at most L^{-1} , yielding inequality 1. Thus we show that the loss is a property of the reduction \mathcal{R} , and not of the particular choice of $\mathcal{E} \in \mathcal{C}$. The results in this paper, using the meta-reduction approach, hold when \mathcal{C} is any non-empty subset of the class of secure schemes that satisfy the key uniqueness assumption. Since \mathcal{C} could contain just one element \mathcal{E} , our results show that even a reduction that is tailored to the specific details of \mathcal{E} cannot be tight. On the other hand, our results are not directly comparable to those of Bellare et al. [6], since the artificially insecure scheme used in their counterexample does not belong to any class \mathcal{C} we consider here.

Remark 6. An alternative definition of tightness might consider only ‘reasonable’ adversaries \mathcal{A} in the supremum, namely those for which $t_{\mathcal{A}}$ is not too large. Our meta-reduction approach would not work in this setting, since the idealised adversary \mathcal{A} we construct has an extremely large (and wholly unfeasible) runtime as it performs an exhaustive search over all possible keys. Nevertheless, reductions \mathcal{R} that are black-box with respect to \mathcal{A} have no way of ‘excluding’ such unrealistic adversaries and so we feel it is not reasonable to exclude them in the definition of tightness. We remark that unrealistic adversaries are not uncommon in the meta-reduction literature [3].

3 Multi-Key Security Notions

Multi-Key Security with Adaptive Corruptions. In the single-key case, the challenge oracles depend on a single hidden bit b and it is the job of the adversary to try and learn b . The straightforward generalization [6] to a multi-key setting (with n keys) is to enrich all the oracles to include the index $i \in [n]$ of the key K_i that will then be used by the oracle. Thus the challenge oracles for distinct keys will all depend on the same single hidden bit b .

However, in a realistic multi-key setting, an adversary might well learn some of the keys. For instance, consider the situation where an attacker passively monitors millions of TLS connections and adaptively implants malware on particular endpoint devices in order to recover the session keys for those devices. We still want security for those keys that have not been compromised; the question is how to appropriately model multi-key security.

There are two natural approaches to model multi-key security games in the presence of an adaptive corruption oracle Cor that, on input $i \in [n]$, returns the key K_i . The approaches differ in how they avoid trivial wins that occur when the adversary corrupts a key that was used for a challenge query. In one approach, the same bit is used for the challenge queries throughout, but the adversary is prohibited from using the same index i for both a corruption and challenge query (cf. [37]). In another approach, for each index i there is an independent hidden bit b_i to guess and the adversary has to specify for which uncorrupted index its guess b' is intended (cf. [9]).

As far as we are aware, these two approaches have not been formally compared; moreover we could not easily establish a tight relationship between them. However, as we show, both options lead to a reduction loss linear in n . To do so, we will use a novel way of formalizing a multi-key security game with adaptive corruptions that encompasses *both* options mentioned above.

In our generalised game (Definition 7) there are n independently, uniformly sampled random bits b_1, \dots, b_n . Each challenge query from the adversary must specify two indices, $i, j \in [n]$, such that the response to the query depends on key K_i and hidden bit b_j . The two ‘natural’ multi-key games are special cases of this general game: in the single-bit game the adversary is restricted to challenge queries with $j = 1$, whereas in the multi-bit game only challenge queries with $i = j$ are allowed.

Our impossibility results hold regardless how the hidden bits are used: we only require that for any $i \in [n]$ there exists *some* $j \in [n]$ such that the adversary can make a challenge query corresponding to K_i and b_j . In other words, our impossibility results hold provided that the adversary can win the game by ‘attacking’ any of the n keys in the game, not just some subset of the keys.

Experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$:

$K_1, \dots, K_n \leftarrow \mathfrak{K}$
 $b_1, \dots, b_n \leftarrow \mathfrak{s}\{0, 1\}$
 $(j, b'_j) \leftarrow \mathcal{A}^{\mathcal{O}}$
 Return $(b'_j = b_j)$

Fig. 3. The $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ games, where $X \in \{\text{IV}, \text{NR}, \text{MR}\}$, $n \geq 1$, $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$ and $\text{POWER} \in \{\text{CCA}, \text{CPA}, \text{CDA}, \text{PAS}\}$. The oracles \mathcal{O} available to the adversary always include the corruption oracle Cor ; the other oracles depend on GOAL and POWER , as indicated in Fig. 2.

<p>Oracle $\text{Enc}(i, j, M, N)$: if $b_j = 0$, $C \leftarrow \mathcal{E}_{K_i}^N(M)$ else $C \leftarrow \mathfrak{s}\{0, 1\}^{\text{encLen}(M)}$ return C</p>	<p>Oracle $\mathcal{E}(i, M, N)$: return $\mathcal{E}_{K_i}^N(M)$</p>
<p>Oracle $\text{LR}(i, j, M_0, M_1, N)$: $C \leftarrow \mathcal{E}_{K_i}^N(M_{b_j})$ return C</p>	<p>Oracle $\mathcal{D}(i, C, N)$: return $\mathcal{D}_{K_i}^N(C)$</p>
<p>Oracle $\text{Dec}(i, j, C, N)$: if $b_j = 0$, $M \leftarrow \mathcal{D}_{K_i}^N(C)$ else $M \leftarrow \perp$ return M</p>	<p>Oracle $\text{Cor}(i)$: return K_i</p>

Fig. 4. Oracles for the $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ security games. Without loss of generality, we assume that all oracles return \perp if the input arguments do not belong to the relevant sets. For example, the \mathcal{E} oracle will return \perp on any input (i, M, N) that is not a member of $[n] \times \mathfrak{M} \times \mathfrak{N}$.

Definition 7 (Security of AE). *Let*

$$\begin{aligned}
 &\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}, \\
 &\text{POWER} \in \{\text{CCA}, \text{CPA}, \text{CDA}, \text{PAS}\}, \\
 &X \in \{\text{IV}, \text{NR}, \text{MR}\}, \\
 &n \geq 1.
 \end{aligned}$$

Then for any authenticated encryption scheme \mathcal{E} and adversary \mathcal{A} , the advantage of \mathcal{A} against \mathcal{E} with respect to $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ is defined as

$$\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER}_{\mathcal{E}}^{X,n}}(\mathcal{A}) := 2 \cdot \Pr \left[\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A}) = 1 \right] - 1,$$

where the experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$ is defined in Fig. 3, with the oracles' behaviour shown in Fig. 4 and their GOAL-POWER -dependent availability in Fig. 2 (all games have access to Cor).

Whenever the experiment $G = \text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$ is clear from the context, we write $\text{Adv}^G(\mathcal{A})$ for the advantage of \mathcal{A} in experiment G .

The outline games are deliberately kept simple, but are trivial to win: if \mathcal{A} corrupts a key K_i and then issues a challenge query corresponding to K_i and a hidden bit b_j , then it is trivial for \mathcal{A} to compute b_j from the response to the query; successfully 'guessing' b_j does not represent a meaningful attack. In our formal syntax, we say j is *compromised* iff there is some $i \in [n]$ such that \mathcal{A} has issued a query $\text{Cor}(i)$ and \mathcal{A} has also issued some challenge query of the form $\text{Enc}(i, j, -, -)$, $\text{LR}(i, j, -, -, -)$ or $\text{Dec}(i, j, -, -)$. We disallow such trivial wins.

Relatedly, we follow the AE literature in disallowing certain combinations of queries that lead to trivial wins (*prohibited queries*), or that are inconsistent with the nonce notion under consideration.

Experiment $\text{KEYREC}_{\mathcal{E}}^{M,n}(\mathcal{A})$:

```

 $K_1, \dots, K_n \leftarrow \mathcal{K}$ 
for  $i$  in  $1, \dots, n$ ,
  for  $j$  in  $1, \dots, l$ ,
     $N_{i,j} \leftarrow \mathcal{N}$ 
     $C_{i,j} \leftarrow \mathcal{E}_{K_i}^{N_{i,j}}(M_j)$ 
 $(i^*, st) \leftarrow \mathcal{A}_1 \left( (C_{i,j}, N_{i,j})_{i \in [n], j \in [l]} \right)$ 
 $K^* \leftarrow \mathcal{A}_2 \left( (K_i)_{i \in [n \setminus i^*]}, st \right)$ 
Return  $(K_{i^*} = K^*)$ 

```

Fig. 5. Key recovery game with n keys and the hard-coded messages M_1, \dots, M_l . Without loss of generality, we separate the adversary \mathcal{A} into two components \mathcal{A}_1 and \mathcal{A}_2 .

Without loss of generality, we also disallow queries where the response from the oracle can be computed by the adversary directly without making the query, e.g. using correctness (*pointless queries*). The relevant—and standard—definitions are given in Appendix A. Combining the various restrictions leads to the notion of valid adversaries (cf. Remark 5), as summarised in Definition 8 below.

Definition 8 (Valid Adversaries). *An adversary against $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ is valid iff:*

1. *it does not output (j, b'_j) where j was compromised;*
2. *it does not make pointless or prohibited queries;*
3. *it uses nonces correctly with respect to X .*

Remark 9 (Recovering the Single-Key Security Notions). Setting $n = 1$ in Definition 7 yields formal definitions of the *single-key* security games for authenticated encryption, albeit with a more complicated interface than one is used to: the specification of i and j becomes redundant, as does the corruption oracle for valid adversaries. Indeed, to simplify notation in the case $n = 1$, we often omit i and j from the queries made, refer to the hidden bit b_1 as b , and only expect a simple guess b' by an adversary.

Relations among Multi-Key Notions. We discuss the relations between different single-user and multi-user security notions in Appendix B.

Key Recovery Notions. For our meta-reduction, we use an auxiliary, key recovery game $\text{KEYREC}_{\mathcal{E}}^{M,n}$ (Definition 10). Here there are n unknown keys and the adversary is provided with encryptions under each of the keys of the hard-coded (and hence known) messages $M \in \mathcal{M}^l$, using known, yet random, nonces. Then the adversary provides an index $i^* \in [n]$, learns the $n - 1$ keys $(K_i)_{i \in [n \setminus i^]}$ and tries to guess the uncorrupted key.

Definition 10. *For any integers $n, \ell \geq 1$, messages $M = (M_1, \dots, M_\ell) \in \mathcal{M}^\ell$, AE scheme \mathcal{E} and any adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, the advantage of \mathcal{A} against $\text{KEYREC}_{\mathcal{E}}^{M,n}$ is defined as*

$$\text{Adv}_{\mathcal{E}}^{\text{KEYREC}, M, n}(\mathcal{A}) := \Pr \left[\text{KEYREC}_{\mathcal{E}}^{M, n}(\mathcal{A}) = 1 \right],$$

where the experiment $\text{KEYREC}_{\mathcal{E}}^{M, n}(\mathcal{A})$ is given in Fig. 5.

Of course, it might be the case that it is impossible to win the key recovery game with certainty, since there could be more than one key that ‘matches’ the messages, nonces and ciphertexts. For our tightness results, we need to assume that there is some reasonably small l and some messages M_1, \dots, M_l such that the key recovery game corresponding to M_1, \dots, M_l can be won with certainty; we call this the *key uniqueness* property; its definition is below.

Definition 11. Let \mathcal{E} be an authenticated encryption scheme. Suppose there is some integer $l \geq 1$ and certain messages $M_1, \dots, M_l \in \mathbb{M}$ such that, for all keys $K \in \mathbb{K}$ and all nonces $N_1, \dots, N_l \in \mathbb{N}$,

$$\left\{ K' \in \mathbb{K} : \mathcal{E}_{K'}^{N_i}(M_i) = \mathcal{E}_K^{N_i}(M_i) \text{ for all } i \in 1, \dots, l \right\} = \{K\}.$$

Then we say \mathcal{E} is M -key-unique, where $M = (M_1, \dots, M_l) \in \mathbb{M}^l$. This means that encryptions of M_1, \dots, M_l under the same key uniquely determine the key, regardless of the nonces used.

As mentioned above, $\text{KEYREC}_{\mathcal{E}}^{M,n}$ corresponds to a very weak notion of security. In the following Lemma, we prove that this weak notion of security is implied, with only a small loss, by many of the more reasonable n -key security notions given in Definition 7. For succinctness we present the reduction in a compact way, but split the analysis in different cases (depending on the adversary goal and on the requirements to respect uniqueness or not).

Lemma 12. Let $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$, $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$ and suppose \mathcal{E} is M -key-unique. Then there exists an (\mathbb{S}, \mathbb{T}) simple reduction from $\text{KEYREC}_{\mathcal{E}}^{M,n}$ to $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ with $\mathbb{T}(t_{\mathcal{A}}) = t_{\mathcal{A}} + (l + m_{\text{GOAL}})t_{\mathcal{E}}$ and $\mathbb{S}(\epsilon_{\mathcal{A}}) = \delta_X \cdot \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}}$, where $m_{\text{IND}} = m \geq 1$, an arbitrary integer; $m_{\text{GOAL}} = 1$ if $\text{GOAL} \neq \text{IND}$; $t_{\mathcal{E}}$ is a bound on the runtime of a single encryption with \mathcal{E} ;

$$\delta_X = \begin{cases} 1 - \frac{nl(l-1) + m_{\text{GOAL}}(m_{\text{GOAL}} + 2l - 1)}{2|\mathbb{N}|}, & \text{if } X = \text{NR} \\ 1, & \text{if } X \neq \text{NR} \end{cases}$$

and

$$\delta_{\text{GOAL}} = \begin{cases} 1 - \frac{1}{2^m}, & \text{if } \text{GOAL} = \text{IND} \\ 1, & \text{if } \text{GOAL} \neq \text{IND}. \end{cases}$$

Note that δ_X and δ_{GOAL} are both close to 1: m can be set arbitrarily large and, for useful encryption schemes, the nonce space \mathbb{N} is very large.

Remark 13. We are unable to show a corresponding result for $\text{POWER} = \text{CDA}$ or $\text{POWER} = \text{PAS}$. This is because we need the ‘always real’ encryption oracle \mathcal{E} to simulate the environment of \mathcal{A} in the key recovery game. As a consequence, looking forward, our lower bounds for tightness of simple reductions hold only for n -key games with such an oracle. Nevertheless, we feel it is natural to give the n -key adversary access to the \mathcal{E} oracle so that, for example, the adversary can use queries to this oracle to determine which keys to corrupt and which to challenge.

Proof. We construct $\mathcal{R}_{\mathcal{A}}$ that runs the key recovery adversary \mathcal{A} to obtain the key used by the challenge oracle(s) and then uses it to guess the hidden bit b_1 . Therefore $\mathcal{R}_{\mathcal{A}}$ will return $(1, b'_1)$ and wins if $b'_1 = b_1$.

For each $i \in [n]$ and $j \in [l]$, $\mathcal{R}_{\mathcal{A}}$ samples $N_{i,j} \leftarrow_{\mathbb{S}} \mathbb{N}$ and then queries the encryption oracle \mathcal{E} on input $(i, M_j, N_{i,j})$, receiving $C_{i,j}$ (unless $\mathcal{R}_{\mathcal{A}}$ has made this query before, since this is a pointless query, in which case it just sets $C_{i,j}$ to be the response from the last time the query was made). Then $\mathcal{R}_{\mathcal{A}}$ passes $(C_{i,j}, N_{i,j})_{i \in [n], j \in [l]}$ to the key recovery adversary \mathcal{A} .

When \mathcal{A} returns an index i^* , $\mathcal{R}_{\mathcal{A}}$ queries Cor on each $i \in [n \setminus i^*]$ and passes $(K_i)_{i \in [n \setminus i^*]}$ to \mathcal{A} .

When \mathcal{A} returns a key K^* , $\mathcal{R}_{\mathcal{A}}$ checks if $\mathcal{E}_{K^*}^{N_{i^*,j}}(M_j) = C_{i^*,j}$ for each $j \in [l]$. If not, then \mathcal{A} has been unsuccessful, so $\mathcal{R}_{\mathcal{A}}$ samples a random bit $b'_1 \leftarrow_{\mathbb{S}} \{0, 1\}$ and returns $(1, b'_1)$. If the tests all succeed, then by M -key-uniqueness, $K^* = K_{i^*}$. Then $\mathcal{R}_{\mathcal{A}}$ does the following:

- If $\text{GOAL} = \text{IND}$, for $i = 1, 2, \dots, m$ (for some ‘‘large’’ m), $\mathcal{R}_{\mathcal{A}}$ chooses random $M_i^* \leftarrow_{\mathbb{S}} \mathbb{M}$ and $N_i^* \leftarrow_{\mathbb{S}} \mathbb{N}$ such that $M_i^* \neq M_j$ for all $j \in [l]$. Then $\mathcal{R}_{\mathcal{A}}$ queries Enc on input $(i^*, 1, M_i^*, N_i^*)$, receiving C_i^* . If for all $i = 1, 2, \dots, m$ it holds that $\mathcal{E}_{K^*}^{N_i^*}(M_i^*) = C_i^*$ then $\mathcal{R}_{\mathcal{A}}$ returns $(1, 0)$. Else, $\mathcal{R}_{\mathcal{A}}$ returns $(1, 1)$.
- If $\text{GOAL} = \text{LRIND}$, $\mathcal{R}_{\mathcal{A}}$ chooses random $M_0^*, M_1^* \leftarrow_{\mathbb{S}} \mathbb{M}$ and $N^* \leftarrow_{\mathbb{S}} \mathbb{N}$ such that $|M_0^*| = |M_1^*|$, $M_0^* \neq M_j$ and $M_1^* \neq M_j$ for all $j \in [l]$. Then $\mathcal{R}_{\mathcal{A}}$ queries LR on input $(i^*, 1, M_0^*, M_1^*, N^*)$, receiving C^* . If $\mathcal{E}_{K^*}^{N^*}(M_0^*) = C^*$, $\mathcal{R}_{\mathcal{A}}$ returns $(1, 0)$. Else, $\mathcal{R}_{\mathcal{A}}$ returns $(1, 1)$.

- If $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{CTI}\}$, $\mathcal{R}_{\mathcal{A}}$ chooses random $M^* \leftarrow_{\$} \mathbb{M}$ and $N^* \leftarrow_{\$} \mathbb{N}$ such that $M^* \neq M_j$ for all $j \in [l]$. Then $\mathcal{R}_{\mathcal{A}}$ computes $C^* \leftarrow \mathcal{E}_{K^*}^{N^*}(M^*)$ and queries Dec on input $(i^*, 1, C^*, N^*)$, receiving M . If $M \neq \perp$, $\mathcal{R}_{\mathcal{A}}$ returns $(1, 0)$. Else, $\mathcal{R}_{\mathcal{A}}$ returns $(1, 1)$.

For $\text{GOAL} \in \{\text{LRIND}, \text{AE}, \text{LRAE}, \text{CTI}\}$, the adversary $\mathcal{R}_{\mathcal{A}}$ returns $(1, b)$ with $b = b_1$ whenever the adversary \mathcal{A} against key recovery is successful.

For $\text{GOAL} \in \{\text{IND}\}$, the adversary $\mathcal{R}_{\mathcal{A}}$ always returns the correct bit if $b_1 = 1$. It also returns the correct bit $b_1 = 0$, provided that the random ciphertexts $(C_i^*)_{i \in [m]}$ that oracle Enc returns do not all collide with the true ciphertexts $\mathcal{E}_{K^*}^{N_i^*}(M_i^*)$. This collision event occurs with probability at most $\frac{1}{2^m}$.

In other words, for $\text{GOAL} \in \{\text{IND}, \text{LRIND}, \text{AE}, \text{LRAE}, \text{CTI}\}$, \mathcal{R} succeeds whenever \mathcal{A} succeeds if $b_1 = 0$, while, if $b_1 = 1$, then \mathcal{R} succeeds with the same probability that \mathcal{A} succeeds multiplied by δ_{GOAL} , where $\delta_{\text{GOAL}} = 1$ for $\text{GOAL} \in \{\text{LRIND}, \text{AE}, \text{LRAE}, \text{CTI}\}$ and $\delta_{\text{GOAL}} = (1 - \frac{1}{2^m})$ for $\text{GOAL} = \text{IND}$.

Whenever \mathcal{A} does not recover K^* , $\mathcal{R}_{\mathcal{A}}$ guesses correctly with probability $\frac{1}{2}$. Putting it all together we get the following:

$$\begin{aligned} & \Pr \left[\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{R}_{\mathcal{A}}) = 1 \right] \\ &= \Pr \left[\text{KEYREC}_{\mathcal{E}}^{M,n}(\mathcal{A}) = 1 \right] \cdot \left(1 - \frac{1 - \delta_{\text{GOAL}}}{2} \right) + \frac{1}{2} \cdot \left(1 - \Pr \left[\text{KEYREC}_{\mathcal{E}}^{M,n}(\mathcal{A}) = 1 \right] \right) \\ &= \Pr \left[\text{KEYREC}_{\mathcal{E}}^{M,n}(\mathcal{A}) = 1 \right] \cdot \left(\frac{1}{2} - \frac{1 - \delta_{\text{GOAL}}}{2} \right) + \frac{1}{2}, \end{aligned}$$

from which we obtain

$$\begin{aligned} \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{R}_{\mathcal{A}}) &= 2 \left(\Pr \left[\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{R}_{\mathcal{A}}) = 1 \right] - \frac{1}{2} \right) \\ &= \delta_{\text{GOAL}} \cdot \Pr \left[\text{KEYREC}_{\mathcal{E}}^{M,n}(\mathcal{A}) = 1 \right] \\ &= \delta_{\text{GOAL}} \cdot \text{Adv}_{\mathcal{E}}^{\text{KEYREC},M,n}(\mathcal{A}) \\ &= \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}}. \end{aligned}$$

Ignoring the time taken for random sampling, the runtime of $\mathcal{R}_{\mathcal{A}}$ is precisely the runtime of \mathcal{A} , plus the time taken for additional encryptions using K^* : if $\text{GOAL} = \text{IND}$, there are $l + m$ additional encryptions and, if $\text{GOAL} \neq \text{IND}$, there are $l + 1$ additional encryptions. It follows that

$$t_{\mathcal{R}_{\mathcal{A}}} = t_{\mathcal{A}} + (l + m_{\text{GOAL}})t_{\mathcal{E}},$$

where $m_{\text{IND}} = m$ and $m_{\text{GOAL}} = 1$ for $\text{GOAL} \neq \text{IND}$.

Moreover, $\mathcal{R}_{\mathcal{A}}$ doesn't compromise b_1 and makes no pointless or prohibited queries: no queries are repeated, the messages used to generate the challenge queries do not appear in any of the previous encryption queries under key K_{i^*} and, in the LRIND case, the challenge messages are of equal length. It follows that $\mathcal{R}_{\mathcal{A}}$ is a valid adversary against $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ for $X \in \{\text{IV}, \text{MR}\}$, since nonces are always chosen uniformly at random.

If $X = \text{NR}$, $\mathcal{R}_{\mathcal{A}}$ might not be a valid adversary, since the randomly chosen nonces might accidentally collide. So we modify $\mathcal{R}_{\mathcal{A}}$ to abort and output a random bit whenever there is a collision among the l randomly chosen nonces $(N_{i,j})_{j \in [l]}$ for each $i \in [n \setminus i^*]$, or among the $l + m_{\text{GOAL}}$ randomly chosen nonces for encryptions under K_{i^*} : the $l + m$ nonces $(N_{i^*,j})_{j \in [l]}$ and $(N_i^*)_{i \in [m]}$, if $\text{GOAL} = \text{IND}$, and the $l + 1$ nonces $(N_{i^*,j})_{j \in [l]}$ and N^* , if $\text{GOAL} \neq \text{IND}$. Then $\mathcal{R}_{\mathcal{A}}$ is a valid adversary and its advantage is $\epsilon_{\mathcal{A}}$ multiplied by the probability that no such nonce collisions happen. By a simple union bound the probability of a collision among the l randomly chosen nonces $(N_{i,j})_{j \in [l]}$ is at most $\frac{l(l-1)}{2|\mathbb{N}|}$ for each $i \in [n \setminus i^*]$ and the probability of a collision among the $l + m_{\text{GOAL}}$ randomly chosen nonces for i^* is at

most $\frac{(l+m_{\text{GOAL}})(l+m_{\text{GOAL}}-1)}{2^{|\mathbf{N}|}}$. Thus the probability of a collision among the nonces for any of the n keys is at most

$$(n-1) \frac{l(l-1)}{2^{|\mathbf{N}|}} + \frac{l+m_{\text{GOAL}}(l+m_{\text{GOAL}}-1)}{2^{|\mathbf{N}|}} = \frac{nl(l-1) + m_{\text{GOAL}}(m_{\text{GOAL}} + 2l - 1)}{2^{|\mathbf{N}|}} = 1 - \delta_{\text{NR}}.$$

Thus the advantage of $\mathcal{R}_{\mathcal{A}}$ is $\epsilon_{\mathcal{R}_{\mathcal{A}}} \geq \delta_{\text{NR}} \cdot \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}}$, as desired. \square

Remark 14. In the proof, we assumed that the adversary is allowed to associate the bit b_1 with any of the n keys K_1, \dots, K_n . While this is permitted according to our definition of the $\text{GOAL-POWER}_{\mathcal{E}}^{n,X}$ game, in fact the result holds for more restrictive games: we only require that for all $i \in [n]$ there exists *some* $j \in [n]$ such that the adversary can associate the bit b_j with the key K_i . In this case, $\mathcal{R}_{\mathcal{A}}$ uses the recovered key K^* from \mathcal{A} to determine the value of any hidden bit b_j that can be associated with K_{i^*} .

4 Multi-Key to Single-Key Reductions are Lossy

In this section we present our main results: any simple black-box reduction from multi-key security (in its many definitional variants) to single-key security loses a linear factor in the number of keys. Two remarks are in order. First, we show the lower bound for reductions from the security of an arbitrary construction of an (authenticated) encryption scheme $C[\mathcal{E}]$ to that of \mathcal{E} (and in particular for the case where $C[\mathcal{E}] = \mathcal{E}$). This more general setting encompasses interesting cases, e.g. where $C[\mathcal{E}]$ is *double encryption* with \mathcal{E} , i.e.

$$C[\mathcal{E}]_{(K_1, K_2)}^{(N_1, N_2)}(M) = \mathcal{E}_{K_2}^{N_2} \left(\mathcal{E}_{K_1}^{N_1}(M) \right),$$

which has been shown to have desirable multi-key properties [27]. Furthermore, showing the separation for $C[\mathcal{E}]$ and \mathcal{E} also suggests how to circumvent the lower bound for the loss that we provide. Our lower bound requires that $C[\mathcal{E}]$ satisfies key-uniqueness. It may therefore be possible to start from a secure single-key security that satisfies key-uniqueness, and show a tight reduction from multi-key security of a *variant* $C[\mathcal{E}]$ of \mathcal{E} , provided that $C[\mathcal{E}]$ somehow avoids key uniqueness.

We consider separately reductions between different security flavours (authenticated encryption, privacy, integrity). For each case in turn, we proceed in two steps. First, we establish that if \mathcal{E} is a (single-key) secure encryption scheme and $C[\mathcal{E}]$ is a key-unique encryption scheme, then all simple reductions from the multi-key key recovery game for $C[\mathcal{E}]$ to the single-key security game for \mathcal{E} are lossy. Since by Lemma 12 there is a tight reduction from multi-key key recovery to multi-key security, it is an immediate corollary that there is no tight reduction from the multi-key security of $C[\mathcal{E}]$ to the single-key security of \mathcal{E} .

An interesting remark is that the bound on the inherent loss of simple reductions depends on the security of the scheme \mathcal{E} : the more secure the scheme, the tighter the bound. While our bound is therefore not meaningful for insecure schemes, this case is of little interest in practice.

Authenticated Encryption. We give the formal results for the case of authenticated encryption below.

Theorem 15. *Let \mathcal{E} and $C[\mathcal{E}]$ be AE schemes such that $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathbb{M}^l$. Then, for $X \in \{\text{IV}, \text{NR}, \text{MR}\}$, any simple reduction \mathcal{R} from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ loses at least $(\frac{1}{n} + 2\epsilon)^{-1}$, where ϵ is the maximum advantage for a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ running in time at most $nt_{\mathcal{R}} + 2l(n-1)t_{C[\mathcal{E}]}$ (where $t_{C[\mathcal{E}]}$ is an upper-bound on the runtime of a single encryption with $C[\mathcal{E}]$).*

We sketch the proof before giving its details below. The crucial idea, following [3], is to construct a *meta-reduction* \mathcal{M} that *rewinds* the reduction \mathcal{R} in order to simulate its interaction with an ideal adversary \mathcal{A} against $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$. If the simulation works correctly, then the output of \mathcal{R} can be used

by \mathcal{M} to win the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game with probability $\epsilon_{\mathcal{R}}$. Then the (single-key) security of \mathcal{E} yields an upper-bound on the success probability of \mathcal{M} , i.e. an upper-bound on $\epsilon_{\mathcal{R}}$.

We view \mathcal{R} as a collection of three algorithms, $\mathcal{R} = (\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3)$. The first, \mathcal{R}_1 , makes oracle queries in the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game, then produces the ciphertexts and nonces that \mathcal{A} expects to receive in the $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ game. The second, \mathcal{R}_2 , receives an index i^* from \mathcal{A} and the state st_1 of the previous algorithm, \mathcal{R}_1 . Then \mathcal{R}_2 makes oracle queries and eventually produces the vector of keys that \mathcal{A} expects to receive in the $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ game. Finally, \mathcal{R}_3 receives a guessed key K^* from \mathcal{A} and the state st_2 of \mathcal{R}_2 . Then \mathcal{R}_3 makes oracle queries and outputs a guessed bit b' .

\mathcal{M} only rewinds \mathcal{R}_2 : \mathcal{M} executes \mathcal{R}_2 on each of the n possible indices i^* that could be returned by \mathcal{A} and each \mathcal{R}_2 then returns a set of keys. Then \mathcal{M} uses the keys returned by one execution of \mathcal{R}_2 to construct the input to a *different* execution of \mathcal{R}_3 , i.e. st_2 given to \mathcal{R}_3 will not be from the same execution of \mathcal{R}_2 used to construct the ‘guessed’ key K^* .

The main obstacle in arguing that the above strategy works is that \mathcal{M} needs to break $\text{AE-PAS}_{\mathcal{E}}^{X,1}$, which is an *interactive* assumption. This is in contrast to the meta-reductions from [3], which are designed to violate a non-interactive complexity assumption. In our case, \mathcal{M} needs to simulate an appropriate environment towards multiple copies of \mathcal{R} , each of which may make oracle queries, yet \mathcal{M} has access to a single set of oracles for the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game. It is not obvious that \mathcal{M} can simply forward queries from all copies of \mathcal{R} to these oracles, since queries across different invocations of \mathcal{R} may interfere with one-another and render \mathcal{M} invalid. The key observation is that we can leverage the single-key security of \mathcal{E} : instead of forwarding queries, \mathcal{M} simply simulates the Enc and Dec oracles by sampling random ciphertexts and returning \perp , respectively. We argue, based on the security of \mathcal{E} , that \mathcal{R} cannot distinguish this simulation from the real oracles in its game.

Proof. For ease of notation, let K , M , N and C be the sets of keys, messages, nonces and ciphertexts, respectively, for the *construction* $C[\mathcal{E}]$ (even though they may differ from the corresponding sets for \mathcal{E} , but we shall not need to refer to those in the proof).

Consider the following (inefficient) adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ in the game $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$. On input

$$(C_{i,j}, N_{i,j})_{i \in [n], j \in [l]},$$

\mathcal{A}_1 first checks that each $C_{i,j} \in \mathsf{C}$ and each $N_{i,j} \in \mathsf{N}$. If this check fails, then \mathcal{A}_1 aborts (by outputting a random index $i^* \in [n]$ and recording an abort message in the state $st_{\mathcal{A}}$ for \mathcal{A}_2 , triggering the latter to output \perp). If the check succeeds, then \mathcal{A}_1 chooses $i^* \in [n]$ uniformly at random, sets

$$st_{\mathcal{A}} \leftarrow \left(i^*, (C_{i,j}, N_{i,j})_{i \in [n], j \in [l]} \right)$$

and outputs $(i^*, st_{\mathcal{A}})$. On input $\left((K_i)_{i \in [n \setminus i^*]}, st_{\mathcal{A}} \right)$, \mathcal{A}_2 checks that K_i is *valid* for each $i \in [n \setminus i^*]$, that is:

1. $K_i \in \mathsf{K}$
2. For each $j \in [l]$, $C[\mathcal{E}]_{K_i}^{N_{i^*,j}}(M_j) = C_{i,j}$.

If this check fails, then \mathcal{A}_2 outputs \perp . If the check succeeds, then \mathcal{A}_2 uses exhaustive search to find some $K^* \in \mathsf{K}$ such that $C[\mathcal{E}]_{K^*}^{N_{i^*,j}}(M_j) = C_{i^*,j}$ for each $j \in [l]$. Since $C[\mathcal{E}]$ is M -key-unique, either K^* exists and is unique, or the ciphertexts $C_{i^*,j}$ were not all encryptions of the messages M_j with the nonces $N_{i^*,j}$ under the same key. So if \mathcal{A}_2 does not find a K^* with this property, it outputs \perp . Otherwise it outputs K^* .

It is clear that the advantage of \mathcal{A} is $\epsilon_{\mathcal{A}} = 1$ since, in the real $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ game, all the checks performed by \mathcal{A} will succeed and K^* is uniquely defined.

We construct a meta-reduction \mathcal{M} that simulates the environment of \mathcal{R} in its interaction with this ideal adversary \mathcal{A} . Then \mathcal{M} will use the output of \mathcal{R} to play the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game. In what follows, we describe \mathcal{M} in detail. A diagram showing the overall structure of the interaction between \mathcal{M} and \mathcal{R} is given in Fig. 6.

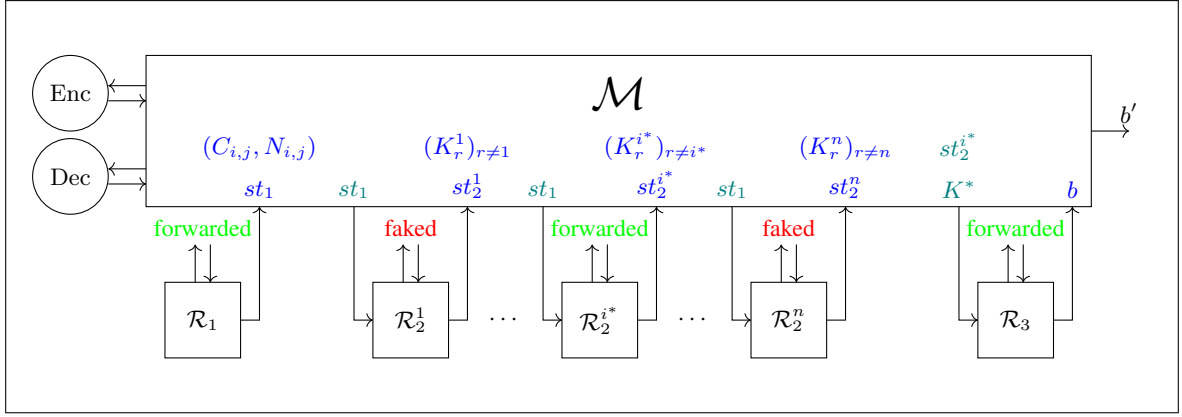


Fig. 6. An overview of the meta-reduction \mathcal{M} , which rewinds a reduction $\mathcal{R} = (\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3)$. The inputs to each component of the reduction \mathcal{R} are shown in teal, while the outputs are shown in blue. Some oracle queries from \mathcal{R} are forwarded to the oracles in the game played by \mathcal{M} and the responses sent back, while other oracle queries from \mathcal{R} are faked by \mathcal{M} .

First, K^* is initialised to \perp . Then, \mathcal{M} uses its oracles to simulate the oracles used by \mathcal{R}_1 by simply forwarding the queries from \mathcal{R}_1 and the responses from the oracles, until \mathcal{R}_1 returns

$$\left((C_{i,j}, N_{i,j})_{i \in [n], j \in [l]}, st_1 \right).$$

Then \mathcal{M} checks that each $C_{i,j} \in \mathcal{C}$ and each $N_{i,j} \in \mathcal{N}$. If this check fails, \mathcal{M} ‘aborts’ just as \mathcal{A} would. That is, \mathcal{M} runs \mathcal{R}_2 on input (i, st_1) for a random index $i^* \in [n]$, forwarding oracle queries and responses, receives $\left((K_i)_{i \in [n \setminus i^*]}, st_2 \right)$ from \mathcal{R}_2 , runs \mathcal{R}_3 on input (\perp, st_2) , receives a bit b' and outputs this in its game. If, on the other hand, the check succeeds, then \mathcal{M} chooses i^* uniformly at random from $[n]$ and does the following for each $i \in [n]$:

1. \mathcal{M} runs \mathcal{R}_2 on input (i, st_1) , which we call \mathcal{R}_2^i for ease of readability.
2. When \mathcal{R}_2^i makes oracle queries:
 - (a) If $i = i^*$, \mathcal{M} uses its oracles to honestly answer all oracle queries; forwarding the queries to its oracles and then forwarding the replies to $\mathcal{R}_2^{i^*}$.
 - (b) If $i \neq i^*$, \mathcal{M} simulates the ‘fake’ oracles, i.e. the oracles Enc and Dec in the case $b = 1$. Concretely, when \mathcal{R}_2^i makes an encryption query (M, N) , \mathcal{M} samples $C \leftarrow_{\mathcal{S}} \{0, 1\}^{\text{enclen}(|M|)}$ and returns this to \mathcal{R}_2^i .³ When \mathcal{R}_2^i makes a decryption query (C, N) , \mathcal{M} returns \perp to \mathcal{R}_2^i .
3. When \mathcal{R}_2^i outputs $\left((K_r^i)_{r \in [n \setminus i]}, st_2^i \right)$, if $i \neq i^*$ then \mathcal{M} checks if $K_{i^*}^i$ is valid, i.e.
 - (a) $K_{i^*}^i \in \mathcal{K}$,
 - (b) For each $j \in [l]$, $C[\mathcal{E}]_{K_{i^*}^i}^{N_{i^*,j}}(M_j) = C_{i^*,j}$.

If $K_{i^*}^i$ is valid, then $K^* \leftarrow K_{i^*}^i$. By \mathcal{M} -key-uniqueness, $K_{i^*}^i$ is the only key with this property.

At the end of these runs of \mathcal{R}_2 , if $\mathcal{R}_2^{i^*}$ did not provide a full set of valid keys, i.e. $K_r^{i^*}$ is not valid for some $r \in [n \setminus i^*]$, then \mathcal{M} sets $K^* \leftarrow \perp$ (mirroring the check performed by \mathcal{A}_2).

If $\mathcal{R}_2^{i^*}$ did provide a full set of valid keys, but $K^* = \perp$, (so none of the $\mathcal{R}_2^i, i \neq i^*$ provided a valid key $K_{i^*}^i$), \mathcal{M} aborts the simulation and returns a random bit. We call this event BAD.

Otherwise, \mathcal{M} runs \mathcal{R}_3 on input $(K^*, st_2^{i^*})$, forwarding oracle queries from \mathcal{R}_3 to its oracles and sending back the responses.

When \mathcal{R}_3 outputs a bit b' , \mathcal{M} returns this bit in its game.

Now we consider the resources of \mathcal{M} and its advantage in the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game.

\mathcal{M} performs n runs of (part of) \mathcal{R} and carries out $2(n-1)l$ encryptions with $C[\mathcal{E}]$ (checking validity of $K_{i^*}^i$ for each $i \neq i^*$ and checking validity of $K_r^{i^*}$ for each $r \neq i^*$), so if we ignore the time taken

³ Of course, here enclen refers to the lengths of ciphertexts from \mathcal{E} , not $C[\mathcal{E}]$.

for random sampling and checking set membership, the runtime of \mathcal{M} is at most $nt_{\mathcal{R}} + 2l(n-1)t_{C[\mathcal{E}]}$. Moreover, \mathcal{M} makes at most $q_{\mathcal{R}}$ oracle queries, since it only forwards the queries from $\mathcal{R}_1, \mathcal{R}_2^{i^*}$ and \mathcal{R}_3 .

Now consider the advantage $\epsilon_{\mathcal{M}}$ of \mathcal{M} in $\text{AE-PAS}_{\mathcal{E}}^{X,1}$. From the definition of a simple reduction, \mathcal{R} must be a valid adversary in $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ whenever \mathcal{A} is a valid adversary in $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$. But all adversaries are automatically valid in $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$, so \mathcal{R} must always be a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$. Now the oracle queries \mathcal{M} makes are exactly the same queries as $(\mathcal{R}_1, \mathcal{R}_2^{i^*}, \mathcal{R}_3)$ makes in the same game. Since \mathcal{R} is a valid adversary, this shows that \mathcal{M} does not make pointless or prohibited queries and uses nonces correctly with respect to X . Therefore \mathcal{M} is a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ and so $\epsilon_{\mathcal{M}} \leq \epsilon$.

Note that for $\mathcal{R}_1, \mathcal{R}_2^{i^*}$ and \mathcal{R}_3 , \mathcal{M} answers the oracle queries honestly with its own oracles. Therefore \mathcal{M} correctly simulates the view of $(\mathcal{R}_1, \mathcal{R}_2^{i^*}, \mathcal{R}_3)$ in the game $\text{AE-PAS}_{\mathcal{E}}^{X,1}$. However, \mathcal{M} might not correctly simulate the responses from \mathcal{A} . Indeed, to correctly simulate \mathcal{A} , \mathcal{M} requires that some $\mathcal{R}_2^i, i \neq i^*$ provides a valid key $K_{i^*}^i$, but the oracle queries from $\mathcal{R}_2^i, i \neq i^*$ are not handled honestly. The imperfect simulation of the view of \mathcal{R}_2^i might make it less likely to provide a valid key $K_{i^*}^i$. We will therefore need to show that the change in behaviour of the \mathcal{R}_2^i due to the imperfect simulation is small. The intuition for this claim is that if \mathcal{R}_2^i could distinguish between the honest and the simulated oracles (having only received an index i from the key-recovery adversary \mathcal{A} , not a key), then one can use $(\mathcal{R}_1, \mathcal{R}_2^i)$ directly, without \mathcal{A} , to win the single-key game $\text{AE-PAS}_{\mathcal{E}}^{X,1}$.

Consider the three possible scenarios:

1. $\mathcal{R}_2^{i^*}$ did not provide a full set of valid keys.
2. $\mathcal{R}_2^{i^*}$ did provide a full set of valid keys and, for some $i \neq i^*$, \mathcal{R}_2^i provided a valid key $K_{i^*}^i$.
3. $\mathcal{R}_2^{i^*}$ did provide a full set of valid keys, but, for each $i \neq i^*$, \mathcal{R}_2^i did not provide a valid key $K_{i^*}^i$.

In the first case, both \mathcal{M} and \mathcal{A} submit \perp to \mathcal{R}_3 as their ‘key’, so the simulation is correct. In the second case, both \mathcal{M} and \mathcal{A} submit a key K^* to \mathcal{R}_3 that satisfies $C[\mathcal{E}]_{K^*}^{N_{i^*,j}}(M_j) = C_{i^*,j}$ for all $j \in [l]$, and K^* is the only key with this property by the M -key-uniqueness of $C[\mathcal{E}]$. So the simulation is correct in this case too.

The third case is the event BAD and is where the simulation fails. By construction \mathcal{M} aborts the simulation if BAD occurs and outputs a random bit. Given that BAD does not occur, the view of $(\mathcal{R}_1, \mathcal{R}_2^{i^*}, \mathcal{R}_3)$ in its interaction with \mathcal{A} and the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ oracles is identical to its view in its interaction with \mathcal{M} and \mathcal{M} returns the bit b' returned by \mathcal{R}_3 . This shows that

$$\Pr \left[\text{AE-PAS}_{\mathcal{E}}^{X,1}(\mathcal{R}) = 1 \right] = \Pr \left[\text{AE-PAS}_{\mathcal{E}}^{X,1}(\mathcal{M}) = 1 \mid \neg \text{BAD} \right].$$

Write $W^X(\mathcal{M})$ (‘Win’) for the event $\text{AE-PAS}_{\mathcal{E}}^{X,1}(\mathcal{M}) = 1$. Then, as \mathcal{M} outputs a random bit if BAD occurs, we have $\Pr [W^X(\mathcal{M}) \mid \text{BAD}] = \frac{1}{2}$ and it follows that:

$$\begin{aligned} & \Pr [W^X(\mathcal{M})] \\ &= \Pr [W^X(\mathcal{M}) \cap \neg \text{BAD}] + \Pr [W^X(\mathcal{M}) \cap \text{BAD}] \\ &= \Pr [W^X(\mathcal{M}) \mid \neg \text{BAD}] (1 - \Pr [\text{BAD}]) + \Pr [W^X(\mathcal{M}) \mid \text{BAD}] \Pr [\text{BAD}] \\ &= \Pr [W^X(\mathcal{M}) \mid \neg \text{BAD}] - \Pr [\text{BAD}] \left(\Pr [W^X(\mathcal{M}) \mid \neg \text{BAD}] - \frac{1}{2} \right). \end{aligned}$$

Then,

$$\begin{aligned}
\text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{M}) &= 2 \left(\Pr [W^X(\mathcal{M})] - \frac{1}{2} \right) \\
&= 2 \left[\Pr [W^X(\mathcal{M}) \mid \neg\text{BAD}] - \Pr [\text{BAD}] \left(\Pr [W^X(\mathcal{M}) \mid \neg\text{BAD}] - \frac{1}{2} \right) - \frac{1}{2} \right] \\
&= 2 \left(\Pr [W^X(\mathcal{M}) \mid \neg\text{BAD}] - \frac{1}{2} \right) - \Pr [\text{BAD}] \cdot 2 \left(\Pr [W^X(\mathcal{M}) \mid \neg\text{BAD}] - \frac{1}{2} \right) \\
&= (1 - \Pr [\text{BAD}]) \text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{R}).
\end{aligned}$$

It follows that:

$$\text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{M}) \geq \text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{R}) - \Pr [\text{BAD}].$$

To complete the proof we bound the probability of BAD (see the next lemma) by $\Pr [\text{BAD}] \leq \frac{1}{n} + \epsilon$. We therefore get that

$$\epsilon \geq \epsilon_{\mathcal{M}} \geq \epsilon_{\mathcal{R}} - \Pr [\text{BAD}] \geq \epsilon_{\mathcal{R}} - \frac{1}{n} - \epsilon.$$

So, $\epsilon_{\mathcal{R}} \leq (\frac{1}{n} + 2\epsilon)$. Since $\epsilon_{\mathcal{A}} = 1$, we get that

$$\frac{\epsilon_{\mathcal{A}}}{\epsilon_{\mathcal{R}}} \geq \left(\frac{1}{n} + 2\epsilon \right)^{-1},$$

as required to show that \mathcal{R} loses $(\frac{1}{n} + 2\epsilon)^{-1}$. \square

Lemma 16.

$$\Pr [\text{BAD}] \leq \frac{1}{n} + \epsilon.$$

Proof. Consider a meta-reduction \mathcal{M}' in the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game that executes \mathcal{R}_1 and each $\mathcal{R}_2^i, i \in [n]$ exactly as \mathcal{M} does, but without treating $\mathcal{R}_2^{i^*}$ differently. That is, encryption and decryption queries from $\mathcal{R}_2^{i^*}$ are ‘faked’ in the same way as for the other $\mathcal{R}_2^i, i \neq i^*$. Such an \mathcal{M}' could have chosen $i^* \leftarrow_{\$} [n]$ after executing each \mathcal{R}_2^i , simply by storing all the keys output by each \mathcal{R}_2^i , and then, once i^* had been chosen, checking if $\mathcal{R}_2^{i^*}$ returned a full set of valid keys and if each $K_{i^*}^i$ was valid for $i \neq i^*$.

Note that the probability of BAD occurring for \mathcal{M}' does not depend on whether i^* was chosen at the start of executing the \mathcal{R}_2^i , or at the end, since \mathcal{M}' runs each \mathcal{R}_2^i in the same way. Moreover, after executing each \mathcal{R}_2^i , there can be at most one $j \in [n]$ such that \mathcal{R}_2^j returned a full set of valid keys but for each $i \neq j$, \mathcal{R}_2^i did not provide a full set of valid keys. Therefore there can be at most one $j \in [n]$ such that \mathcal{R}_2^j returned a full set of valid keys but for each $i \neq j$, \mathcal{R}_2^i did not provide a valid key K_j^i . Since i^* was sampled uniformly from $[n]$, the probability that i^* has the latter property, i.e. that BAD occurs for \mathcal{M}' , is at most $\frac{1}{n}$.

Now we compare the probability that BAD occurs for the two meta-reductions \mathcal{M} and \mathcal{M}' . Let $\text{BAD}_{\mathcal{M}} = \text{BAD}$ and let $\text{BAD}_{\mathcal{M}'}$ be the event that BAD occurs in the game played by \mathcal{M}' .

Consider the hidden bit b in the game played by \mathcal{M} and \mathcal{M}' . If $b = 1$, then the views of \mathcal{R}_1 and each \mathcal{R}_2^i are identically distributed in their interactions with \mathcal{M} and \mathcal{M}' (since $\mathcal{R}_2^{i^*}$ receives ‘fake’ responses to its queries, regardless of whether the meta-reduction forwards them to its own oracles or simulates the responses.) It follows that $\Pr [\text{BAD}_{\mathcal{M}'} \mid b = 1] = \Pr [\text{BAD}_{\mathcal{M}} \mid b = 1]$.

Then

$$\begin{aligned}
\Pr [\text{BAD}] &= \Pr [\text{BAD}_{\mathcal{M}}] - \Pr [\text{BAD}_{\mathcal{M}'}] + \Pr [\text{BAD}_{\mathcal{M}'}] \\
&\leq \Pr [\text{BAD}_{\mathcal{M}}] - \Pr [\text{BAD}_{\mathcal{M}'}] + \frac{1}{n} \\
&= \frac{1}{2} (\Pr [\text{BAD}_{\mathcal{M}} \mid b = 0] - \Pr [\text{BAD}_{\mathcal{M}'} \mid b = 0]) + \frac{1}{n}.
\end{aligned}$$

Now we construct an adversary \mathcal{B} that simulates the environment of \mathcal{R}_1 and the \mathcal{R}_2^i in their interaction with either \mathcal{M} or \mathcal{M}' , depending on the hidden bit b' in the game played by \mathcal{B} . If BAD occurs, \mathcal{B} will output 0. Otherwise \mathcal{B} will output 1.

Consider \mathcal{B} in the $\text{AE-CCA}_{\mathcal{E}}^{X,1}$ game. That is, \mathcal{B} has access to the usual challenge oracles Enc and Dec, but can also query the ‘always real’ oracles \mathcal{E} and \mathcal{D} (provided it does not make pointless or prohibited queries). But if \mathcal{B} has significant advantage in this game, then there is another adversary, with the same resources as \mathcal{B} , that has significant advantage against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$:

Lemma 17. *Suppose \mathcal{A} is a valid adversary against $\text{AE-CCA}_{\mathcal{E}}^{X,1}$, where $X \in \{\text{IV}, \text{NR}, \text{MR}\}$. Then*

$$\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{A}) \leq 2\epsilon,$$

where ϵ is the maximum advantage of a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ that runs in the same time as \mathcal{A} and makes the same number of oracle queries as \mathcal{A} .

The proof of Lemma 17 is in Appendix C. We remark that a similar statement can be easily derived by combining results from an existing work [5]. However, this approach only shows that the advantage in $\text{AE-CCA}_{\mathcal{E}}^{X,1}$ is at most *four times* the maximum advantage in $\text{AE-PAS}_{\mathcal{E}}^{X,1}$, whereas proving the statement directly gives a tighter bound.

Now we describe the adversary \mathcal{B} in the $\text{AE-CCA}_{\mathcal{E}}^{X,1}$ game. First, \mathcal{B} runs \mathcal{R}_1 , but all queries are forwarded to the *genuine* oracles \mathcal{E} and \mathcal{D} . Then \mathcal{B} carries out the same checks as \mathcal{M} (or \mathcal{M}') and, if the checks succeed, \mathcal{B} samples $i^* \leftarrow_{\$}[n]$ and, for each $i \in [n]$, \mathcal{B} runs \mathcal{R}_2 on input (i, st_1) .

When \mathcal{R}_2^i makes oracle queries:

1. If $i = i^*$, \mathcal{B} uses its *challenge* oracles Enc and Dec to honestly answer all oracle queries; forwarding the queries to its oracles and then forwarding the replies to $\mathcal{R}_2^{i^*}$.
2. If $i \neq i^*$, \mathcal{B} simulates the ‘fake’ oracles, i.e. the oracles Enc and Dec with $b = 1$, just as \mathcal{M} (or \mathcal{M}') does.

Finally, \mathcal{B} checks if BAD has occurred. If so \mathcal{B} outputs 0. Otherwise, \mathcal{B} outputs 1.

Let b' be the hidden bit in the game played by \mathcal{B} . So the oracle queries from \mathcal{R}_1 will always be ‘real’ (as they are for \mathcal{M} and \mathcal{M}' , given that $b = 0$), the oracle queries from \mathcal{R}_2^i for $i \neq i^*$ will always be ‘fake’ (as they are for \mathcal{M} and \mathcal{M}') and, depending on b' , the oracle queries from $\mathcal{R}_2^{i^*}$ will be real (like \mathcal{M} , given that $b = 0$), or fake (like \mathcal{M}'). It follows that $\Pr[0 \leftarrow \mathcal{B} \mid b' = 0] = \Pr[\text{BAD}_{\mathcal{M}} \mid b = 0]$ and $\Pr[0 \leftarrow \mathcal{B} \mid b' = 1] = \Pr[\text{BAD}_{\mathcal{M}'} \mid b = 0]$. Now,

$$\begin{aligned} \Pr[\text{AE-CCA}_{\mathcal{E}}^{X,1}(\mathcal{B}) = 1] &= \frac{1}{2} (\Pr[0 \leftarrow \mathcal{B} \mid b' = 0] + \Pr[1 \leftarrow \mathcal{B} \mid b' = 1]) \\ &= \frac{1}{2} (\Pr[0 \leftarrow \mathcal{B} \mid b' = 0] - \Pr[0 \leftarrow \mathcal{B} \mid b' = 1]) + \frac{1}{2} \end{aligned}$$

and so

$$\begin{aligned} \text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{B}) &= 2 \left(\Pr[\text{AE-CCA}_{\mathcal{E}}^{X,1}(\mathcal{B}) = 1] - \frac{1}{2} \right) \\ &= \Pr[0 \leftarrow \mathcal{B} \mid b' = 0] - \Pr[0 \leftarrow \mathcal{B} \mid b' = 1] \\ &= \Pr[\text{BAD}_{\mathcal{M}} \mid b = 0] - \Pr[\text{BAD}_{\mathcal{M}'} \mid b = 0] \\ &\geq 2 \left(\Pr[\text{BAD}] - \frac{1}{n} \right). \end{aligned}$$

Like \mathcal{M} (or \mathcal{M}'), \mathcal{B} performs n runs of (part of) \mathcal{R} and carries out $2(n-1)l$ encryptions to check if BAD has occurred. So the runtime of \mathcal{B} is at most $nt_{\mathcal{R}} + 2l(n-1)t_{C[\mathcal{E}]}$. Moreover \mathcal{B} makes at most $q_{\mathcal{R}}$ oracle queries (only forwarding queries from \mathcal{R}_1 and $\mathcal{R}_2^{i^*}$).

Consider $\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{B})$. Firstly, note that \mathcal{B} uses nonces correctly with respect to X , since any query to Enc or \mathcal{E} is a query made to Enc by $(\mathcal{R}_1, \mathcal{R}_2^{i^*}, \mathcal{R}_3)$ and \mathcal{R} is a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$. Also, \mathcal{B} will not make pointless queries:

- A repeated query to \mathcal{E} or \mathcal{D} by \mathcal{B} would be a repeated query to Enc or Dec from \mathcal{R}_1 , which is a pointless or prohibited query in the game played by \mathcal{R} .
- A repeated query to Dec by \mathcal{B} would be a repeated query to Dec from \mathcal{R}_2^{i*} , which is a pointless query in the game played by \mathcal{R} .
- A query $\mathcal{D}(C, N)$ by \mathcal{B} , where C was the response to a query $\mathcal{E}(M, N)$, would be a query Dec(C, N) from \mathcal{R}_1 , where C was the response to a query Enc(M, N), which is a prohibited query in the game played by \mathcal{R} .
- A query $\mathcal{E}(M, N)$ by \mathcal{B} , where $M \neq \perp$ was the response to a query $\mathcal{D}(C, N)$, would be a query Enc(M, N) from \mathcal{R}_1 , where $M \neq \perp$ was the response to a query Dec(C, N), which is a pointless query in the game played by \mathcal{R} .
- Finally, suppose \mathcal{B} makes a query $\mathcal{E}(M, N)$ or Enc(M, N), where $M \neq \perp$ was the response to a query Dec(C, N). The query Dec(C, N) from \mathcal{B} would correspond to a query Dec(C, N) from \mathcal{R}_2^{i*} and so the subsequent encryption query would correspond to a query Enc(M, N) from \mathcal{R}_2^{i*} . But as $M \neq \perp$ this is a pointless query for \mathcal{R} .

Moreover, \mathcal{B} will not make prohibited queries:

- A repeated query to Enc by \mathcal{B} would be a repeated query to Enc from \mathcal{R}_2^{i*} , which is a prohibited query in the game played by \mathcal{R} .
- Suppose \mathcal{B} makes two queries of the form Enc(M, N) and $\mathcal{E}(M, N)$. Each of these queries would correspond to the same query Enc(M, N) from \mathcal{R} , which is prohibited in the game played by \mathcal{R} .
- A query $\mathcal{D}(C, N)$ from \mathcal{B} , where C was the response to a query Enc(M, N), is impossible since \mathcal{B} only queries Enc and Dec after querying \mathcal{E} and \mathcal{D} .
- A query Dec(C, N) from \mathcal{B} , where C was the response to a query $\mathcal{E}(M, N)$ or Enc(M, N), would correspond to a query Dec(C, N) from \mathcal{R}_2^{i*} , where C was the response to a query Enc(M, N) from \mathcal{R}_1 or \mathcal{R}_2^{i*} , which is a prohibited query in the game played by \mathcal{R} .
- A query Enc(M, N) from \mathcal{B} , where $M \neq \perp$ was the response to a query $\mathcal{D}(C, N)$, would correspond to a query Enc(M, N) from \mathcal{R}_2^{i*} , where $M \neq \perp$ was the response to a query Dec(C, N) from \mathcal{R}_1 , which is a pointless query in the game played by \mathcal{R} .

It follows that \mathcal{B} is a valid adversary against $\text{Adv}_{\mathcal{E}}^{\text{AE-CCA}, X, 1}(\mathcal{B})$. Then, by Lemma 17, we have

$$2 \left(\Pr[\text{BAD}] - \frac{1}{n} \right) \leq \text{Adv}_{\mathcal{E}}^{\text{AE-CCA}, X, 1}(\mathcal{B}) \leq 2\epsilon,$$

from which the result follows. \square

Theorem 15 establishes that reductions from $\text{KEYREC}_{C[\mathcal{E}]}^{M, n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X, 1}$ are lossy. By Lemma 12, there exists a tight reduction from $\text{KEYREC}_{C[\mathcal{E}]}^{M, n}$ to $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X', n}$ (for $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$); it immediately follows that reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X', n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X, 1}$ must be lossy (for $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$). We formalise this intuition in the following corollary:

Corollary 18. *Let \mathcal{E} and $C[\mathcal{E}]$ be AE schemes such that $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathbb{M}^l$. Then for $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$, $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$, $X, X' \in \{\text{IV}, \text{NR}, \text{MR}\}$ and $n > 1$, all simple reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X', n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X, 1}$ must lose*

$$L = \delta_{\text{GOAL}} \cdot \delta_{X'} \cdot \left(\frac{1}{n} + 2\epsilon \right)^{-1},$$

where δ_{GOAL} and $\delta_{X'}$ are as in Lemma 12 and ϵ is as given in Theorem 15.

We emphasise that the ‘nonce use’ parameters $X', X \in \{\text{IV}, \text{NR}, \text{MR}\}$ can differ between the n -key game and the single key game. While it is natural to consider $X' = X$ we prefer to state the result in full generality and show that a very large class of reductions are necessarily lossy. Note that

multi-key games for $\text{POWER} \in \{\text{PAS}, \text{CDA}\}$ are not known to be (tightly) equivalent to those where $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$ (see Appendix B). It therefore remains an open problem to obtain tightness lowerbounds for $\text{POWER} \in \{\text{PAS}, \text{CDA}\}$.

Proof. Recall from Lemma 12 the (\mathbb{S}, \mathbb{T}) -simple reduction from $\text{KEYREC}_{\mathcal{E}}^{M,n}$ to $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$, where $\mathbb{S}(\epsilon_{\mathcal{A}}) = \delta_X \cdot \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}}$ and $\mathbb{T}(t_{\mathcal{A}}) = t_{\mathcal{A}} + (l + m_{\text{GOAL}})t_{\mathcal{E}}$. Relabelling, we obtain a $(\mathbb{S}', \mathbb{T}')$ -simple reduction from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$, where $\mathbb{S}'(\epsilon_{\mathcal{A}}) = \delta_{X'} \cdot \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}}$ and $\mathbb{T}'(t_{\mathcal{A}}) = t_{\mathcal{A}} + (l + m_{\text{GOAL}})t_{C[\mathcal{E}]}$, which we call \mathcal{R} .

We argue by contradiction. Suppose there is a simple reduction \mathcal{R}' from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ such that, for all valid adversaries \mathcal{B} against $\text{AE-PAS}_{\mathcal{E}}^{X,n}$, $\epsilon_{\mathcal{R}'} > L^{-1}\epsilon_{\mathcal{B}}$.

Then we can form a simple reduction \mathcal{R}'' from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X,1}$: for any adversary \mathcal{A} against $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$, running \mathcal{R} with \mathcal{A} provides a valid adversary \mathcal{B} against $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ for \mathcal{R}' to turn into a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$.

By construction, the advantage $\epsilon_{\mathcal{R}''}$ of \mathcal{R}'' is equal to the advantage of \mathcal{R}' with access to an adversary with advantage $\epsilon_{\mathcal{R}}$, i.e. $\epsilon_{\mathcal{R}''} > L^{-1}\epsilon_{\mathcal{R}}$. Since $\epsilon_{\mathcal{R}} \geq \delta_{X'} \cdot \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}}$ for all adversaries \mathcal{A} against $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$, we have

$$\epsilon_{\mathcal{R}''} > L^{-1} \cdot \delta_{X'} \cdot \delta_{\text{GOAL}} \cdot \epsilon_{\mathcal{A}} = \left(\frac{1}{n} + 2\epsilon\right) \epsilon_{\mathcal{A}}.$$

But this is a contradiction since, by Theorem 15, for any simple reduction \mathcal{R}'' from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X,1}$, there exists an adversary \mathcal{A} against $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ such that

$$\epsilon_{\mathcal{R}''} \leq \left(\frac{1}{n} + 2\epsilon\right) \epsilon_{\mathcal{A}}.$$

Thus for any simple reduction \mathcal{R}' from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ there exists a valid adversary \mathcal{B} against $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ such that $\epsilon_{\mathcal{R}'} \leq L^{-1}\epsilon_{\mathcal{B}}$, i.e. \mathcal{R}' loses L . □

Privacy and Integrity. The above results hold for notions of authenticated encryption schemes. It is natural to ask whether the loss for simple reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ is an artefact of considering the two orthogonal single-key security properties of secrecy and authenticity at the same time. Perhaps it is possible to circumvent the loss when looking at these properties separately, e.g. there could there be non-lossy simple reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ and from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$. We show that this is not the case.

We proceed as for the authenticated encryption case. For privacy and integrity, in turn, we show that reductions from multi-key key recovery to single-key security are inherently lossy; the lower bound then follows by Lemma 12. We give the details in Appendix D.

Other Single-Key Security Notions. Given the results above concerning reductions from multi-key AE security notions to the single-key notions AE-PAS , IND-PAS and CTI-CPA , one can obtain analogous results for *equivalent or weaker* single key notions, such as where ciphertexts being indistinguishable from random strings (IND , AE) is replaced by (weaker) left-or-right indistinguishability (LRIND , LRAE). The idea is that if there were a tight reduction from an n -key game to single-key LRAE-PAS , say, then this reduction could be combined with the tight reduction from LRAE-PAS to AE-PAS to obtain a tight reduction from the n -key game to AE-PAS that contradicts Corollary 18. However, “tight” is defined here with respect to a number of parameters including, crucially, ϵ : the maximum advantage *in the AE-PAS game*. If ϵ is close to 1, then so is the “loss”. In other words, the tightness lowerbounds that

one can prove using our existing results for strictly weaker single-key security notions are only meaningful for schemes that are secure according to the stronger notions. This leaves open the possibility that tight multi-key to single-key reductions exist for schemes that achieve the weaker single-key security notions, but *not* the stronger ones. Moreover, our meta-reduction techniques cannot be directly applied to left-or-right indistinguishability, since the meta-reduction cannot correctly simulate left-or-right encryption queries during the rewinding phase without making its own (possibly prohibited) oracle queries (unlike for IND when the meta-reduction simply samples random strings of the appropriate length).

Public Key Encryption. It should be possible to adapt our existing techniques to the public key setting. Let LRIND-CPA be the standard game in which the adversary is given the public key and can query a left-or-right encryption oracle. Note that the honest encryption oracle is omitted as it is rendered superfluous by the public key. Since public key encryption is typically randomised rather than nonce-based, repeated left-or-right encryption queries are not prohibited, so a meta-reduction \mathcal{M} can use its own left-or-right challenge oracle to correctly simulate left-or-right queries from the reduction \mathcal{R} during the rewinding phase, without \mathcal{M} becoming an invalid adversary. However, if \mathcal{R} can also make decryption queries, then simulating these queries during the rewinding phase might force \mathcal{M} to be invalid (such as if one instance of \mathcal{R} attempts to decrypt the output of the left-or-right encryption oracle sent to an earlier instance of \mathcal{R}). In summary, it should be possible to show reductions from multi-key games to single-key LRIND-CPA are lossy for public key encryption schemes secure according to LRIND-CPA, but to show an analogous result for LRIND-CCA one needs to additionally assume that ciphertexts are indistinguishable from random strings (which is a rather strong assumption in the public key setting). We leave formally proving these claims for future work.

5 Conclusion

We have presented a general family of multi-key security definitions for authenticated encryption, where the adversary can adaptively corrupt keys. We have shown, for a very large class of authenticated encryption schemes, for most members of our family of definitions and for widely-accepted single-key security definitions, that any black-box reduction from the n -key security of an encryption scheme to its single-key security will incur a loss close to n .

For practitioners who set security parameters based on provable guarantees, this shows that security reductions have an inherent shortcoming. Since keys are sampled independently, the corruption of one key should not affect the security of another, yet it is impossible in many cases to prove that security does not degrade from the single-key setting to the n -key setting. It appears that the loss of n is an unfortunate, unavoidable artefact of the proof.

We have shown that the loss of reductions is inevitable for multi-key definitions where the adversary has access to an honest encryption oracle. We therefore left open the possibility that for security notions without such an oracle, tight reductions may be found. Furthermore, our impossibility results apply to schemes where ciphertexts are indistinguishable from random strings. It may be possible that tight reductions exist for schemes that achieve weaker forms of confidentiality, such as left-or-right indistinguishability. Historically, the community has tended to opt for stronger and stronger security notions, but perhaps a slightly weaker single-key notion would be preferred if it tightly implied a meaningful multi-key notion. Finally, it was pointed out by an anonymous reviewer that, in practice, the number of keys an adversary can corrupt is likely to be much smaller than the number of keys in use; it might be possible to find tighter multi-key to single-key reductions for multi-key games where the adversary can corrupt at most q_c keys (with $q_c \ll n$). We leave these interesting open questions for future work.

Acknowledgements. This work was supported by an EPSRC Industrial CASE award and DFG grant JA 2445/1-1. The authors would also like to thank the anonymous TCC reviewers for their constructive comments on our paper.

References

1. Abdalla, M., Fouque, P.A., Lyubashevsky, V., Tibouchi, M.: Tightly-secure signatures from lossy identification schemes. In: Pointcheval and Johansson [38], pp. 572–590
2. Bader, C., Hofheinz, D., Jager, T., Kiltz, E., Li, Y.: Tightly-secure authenticated key exchange. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 629–658. Springer, Heidelberg (Mar 2015)
3. Bader, C., Jager, T., Li, Y., Schäge, S.: On the impossibility of tight cryptographic reductions. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 273–304. Springer, Heidelberg (May 2016)
4. Baecker, P., Brzuska, C., Fischlin, M.: Notions of black-box reductions, revisited. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 296–315. Springer, Heidelberg (Dec 2013)
5. Barwell, G., Page, D., Stam, M.: Rogue decryption failures: Reconciling AE robustness notions. In: Groth, J. (ed.) 15th IMA International Conference on Cryptography and Coding. LNCS, vol. 9496, pp. 94–111. Springer, Heidelberg (Dec 2015)
6. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: Security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (May 2000)
7. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among notions of security for public-key encryption schemes. In: Krawczyk, H. (ed.) CRYPTO’98. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (Aug 1998)
8. Bellare, M., Hofheinz, D., Kiltz, E.: Subtleties in the definition of IND-CCA: When and how should challenge decryption be disallowed? *Journal of Cryptology* 28(1), 29–48 (Jan 2015)
9. Bellare, M., Ristenpart, T., Tessaro, S.: Multi-instance security and its application to password-based cryptography. In: Safavi-Naini and Canetti [40], pp. 312–329
10. Bellare, M., Tackmann, B.: The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 247–276. Springer, Heidelberg (Aug 2016)
11. Biham, E.: How to decrypt or even substitute DES-encrypted messages in 2^{28} steps. *Inf. Process. Lett.* 84(3), 117–124 (2002), [http://dx.doi.org/10.1016/S0020-0190\(02\)00269-7](http://dx.doi.org/10.1016/S0020-0190(02)00269-7)
12. Biryukov, A., Mukhopadhyay, S., Sarkar, P.: Improved time-memory trade-offs with multiple data. In: Preneel, B., Tavares, S. (eds.) SAC 2005. LNCS, vol. 3897, pp. 110–127. Springer, Heidelberg (Aug 2006)
13. Blazy, O., Kakvi, S.A., Kiltz, E., Pan, J.: Tightly-secure signatures from chameleon hash functions. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 256–279. Springer, Heidelberg (Mar / Apr 2015)
14. Blazy, O., Kiltz, E., Pan, J.: (Hierarchical) identity-based encryption from affine message authentication. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 408–425. Springer, Heidelberg (Aug 2014)
15. Boneh, D., Venkatesan, R.: Breaking RSA may not be equivalent to factoring. In: Nyberg, K. (ed.) EUROCRYPT’98. LNCS, vol. 1403, pp. 59–71. Springer, Heidelberg (May / Jun 1998)
16. Chatterjee, S., Kobitz, N., Menezes, A., Sarkar, P.: Another look at tightness II: Practical issues in cryptography. *Cryptology ePrint Archive, Report 2016/360* (2016), <http://eprint.iacr.org/2016/360>
17. Chatterjee, S., Menezes, A., Sarkar, P.: Another look at tightness. In: Miri, A., Vaudenay, S. (eds.) SAC 2011. LNCS, vol. 7118, pp. 293–319. Springer, Heidelberg (Aug 2012)
18. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg (Aug 2013)
19. Coron, J.S.: Security proof for partial-domain hash signature schemes. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 613–626. Springer, Heidelberg (Aug 2002)
20. Fischlin, M., Fleischhacker, N.: Limitations of the meta-reduction technique: The case of Schnorr signatures. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 444–460. Springer, Heidelberg (May 2013)
21. Fischlin, M., Lehmann, A., Ristenpart, T., Shrimpton, T., Stam, M., Tessaro, S.: Random oracles with(out) programmability. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 303–320. Springer, Heidelberg (Dec 2010)
22. Fleischhacker, N., Jager, T., Schröder, D.: On tight security proofs for Schnorr signatures. In: Sarkar and Iwata [41], pp. 512–531
23. Fouque, P.A., Joux, A., Mavromati, C.: Multi-user collisions: Applications to discrete logarithm, Even-Mansour and PRINCE. In: Sarkar and Iwata [41], pp. 420–438
24. Garg, S., Bhaskar, R., Lokam, S.V.: Improved bounds on security reductions for discrete log based signatures. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 93–107. Springer, Heidelberg (Aug 2008)
25. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part I. LNCS, vol. 9665, pp. 1–27. Springer, Heidelberg (May 2016)
26. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption AEZ and the problem that it solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 15–44. Springer, Heidelberg (Apr 2015)
27. Hoang, V.T., Tessaro, S.: The multi-user security of double encryption. In: Coron, J.S., Nielsen, J.B. (eds.) *Advances in Cryptology – EUROCRYPT 2017: 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 – May 4, 2017, Proceedings, Part II*. pp. 381–411. Springer International Publishing, Cham (2017)
28. Hofheinz, D.: Algebraic partitioning: Fully compact and (almost) tightly secure cryptography. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016-A, Part I. LNCS, vol. 9562, pp. 251–281. Springer, Heidelberg (Jan 2016)
29. Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. In: Safavi-Naini and Canetti [40], pp. 590–607

30. Hofheinz, D., Jager, T., Knapp, E.: Waters signatures with optimal security reduction. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 66–83. Springer, Heidelberg (May 2012)
31. Hsiao, C.Y., Reyzin, L.: Finding collisions on a public road, or do secure hash functions need secret coins? In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 92–105. Springer, Heidelberg (Aug 2004)
32. Kakvi, S.A., Kiltz, E.: Optimal security proofs for full domain hash, revisited. In: Pointcheval and Johansson [38], pp. 537–553
33. Katz, J., Wang, N.: Efficiency improvements for signature schemes with tight security reductions. In: Jajodia, S., Atluri, V., Jaeger, T. (eds.) ACM CCS 03. pp. 155–164. ACM Press (Oct 2003)
34. Mouha, N., Luykx, A.: Multi-key security: The Even-Mansour construction revisited. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 209–223. Springer, Heidelberg (Aug 2015)
35. Namprempe, C., Rogaway, P., Shrimpton, T.: Reconsidering generic composition. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 257–274. Springer, Heidelberg (May 2014)
36. Paillier, P., Vergnaud, D.: Discrete-log-based signatures may not be equivalent to discrete log. In: Roy, B.K. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 1–20. Springer, Heidelberg (Dec 2005)
37. Panjwani, S.: Tackling adaptive corruptions in multicast encryption protocols. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 21–40. Springer, Heidelberg (Feb 2007)
38. Pointcheval, D., Johansson, T. (eds.): EUROCRYPT 2012, LNCS, vol. 7237. Springer, Heidelberg (Apr 2012)
39. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (May / Jun 2006)
40. Safavi-Naini, R., Canetti, R. (eds.): CRYPTO 2012, LNCS, vol. 7417. Springer, Heidelberg (Aug 2012)
41. Sarkar, P., Iwata, T. (eds.): ASIACRYPT 2014, Part I, LNCS, vol. 8873. Springer, Heidelberg (Dec 2014)
42. Schäge, S.: Tight proofs for signature schemes without random oracles. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 189–206. Springer, Heidelberg (May 2011)
43. Seurin, Y.: On the exact security of Schnorr-type signatures in the random oracle model. In: Pointcheval and Johansson [38], pp. 554–571
44. Tessaro, S.: Optimally secure block ciphers from ideal primitives. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part II. LNCS, vol. 9453, pp. 437–462. Springer, Heidelberg (Nov / Dec 2015)

A Valid Adversarial Behaviour for AE Games

Pointless and Prohibited Queries. Since encryption is deterministic, the response to certain oracle queries can be predicted in advance. Therefore the adversary learns nothing from these queries; we call them *pointless*. Without loss of generality we assume that valid adversaries do not make such queries. The following queries are pointless:

- Repeat a query to any oracle other than Enc(the Enc oracle sometimes samples random ciphertexts, but all other oracles are deterministic).
- Make a query $\mathcal{D}(i, C, N)$, where C was the response to a query $\mathcal{E}(i, M, N)$ (since the response will be M , by correctness).
- Make a query $\mathcal{E}(i, M, N)$, where $M \neq \perp$ was the response to a query $\mathcal{D}(i, C, N)$ (since the response will be C , by tidiness).
- Make a query $\mathcal{E}(i, M, N)$ or $\text{Enc}(i, j, M, N)$, where a query $\text{Dec}(i, j, C, N)$ was made with response $M \neq \perp$ (since the response $M \neq \perp$ reveals $b_j = 0$ and $\mathcal{E}_{K_i}^N(M) = C$ by tidiness).

Some other queries lead to hidden bits being trivial to recover (without having to corrupt a key); we call these queries *prohibited*, since valid adversaries are not permitted to make them. The following queries are prohibited:

- Repeat a query $\text{Enc}(i, j, M, N)$ (if the response to both queries is the same, then with very high probability $b_j = 0$ and otherwise $b_j = 1$).
- Make a query of the form $\text{LR}(i, j, M_0, M_1, N)$ with $|M_0| \neq |M_1|$ (since the length of the ciphertext reveals the length of the plaintext, trivially revealing which of M_0 or M_1 was encrypted).
- Make two queries of the form $\text{LR}(i, j, M_0, M_1, N)$, $\text{LR}(i, j, M'_0, M'_1, N)$ such that $M_b = M'_b$ and $M_{1-b} \neq M'_{1-b}$ for some $b \in \{0, 1\}$ (if the response to both queries is the same, then $b_j = b$ by correctness, and otherwise $b_j = 1 - b$).
- Make two queries of the form $\text{Enc}(i, j, M, N)$ and $\mathcal{E}(i, M, N)$, in any order (which trivially reveals b_j).

- Make two queries of the form $\text{LR}(i, j, M_0, M_1, N)$ and $\mathcal{E}(i, M_b, N)$, in any order, for some $b \in \{0, 1\}$ (which trivially reveals b_j).
- Make a query $\mathcal{D}(i, C, N)$, where a query $\text{Enc}(i, j, M, N)$ or $\text{LR}(i, j, M_0, M_1, N)$ was made with response C (which trivially reveals b_j , by correctness).
- Make a query $\text{Dec}(i, j, C, N)$, where a query $\mathcal{E}(i, M, N)$, $\text{Enc}(i, j, M, N)$ or $\text{LR}(i, j, M_0, M_1, N)$ was previously made with response C (which trivially reveals b_j , by correctness).
- Make a query $\text{Enc}(i, j, M, N)$, $\text{LR}(i, j, M, M_1, N)$ or $\text{LR}(i, j, M_0, M, N)$, where $M \neq \perp$ was the response to a query $\mathcal{D}(i, C, N)$ (which trivially reveals b_j , by tidiness).

It is not necessary to prohibit queries being forwarded between the Enc and LR oracles, since we do not consider games where both these challenge oracles are present.

Correct Nonce Use. The parameter $X \in \{\text{IV}, \text{NR}, \text{MR}\}$ determines how the adversary may use nonces in encryption queries. We say \mathcal{A} uses nonces correctly with respect to X if the following statements hold:

- If $X = \text{IV}$, then for each query of the form $\mathcal{E}(-, -, N)$, $\text{Enc}(-, -, -, N)$, or $\text{LR}(-, -, -, -, N)$, N is sampled uniformly at random from \mathcal{N} .
- If $X = \text{NR}$, then each nonce appears in at most one encryption query *under the same key*. That is, for each $i \in [n]$, each nonce N appears in at most one query of the form $\text{Enc}(i, -, -, N)$, $\text{LR}(i, -, -, -, N)$ or $\mathcal{E}(i, -, N)$.
- If $X = \text{MR}$, then nonces may be chosen arbitrarily and repeated in different queries (modulo the pointless and prohibited queries specified above).

B Relations between Security Notions

Reducing Multi-Key Security to Single-Key Security. To show that our lower tightness bounds are optimal, in the sense that they match the upper bound provided by simple generic reductions from single-key security to n -key security with a loss linear in n , we consider the two natural formulations of multi-key security discussed at the beginning of Sect. 3.

The Multi-Bit Setting. We start with considering the variant of the multi-key game of Definition 7 where the adversary is restricted to challenge queries with $i = j$. This corresponds to the natural ‘multi-bit’ game, where there is one hidden bit b_i for each key K_i and the adversary wins the game if it can guess b_i without corrupting K_i .

Theorem 19. *Let $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$, $\text{POWER} \in \{\text{CCA}, \text{CPA}, \text{CDA}, \text{PAS}\}$, $n \geq 1$, $X \in \{\text{IV}, \text{NR}, \text{MR}\}$ and let \mathcal{E} be an AE scheme. Suppose \mathcal{A} is a valid adversary in $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ such that, for all queries to $\text{Enc}(i, j, M, N)$, $\text{LR}(i, j, M_0, M_1, N)$, and $\text{Dec}(i, j, C, N)$, it holds that $i = j$. Then we can construct a valid adversary \mathcal{B} against $\text{GOAL-POWER}_{\mathcal{E}}^{X,1}$, running in approximately the same time as \mathcal{A} , such that*

$$\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) = \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}).$$

Proof. The adversary \mathcal{B} simulates the environment of \mathcal{A} in the $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ game. First, \mathcal{B} chooses a random index $i^* \leftarrow_{\$} [n]$, and then, for each $i \in [n \setminus i^*]$, \mathcal{B} generates $n - 1$ keys $K_i \leftarrow_{\$} \mathcal{KG}$ and chooses $n - 1$ random bits $b_i \leftarrow_{\$} \{0, 1\}$. Implicitly, K_{i^*} is the secret key in the game played by \mathcal{B} and b_{i^*} is the hidden bit in the game played by \mathcal{B} . Formally, the simulation of the environment of \mathcal{A} works as follows:

- If \mathcal{A} makes the query $\text{Cor}(i^*)$, then \mathcal{B} outputs a random bit and terminates.
- If \mathcal{A} makes an oracle query with $i = i^*$, either to one of the challenge oracles Enc, Dec and LR, or to one of the honest oracles \mathcal{E} and \mathcal{D} , then \mathcal{B} forwards this query to its own oracle in the $\text{GOAL-POWER}_{\mathcal{E}}^{X,1}$ experiment and returns the response to \mathcal{A} .

- If \mathcal{A} makes an oracle query with $i \neq i^*$, then \mathcal{B} simulates the response to the query using its own K_i and b_i , exactly as in the $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ game.
- When \mathcal{A} terminates and outputs a tuple (j, b') , if $j \neq i^*$, then \mathcal{B} outputs a random bit. Otherwise \mathcal{B} outputs b' .

Since all the queries made by \mathcal{B} are also made by \mathcal{A} and since \mathcal{A} is a valid adversary, it follows that \mathcal{B} uses nonces correctly with respect to X and makes no pointless or prohibited queries. Thus \mathcal{B} is a valid adversary. Moreover, the running time of \mathcal{B} essentially consists of the running time of \mathcal{A} , plus a minor overhead to simulate the $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ experiment once. The precise runtime depends on the queries made by \mathcal{A} . It remains to show that

$$\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) = \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}).$$

Let GUESS be the event that \mathcal{A} returns a tuple (j, b') in the real experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$ such that $j = i^*$. Since i^* is chosen by \mathcal{B} uniformly at random from $[n]$ and does not appear in the real experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$, we must have

$$\Pr[\text{GUESS}] = \frac{1}{n}.$$

Also, let SMART be the event that \mathcal{A} does not corrupt the key K_j corresponding to the bit b_j it guesses in the real experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$. This event is so named because corrupting K_j would not be ‘smart’ at all: for \mathcal{A} to be valid and also exhibit this behaviour, \mathcal{A} must make no challenge queries whatsoever for K_j and b_j and therefore its guess b' is independent of b_j , yielding 0 advantage. We remark that GUESS and SMART are independent, since i^* is chosen uniformly from $[n]$ regardless of the behaviour of \mathcal{A} .

Given that SMART and GUESS both occur, the simulation by \mathcal{B} of the environment of \mathcal{A} is perfect:

- \mathcal{B} is able to answer all of \mathcal{A} ’s queries to Cor, since it knows the keys K_i for $i \neq i^*$ and \mathcal{A} does not make a query $\text{Cor}(i^*)$.
- When \mathcal{A} makes a query of the form $\mathcal{E}(i, M, N)$ or $\mathcal{D}(i, C, N)$ with $i \neq i^*$, \mathcal{B} is able to respond correctly since it knows K_i .
- When \mathcal{A} makes a query of the form $\mathcal{E}(i^*, M, N)$ or $\mathcal{D}(i^*, C, N)$, \mathcal{B} is able to respond correctly using its own \mathcal{E} and \mathcal{D} oracles.
- When \mathcal{A} makes a query of the form $\text{Enc}(i, j, M, N)$, $\text{LR}(i, j, M_0, M_1, N)$, or $\text{Dec}(i, j, C, N)$, we have $i = j$ and:
 - either $i \neq i^*$, in which case \mathcal{B} responds correctly using its own K_i and b_i ,
 - or $i = i^*$, in which case \mathcal{B} responds correctly using its own challenge oracles, with K_{i^*} and b_{i^*} being the secret key and the hidden bit, respectively, used by the challenge oracles in the game played by \mathcal{B} .

Let $W_{\mathcal{A}}$ be the event $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A}) = 1$ and let $W_{\mathcal{B}}$ be the event $\text{GOAL-POWER}_{\mathcal{E}}^{X,1}(\mathcal{B}) = 1$. By the preceding observation, we have

$$\Pr[W_{\mathcal{B}} \mid \text{SMART} \cap \text{GUESS}] = \Pr[W_{\mathcal{A}} \mid \text{SMART} \cap \text{GUESS}] = \Pr[W_{\mathcal{A}} \mid \text{SMART}],$$

where the second equality holds since the choice of i^* made by \mathcal{B} is independent of the output of the real experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$. Furthermore,

$$\Pr[W_{\mathcal{A}} \mid \neg \text{SMART}] = \Pr[W_{\mathcal{B}} \mid \neg \text{SMART}] = \frac{1}{2}.$$

This is because, given that $\neg \text{SMART}$ occurs, \mathcal{A} makes no challenge queries corresponding to the bit it tries to guess and \mathcal{B} either returns a random bit or, if \mathcal{A} tries to guess b_{i^*} , then \mathcal{B} returns the bit returned

by \mathcal{A} , but will have made no challenge queries. It follows that

$$\begin{aligned}\Pr[W_{\mathcal{A}}] &= \Pr[W_{\mathcal{A}} \mid \text{SMART}] \cdot \Pr[\text{SMART}] + \frac{1}{2} \cdot (1 - \Pr[\text{SMART}]) \\ &= \Pr[\text{SMART}] \cdot \left(\Pr[W_{\mathcal{A}} \mid \text{SMART}] - \frac{1}{2} \right) + \frac{1}{2}\end{aligned}$$

and hence

$$\begin{aligned}\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}) &= 2 \cdot \Pr[W_{\mathcal{A}}] - 1 \\ &= \Pr[\text{SMART}] \cdot (2 \cdot \Pr[W_{\mathcal{A}} \mid \text{SMART}] - 1).\end{aligned}$$

Similarly,

$$\Pr[W_{\mathcal{B}}] = \Pr[\text{SMART}] \cdot \left(\Pr[W_{\mathcal{B}} \mid \text{SMART}] - \frac{1}{2} \right) + \frac{1}{2}$$

and

$$\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) = \Pr[\text{SMART}] \cdot (2 \cdot \Pr[W_{\mathcal{B}} \mid \text{SMART}] - 1).$$

Now, given that SMART occurs, if GUESS does not occur, then either \mathcal{A} tries to corrupt K_{i^*} and \mathcal{B} aborts the simulation – since \mathcal{A} is ‘smart’, this means \mathcal{A} would have guessed b_j for some $j \neq i^*$ in the real experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}(\mathcal{A})$ – or \mathcal{B} completes the simulation but \mathcal{A} returns (j, b') for some $j \neq i^*$. In either case, \mathcal{B} returns a random bit, so $\Pr[W_{\mathcal{B}} \mid \neg\text{GUESS} \cap \text{SMART}] = \frac{1}{2}$. It follows that

$$\begin{aligned}\Pr[W_{\mathcal{B}} \mid \text{SMART}] &= \Pr[W_{\mathcal{B}} \cap \text{GUESS} \mid \text{SMART}] + \Pr[W_{\mathcal{B}} \cap \neg\text{GUESS} \mid \text{SMART}] \\ &= \Pr[W_{\mathcal{B}} \mid \text{GUESS} \cap \text{SMART}] \Pr[\text{GUESS}] \\ &\quad + \Pr[W_{\mathcal{B}} \mid \neg\text{GUESS} \cap \text{SMART}] (1 - \Pr[\text{GUESS}]) \\ &= \Pr[W_{\mathcal{A}} \mid \text{SMART}] \cdot \frac{1}{n} + \frac{1}{2} \cdot \left(1 - \frac{1}{n} \right) \\ &= \frac{1}{n} \cdot \left(\Pr[W_{\mathcal{A}} \mid \text{SMART}] - \frac{1}{2} \right) + \frac{1}{2},\end{aligned}$$

so that

$$\begin{aligned}\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) &= \frac{1}{n} \cdot \Pr[\text{SMART}] \cdot (2 \cdot \Pr[W_{\mathcal{A}} \mid \text{SMART}] - 1) \\ &= \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}),\end{aligned}$$

as required. \square

The Single-Bit Setting. Since Theorem 19 covers only one of the two natural formulations of multi-key security discussed at the beginning of Sect. 3, it is natural to ask whether a similar result can be proven for the other natural variant: the single-bit setting, where $j = 1$ for all queries made by the adversary. Indeed, this is the case.

Theorem 20. *Let $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$, $\text{POWER} \in \{\text{CCA}, \text{CPA}, \text{CDA}, \text{PAS}\}$, $n \geq 1$, $X \in \{\text{IV}, \text{NR}, \text{MR}\}$ and let \mathcal{E} be an AE scheme. Suppose \mathcal{A} is a valid adversary in $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ such that, for all queries to $\text{Enc}(i, j, M, N)$, $\text{LR}(i, j, M_0, M_1, N)$, and $\text{Dec}(i, j, C, N)$, it holds that $j = 1$ and, at the end of the game, \mathcal{A} returns the guess $(1, b')$.⁴ Then we can construct a valid adversary \mathcal{B} against $\text{GOAL-POWER}_{\mathcal{E}}^{X,1}$, running in approximately the same time as \mathcal{A} , such that*

$$\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) = \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}).$$

⁴ We exclude the uninteresting case in which \mathcal{A} guesses a hidden bit other than b_1 , since such an \mathcal{A} must have 0 advantage. It is straightforward to adapt the existing argument to take care of this case, similarly to the use of the event SMART in the proof of Theorem 19.

Proof. This is essentially a standard hybrid argument. The adversary \mathcal{B} computes as follows. First, \mathcal{B} chooses random $i^* \leftarrow_{\$} [n]$ and generates $n - 1$ keys $K_i \leftarrow_{\$} \mathcal{KG}$ for $i \neq i^*$. Implicitly, K_{i^*} and b_1 are the secret key and the hidden bit, respectively, in the game played by \mathcal{B} .

For all $1 \leq i < i^*$, if \mathcal{A} makes a challenge query for K_i , then \mathcal{B} answers as if $b = 0$. For all $i^* < i \leq n$, if \mathcal{A} makes a challenge query for K_i , then \mathcal{B} answers as if $b = 1$. For honest encryption and decryption queries under K_i ($i \neq i^*$), \mathcal{B} simulates the oracles honestly. If \mathcal{A} makes challenge or honest queries under K_{i^*} , then \mathcal{B} passes these to its own oracles. If \mathcal{A} queries $\text{Cor}(i)$ for $i \neq i^*$, then \mathcal{B} returns K_i to \mathcal{A} . If \mathcal{A} queries $\text{Cor}(i^*)$, then \mathcal{B} aborts and outputs 0. Finally, when \mathcal{A} returns $(1, b')$, \mathcal{B} returns b' . To simplify notation, we write $(b' \leftarrow \mathcal{A})$ for the event that \mathcal{A} returns $(1, b')$.

We remark that \mathcal{B} is valid: every oracle query from \mathcal{B} corresponds to a query for the key K_{i^*} made by \mathcal{A} and \mathcal{A} is valid, so \mathcal{B} makes no prohibited or pointless queries and uses nonces correctly with respect to X . Moreover, the runtime of \mathcal{B} is essentially the same as \mathcal{A} , plus the small overhead for simulating the experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$ once. It remains to show that

$$\text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) = \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}).$$

Let $b_{\mathcal{A}}$ (respectively, $b_{\mathcal{B}}$) be the hidden bit in the game played by \mathcal{A} (resp., \mathcal{B}). For any $i \in [n] \cup \{0\}$, let i be the event that all challenge queries from \mathcal{A} for the keys K_1, \dots, K_i are answered as if $b = 0$ and all challenge queries for the keys K_{i+1}, \dots, K_n are answered as if $b = 1$. We remark that the event n is where *all* challenge queries are answered as if $b = 0$ and the event 0 is where *all* challenge queries are answered as if $b = 1$. Finally, let BAD_i be the event that \mathcal{A} issues the query $\text{Cor}(i)$ in the real experiment $\text{GOAL-POWER}_{\mathcal{E}}^{X,n}$. Then,

$$\begin{aligned} & \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) \\ &= \Pr[1 \leftarrow \mathcal{B} \mid b_{\mathcal{B}} = 1] - \Pr[1 \leftarrow \mathcal{B} \mid b_{\mathcal{B}} = 0] \\ &= \sum_{i=1}^n \Pr[1 \leftarrow \mathcal{B} \cap i^* = i \mid b_{\mathcal{B}} = 1] - \Pr[1 \leftarrow \mathcal{B} \cap i^* = i \mid b_{\mathcal{B}} = 0] \\ &= \sum_{i=1}^n \Pr[i^* = i] (\Pr[1 \leftarrow \mathcal{B} \mid i^* = i \cap b_{\mathcal{B}} = 1] - \Pr[1 \leftarrow \mathcal{B} \mid i^* = i \cap b_{\mathcal{B}} = 0]) \\ &= \frac{1}{n} \cdot \left(\sum_{i=1}^n \Pr[1 \leftarrow \mathcal{B} \mid i^* = i \cap b_{\mathcal{B}} = 1] - \Pr[1 \leftarrow \mathcal{B} \mid i^* = i \cap b_{\mathcal{B}} = 0] \right). \end{aligned}$$

Now, given that \mathcal{B} chooses $i^* = i$, the event $(1 \leftarrow \mathcal{B})$ occurs exactly when the event $(1 \leftarrow \mathcal{A} \cap \neg \text{BAD}_i)$ occurs. Thus, for $b \in \{0, 1\}$,

$$\begin{aligned} \Pr[1 \leftarrow \mathcal{B} \mid i^* = i \cap b_{\mathcal{B}} = b] &= \Pr[1 \leftarrow \mathcal{A} \cap \neg \text{BAD}_i \mid i^* = i \cap b_{\mathcal{B}} = b] \\ &= \Pr[1 \leftarrow \mathcal{A} \mid i^* = i \cap b_{\mathcal{B}} = b] - \Pr[1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i^* = i \cap b_{\mathcal{B}} = b]. \end{aligned}$$

But $\Pr[1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i^* = i \cap b_{\mathcal{B}} = 1] = \Pr[1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i^* = i \cap b_{\mathcal{B}} = 0]$. This is because

$$\begin{aligned} (1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i^* = i \cap b_{\mathcal{B}} = 0) &= (1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i), \\ (1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i^* = i \cap b_{\mathcal{B}} = 1) &= (1 \leftarrow \mathcal{A} \cap \text{BAD}_i \mid i-1) \end{aligned}$$

and \mathcal{A} is a valid adversary, so if \mathcal{A} corrupts K_i , then \mathcal{A} issues no challenge queries for K_i and hence $(1 \leftarrow \mathcal{A})$ has the same probability regardless of whether i_{-1} or i occurs. It follows that

$$\begin{aligned}
& \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,1}(\mathcal{B}) \\
&= \frac{1}{n} \cdot \left(\sum_{i=1}^n \Pr[1 \leftarrow \mathcal{A} \mid i^* = i \cap b_{\mathcal{B}} = 1] - \Pr[1 \leftarrow \mathcal{A} \mid i^* = i \cap b_{\mathcal{B}} = 0] \right) \\
&= \frac{1}{n} \cdot \left(\sum_{i=1}^n \Pr[1 \leftarrow \mathcal{A} \mid i_{-1}] - \Pr[1 \leftarrow \mathcal{A} \mid i] \right) \\
&= \frac{1}{n} \cdot (\Pr[1 \leftarrow \mathcal{A} \mid 0] - \Pr[1 \leftarrow \mathcal{A} \mid n]) \\
&= \frac{1}{n} \cdot (\Pr[1 \leftarrow \mathcal{A} \mid b_{\mathcal{A}} = 1] - \Pr[1 \leftarrow \mathcal{A} \mid b_{\mathcal{A}} = 0]) \\
&= \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{GOAL-POWER},X,n}(\mathcal{A}),
\end{aligned}$$

as required. \square

The General Case. One may wonder why we had to distinguish between the multi-bit and the single-bit settings in Theorems 19 and 20, instead of directly proving a result for the generalized security notion from Definition 7. One may be tempted to think that the latter should be possible, again with a linear loss n . We remark that this seems not to be the case; removing the assumptions on the adversary in Theorems 19 and 20 (that is, the restriction that \mathcal{A} always chooses $i = j$, or $j = 1$, respectively) appears difficult.

In the game played by the single-key adversary \mathcal{B} , there is one secret key K and one hidden bit b ; in the proof for the multi-bit case (Theorem 19) \mathcal{B} chooses random $i^* \leftarrow_{\$}[n]$ and associates b_{i^*} and K_{i^*} with b and K , respectively, while in the proof for the single-bit case (Theorem 20) \mathcal{B} chooses random $i^* \leftarrow_{\$}[n]$ and associates b_1 and K_{i^*} with b and K , respectively. In the general case, the multi-key adversary \mathcal{A} can associate the same secret key with multiple bits and even corrupt this key as long as none of those bits are subsequently guessed; it is unclear how \mathcal{B} can simulate this consistently with just one secret key K and one hidden bit b . To answer queries from \mathcal{A} , \mathcal{B} may have to ‘combine’ keys that it has generated with the hidden bit b it does not know, which does not seem possible using the oracles available to \mathcal{B} .

Nevertheless, our main objective behind the introduction of the generalised multi-key security definition (Definition 7) was to capture both natural variants of multi-key security – the single-bit setting and the multi-bit setting – within a unified framework. Theorems 19 and 20 show that security with respect to both these natural notions is implied by security with respect to a corresponding single-key security notion, albeit with a loss of n , the number of keys.

Relations among Multi-Key Notions. In the single-key case, many security notions that at first appear different are, in fact, equivalent [5,7]. We already mentioned the equivalence between left-or-right versus real-or-random challenge oracles. In a different dimension, some oracle access can become redundant, for instance LRIND-PAS and LRIND-CPA are equivalent: a valid adversary can use its challenge oracle LR to perfectly emulate access to a ‘true’ encryption oracle.

However, if the adversary plays a game with $n > 1$ keys and is allowed to adaptively corrupt some of these keys, achieving tight reductions between the various notions becomes considerably harder. For instance, the reduction used to show equivalence between LRIND-PAS and LRIND-CPA above no longer works for a valid LRIND-CPA adversary that made both a corrupt and a true encryption query to the same key: the reduction would not yield a valid LRIND-PAS adversary.

Nevertheless, we show in Theorem 21 that relations among the single-key security notions induce relations among the n -key security notions, but this approach results in an additional loss factor of n in any relation, compared to the single-key case. Therefore it remains an open problem to determine if any of the multi-key security notions are (tightly) equivalent.

The disappearance of known equivalences was previously observed by Bellare et al. [9] in the context of multi-instance security (here an adversary has to guess *all* the challenge bits). They argued why left-or-right security was the most ‘robust’ notion in their context. Rather than trying to determine the most robust notion in our context, we ensure that our negative results hold for almost all conceivable security notions covered by Definition 7. For instance, we treat *both* the real-or-random and left-or-right formulations of confidentiality, without having to worry how tightly they relate to one another in our multi-key setting. Thus we rule out bypassing our negative results simply by tweaking the multi-key security definition.

Theorem 21. *Suppose $G = \text{GOAL-POWER}$ and $H = \text{GOAL'-POWER'}$, where GOAL , POWER , GOAL' and POWER' are such that, for any valid adversary \mathcal{A}_1 in $\mathbb{G}_{\mathcal{E}}^{X,1}$ there is a valid adversary \mathcal{B}_1 in $\mathbb{H}_{\mathcal{E}}^{X,1}$ with close to the same runtime as \mathcal{A}_1 and*

$$\text{Adv}_{\mathcal{E}}^{\text{H},X,1}(\mathcal{B}_1) \geq \frac{1}{L} \cdot \text{Adv}_{\mathcal{E}}^{\text{G},X,1}(\mathcal{A}_1)$$

for some L .

Then, for each valid adversary \mathcal{A}_n in $\mathbb{G}_{\mathcal{E}}^{X,n}$ where, for all queries of the form $\text{Enc}(i, j, M, N)$, $\text{Dec}(i, j, C, N)$ or $\text{LR}(i, j, M_0, M_1, N)$ it holds that $i = j$ (or $i = 1$, respectively), there is a valid adversary \mathcal{B}_n in $\mathbb{H}_{\mathcal{E}}^{X,n}$ satisfying the same constraint on its challenge queries as \mathcal{A}_n , with close to the same runtime as \mathcal{A}_n and

$$\text{Adv}_{\mathcal{E}}^{\text{H},X,n}(\mathcal{B}_n) \geq \frac{1}{n \cdot L} \cdot \text{Adv}_{\mathcal{E}}^{\text{G},X,n}(\mathcal{A}_n).$$

Proof. We first use Theorem 19 (or Theorem 20, respectively) to build a valid adversary \mathcal{A}_1 from \mathcal{A}_n with

$$\text{Adv}_{\mathcal{E}}^{\text{G},X,1}(\mathcal{A}_1) \geq \frac{1}{n} \cdot \text{Adv}_{\mathcal{E}}^{\text{G},X,n}(\mathcal{A}_n).$$

Next, we use the hypothesis of the theorem to build a valid adversary \mathcal{B}_1 from \mathcal{A}_1 , with

$$\text{Adv}_{\mathcal{E}}^{\text{H},X,1}(\mathcal{B}_1) \geq \frac{1}{L} \cdot \text{Adv}_{\mathcal{E}}^{\text{G},X,1}(\mathcal{A}_1).$$

Finally, we apply the trivial tight implication that, from a valid adversary \mathcal{B}_1 in $\mathbb{H}_{\mathcal{E}}^{X,1}$, we can construct a valid adversary \mathcal{B}_n in $\mathbb{H}_{\mathcal{E}}^{X,n}$ with identical running time and success probability, i.e.

$$\text{Adv}_{\mathcal{E}}^{\text{H},X,n}(\mathcal{B}_n) \geq \text{Adv}_{\mathcal{E}}^{\text{H},X,1}(\mathcal{B}_1).$$

Note that \mathcal{B}_n can easily satisfy the same constraint on its challenge queries as \mathcal{A}_n , since \mathcal{B}_n can associate any one of its n keys K_1, \dots, K_n and any one of its n bits b_1, \dots, b_n with the single key K and hidden bit b , respectively, in the game played by \mathcal{B}_1 . This yields the claim. \square

C Proof of Lemma 17

We prove that access to honest encryption and decryption oracles does not (meaningfully) help an adversary in the security game for authenticated encryption. More precisely, the maximum advantage of a valid adversary in the AE-CCA game is at most twice the maximum advantage of a valid adversary with the same resources in the AE-PAS game.

Proof. It is easy to show that

$$\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{A}) = \Pr [1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1] - \Pr [1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 0].$$

Now let $\mathbb{G}_{b,b'}$ be the game played by \mathcal{A} where the Enc and Dec oracles respond according to b and, if $b' = 0$ the \mathcal{E} and \mathcal{D} oracles respond as normal, but if $b' = 1$ these oracles respond exactly as Enc and Dec respond when $b = 1$, i.e. with random strings of the appropriate length and \perp , respectively.

It follows that

$$\begin{aligned}
& \text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{A}) \\
&= \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 0] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 0, b' = 0] \\
&= (\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 0] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 1]) \\
&\quad + (\Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 0, b' = 0]).
\end{aligned}$$

Now we construct adversaries \mathcal{A}_1 and \mathcal{A}_2 against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ that run in the same time as \mathcal{A} and make the same number of oracle queries as \mathcal{A} . The adversaries \mathcal{A}_1 and \mathcal{A}_2 use their Enc and Dec oracles to simulate the environment of \mathcal{A} . Then when \mathcal{A} returns a bit, this is the bit that \mathcal{A}_1 and \mathcal{A}_2 output in their games.

When \mathcal{A} makes queries to Enc or Dec, \mathcal{A}_1 simulates the ‘fake oracles’, i.e. responds with random strings and \perp , respectively. When \mathcal{A} makes queries to \mathcal{E} or \mathcal{D} , \mathcal{A}_1 forwards these to its own Enc and Dec oracles and returns the responses to \mathcal{A} . Thus \mathcal{A}_1 correctly simulates the game $G_{1,b'}$ for \mathcal{A} , where b' is the hidden bit in the game played by \mathcal{A}_1 .

On the other hand, \mathcal{A}_2 forwards all queries from \mathcal{A} , whether to the challenge oracles or the ‘always real’ oracles, to its challenge oracles Enc and Dec and returns the responses to \mathcal{A} . Thus \mathcal{A}_2 correctly simulates the game $G_{b'',b''}$ for \mathcal{A} , where b'' is the hidden bit in the game played by \mathcal{A}_2 .

It follows that

$$\begin{aligned}
\text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{A}_1) &= \Pr[1 \leftarrow \mathcal{A}_1^{\mathcal{O}} \mid b'' = 1] - \Pr[1 \leftarrow \mathcal{A}_1^{\mathcal{O}} \mid b'' = 0] \\
&= \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 0] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 1],
\end{aligned}$$

and similarly,

$$\begin{aligned}
\text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{A}_2) &= \Pr[1 \leftarrow \mathcal{A}_2^{\mathcal{O}} \mid b'' = 1] - \Pr[1 \leftarrow \mathcal{A}_2^{\mathcal{O}} \mid b'' = 0] \\
&= \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 1, b' = 1] - \Pr[1 \leftarrow \mathcal{A}^{\mathcal{O}} \mid b = 0, b' = 0].
\end{aligned}$$

Thus

$$\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{A}) = \text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{A}_1) + \text{Adv}_{\mathcal{E}}^{\text{AE-PAS},X,1}(\mathcal{A}_2)$$

and the result follows if we can show that \mathcal{A}_1 and \mathcal{A}_2 are valid in $\text{AE-PAS}_{\mathcal{E}}^{X,1}$.

Since \mathcal{A} is a valid adversary in $\text{AE-CCA}_{\mathcal{E}}^{X,1}$, \mathcal{A} uses nonces correctly with respect to X . Therefore the same is true of \mathcal{A}_1 and \mathcal{A}_2 . It remains to check that \mathcal{A}_1 and \mathcal{A}_2 make no pointless or prohibited queries, given that \mathcal{A} makes no pointless or prohibited queries.

If \mathcal{A}_1 repeated a query to Dec or Enc, these would be repeated queries to \mathcal{E} or \mathcal{D} from \mathcal{A} , which would be pointless or prohibited queries for \mathcal{A} . If \mathcal{A}_1 made a query $\text{Enc}(M, N)$, where $M \neq \perp$ was the response to a query $\text{Dec}(C, N)$, then this would be a query $\mathcal{E}(M, N)$ from \mathcal{A} , where $M \neq \perp$ was the response to a query $\mathcal{D}(C, N)$, which would be a pointless query for \mathcal{A} . Finally if \mathcal{A}_1 made a query $\text{Dec}(C, N)$, where C was the response to a query $\text{Enc}(M, N)$, then this would be a query $\mathcal{D}(C, N)$ from \mathcal{A} , where C was the response to a query $\mathcal{E}(M, N)$, which would be a pointless query for \mathcal{A} . Thus, \mathcal{A}_1 makes no pointless or prohibited queries.

Similarly, if \mathcal{A}_2 repeated a query to Dec or Enc, then either: \mathcal{A} repeated a query to \mathcal{E} , \mathcal{D} , Enc or Dec, which would be a pointless or prohibited query for \mathcal{A} ; or \mathcal{A} made the same query to \mathcal{E} and Enc, in either order, which would be a prohibited query for \mathcal{A} ; or \mathcal{A} made the same query to \mathcal{D} and Dec, in either order. The latter repeated query is valid for \mathcal{A} , but leads to a pointless repeated query to Dec for \mathcal{A}_2 . To avoid this pointless query, \mathcal{A}_2 can just return the response from the first query (and correctly simulate \mathcal{A} 's environment in the game where \mathcal{D} has been replaced by Dec). If \mathcal{A}_2 made a query $\text{Enc}(M, N)$, where $M \neq \perp$ was the response to a query $\text{Dec}(C, N)$, then this would be a query $\mathcal{E}(M, N)$ or $\text{Enc}(M, N)$ from \mathcal{A} , where $M \neq \perp$ was the response to a query $\mathcal{D}(C, N)$ or $\text{Dec}(C, N)$, which is either a pointless or prohibited query for \mathcal{A} . Finally if \mathcal{A}_2 made a query $\text{Dec}(C, N)$, where C was the response to a query $\text{Enc}(M, N)$, then this would be a query $\mathcal{D}(C, N)$ or $\text{Dec}(C, N)$ from \mathcal{A} , where C was the response to a query $\mathcal{E}(M, N)$ or $\text{Enc}(M, N)$, which is either a pointless or prohibited query for \mathcal{A} . Thus, \mathcal{A}_2 makes no pointless or prohibited queries. \square

D Lower Bounds for Simple Reductions for Secrecy and Integrity Notions

In this section we provide lower bounds on simple reductions to single-key security notions for privacy. Informally, we show that for all $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathcal{M}^l$ then: if all reasonable adversaries have small $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ advantage, then any simple reduction from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ loses approximately n .

We proceed in the same way as for $\text{AE-PAS}_{C[\mathcal{E}]}^{X',n}$: we first prove an impossibility result for reductions from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to the relevant single-key game, then combine this with the non-lossy reduction from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ from Lemma 12 (after relabelling X and \mathcal{E}). The proofs are similar to those of Theorem 15 and Corollary 18. We omit some of the details and highlight the main differences.

Secrecy Notions. We start with lowerbounds for simple reductions from multi-key to single-key secrecy properties.

Theorem 22. *Let \mathcal{E} and $C[\mathcal{E}]$ be authenticated encryption schemes such that $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathcal{M}^l$. Then, for $X \in \{\text{IV}, \text{NR}, \text{MR}\}$, any simple reduction \mathcal{R} from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ loses $L = \left(\frac{1}{n} + 2\epsilon\right)^{-1}$. Here ϵ is the maximum advantage of a valid adversary against $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ that runs in time at most $nt_{\mathcal{R}} + 2l(n-1)t_{C[\mathcal{E}]}$ and $t_{C[\mathcal{E}]}$ is an upper-bound on the runtime of a single encryption with $C[\mathcal{E}]$.*

Proof. We reuse the hypothetical adversary \mathcal{A} against $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ that was described in the proof of Theorem 15. We construct a meta-reduction \mathcal{M} that simulates the environment of \mathcal{R} in its interaction with \mathcal{A} and uses the output of \mathcal{R} to play the $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ game.

The meta-reduction \mathcal{M} is exactly as is described in the proof of Theorem 15, except that now \mathcal{M} and \mathcal{R} may only query an encryption oracle, not a decryption oracle. As before, queries from $\mathcal{R}_2^{i^*}$ are forwarded, while queries from $\mathcal{R}_2^i, i \neq i^*$ are simulated: when \mathcal{R}_2^i makes an encryption query (M, N) , \mathcal{M} samples $C \leftarrow_{\$} \{0, 1\}^{\text{enclen}(|M|)}$ and returns this to \mathcal{R}_2^i . How \mathcal{M} computes K^* and the conditions under which \mathcal{M} aborts and returns a random bit are exactly as before.

We remark that, since \mathcal{R} is a valid adversary against $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ and all the oracle queries made by \mathcal{M} were chosen by $\mathcal{R}_1, \mathcal{R}_2^{i^*}$ or \mathcal{R}_3 , \mathcal{M} uses nonces correctly with respect to X and does not make pointless or prohibited queries. So \mathcal{M} is a valid adversary against $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ and $\text{Adv}_{\mathcal{E}}^{\text{IND-PAS}, X, 1}(\mathcal{M}) \leq \epsilon$.

As before, \mathcal{M} correctly simulates the view of $(\mathcal{R}_1, \mathcal{R}_2^{i^*}, \mathcal{R}_3)$ in the game $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ and correctly simulates the responses from \mathcal{A} , provided that the event BAD does not occur.

This shows that

$$\Pr \left[\text{IND-PAS}_{\mathcal{E}}^{X,1}(\mathcal{R}) = 1 \right] = \Pr \left[\text{IND-PAS}_{\mathcal{E}}^{X,1}(\mathcal{M}) = 1 \mid \neg \text{BAD} \right]$$

and hence:

$$\text{Adv}_{\mathcal{E}}^{\text{IND-PAS}, X, 1}(\mathcal{M}) \geq \text{Adv}_{\mathcal{E}}^{\text{IND-PAS}, X, 1}(\mathcal{R}) - \Pr[\text{BAD}].$$

We now show that $\Pr[\text{BAD}] \leq \frac{1}{n} + \epsilon$, which proves the theorem.

As in the proof of Lemma 16, we consider an alternative meta-reduction \mathcal{M}' that simulates the encryption queries from all \mathcal{R}_2^i , including $\mathcal{R}_2^{i^*}$. With $\text{BAD}_{\mathcal{M}}$ and $\text{BAD}_{\mathcal{M}'}$ defined as before, we have

$$\Pr[\text{BAD}] = \frac{1}{2} (\Pr[\text{BAD}_{\mathcal{M}} \mid b = 0] - \Pr[\text{BAD}_{\mathcal{M}'} \mid b = 0]) + \frac{1}{n}.$$

We construct an adversary \mathcal{B} that simulates the environment of \mathcal{R}_1 and the \mathcal{R}_2^i in their interaction with either \mathcal{M} or \mathcal{M}' , depending on the hidden bit b' in the game played by \mathcal{B} . If BAD occurs, \mathcal{B} will output 0. Otherwise \mathcal{B} will output 1. In this case, \mathcal{B} plays the $\text{IND-CPA}_{\mathcal{E}}^{X,1}$ game where, in addition to the ‘challenge’ oracle Enc , one can query a genuine encryption oracle \mathcal{E} , provided that one does

not make pointless or prohibited queries. It has been shown that the advantage of an adversary in this game is at most twice the maximum advantage of an adversary, with the same resources, in the standard $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ game [5].

As in the proof of Theorem 15, \mathcal{B} uses its genuine encryption oracle to answer queries from \mathcal{R}_1 , simulates the fake encryption oracle to answer queries from $\mathcal{R}_2^i, i \neq i^*$ and uses its challenge encryption oracle to answer queries from $\mathcal{R}_2^{i^*}$.

Since \mathcal{R} is a valid adversary in $\text{IND-PAS}_{\mathcal{E}}^{X,1}$, it is clear that \mathcal{B} uses nonces correctly with respect to X . Moreover, \mathcal{B} makes no pointless or prohibited queries:

- A repeated query to \mathcal{E} by \mathcal{B} would correspond to a repeated query to Enc by \mathcal{R}_1 , which would be a prohibited query in the game played by \mathcal{R} .
- A repeated query to Enc by \mathcal{B} would correspond to a repeated query to Enc by $\mathcal{R}_2^{i^*}$, which would be a prohibited query in the game played by \mathcal{R} .
- If \mathcal{B} were to make two queries of the form $\text{Enc}(M, N)$ and $\mathcal{E}(M, N)$, then these would correspond to a query $\text{Enc}(M, N)$ by \mathcal{R}_1 and a query $\text{Enc}(M, N)$ by $\mathcal{R}_2^{i^*}$, but the latter would be a prohibited query in the game played by \mathcal{R} .

Thus \mathcal{B} is a valid adversary in $\text{IND-CPA}_{\mathcal{E}}^{X,1}$ and $\text{Adv}_{\mathcal{E}}^{\text{IND-CPA},X,1}(\mathcal{B}) \leq 2\epsilon$.

It follows that:

$$\begin{aligned} 2\epsilon &\geq \text{Adv}_{\mathcal{E}}^{\text{IND-CPA},X,1}(\mathcal{B}) = 2 \left(\Pr \left[\text{IND-CPA}_{\mathcal{E}}^{X,1}(\mathcal{B}) = 1 \right] - \frac{1}{2} \right) \\ &= \Pr [0 \leftarrow \mathcal{B} \mid b' = 0] - \Pr [0 \leftarrow \mathcal{B} \mid b' = 1] \\ &= \Pr [\text{BAD}_{\mathcal{M}} \mid b = 0] - \Pr [\text{BAD}_{\mathcal{M}'} \mid b = 0] \\ &= 2 \left(\Pr [\text{BAD}] - \frac{1}{n} \right), \end{aligned}$$

which yields the result. \square

Corollary 23. *Let \mathcal{E} and $C[\mathcal{E}]$ be authenticated encryption schemes such that $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathcal{M}^l$. Let $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$, $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$, $X, X' \in \{\text{IV}, \text{NR}, \text{MR}\}$ and $n > 1$. Then, all simple reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ must lose*

$$L = \delta_{X'} \cdot \delta_{\text{GOAL}} \cdot \left(\frac{1}{n} + 2\epsilon \right)^{-1},$$

with $\delta_{X'}$ as given in Lemma 12 and ϵ as given in Theorem 22.

Proof. Recall, from the proof of Lemma 12, the reduction \mathcal{R} from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ (after relabelling \mathcal{E} and X). Suppose there is a simple reduction \mathcal{R}' from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{IND-PAS}_{\mathcal{E}}^{X,1}$, where the loss is less than L for all valid adversaries against $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$. Then one can combine \mathcal{R} and \mathcal{R}' to obtain a reduction $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{IND-PAS}_{\mathcal{E}}^{X,1}$ that contradicts Theorem 22. \square

Integrity Notions. We prove similar results for integrity. Specifically we show that if all reasonable adversaries have small $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ advantage, then any simple reduction from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ loses approximately n .

Notice that to show the lower bound for reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$, we have to assume \mathcal{E} has both authenticity *and* secrecy. While this suffices for authenticated encryption schemes, it does not rule out non-lossy reductions for schemes that have authenticity but not secrecy. The technical reason for this additional assumption is that, to make sure that the meta-reduction \mathcal{M} is a valid adversary against $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$, i.e. does not attempt to decrypt a ciphertext that was output

of the encryption oracle, \mathcal{M} must suppress both the encryption and decryption queries made during the rewinding of the reduction \mathcal{R} . For the case $X = \text{IV}$, it should be possible to argue that prohibited decryption queries are unlikely to occur, even if \mathcal{R} can make encryption queries during the rewinding phase, as the nonces are chosen independently at random for each encryption query. We leave this for future work. In any case, if nonces are chosen (possibly deterministically) by \mathcal{R} and \mathcal{M} does not suppress the encryption queries made by \mathcal{R} during the rewinding phase, then it is impossible to avoid \mathcal{M} making prohibited queries.

Theorem 24. *Let \mathcal{E} and $C[\mathcal{E}]$ be authenticated encryption schemes such that $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathcal{M}^l$. Then, for $X \in \{\text{IV}, \text{NR}, \text{MR}\}$, any simple reduction \mathcal{R} from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ loses $L = (\frac{1}{n} + 3\epsilon)^{-1}$. Here ϵ is the maximum advantage of a valid adversary against $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ running in time at most $nt_{\mathcal{R}} + 2l(n-1)t_{C[\mathcal{E}]}$ and $t_{C[\mathcal{E}]}$ is an upper-bound on the runtime of a single encryption with $C[\mathcal{E}]$.*

Proof. The proof is similar to that of Theorem 15. Indeed, we do not modify the meta-reduction \mathcal{M} or its variant \mathcal{M}' at all, since \mathcal{R} , \mathcal{M} and \mathcal{M}' have access to both an encryption oracle and a decryption oracle in the $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ game, just as in the $\text{AE-PAS}_{\mathcal{E}}^{X,1}$ game.⁵

However, whereas before we had $\Pr[\text{BAD}] \leq \frac{1}{n} + \epsilon$, when \mathcal{M} plays the $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ game we can only obtain the looser bound $\Pr[\text{BAD}] \leq \frac{1}{n} + 2\epsilon$. When $b = 1$ in the game played by \mathcal{M} or \mathcal{M}' , the meta-reductions have access to a *real* encryption oracle and a fake decryption oracle, so \mathcal{M} responds to the encryption queries from \mathcal{R}_2^{i*} with real ciphertexts, while \mathcal{M}' responds with random strings. Thus there is a difference in the view of \mathcal{R}_2^{i*} between \mathcal{M} and \mathcal{M}' even when $b = 1$, which was not the case in the proof of Theorem 15.

Accordingly, $\Pr[\text{BAD}] \leq \frac{1}{n} + \gamma$, where γ is given by:

$$\frac{1}{2} (\Pr[\text{BAD}_{\mathcal{M}} | b = 0] - \Pr[\text{BAD}_{\mathcal{M}'} | b = 0]) + \frac{1}{2} (\Pr[\text{BAD}_{\mathcal{M}} | b = 1] - \Pr[\text{BAD}_{\mathcal{M}'} | b = 1]).$$

We construct adversaries \mathcal{B}_0 and \mathcal{B}_1 , running in time at most $nt_{\mathcal{R}} + 2l(n-1)t_{C[\mathcal{E}]}$ and making at most $q_{\mathcal{R}}$ oracle queries, such that

$$\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{B}_0) = \Pr[\text{BAD}_{\mathcal{M}} | b = 0] - \Pr[\text{BAD}_{\mathcal{M}'} | b = 0]$$

and

$$\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{B}_1) = \Pr[\text{BAD}_{\mathcal{M}} | b = 1] - \Pr[\text{BAD}_{\mathcal{M}'} | b = 1]$$

and we show that both \mathcal{B}_0 and \mathcal{B}_1 are valid adversaries. By Lemma 17, this shows that $\Pr[\text{BAD}] \leq \frac{1}{n} + 2\epsilon$ and hence

$$\epsilon_{\mathcal{R}} \leq \frac{1}{n} + 3\epsilon$$

as required.

For \mathcal{B}_0 we simply use the adversary \mathcal{B} described in the proof of Theorem 15. The adversary \mathcal{B}_1 is as follows: first, \mathcal{B}_1 runs \mathcal{R}_1 , forwarding all encryption queries to the genuine encryption oracle and responding to all decryption queries with \perp . Then \mathcal{B}_1 carries out the same checks as \mathcal{M} (or \mathcal{M}') and, if the checks succeed, \mathcal{B}_1 samples $i^* \leftarrow_{\$} [n]$ and, for each $i \in [n]$, \mathcal{B}_1 runs \mathcal{R}_2 on input (i, st_1) .

When \mathcal{R}_2^i makes oracle queries:

1. If $i = i^*$, \mathcal{B}_1 answers all *encryption* queries with its challenge encryption oracle, but answers all *decryption* queries with \perp .
2. If $i \neq i^*$, \mathcal{B}_1 simulates the ‘fake’ oracles, Enc and Dec in the case $b = 1$, just as \mathcal{M} (or \mathcal{M}') does.

⁵ Even though \mathcal{M} and \mathcal{M}' now query the oracles \mathcal{E} and Dec instead of Enc and Dec.

Finally, \mathcal{B}_1 outputs 0 if BAD has occurred and 1 otherwise.

Let b' be the hidden bit in the $\text{AE-CCA}_{\mathcal{E}}^{X,1}$ game played by \mathcal{B}_1 and recall that b is the hidden bit in the $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ game played by \mathcal{M} or \mathcal{M}' . Then, given that $b = 1$, \mathcal{B}_1 correctly simulates the view of \mathcal{R}_1 and the \mathcal{R}_2^i in their interaction with \mathcal{M} (respectively, \mathcal{M}') if $b' = 0$ (respectively, $b' = 1$). This shows that

$$\begin{aligned}\Pr[0 \leftarrow \mathcal{B}_1 \mid b' = 0] &= \Pr[\text{BAD}_{\mathcal{M}} \mid b = 1], \\ \Pr[0 \leftarrow \mathcal{B}_1 \mid b' = 1] &= \Pr[\text{BAD}_{\mathcal{M}'} \mid b = 1].\end{aligned}$$

Then we have:

$$\begin{aligned}\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{B}_1) &= 2 \left(\Pr[\text{AE-CCA}_{\mathcal{E}}^{X,1}(\mathcal{B}_1) = 1] - \frac{1}{2} \right) \\ &= \Pr[0 \leftarrow \mathcal{B}_1 \mid b' = 0] - \Pr[0 \leftarrow \mathcal{B}_1 \mid b' = 1] \\ &= \Pr[\text{BAD}_{\mathcal{M}} \mid b = 1] - \Pr[\text{BAD}_{\mathcal{M}'} \mid b = 1].\end{aligned}$$

Now consider $\text{Adv}_{\mathcal{E}}^{\text{AE-CCA},X,1}(\mathcal{B}_1)$. Clearly \mathcal{B}_1 uses nonces correctly with respect to X , since \mathcal{R} is a valid adversary against $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$. Moreover, \mathcal{B}_1 does not make pointless or prohibited queries: there are no queries to Dec or \mathcal{D} ; repeated queries to \mathcal{E} or Enc by \mathcal{B}_1 would correspond to repeated queries to \mathcal{E} by \mathcal{R} (which are pointless); and if \mathcal{B} were to make two queries of the form $\text{Enc}(M, N)$ and $\mathcal{E}(M, N)$, then these would correspond to a query $\mathcal{E}(M, N)$ by \mathcal{R}_1 and a query $\mathcal{E}(M, N)$ by \mathcal{R}_2^{i*} , but the latter would be a pointless query in the game played by \mathcal{R} . Therefore \mathcal{B}_1 is a valid adversary against $\text{AE-CCA}_{\mathcal{E}}^{X,1}$. (In fact, \mathcal{B}_1 can be viewed as a valid adversary against $\text{IND-CPA}_{\mathcal{E}}^{X,1}$, but we have to assume both integrity and secrecy for the rest of the proof, so it is simpler to view \mathcal{B}_1 as an adversary against $\text{AE-CCA}_{\mathcal{E}}^{X,1}$.) \square

Corollary 25. *Let \mathcal{E} and $C[\mathcal{E}]$ be authenticated encryption schemes such that $C[\mathcal{E}]$ is M -key-unique for some $M \in \mathcal{M}^l$. Let $\text{GOAL} \in \{\text{AE}, \text{LRAE}, \text{IND}, \text{LRIND}, \text{CTI}\}$, $\text{POWER} \in \{\text{CCA}, \text{CPA}\}$, $X, X' \in \{\text{IV}, \text{NR}, \text{MR}\}$ and $n > 1$. Then all simple reductions from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ must lose*

$$L = \delta_{X'} \cdot \delta_{\text{GOAL}} \cdot \left(\frac{1}{n} + 3\epsilon \right)^{-1},$$

with $\delta_{X'}$ as given in Lemma 12 and ϵ as given in Theorem 24.

Proof. Again, we use the reduction \mathcal{R} from $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ given in the proof of Lemma 12 (after relabelling X and \mathcal{E}). If there were a simple reduction \mathcal{R}' from $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$, where the loss was less than L for all valid adversaries against $\text{GOAL-POWER}_{C[\mathcal{E}]}^{X',n}$, one could combine \mathcal{R} and \mathcal{R}' to obtain a reduction $\text{KEYREC}_{C[\mathcal{E}]}^{M,n}$ to $\text{CTI-CPA}_{\mathcal{E}}^{X,1}$ that contradicted Theorem 24. \square