Tightness Subtleties for Multi-user PKE Notions

Hans Heum[ID] and Martijn Stam[(✉)][ID]

Simula UiB, Bergen, Norway
`{hansh,martijn}@simula.no`

Abstract. Public key encryption schemes are increasingly being studied concretely, with an emphasis on tight bounds even in a multi-user setting. Here, two types of formalization have emerged, one with a single challenge bit and one with multiple challenge bits. Another modelling choice is whether to allow key corruptions or not. How tightly the various notions relate to each other has hitherto not been studied in detail. We show that in the absence of corruptions, single-bit left-or-right indistinguishability is the preferred notion, as it tightly implies the other (corruption-less) notions. However, in the presence of corruptions, this implication no longer holds; we suggest the use of a more general notion that tightly implies both existing options. Furthermore, for completeness we study how the relationship between left-or-right versus real-or-random evolves in the multi-user PKE setting.

Keywords: Indistinguishability · Public key encryption · Multi-user security · Adaptive corruptions

1 Introduction

Historically, a primitive like public key encryption (PKE) is often studied in a setting where a single key-pair is generated for an adversary to attack, often based on a single challenge ciphertext only [27]. Yet, in reality there will be many users, each generating their own key pairs, to be used repeatedly. To study the concrete security risk of having very many keys in play simultaneously, Bellare et al. [5] introduced the multi-user setting. They considered an adversary with access to n different public keys and the ability to challenge (in an indistinguishability fashion) each of them, and concluded that the security loss is at worst linear in the total number challenge queries. Loosely speaking, such a linear security loss implies that a scheme that is believed to offer, say, 128-bit security in the single user setting, may only guarantee 80-bit security if there are 2^{20} users each receiving 2^{28} messages (based on the same hardness assumption).

Unfortunately, there have been ample examples of schemes where practical attacks can indeed exploit the increased attack surface, demonstrating that these theoretical security losses can be realized. Consequently, the generic tightness losses to move from a single-user, single-challenge setting to a more realistic

multi-user, multi-challenge setting are problematic as, conservatively, one would have to increase key sizes to compensate. Alternatively, a growing number of works have looked at schemes with tighter security guarantees, either if the number of users goes up, the number of challenge encryptions per key goes up, or both [2,5,12,16,21,22,28].

Moreover, in a system with many users, it is not inconceivable that some private keys eventually become available to an adversary, which can be modelled using key corruptions. An adversary learning a private key can obviously decrypt all ciphertexts that were encrypted under the corresponding public key, thus some care has to be taken to avoid trivial wins when allowing key corruptions. The two simplest mechanisms are either using independent challenge bits for each key or disallowing an adversary to both challenge and corrupt a single key. As we detail in Appendix A, both these mechanisms have been used, also in related contexts such as key encapsulation mechanisms (KEMs), authenticated encryption (AE), and authenticated key exchange (AKE), raising the inevitable question which notion should be the preferred one.

In the context of lower bounding tightness losses for multi-user AE, Jager et al. [25] employed a novel multi-key, multi-challenge-bit notion that generalizes both mechanisms; however, the main motivation of this generalized mechanism was universality of their impossibility result, allowing them to side-step the question which mechanism to focus on. Recently, in the context of AKE, Jager et al. [24] argued in favour of the single-bit notion, primarily as it composes more easily. For KEMs a similar argument holds, yet for PKE composition is arguably less relevant. Instead, a more direct interpretation of what the various notions entail might well be preferable.

Our Contribution. Both the single-bit and multi-bit approaches are implied by the single user notion at the cost of a tightness loss linear in the number of users. Consequently, the two multi-user notions are also within that linear factor in the number of user. As our goal is to avoid such tightness losses, we are interested in identifying the most suitable, general notion as possible, guaranteeing that there are no "hidden" linear losses in the choice of notion—an issue already pointed to by Jager et al. [24]

To this end, we adapt the multi-key, multi-bit notion of Jager et al. [25] to the PKE setting, slightly generalizing it in the process. We show how it tightly implies, and therefore unifies, the previous multi-user notions, and give novel interpretations of each (see Sect. 3).

We then shift our focus to how tightly the different notions relate to each other, with the goal of identifying the strongest, and therefore preferred, multi-user notions. We find that the answer depends on whether or not corruptions are present: in the absence of corruptions, we find that the single-challenge-bit notion is *as strong or stronger* than any of the other (see Sect. 4.2). Given that this notion is significantly simpler than the fully general game, this makes the single-bit notion the preferred one in the absence of corruptions. With corruptions, this

relation breaks down, and the general "free-bit" game indeed seems the stronger, and therefore preferred, notion (see Sect. 4.3).

Finally, we fill some holes largely left as folklore until now regarding how the well-known factor-2 reduction from real-or-random to left-or-right indistinguishability, as shown by Bellare et al. [7] for the single-user, single-challenge setting, generalizes to the multi-user setting. We find that, as expected, the relation remains intact in the single-bit setting, regardless of whether corruptions are present (see Sect. 4.4). In contrast, with multiple challenge bits the best-known reductions turn lossy. Whether these losses are inevitable remains open; however, it reinforces the by now established notion that left-or-right indistinguishability is to be preferred over its real-or-random counterpart whenever possible.

The appendices contain much additional material: highlights include Appendix A giving context to the present work by presenting a more complete history of multi-user indistinguishability than that presented here, and Appendix F, illustrating the difficulty of achieving tight composition in multi-bit settings, as alluded to by Jager et al. [24], by giving an overview of how additional losses can appear in PKE schemes built using the widely used KEM/DEM paradigm.

2 Preliminaries

2.1 Notation

For an integer n, we will write $[n]$ for the set $\{1, \ldots, n\}$. We will also use the abbreviation $X \xleftarrow{\cup} x$ for the operation $X \leftarrow X \cup \{x\}$. The event of an adversary $\mathbb{A}$ outputting 0 is denoted $0 \leftarrow \mathbb{A}$. We use $\Pr[Code : Event \,|\, Condition]$ to denote the conditional probability of $Event$ occurring when $Code$ is executed, conditioned on $Condition$. We omit $Code$ when it is clear from the context and $Condition$ when it is not needed.

2.2 PKE Syntax

A public key encryption scheme PKE consists of three algorithms: the probabilistic *key generation* algorithm Pk.Kg, which takes as input some system parameter pm and outputs a public/private key pair $(\mathsf{pk}, \mathsf{sk}) \in (\mathcal{PK}, \mathcal{SK})$; the probabilistic *encryption* algorithm Pk.Enc, which on input a public key $\mathsf{pk} \in \mathcal{PK}$ and a message $m \in \mathcal{M}$, outputs a ciphertext c; and the deterministic *decryption* algorithm Pk.Dec, which on input of a secret key $\mathsf{sk} \in \mathcal{SK}$ and a ciphertext c, outputs either the message m, or a special symbol $\bot$ denoting failure.

We allow the message space $\mathcal{M}$ to depend on the parameters pm, but insist it is independent of the public key pk. We furthermore assume that there exists an equivalence relation $\sim$ on the message space that partitions $\mathcal{M}$ into finite equivalence classes. For $m \in \mathcal{M}$, we let $\llbracket m \rrbracket$ denote its equivalence class, so $\llbracket m \rrbracket = \{\tilde{m} \in \mathcal{M} : m \sim \tilde{m}\}$. Often $\mathcal{M}$ consists of arbitrary length bitstrings, or at least all bitstrings up to some large length (e.g. 2^{64}), and two messages are equivalent

iff they have the same length, so $[\![m]\!] = \{0,1\}^{|m|}$; for other cryptosystems, such as ElGamal, messages are group elements that are essentially all equivalent, so $[\![m]\!] = \mathcal{M}$. (Note that the case where $[\![m]\!] = \{m\}$ for all m is degenerate and 'security' is often trivially satisfied.)

The scheme must satisfy ϵ-correctness [20], namely that for any pm:

$$\mathbb{E}_{(\mathsf{pk},\mathsf{sk}) \leftarrow\$ \, \mathsf{Pk.Kg}(\mathsf{pm})} \left[\max_{m \in \mathcal{M}} \Pr[c \leftarrow\$ \, \mathsf{Pk.Enc}_{\mathsf{pk}}(m)) \,:\, \mathsf{Pk.Dec}_{\mathsf{sk}}(c) \neq m] \right] \leq \epsilon \,.$$

If $\epsilon = 0$ we speak of perfect correctness; the case $\epsilon > 0$ is especially useful to model decryption errors typical to lattice-based schemes.

Remark 1. The system parameters pm are implicitly input to $\mathsf{Pk.Enc}$ and $\mathsf{Pk.Dec}$ as well; for concreteness, they can for instance be the description of an elliptic curve group with generator for an ECDLP-based system or the dimensions and noise sampling algorithm for an LWE-based system. When one is interested in re-phrasing our results in an asymptotic setting, the parameters pm will be generated by a probabilistic, polynomial-time algorithm that only takes the security parameter as input.

2.3 Concrete Security

Indistinguishability. The standard notion of security for encryption systems has become that of indistinguishability. Here the adversary is given access to a challenge encryption oracle implementing one of two "worlds"; the adversary needs to find out which. Several choices appear regarding the exact nature of these worlds, leading to different notions of indistinguishability such as real-or-random and left-or-right. Henceforth we refer to those two notions ROR and LOR, respectively, and we will refer to them collectively as IND. We will flesh out the details in Sect. 3.

Security definitions furthermore take into account the POWER given to the adversary, for example that of chosen plaintext attacks (CPA), or chosen ciphertext attacks (CCA). The distinguishing advantage of an adversary $\mathbb{A}$ against a scheme relative to some notion will then be IND-POWER$_{\mathrm{PKE}}(\mathbb{A})$, see Definition 1. As randomly guessing a world is correct half of the time, the distinguishing advantage is of course suitably offset.

Definition 1. *The* distinguishing advantage *of an adversary* $\mathbb{A}$ *against an encryption scheme* PKE *is*

$$\mathrm{IND\text{-}POWER}_{\mathrm{PKE}}(\mathbb{A}) := 2 \cdot \Pr\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ind\text{-}power}}(\mathbb{A}) = 1 \right] - 1 \,.$$

Implications and Separations. Our main focus will be comparing different notions of security, especially showing that if security is met under one notion, then it is also met under another one. We will prove these implication using fully black box reductions [4,31] that are furthermore simple [29]. A fully black box

reduction works for all schemes and adversaries, and only accesses them in a black box manner. Moreover, if the reduction only runs its adversary once and without rewinding, then the reduction is simple.

To allow for black-box access to the scheme, we will add an auxiliary oracle for the PKE to operate on the message space and the key space. A simple fully-black box (SFBB) reduction has access to this auxiliary oracle, as well as to the oracles corresponding to the PKE's algorithms, the oracles provided to the reduction by the game it is playing, and finally its single straight copy of the adversary. We will insist that the overhead of such a reduction, namely the number of oracle calls it makes more than the adversary it is running, is not undue: it can be upper bounded in terms of the parameters that define the security game(s) at hand, such as the number of keys in the system.

Definition 2 (Tightness). *Let* IND_1 *and* IND_2 *be two indistinguishability notions for PKE schemes, let c be a positive real number, then* $\mathrm{IND}_1 \overset{\leq c}{\Longrightarrow} \mathrm{IND}_2$ *iff there exists a simple fully-black box reduction* $\mathbb{B}_1$ *such that for all PKE schemes* PKE *and adversaries* $\mathbb{A}_2$,

$$\mathrm{IND}_2(\mathbb{A}_2) \leq c \cdot \mathrm{IND}_1(\mathbb{B}_1^{\mathbb{A}_2,\mathrm{PKE}})$$

and the overhead of $\mathbb{A}_2$ *is not undue.*

Refer also to Jager et al. [25] for a discussion on how to express tightness for more general reductions. They also formalize the folklore that simple reductions compose neatly; in our case if $\mathrm{IND}_1 \overset{\leq c}{\Longrightarrow} \mathrm{IND}_2$ and $\mathrm{IND}_2 \overset{\leq d}{\Longrightarrow} \mathrm{IND}_3$ then also $\mathrm{IND}_1 \overset{\leq c \cdot d}{\Longrightarrow} \mathrm{IND}_3$.

If $c = 1$, the reduction is called tight; if $c > 1$ we call the reduction lossy. Note that our notion of tightness is stricter than in some other works where a constant factor of say 2 will still be considered tight [18]; our convention has the benefit of not depending on any (security) parameter. A natural question for lossy reductions is whether the loss is inevitable or not—if it is, the bound is called sharp. Questions of sharpness are not the focus of our work, although we do remark upon it in more detail in Appendix B.

3   A General Definition of PKE Multi-user Security

3.1   A General Game

In order to compare various flavours of multi-user notions for PKE, we take Jager et al.'s framework for multi-user AE notions [25] and port it to the PKE setting, using some slightly different game-mechanics in the process. A multi-user security game is parametrized by the number of keys κ and the number of bits β. Usually one can imagine $\beta \leq \kappa$ and in fact Jager et al. only considered $\beta = \kappa$. However, keeping κ and β distinct helps when expressing and interpreting security losses.

$$\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ind\text{-}cca},\not\kappa,\beta}(\mathbb{A})$$

$(\mathsf{pk}_1, \mathsf{sk}_1), \ldots, (\mathsf{pk}_\kappa, \mathsf{sk}_\kappa) \leftarrow\!\!\$\ \mathsf{Pk.Kg}$

$b_1, \ldots, b_\beta, \delta \leftarrow\!\!\$\ \{0, 1\}$

$\mathsf{C}_1, \ldots, \mathsf{C}_\kappa, \mathsf{I}_1^{\mathcal{E}}, \ldots, \mathsf{I}_\beta^{\mathcal{E}}, \mathsf{I}^{\mathcal{K}} \leftarrow \emptyset$

$(j, \hat{b}_j) \leftarrow\!\!\$\ \mathbb{A}^{\mathcal{E},\mathcal{D},\mathcal{K}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$

if $\mathsf{I}_j^{\mathcal{E}} \cap \mathsf{I}^{\mathcal{K}} \neq \emptyset$ **then return** δ

else return $b_j = \hat{b}_j$

$$\mathcal{E}_{\mathrm{LOR}}(i, j, m_0, m_1)$$

if $m_0 \not\sim m_1$ **then return** $\not$

$\mathsf{I}_j^{\mathcal{E}} \xleftarrow{\cup} i$

$c^* \leftarrow\!\!\$\ \mathsf{Pk.Enc}_{\mathsf{pk}_i}(m_{b_j})$

$\mathsf{C}_i \xleftarrow{\cup} c^*$

return c^*

$$\mathcal{E}_{\mathrm{ROR}}(i, j, m)$$

$m' \leftarrow\!\!\$\ [\![m]\!]$

return $\mathcal{E}_{\mathrm{LOR}}(i, j, m, m')$

$$\mathcal{D}(i, c)$$

if $c \in \mathsf{C}_i$ **then return** $\not$

$m \leftarrow \mathsf{Pk.Dec}_{\mathsf{sk}_i}(c)$

return m

$$\mathcal{K}(i)$$

$\mathsf{I}^{\mathcal{K}} \xleftarrow{\cup} i$

return sk_i

Fig. 1. The generalised multi-user distinguishing experiment $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ind\text{-}cca},\not\kappa,\beta}(\mathbb{A})$; the adversary has access to either the left-or-right $\mathcal{E}_{\mathrm{LOR}}$ or the real-or-random $\mathcal{E}_{\mathrm{ROR}}$ challenge oracle.

Given a public key encryption scheme PKE, let $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ind\text{-}cca},\not\kappa,\beta}(\mathbb{A})$ be the experiment given in Fig. 1, where $\mathbb{A}$ is the adversary. The corresponding distinguishing advantage (see Definition 1) is denoted by $\mathrm{IND\text{-}CCA}_{\mathrm{PKE}}^{\not\kappa,\beta}(\mathbb{A})$. The κ is slashed to denote the presence of a key corruption oracle; the corresponding notion without corruptions is $\mathrm{IND\text{-}CCA}_{\mathrm{PKE}}^{\kappa,\beta}$. Without the decryption oracle the notion becomes a chosen-plaintext attack (CPA) instead. Often our results are oblivious of whether the power is CPA or CCA; we will then use CXA to refer to them collectively.

In the game, an adversary is given κ public keys, and a choice of β bits to try and attack through one of the two challenge oracles depending on the flavour of indistinguishability: for left-or-right indistinguishability, it gains access to $\mathcal{E}_{\mathrm{LOR}}$, whereas for real-or-random, it instead gains access to $\mathcal{E}_{\mathrm{ROR}}$. Both oracles have the usual interface, augmented by a key handle i and a bit handle j. For instance, for $\mathcal{E}_{\mathrm{LOR}}$ an adversary picks handles i and j as well as two equivalent messages m_0 and m_1 to receive the encryption of m_{b_j} under public key pk_i. For $\mathcal{E}_{\mathrm{ROR}}$ only a single message m is provided in addition to the two handles and, depending on the value of b_j, $\mathbb{A}$ receives the encryption of either the message or of a uniformly chosen equivalent message.

The adversary has possible access to two additional powers: a decryption oracle $\mathcal{D}$ and a corruption oracle $\mathcal{K}$. The former takes as input a ciphertext c together with a key handle i, and returns the decryption of c under private key sk_i. The latter takes as input a key handle i and directly returns said sk_i.

The adversary has in principle unlimited adaptive access to the available oracles, necessitating some admin in the game to deal with trivial wins. Firstly, if $m_0 \not\sim m_1$ for $\mathcal{E}_{\mathrm{LOR}}$, or if a challenge ciphertext is submitted to the decryption oracle under its handle of creation, then the adversary receives the special symbol $\not$ instead. Secondly, once the adversary outputs a bit handle j and a guess $\hat{b}_j$,

the game checks through $I_j^{\mathcal{E}} \cap I^{\mathcal{K}} = \emptyset$ whether the challenge bit has become compromised by virtue of being challenged together with a corrupted key. If so, the game outputs the uniformly random bit δ, yielding the adversary no advantage; otherwise, the game outputs whether $\hat{b}_j = b_j$.

Unlike Jager et al., we do not consider valid or invalid adversaries, but rather deal with bad behaviour in-game. Specifically, we want the adversary to be able to challenge on a key both before and after it becomes corrupted, but trying to win by attacking any of the corrupted challenge bits must of course be disallowed, regardless of the order of the queries. Thus, for problematic combinations of challenge/corrupt/target we necessarily had to wait until the adversary announced its target j before, if need be, penalizing. For bad decryption queries, penalizing at the end is discouraged [8], moreover it is easy to check on-the-fly.

Finally, we use $q_i^{\mathcal{E}}$ to refer to the number of challenge queries on public key pk_i; $q_{\Sigma}^{\mathcal{E}}$ for the total number of challenge oracle calls; and $q_{\max}^{\mathcal{E}}$ for the maximum number of challenge queries per key. Similarly, $q_i^{\mathcal{D}}$ is the number of decryption calls on private key sk_i and $q^{\mathcal{K}}$ the number of corruption calls.

## 3.2 Notational Conventions

Jager et al. [25] introduced their unified game in order to show that, for authenticated encryption, tightness losses are inevitable in a multi-key with corruption setting, irrespective of certain definitional choices. Thus they can avoid having to choose one notion over the other. We are interested in finding out, for public key encryption, whether some notion is preferred over the other. To that end, we will introduce some notation to more easily identify known notions and express relationships between them.

One can visualize the IND-CXA$^{\not{\kappa},\beta}$ experiment using a binary matrix of dimension $\kappa \times \beta$, where an entry be set wherever a key and a bit may be called together. For the general game, the matrix has every entry filled (see the leftmost matrix of Fig. 2). We will refer to this as the free-bit notion. By restricting the matrix, we can easily express existing notions.

Bellare et al.'s original single-challenge-bit notion [5] corresponds to a $\kappa \times \beta$-matrix (for arbitrary β) with only a single set row to force all challenge queries to the same bit handle (see the middle matrix of Fig. 2). If $\beta = 1$, the notion matches the free-bit notion, so we may write IND-CXA$^{\kappa,1}$, or IND-CXA$^{\not{\kappa},1}$ if corruptions are present, for the single-bit notion.

On the other hand, for the one-challenge-bit-per-key notion we have that $\beta = \kappa$ and the restriction $i = j$ for all challenge queries. These restrictions correspond to a square matrix in which only the diagonal is set (see the rightmost matrix of Fig. 2), inspiring us to refer to this notion as *diagonal-bit*, or just diagonal, and denote it by IND-CXA$^{\kappa,\boxtimes}$, or IND-CXA$^{\not{\kappa},\boxtimes}$ with corruptions.

The single-bit and diagonal-bit notions we will collectively refer to as the simple notions. Our notation and terminology differs from prior art, which is to some extent inevitable. The distinction between the various notions has only recently received explicit attention [24, 25] and no clear terminology has yet been set. For instance, we drop the prefix MU (for multi-user, to contrast with the

$$
\begin{array}{cc}
\begin{array}{cccccc}
 & b_1 & b_2 & b_3 & \dots & b_\beta \\
\mathsf{pk}_1 & \circ & \circ & \circ & \dots & \circ \\
\mathsf{pk}_2 & \circ & \circ & \circ & \dots & \circ \\
\mathsf{pk}_3 & \circ & \circ & \circ & \dots & \circ \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
\mathsf{pk}_\kappa & \circ & \circ & \circ & \dots & \circ
\end{array}
&
\begin{array}{cccccc}
 & b_1 & b_2 & b_3 & \dots & b_\beta \\
 & \circ & \times & \times & \dots & \times \\
 & \circ & \times & \times & \dots & \times \\
 & \circ & \times & \times & \dots & \times \\
 & \vdots & \vdots & \vdots & \ddots & \vdots \\
 & \circ & \times & \times & \dots & \times
\end{array}
&
\begin{array}{cccccc}
 & b_1 & b_2 & b_3 & \dots & b_\kappa \\
 & \circ & \times & \times & \dots & \times \\
 & \times & \circ & \times & \dots & \times \\
 & \times & \times & \circ & \dots & \times \\
 & \vdots & \vdots & \vdots & \ddots & \vdots \\
 & \times & \times & \times & \dots & \circ
\end{array}
\\
\mathrm{IND}^{\kappa,\beta} & \mathrm{IND}^{\kappa,1} & \mathrm{IND}^{\kappa,\boxslash}
\end{array}
$$

Fig. 2. Matrices of allowed key/bit combinations in challenge oracle calls for the free-bit, single-bit, and diagonal-bit multi-user notion, respectively; circles mark allowed queries, while crosses mark disallowed ones. The visualization highlights that the free-bit notion is a strict generalization of the two other, simple notions.

older single user notions) as on the one hand we believe that these days multi-user security should be the default from which single user notions can be derived if needed, and on the other hand we wish to maintain a clean GOAL–POWER nomenclature: having multiple users to target primarily modifies an adversary's power, not its goal.

3.3   Interpretation

Both simple notions with corruptions have appeared in the literature, both in a PKE setting but also in related KEM, AKE, and to a lesser extent AE settings. One key question is which notion to opt for when. Establishing relationships between the notions, as in the next section, helps answer this question. Here, we want to address the meaning and usefulness of the notions as they are.

In the context of AKE, Jager et al. [24] discuss the difference between the single-bit notion ("single-bit guess") and the diagonal notion ("multi-bit guess"). Earlier works on tight security for AKE focused on the diagonal setting [2], yet as Cohn-Gorden et al. [13, Section 3] point out, that notion does not lend itself very well for tight composition: when the keys produced by an AKE are subsequently used, in a proof it is convenient to swap out all keys from real to random in one fell swoop. The single-bit notion allows such a massive substitution, but the diagonal notion does not. Moreover, Jager et al. wonder whether the diagonal notion is meaningful, which would "provide a good intuition of what [it] tries to model".

Whereas AKE and KEMs are primarily tools to set up symmetric keys for subsequent use, the situation for PKE is different as it is much closer to the end user. The difference is reflected in the kind of indistinguishability as well: for AKE and KEMs, a ROR-style notion is used where the adversary cannot even control the real world's "message", yet for PKE's LOR-notion, an adversary has full control over the left-versus-right challenge messages. Thus, for PKE the diagonal-LOR notion does seem meaningful, as we explain below.

Suppose we interpret each key to correspond to a *user* and each challenge bit to correspond to a *conversation*. Then the different notions model different scenarios. For instance, the diagonal notion models a scenario where the users

take part in independent conversations, and an adversary can decide which honest conversation to target after corrupting a number of other ones. In contrast, the single-bit notion models a scenario where all users are engaged in the *same* conversation. The latter scenario allows an adversary to accumulate information on the conversation across users, although none of the active parties may be corrupted.

Finally, the free-bit notion models a situation where there are a number of independent conversations, each with their own potentially overlapping set of users. The adversary can adaptively corrupt a number of users, and finally targets a conversation conducted by honest users only.

Of course, there are already existing notions that study PKE security in the presence of corruptions, under the term "selective opening attacks" (SOA, [9, 15]). There are various formalizations of SOA, the most relevant ones to our work are receiver SOA [19] where an adversary can corrupt private keys (as opposed to sender SOA, where an adversary learns how a ciphertext was created). Most of these SOA notions are considerably stronger than the notions we consider: our strongest notion is still implied by the customary single-user single-challenge LOR–CCA (just rather lossy), yet for SOA strong separations, and in some cases impossibility results, are known [23]. The link between multi-user security with corruptions on the one hand and SOA on the other has largely been ignored and appears worth expanding further.

We remark that the multi-bit notion also occurs naturally when studying multi-instance security [10], which has been studied in the context of PKE [1]. We leave the adaptation of our work, and specifically the general free-bit game to that setting as an enticing open problem.

4   Relations Between Indistinguishability Notions

In this section we investigate how tightly the various multi-user notions relate to each other, and how each relates to single-user notions. Some implications are known or folklore and others follow quite naturally from the literature, but not all. As expected, most of the notions are equivalent within a factor linear in the number of users. Yet, some notions turn out to be more, or less, tightly related.

There is for instance the surprising and completely tight reduction from LOR-CXA$_{\mathrm{PKE}}^{\kappa,1}$ to LOR-CXA$_{\mathrm{PKE}}^{\kappa,\boxtimes}$ (Theorem 1). However, the proof technique breaks down for real-or-random indistinguishability and in notions with corruptions. Furthermore, for the latter, there doesn't seem to be a way of relating the notions more tightly than by a linear loss. We conjecture this linear loss to be sharp, yet proving so we leave open.

Shorthand for Unified Implications. Given the large number of notions resulting from the various orthogonal definitional choices, we use shorthand, as presented in Table 1, to state various theorems. The shorthand serves as an implicit quantifier, so a theorem statement in shorthand essentially holds for all notions included in the shorthand. To avoid clutter, we will sometimes abbreviate IND-CXA^{u,c} to

Table 1. A modular framework for multi-user security notions.

Shorthand	Stand-in for	Relates to
IND	$\{\mathrm{LOR}, \mathrm{ROR}\}$	Type of challenge oracle
CXA	$\{\mathrm{CPA}, \mathrm{CCA}\}$	Presence of decryption oracle
u	$\{\kappa, \not\kappa\}$	Number of keys; presence of corruption oracle
c	$\{1, \boxbslash, \beta\}$	Number of challenge bits; relation with keys

just $\mathrm{IND}^{u,c}$, and let it be implied that the result holds for both CPA and CCA. We will refer to single-user, multi-challenge notions by dropping the superscripts, e.g. IND.

As a concrete example, consider the trivial statement

$$\mathrm{IND}^{u,c} \implies \mathrm{IND}.$$

This is then to be read as, "Both in the cpa and the cca setting, and regardless of the nature of the challenge oracle, the presence or absense of corruptions, or the number and structure of the challenge bits, security under a multi-user notion implies security under the corresponding single-user notion." Written out in full, the statement becomes:

Lemma 1. *For all* $\mathrm{IND} \in \{\mathrm{LOR}, \mathrm{ROR}\}$, $\mathrm{CXA} \in \{\mathrm{CPA}, \mathrm{CCA}\}$, $u \in \{\kappa, \not\kappa\}$, *and* $c \in \{1, \boxbslash, \beta\}$, *there is a reduction* $\mathbb{B}$ *such that, for every adversary* $\mathbb{A}$,

$$\mathrm{IND\text{-}CXA}_{\mathrm{PKE}}(\mathbb{A}) \le \mathrm{IND\text{-}CXA}_{\mathrm{PKE}}^{u,c}(\mathbb{B}),$$

where $\mathbb{B}$ *calls* $\mathbb{A}$ *once, with no undue overhead.*

Tight Implications From Strict Generalizations. Security under a multi-user notion tightly implies single-user security under the corresponding notion, and adding helper oracles (like decryption for CCA, or a corruption oracle) yields strictly more general notions; as does increasing the parameters (number of users/number of challenge bits), and for all notions, left-or-right security implies real-or-random security, as can be seen from Fig. 1. For completeness, we summarize these trivial implications in the full version.

4.1 Simple Multi-user Notions Versus Classical Single-Key Notions

Bellare et al. [5] used a hybrid argument to show that single-user single-challenge security implies $\mathrm{LOR}^{\kappa,1}$ with a security loss linear in the total number of challenge encryption queries. They phrased this total as the product of the number of users and the number of challenges per user. As all our notions are explicitly multi-challenge, we will ignore the number of challenge queries, meaning the loss simply becomes linear in the number of users (in line with the original claim).

Fig. 3. Known relations between single-user (but multi-challenge) indistinguishability and the two different generalizations to multi-user indistinguishability, with and without corruptions; refer to Table 1 for an overview of the shorthand. Recall that IND without any superscripts means single-user notions. (Double arrows: trivially tight.)
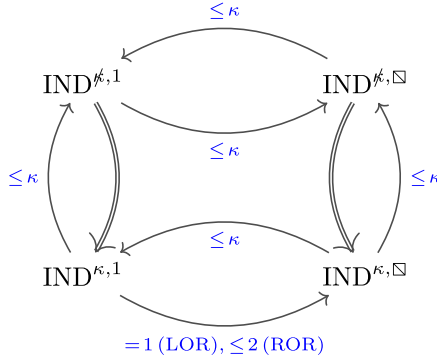
Bellare et al. did not consider the diagonal notion or corruption, however later, when Jager et al. [25] introduced the free-bit notion to the setting of AE, they also showed that the simple notions are implied by the single-user notion, again with a linear loss, even when corruptions are considered. For completeness, we reprove the relevant linear losses in our new PKE context in Appendix C. The resulting relations are summarized in Fig. 3.

As explained in Sect. 3.1, Jager et al. used slightly different game mechanics by prohibiting certain adversarial behaviour. In contrast, we allow such bad behaviour and just ignore the adversary's output instead. We introduce a useful lemma (Lemma 3) that formalizes that, in the single-key setting, our mechanism is sound and corrupting that single-key yields no adversarial advantage. This single-key-with-corruptions game is often easier to use in reductions.

Existing sharpness results can be used to show that linear losses are inevitable, see Appendix B.2 for details.

4.2   Relationship Between Simple Multi-user Notions

Now that we have affirmed that the single-user notion implies any of the four simple multi-user notions with a loss linear in the number of users, a natural question is how the simple multi-user notions relate to each other. As the multi-user notions all tightly imply the single-user notion, one can always just go via the single-user notion. As already noted by Jager et al. [25], this strategy will again lead to a loss linear in the number of users. Lemma 2 formalizes this trivial loss and Fig. 4 provides an overview of the relations. One notable exception from the linear losses is the implication from the single-bit notion to the diagonal notion if there are no corruptions, which is tight for the case of left-or-right indistinguishability and almost tight for real-or-random indistinguishability. We will explain why this is in the next paragraph.

$$\text{IND}^{\not{k},1} \quad\quad\quad \text{IND}^{\not{k},\boxed{N}}$$

$$\leq \kappa$$

$$\text{IND}^{\kappa,1} \quad\quad\quad \text{IND}^{\kappa,\boxed{N}}$$

$$= 1\,(\text{LOR}),\, \leq 2\,(\text{ROR})$$

Fig. 4. Relations between the simple multi-user notions, including the non-trivially tight relation between $\text{LOR}^{\kappa,\beta}$ and $\text{LOR}^{\kappa,\boxed{N}}$ as captured by Corollaries 1 and 2 for LOR and ROR, respectively. (Double arrows: trivially tight.)

Lemma 2 ($\text{IND}^{u,c} \stackrel{\leq \kappa}{\Longrightarrow} \text{IND}^{u,c'}$). *Let $c' \in \{1, \boxed{N}\}$. Then, there is an SFBB reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,*

$$\text{IND-CXA}_{\text{PKE}}^{u',c'}(\mathbb{A}) \leq \kappa \cdot \text{IND-CXA}_{\text{PKE}}^{u,c}(\mathbb{B})\,.$$

Proof (sketch). Trivially, $\text{IND}^{\kappa,c} \implies \text{IND}$. Meanwhile, Theorems 3 and 4 together say that $\text{IND} \stackrel{\leq \kappa}{\Longrightarrow} \text{IND}^{\not{k},c}$. Combining (in the manner discussed in Sect. 2) gives $\text{IND}^{\kappa,c} \implies \text{IND} \stackrel{\leq \kappa}{\Longrightarrow} \text{IND}^{\not{k},c'}$.

A Tight Relation: From Single-Bit to Multi-bit Without Corruptions. Surprisingly, left-or-right indistinguishability allows for a 'bit-hiding' argument that lets an adversary playing a single-bit multi-user game simulate the full free-bit game (and therefore also the diagonal-bit game), by simply exchanging the order in which it forwards its two messages. We formalize this argument in Theorem 1 and its proof. Consequently, $\text{LOR}^{\kappa,1}$ tightly implies of $\text{LOR}^{\kappa,\boxed{N}}$ (Corollary 1), whereas the implication in the other direction appears lossy, this clearly renders $\text{LOR}^{\kappa,1}$ the preferred notion.

Theorem 1 ($\text{LOR}^{\kappa,1} \implies \text{LOR}^{\kappa,\beta}$). *There is an SFBB reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,*

$$\text{LOR-CXA}_{\text{PKE}}^{\kappa,\beta}(\mathbb{A}) \leq \text{LOR-CXA}_{\text{PKE}}^{\kappa,1}(\mathbb{B})\,,$$

where $\mathbb{B}$'s overhead is limited to drawing β uniformly random bits.

Proof. The reduction $\mathbb{B}$, playing $\mathsf{Exp}_{\text{PKE}}^{\text{lor-cxa},\kappa,1}$, simulates $\mathsf{Exp}_{\text{PKE}}^{\text{lor-cxa},\kappa,\beta}$ for $\mathbb{A}$ by drawing β fresh challenge bits, and simply exchanging the order of m_0 and m_1

$\mathbb{B}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$	if $\mathbb{A}$ calls $\mathcal{E}(i, j, m_0, m_1)$	if $\mathbb{A}$ calls $\mathcal{D}(i, c)$
$b_1, \ldots, b_\beta \leftarrow\!\!\$\ \{0, 1\}$	$c^* \leftarrow \mathcal{E}(i, m_{b_j}, m_{\bar{b}_j})$	$m \leftarrow \mathcal{D}(i, c)$
$(j, \hat{b}_j) \leftarrow\!\!\$\ \mathbb{A}^{\mathcal{E}, \mathcal{D}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$	**return** c^*	**return** m
return $\hat{b}_j \oplus b_j$		

Fig. 5. The adversary $\mathbb{B}$, playing $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cxa}, \kappa, 1}$ while simulating $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cxa}, \kappa, \beta}$ for $\mathbb{A}$.

in accordance to the value of the simulated challenge bit when forwarding to its own left-or-right oracle (see Fig. 5). Denoting the challenge bit of $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cxa}, \kappa, 1}(\mathbb{B})$ by b, the ciphertext that $\mathbb{A}$ receives upon the query $\mathcal{E}(i, j, m_0, m_1)$ will be an encryption of the message $m_{b \oplus b_j}$ under pk_i; $\mathbb{B}$ then simply makes sure to undo this xor again before returning its final guess. □

Corollary 1 ($\mathrm{LOR}^{\kappa, 1} \implies \mathrm{LOR}^{\kappa, \boxtimes}$)**.** *There is an SFBB reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,*

$$\mathrm{LOR\text{-}CXA}_{\mathrm{PKE}}^{\kappa, \boxtimes}(\mathbb{A}) \leq \mathrm{LOR\text{-}CXA}_{\mathrm{PKE}}^{\kappa, 1}(\mathbb{B}) \,.$$

In the presence of a corruption oracle, the reduction breaks down as it is no longer able to simulate properly: it cannot both challenge on and corrupt the same key (a behaviour that is allowed in the diagonal and free-bit games). We will return to the free-bit game in the presence of corruptions below, but first we turn our attention to that other indistinguishability notion, real-or-random.

Extending the Argument to Real-or-Random. The proof of Theorem 1 makes use of the fact that the LOR challenge oracle allows both a left and a right message to be input, enabling us to hide the bit in the ordering of the two messages. For ROR, the challenge oracle only accepts a single message, so hiding the bit as above is no longer possible.

However, when Bellare et al. [6] introduced the distinction between LOR versus ROR indistinguishability in the context of single-user probabilistic symmetric encryption, they also showed a factor-2 loss from ROR to LOR. As we will show in Theorem 5 (to be presented shortly), their proof technique is readily adapted to a relation between single-bit multi-user PKE notions. Theorems 1 and 5 can then be combined into the corollary below (which itself implies the equivalent of Corollary 1 for ROR, again with a factor 2 loss).

Corollary 2 ($\mathrm{ROR}^{\kappa, 1} \overset{\leq 2}{\implies} \mathrm{ROR}^{\kappa, \beta}$)**.** *There is an SFBB reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,*

$$\mathrm{ROR\text{-}CXA}_{\mathrm{PKE}}^{\kappa, \beta}(\mathbb{A}) \leq 2 \cdot \mathrm{ROR\text{-}CXA}_{\mathrm{PKE}}^{\kappa, 1}(\mathbb{B}) \,.$$

Proof (Sketch). Theorem 5 states that $\mathrm{ROR}^{\kappa,1} \overset{\leq 2}{\Longrightarrow} \mathrm{LOR}^{\kappa,1}$, while trivially $\mathrm{LOR}^{\kappa,\beta} \implies \mathrm{ROR}^{\kappa,\beta}$. Then, using Theorem 1, we get $\mathrm{ROR}^{\kappa,1} \overset{\leq 2}{\Longrightarrow} \mathrm{LOR}^{\kappa,1} \implies \mathrm{LOR}^{\kappa,\beta} \implies \mathrm{ROR}^{\kappa,\beta}$.

4.3   The Free-Bit Game with Corruptions

In the free-bit game, the adversary can both challenge on and corrupt keys, provided the final targeted bit remains uncorrupted. In the single-bit game, however, challenging on and corrupting a key are mutually exclusive, causing the bit-hiding argument, that tightly related $\mathrm{LOR}^{\kappa,1}$ to $\mathrm{LOR}^{\kappa,\beta}$, to break down in the presence of corruptions. Seemingly, the best we can do is a standard bit-guessing argument, suffering a β loss, as formalized in Theorem 2 below (see Appendix E for a full proof).

Theorem 2 ($\mathrm{IND}^{\not{k},1} \overset{\leq \beta}{\Longrightarrow} \mathrm{IND}^{\not{k},\beta}$)**.** *There is an SFBB reduction* $\mathbb{B}$ *such that, for any adversary* $\mathbb{A}$,

$$\mathrm{IND\text{-}CXA}^{\not{k},\beta}_{\mathrm{PKE}}(\mathbb{A}) \leq \beta \cdot \mathrm{IND\text{-}CXA}^{\not{k},1}_{\mathrm{PKE}}(\mathbb{B}),$$

where $\mathbb{B}$*'s overhead consists of drawing* β *uniformly random bits.*

Combining with $\mathrm{IND} \overset{\leq \kappa}{\Longrightarrow} \mathrm{IND}^{\not{k},1}$ (Theorem 3) yields an upper bound on the free-bit advantage as it relates to single-user advantage, see Corollary 3. Notably, when Jager et al. [25] introduced the free-bit notion (for AE), they observed that proving a linear loss was beyond them, yet they did not provide an alternative, looser bound instead. We therefore plug this gap in the literature. Figure 6 provides an overview of how the single-user and simple multi-user notions relate to the free-bit notions.

Corollary 3 ($\mathrm{IND} \overset{\leq \kappa \beta}{\Longrightarrow} \mathrm{IND}^{\not{k},\beta}$)**.** *There is a reduction* $\mathbb{B}$ *such that, for any adversary* $\mathbb{A}$,

$$\mathrm{IND\text{-}CXA}^{\not{k},\beta}_{\mathrm{PKE}}(\mathbb{A}) \leq \kappa \cdot \beta \cdot \mathrm{IND\text{-}CXA}_{\mathrm{PKE}}(\mathbb{B}).$$

$\mathbb{B}$ *calls* $\mathbb{A}$ *once, and additionally uses the resources needed to draw* κ *fresh key-pairs and* β *uniformly random bits.*

Interestingly, Corollary 3 tightly implies Theorem 3, but not Theorem 4: setting $\kappa = \beta$ in Corollary 3 yields a κ^2 loss. This gives some hope that a tighter relation than that of Corollary 3 might still be possible, one that would imply both Theorems 3 and 4. We leave this an open problem, although present some initial thoughts in Appendix B.

$$\text{IND}^{\kappa,\beta} \qquad\qquad\qquad \text{IND}^{\not\kappa,\beta}$$

$$\begin{array}{ll} =1\,(\text{LOR}) \\ \leq 2\,(\text{ROR}) \end{array} \qquad \begin{array}{ll} \leq\kappa\,(\text{LOR}) \\ \leq 2\kappa\,(\text{ROR}) \end{array} \qquad \leq\beta \qquad\qquad \leq\beta$$

$$\text{IND}^{\kappa,1} \qquad\qquad \text{IND}^{\kappa,\boxed{}} \qquad \text{IND}^{\not\kappa,1} \qquad \leq\kappa\beta \qquad \text{IND}^{\not\kappa,\boxed{}}$$

$$\begin{array}{ll} \leq\kappa\,(\text{LOR}) \\ \leq 2\kappa\,(\text{ROR}) \end{array}$$

$$\text{IND} \qquad\qquad\qquad\qquad \text{IND}$$

Fig. 6. Relations between different multi-user notions, without corruptions (left), and with corruptions (right).

4.4   LOR Versus ROR, or When the Challenge Oracle Matters

Until now, we have for the most part treated the two flavours of indistinguishability as one. However, as we saw for Theorem 1, the choice of challenge oracle can sometimes make a difference. Of course, left-or-right indistinguishability always implies real-or-random indistinguishability. Furthermore, for single-user notions, it has been long been known that ROR implies LOR with only a factor 2 tightness loss [5]. However, for multi-instance security [10], the loss is known to blow up exponentially. Thus, it is a priori unclear what losses one should expect for the multi-user setting, both between corresponding LOR and ROR notions, but also between the ROR notions themselves.

Jager et al. [26, Theorem 21] showed a general result that a loss L in the single user setting can be turned into a loss $L\kappa$ for the simple notions (for AE); the free-bit case is not addressed. We complement their results for the PKE setting, as summarized in Fig. 7 and formalized in Appendix D.

Some relations are worth highlighting. First, note that the same factor 2 reduction still lends itself to the single-bit multi-key setting (with or without corruptions). The argument is very similar to that of the single-user case: either the bit is "real", in which case the simulated game is equivalent to the left-or-right one, or the bit is "random", in which case the simulated challenge bit is information-theoretically hidden from the adversary; the main complication in going to a multi-key setting with corruptions being dealing with disallowed guesses. See Theorem 5. This contrasts to the diagonal-bit setting, in which the tightest known reduction loses a factor 2κ, as achieved via the single-user relation: $\text{ROR}^{u,\boxed{}} \implies \text{ROR} \overset{\leq 2}{\Longrightarrow} \text{LOR} \overset{\leq\kappa}{\Longrightarrow} \text{LOR}^{u,\boxed{}}$.

Second, note that the fact that $\text{LOR}^{\kappa,1} \implies \text{LOR}^{\kappa,\beta}$ (Theorem 1) allows us to conclude that the factor 2 reduction still holds for the free-bit notion absent corruptions: $\text{ROR}^{\kappa,\beta} \implies \text{ROR}^{\kappa,1} \overset{\leq 2}{\Longrightarrow} \text{LOR}^{\kappa,1} \implies \text{LOR}^{\kappa,\beta}$. Compare with the situation in the presence of corruptions, where the corresponding implications yield $\text{ROR}^{\not\kappa,\beta} \overset{\leq 2\beta}{\Longrightarrow} \text{LOR}^{\not\kappa,\beta}$.

Fig. 7. Relations between lor and ror for the different multi-user notions, without corruptions (left) and with corruptions (right). The placement of notions roughly translate to their relative strength, with stronger notions placed higher, (see Fig. 6 for the implications missing from the figure.) As before, double arrows indicate trivially tight.

As before, we leave the question of whether there exist tighter reductions, or these losses really are inevitable, as open questions. Nevertheless, these additional losses serve to reinforce the folklore that left-or-right notions should be preferred over real-or-random whenever possible.

5   Conclusion

In this article, we surveyed several possible notions of multi-user security, showing how they relate to each other, and identifying a unified and general free-bit notion. We also conclusively answered the question of which canonical multi-user notion is the preferred one in the absence of corruptions, namely the single-bit left-or-right notion, as it is as strong or stronger than any of the others. In the presence of corruptions, the situation is less clear, particularly as it is not currently known whether the ability to both challenge and corrupt a key yields the adversary any additional power. What *is* known, however, is that the ability to challenge the same bit on several keys *does* give the adversary extra power. Until these questions have been definitively settled, we therefore suggest aiming for security under a free-bit notion whenever multi-user security with adaptive corruptions is to be considered.

A   A Brief History of Indistinguishability

The traditional 'IND-CPA' security notion for public key encryption (PKE) is an indistinguishability notion (IND) under adaptively chosen plaintext attacks

(CPA). Here an adversary receives a challenge ciphertext either for a plaintext of its choosing or an alternative challenge ciphertext, and needs to decide which one was received. The alternate challenge ciphertext can be generated in different ways, leading to subtly different notions [6]. The two common choices are: left-or-right (LOR), in which the adversary supplies two messages and receives the encryption of one of them; and real-or-random (ROR), in which the adversary supplies a message and either receives its encryption or the encryption of a random bit string. When Bellare et al. [7] considered various PKE security notions they showed that LOR-security tightly implies ROR-security, whereas the other direction incurs a modest security loss of a factor 2.

Stronger, more realistic notions are indistinguishability under adaptively chosen ciphertext attacks (CCA), or IND-CCA (historically also called IND-CCA2 to distinguish it from its non-adaptive counterpart IND-CCA1 [7]). Here, in addition to choosing the plaintexts to be challenged on, the adversary is given access to a decryption oracle, which it can query on any valid ciphertext receiving the corresponding plaintext, or the ciphertext-reject symbol $\perp$. To avoid trivial wins, some care needs to be taken when the challenge ciphertext is submitted to the encryption oracle; there are several mechanisms to deal with this subtlety [8]. Ignoring the decryption oracle gives back IND-CPA, making IND-CCA the stronger notion. Moreover, several real-world attacks not covered by IND-CPA (such as Bleichenbacher's attack [11]) are captured by IND-CCA, making the latter the preferred notion to aim for.

We are concerned with the multi-user setting, leading to further definitional choices. Although it might appear that these choices are largely irrelevant in an asymptotic context, as they are all polynomially equivalent, a concrete security treatment can surface non-trivial differences. These differences are often amplified with the introduction of multiple users, particularly when considering adaptive corruptions (see below).

First of all, while the notions above initially only allowed for a single challenge query, when Bellare et al. [5] investigated multi-user security, they simultaneously generalized the single-user notions by giving each user multiple challenges. Moreover, they showed that security under single-challenge implies security under multi-challenge with an (inevitable) loss linear in the number of challenges (cf. [6]).

In the present work, we consider all notions, including the single-user notions, to be multi-challenge. To adapt our results to a single-challenge setting, simply note that our single-user notions imply the corresponding single-challenge notion with a tightness loss $q^{\mathcal{E}}$, and insert the factor as needed. For instance, writing SC-IND for single-challenge indistinguishability, the analogue to Corollary 3 becomes SC-IND $\overset{\leqslant q^{\mathcal{E}}\kappa\beta}{=\!=\!=\!=\!\Longrightarrow}$ IND$^{\not{k},\beta}$.

Another choice is how to 'multiplex' the challenge oracles: should each user be independent of the others, or should they depend on each other? When multi-user security was introduced [5], the game only had a single challenge bit shared across all users for an adversary to guess. This choice intuitively leads to a stronger notion than if each user was given its own challenge bit as, with a

single shared bit, an adversary can 'gather evidence' for the true value of this challenge bit across the users (we provide evidence to this intuition in Corollary 1). Yet, the notion feels awkward when introducing corruptions, given that both corrupting and challenging on the same key would immediately yield a trivial win. One option is to disallow corrupting 'challenge' keys, and vice versa, leading to the single-bit notion $\text{IND}^{\kappa,1}$ [3,24,28]. Another option is to introduce user-specific bits. This was the approach employed by Bader et al. [2] in their study of authenticated key exchange: they considered a multi-user KEM notion with corruptions where each user was associated with its own challenge bit and the adversary had to declare at the end which uncorrupted bit it was guessing. Thus, even if a user was both challenged on and corrupted, a non-trivial win would still be possible. In the present work, we refer to this notion as *diagonal-bit* $(\text{IND}^{\not\kappa,\boxminus})$, as explained in Sect. 3.

Recently, Jager et al. [24] pointed out that this notion is problematic in the AKE setting, as unlike in the single-bit setting, a KEM secure under the diagonal-bit notion is not known to tightly compose to an AKE. They went on to construct a KEM tightly secure under the single-bit notion instead, which was therefore guaranteed to compose tightly.

Apart from the multi-user setting, the diagonal notion has seen use in the multi-*instance* setting [1,10], in which the adversary is asked to make a guess on *every* bit; in such settings, single-bit notions make little sense.

When Jager et al. [25] investigated the inevitability of multi-user tightness losses in the setting of authenticated encryption, they wanted their result to capture both of the the single-bit and the diagonal-bit notions, without having to provide separate proofs for the distinct cases. They therefore introduced a generalized notion, in which an adversary was free to choose the exact relations between the keys and challenge bits. This notion, which avoids the awkwardness of not being able to both challenge and corrupt the same key without sacrificing the ability to "gather evidence" on a bit over several keys, sits at the centre of much of the present work, and we refer to it as the *free-bit* notion $(\text{IND}^{\not\kappa,\beta})$.

B   Sharpness or When Tightness Losses are Inevitable

B.1   Sharpness and Inevitably Lossy Reductions

A natural question for lossy reductions is whether the loss is inevitable or not. To determine inevitability, we only need to 'invert' Definition 2, as below in Definition 3.

Definition 3 (Lossy). *Let* IND_1 *and* IND_2 *be two indistinguishability notions for PKE schemes, and let c be a positive real number, then* $\text{IND}_1 \overset{\geq c}{\Longrightarrow} \text{IND}_2$ *iff for all simple fully-black box reductions* $\mathbb{B}_1$ *there exist PKE schemes* PKE *and adversary* $\mathbb{A}_2$,

$$\text{IND}_2(\mathbb{A}_2) \geq c \cdot \text{IND}_1(\mathbb{B}_1^{\mathbb{A}_2,\text{PKE}}).$$

If both $\mathrm{IND}_1 \overset{\leq c}{\Longrightarrow} \mathrm{IND}_2$ and $\mathrm{IND}_1 \overset{\geq c}{\Longrightarrow} \mathrm{IND}_2$, then the reduction (for the first term) is sharp and we may write $\mathrm{IND}_1 \overset{\equiv c}{\Longrightarrow} \mathrm{IND}_2$.

B.2   Sharpness of Single-to-Simple Reductions

Below we discuss some relevant methods and results regarding the inevitability of lossy reductions in the context of multi-user PKE, showing that linear losses (in the number of users) is often sharp. Such results are often called impossibility results, yet to contrast with impossibility results that show that no constructions can achieve a notion (irrespective of the lossiness of the reduction), we prefer the term sharpness result when the impossibility is restricted to tightness only. The two main techniques are counterexamples and meta-reductions.

Counterexamples. As already pointed out by Bellare et al. [5], a simple counterexample shows that the bounds are generally sharp. They modified a PKE scheme that was identical to a 'secure' one except that with a small probability its encryption would be trivial and essentially just output the plaintext as the ciphertext (with some additional modifications to ensure correctness and that this event is easily recognizable publicly). Thus, when the challenge encryption oracle hits the trivial encryption, an adversary can trivially win its game; moreover the probability of this event happening at some point during the game is roughly linear with the number of challenge encryption queries.

However, given that we consider all our notions to be multi-challenge, we prefer a counterexample whose security degrades linearly in the number of available *keys*, not challenges. One might therefore instead consider a scheme for which a small-but-nonempty subset of the public keyspace returns messages in the clear. This "weak key" counterexample already works without corruptions for both of the simple multi-user notions, which implies sharpness for the more general notions.

Note that a similar critique of Bellare et al.'s original counterexample and (more refined) link with weak keys was made by Luykx et al. [30].

Meta-reductions. Another line of work has aimed to show sharpness through meta-reduction, thus ruling out tight reductions for a larger class of PKE schemes. The gain in generality is however traded for restrictions on the type of reductions that are ruled out, typically referred to as "simple" reductions (e.g., blackbox, no rewinding, etc.).

Bader et al. [3] showed that, for a large class of PKE systems, any simple reduction from a multi-user notion with corruptions to an underlying non-interactive hardness assumption must be lossy, with the loss linear in the number of keys. Meanwhile, Jager et al. [25] showed a similar result in the setting of authenticated encryption when reducing to single-user notions. In both cases, though, the proof technique crucially relied on the ability to corrupt keys, meaning that sharpness for the corruptionless notions aren't covered by their results.

Meta-reductions also don't rule out tight reductions for schemes outside the class considered; in fact, part of the usefulness of these results is the ability to look for tightly secure constructions outside these classes. This is exactly what Bader et al. [2] did when they constructed a tightly secure authenticated key exchange by deliberately breaking the requirement of public–private key uniqueness.

B.3   Tightening the Single-to-Free Implication?

Corollary 3, IND $\overset{\leq \kappa\beta}{\Longrightarrow}$ IND$^{\not\kappa,\beta}$, tightly implies Theorem 3, but not Theorem 4: setting $\kappa = \beta$ in Corollary 3 yields a κ^2 loss. This gives some hope that we might be able to show a tighter relation than that of Corollary 3, as in order to imply both Theorem 3 and 4, the statement would have to look something like the following.

Conjecture 1 (IND $\overset{\leq \mathtt{I}^{\mathcal{E}}_{\max}\beta}{\Longrightarrow}$ IND$^{\not\kappa,\beta}$). Let $\mathtt{I}^{\mathcal{E}}_{\max}$ be the maximum number of keys called together with any one challenge bit, (i.e., for any run of the game, we now require that $\forall j, |\mathtt{I}_j| \leq \mathtt{I}^{\mathcal{E}}_{\max}$; see Fig. 1). Then, there is a reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,

$$\text{IND-CXA}^{\not\kappa,\beta}_{\text{PKE}}(\mathbb{A}) \leq \mathtt{I}^{\mathcal{E}}_{\max} \cdot \beta \cdot \text{IND-CXA}_{\text{PKE}}(\mathbb{B}),$$

where $\mathbb{B}$ calls $\mathbb{A}$ once, and the overhead of $\mathbb{B}$ is small.

Then, for IND$^{\not\kappa,1}$, we would set $\mathtt{I}^{\mathcal{E}}_{\max} = \kappa$ and $\beta = 1$, while for IND$^{\not\kappa,\boxminus}$, $\mathtt{I}^{\mathcal{E}}_{\max} = 1$ and $\beta = \kappa$. Thus, both theorems are recovered.

To prove the statement, a natural strategy would be to combine the proof techniques of each of the theorems it is generalising, i.e. by first guessing a challenge bit, and then doing a hybrid argument over the keys relating to that bit. However, given that the free-bit game allows the adversary to choose the relations between keys and bits adaptively, this hybrid argument does not work without incurring losses larger than that of Corollary 3. We nevertheless present Conjecture 1 as an interesting open problem.

C   Formalization of Single to Simple Implications

A Single-User Notion with Corruptions. First, let us establish the trivial yet useful Lemma 3. Let $\mathsf{Exp}^{\text{lor-cxa},1,1}_{\text{PKE}}(\mathbb{A})$ be exactly as the single-key game, except that the player now has the option to corrupt the key. In other words, the game will be equivalent to that of Fig. 1, with $\kappa = \beta = 1$ (and with or without decryption oracle). Given that in this game, an adversary that both challenges and corrupts will trigger the game to output the uniformly random value δ, the presence of a corruption oracle should yield it no extra power. We formalize this intuition next.

Lemma 3 (IND $\implies$ IND^{1,1}). *There is an SFBB reduction $\mathbb{B}$ with no additional overhead such that, for every adversary $\mathbb{A}$,*

$$\text{IND-CXA}^{1,1}_{\text{PKE}}(\mathbb{A}) \leq \text{IND-CXA}_{\text{PKE}}(\mathbb{B}).$$

Proof. The following argument works the same whether IND = LOR or ROR, and whether CXA = CCA or CPA. The reduction $\mathbb{B}$, playing the regular single-key game, simulates the game with corruptions to $\mathbb{A}$ by forwarding every oracle call and mimicking $\mathbb{A}$'s output, unless at some point $\mathbb{A}$ asks to corrupt: in that case $\mathbb{B}$ aborts $\mathbb{A}$ and simply returns 0. This works because if $\mathbb{A}$ corrupts, either $\mathbb{A}$ also challenges, in which case the advantage will be forced to 0, or it corrupts the key and outputs a guess without challenging, in which case the challenge bit will be information-theoretically hidden from it, so that its advantage is 0 by necessity. Thus, in the event that $\mathbb{A}$ corrupts at all, its win advantage will be exactly 0; the same that $\mathbb{B}$ gets from simply aborting $\mathbb{A}$ and outputting 0. We provide a formal derivation below.

$$
\begin{aligned}
\Pr\left[\mathsf{Exp}^{\text{ind-cxa}}_{\text{PKE}}(\mathbb{B}) = 1\right] &= \Pr[\mathbb{A} \text{ did not corrupt} \wedge \mathbb{A} \text{ wins}] \\
&\quad + \Pr[\mathbb{A} \text{ did corrupt} \wedge b = 0] \\
&= \Pr[\mathbb{A} \text{ did not corrupt} \wedge \mathbb{A} \text{ wins}] \\
&\quad + \Pr[\mathbb{A} \text{ did corrupt}] \cdot \Pr[b = 0 \mid \mathbb{A} \text{ did corrupt}] \\
&= \Pr[\mathbb{A} \text{ did not corrupt} \wedge \mathbb{A} \text{ wins}] \\
&\quad + \Pr[\mathbb{A} \text{ did corrupt}] \cdot {}^1\!/_2 \\
&= \Pr[\mathbb{A} \text{ did not corrupt} \wedge \mathbb{A} \text{ wins}] \\
&\quad + \Pr[\mathbb{A} \text{ did corrupt}] \cdot \Pr[\mathbb{A} \text{ wins} \mid \mathbb{A} \text{ did corrupt}] \\
&= \Pr[\mathbb{A} \text{ did not corrupt} \wedge \mathbb{A} \text{ wins}] \\
&\quad + \Pr[\mathbb{A} \text{ did corrupt} \wedge \mathbb{A} \text{ wins}] \\
&= \Pr[\mathbb{A} \text{ wins}], \\
\Rightarrow \text{IND-CXA}_{\text{PKE}}(\mathbb{B}) &\geq \text{IND-CXA}^{1,1}_{\text{PKE}}(\mathbb{A}).
\end{aligned}
$$

$\square$

Single-Bit Security with Corruptions. We can then show a reduction from IND-CXA$^{\not{k},1}_{\text{PKE}}$ to IND-CXA$^{1,1}_{\text{PKE}}$, using exactly same hybrid argument that was used by Bellare et al. [5] in the absence of corruptions, and let Lemma 3 imply the result.

Theorem 3 (IND $\overset{\leq \kappa}{\Longrightarrow}$ IND$^{\not{k},1}$). *There is a SFBB reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,*

$$\text{IND-CXA}^{\not{k},1}_{\text{PKE}}(\mathbb{A}) \leq \kappa \cdot \text{IND-CXA}_{\text{PKE}}(\mathbb{B}),$$

where $\mathbb{B}$'s overhead consists of $\kappa - 1$ fresh keypair generations.

Proof (sketch). Through a standard hybrid argument completely analogous to that used to prove the corruptionless version, we can show that there is an adversary $\mathbb{B}$ such that for every adversary $\mathbb{A}$,

$$\text{IND-CXA}_{\text{PKE}}^{\kappa,1}(\mathbb{A}) \leq \kappa \cdot \text{IND-CXA}_{\text{PKE}}^{1,1}(\mathbb{B}) .$$

Then, Lemma 3 implies the result.

See the full version for the complete proof.

Diagonal-Bit Security with Corruptions. For the diagonal notion, showing the relation to the single-user notion is done using a different—and arguably simpler—proof technique: the reduction $\mathbb{B}$ simply guesses which user $\mathbb{A}$ is going to attack, forwarding the oracles called to that user to its own oracles and simulating the rest; it will guess correctly with probability $1/\kappa$, leading to the κ security loss.

Theorem 4 $(\text{IND} \overset{\leq\kappa}{\Longrightarrow} \text{IND}^{\kappa,\boxtimes})$. *There is an SFBB reduction $\mathbb{B}$ such that, for every adversary $\mathbb{A}$,*

$$\text{IND-CXA}_{\text{PKE}}^{\kappa,\boxtimes}(\mathbb{A}) \leq \kappa \cdot \text{IND-CXA}_{\text{PKE}}(\mathbb{B}) ,$$

where $\mathbb{B}$'s overhead consists of $\kappa - 1$ fresh keypair generations.

Proof (sketch). $\mathbb{B}$ draws a key handle $i^* \in [\kappa]$ uniformly at random. Whenever $\mathbb{A}$ calls an oracle using this handle, $\mathbb{B}$ will forward the call to its own oracle. To simulate the rest of the users, $\mathbb{B}$ draws fresh keypairs and challenge bits, simulating the oracles as needed. If $\mathbb{A}$ returns a guess on challenge bit i^*, $\mathbb{B}$ forwards the guess, gaining the winning advantage of $\mathbb{A}$. Given that the value of i^* is information-theoretically hidden from $\mathbb{A}$, this happens with probability exactly $1/\kappa$. Otherwise, $\mathbb{B}$ returns 0, achieving advantage 0.

The full proof can be found in the full version.

D   Formalization of ROR to LOR Implications

Theorem 5 $(\text{ROR}^{u,1} \overset{\leq 2}{\Longrightarrow} \text{LOR}^{u,1})$. *There is a simple, fully black box reduction $\mathbb{B}$ such that, for any adversary $\mathbb{A}$,*

$$\text{LOR-CXA}_{\text{PKE}}^{u,1}(\mathbb{A}) \leq 2 \cdot \text{ROR-CXA}_{\text{PKE}}^{u,1}(\mathbb{B}) ,$$

where $\mathbb{B}$'s overhead consists of drawing one uniformly random bit.

Proof (Sketch). Essentially, there are only two, equally likely cases: either the bit is "real", in which case $\mathbb{B}$ is able to simulate the left-or-right game perfectly; or the bit is "random", in which case the advantage of $\mathbb{A}$ against the simulated game will be exactly 0—and the addition of corruptions does nothing to change this fact.

$\mathbb{B}(\mathsf{pk}_1, \dots, \mathsf{pk}_\kappa)$	if $\mathbb{A}$ calls $\mathcal{E}(i, m_0, m_1)$	if $\mathbb{A}$ calls $\mathcal{D}(i, c)$
$d \leftarrow\!\!\$ \{0, 1\}$	$c^* \leftarrow \mathcal{E}(i, m_d)$	$m \leftarrow \mathcal{D}(i, c)$
$\hat{d} \leftarrow\!\!\$ \mathbb{A}^{\mathcal{E}, \mathcal{D}, \mathcal{K}}(\mathsf{pk}_1, \dots, \mathsf{pk}_\kappa)$	**return** c^*	**return** m
$\hat{b} \leftarrow d \neq \hat{d}$		
return $\hat{b}$		if $\mathbb{A}$ calls $\mathcal{K}(i)$
		$\mathsf{sk}_i \leftarrow \mathcal{K}(i)$
		return sk_i

Fig. 8. The adversary $\mathbb{B}$, playing $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ror\text{-}cca}, \not\kappa, 1}$ while simulating $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca}, \not\kappa, 1}$ for $\mathbb{A}$.

Proof. We will show the theorem for the case $u = \not\kappa$ and CXA = CCA; by inspection, the proof also holds for the cases $u = \kappa$ (by setting $\Pr[1 \in \mathsf{J}^\mathcal{K}] = 0$), and CXA = CPA.

In the following, let b be the challenge bit of $\mathbb{B}$'s game $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ror\text{-}cca}, \not\kappa, 1}$ (see Fig. 1, with $\beta = 1$). Let $\mathsf{J}^\mathcal{K}$ denote the set of compromised bits; note however that there is now only one challenge bit per game, meaning its bit handle is 1, and the event that it was compromised is denoted by $1 \in \mathsf{J}^\mathcal{K}$. Using the strategy of Fig. 8, we then get $\Pr\!\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{ror\text{-}cca}, \not\kappa, 1}(\mathbb{B}) = 1 \right]$

$$
\begin{aligned}
&= \Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d = \hat{d} \wedge b = 0 \right] \\
&\quad + \Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d \neq \hat{d} \wedge b = 1 \right] \\
&\quad + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \right] \\
&= \Pr[b = 0] \left(\Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d = \hat{d} \,\middle|\, b = 0 \right] + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 0 \right] \right) \\
&\quad + \Pr[b = 1] \left(\Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d \neq \hat{d} \,\middle|\, b = 1 \right] + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 1 \right] \right).
\end{aligned}
$$

Note that if $b = 1$, then the value of d is information-theoretically hidden from $\mathbb{A}$, so we have that $\Pr\!\left[d \neq \hat{d} \,\middle|\, b = 1 \right] = \Pr[\delta = 1 \mid b = 1] = 1/2$, allowing us to write

$$
\begin{aligned}
&= \Pr[b = 0] \left(\Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d = \hat{d} \,\middle|\, b = 0 \right] + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 0 \right] \right) \\
&\quad + \Pr[b = 1] \left(\Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 1 \right] + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 1 \right] \right) \\
&= \Pr[b = 0] \left(\Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d = \hat{d} \,\middle|\, b = 0 \right] + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 0 \right] \right) \\
&\quad + \Pr[b = 1] \cdot \Pr[\delta = 1 \mid b = 1] \\
&= \frac{1}{2} \left(\Pr\!\left[1 \notin \mathsf{J}^\mathcal{K} \wedge d = \hat{d} \,\middle|\, b = 0 \right] + \Pr\!\left[1 \in \mathsf{J}^\mathcal{K} \wedge \delta = 1 \,\middle|\, b = 0 \right] + \frac{1}{2} \right) \\
&= \frac{1}{2} \left(\Pr\!\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca}, \not\kappa, 1}(\mathbb{A}) = 1 \right] + \frac{1}{2} \right).
\end{aligned}
$$

Which implies that $\text{ROR-CCA}_{\text{PKE}}^{\not{k},1}(\mathbb{B})$

$$= 2 \cdot \Pr\left[\text{Exp}_{\text{PKE}}^{\text{ror-cca},\not{k},1}(\mathbb{B}) = 1\right] - 1$$

$$= 2 \cdot \frac{1}{2}\left(\Pr\left[\text{Exp}_{\text{PKE}}^{\text{lor-cca},\not{k},1}(\mathbb{A}) = 1\right] + \frac{1}{2}\right) - 1$$

$$= \frac{1}{2}\left(2 \cdot \Pr\left[\text{Exp}_{\text{PKE}}^{\text{lor-cca},\not{k},1}(\mathbb{A}) = 1\right] - 1\right)$$

$$= \frac{1}{2} \cdot \text{LOR-CCA}_{\text{PKE}}^{\not{k},1}(\mathbb{A}),$$

which is what we aimed to show.   □

Taken together with Theorem 1, this implies that the left-or-right free-bit notion without corruptions is separated from the single-bit real-or-random notion by at most a factor 2.

Corollary 4 ($\text{ROR}^{\kappa,1} \overset{\leq 2}{\Longrightarrow} \text{LOR}^{\kappa,\beta}$)**.** *There is a reduction $\mathbb{B}$ such that, for any adversary $\mathbb{A}$,*

$$\text{LOR-CXA}_{\text{PKE}}^{\kappa,\beta}(\mathbb{A}) \leq 2 \cdot \text{ROR-CXA}_{\text{PKE}}^{\kappa,1}(\mathbb{B}).$$

$\mathbb{B}$ *calls $\mathbb{A}$ once, and additionally uses the resources needed to draw β uniformly random bits.*

Proof (Sketch). Theorem 5 states that $\text{ROR}^{\kappa,1} \overset{\leq 2}{\Longrightarrow} \text{LOR}^{\kappa,1}$, while Theorem 1 states that $\text{LOR}^{\kappa,1} \implies \text{LOR}^{\kappa,\beta}$, allowing us to conclude that $\text{ROR}^{\kappa,1} \overset{\leq 2}{\Longrightarrow} \text{LOR}^{\kappa,\beta}$.

Given that the free-bit notion generalises the single-bit notion, this in turn implies that LOR and ROR are separated by at most a factor 2 between the corruptionless free-bit notions, even if the number of challenge bits varies between them.

With corruptions, however, any direct simulation would become trivially recognisable—meaning that in order to do a faithful simulation, the reduction would have to guess which bit the adversary is going to attack, leading to a loss linear in β. Instead of reformulating this argument, we let it follow as a corollary to previous results, yielding a slightly tighter statement by letting $\mathbb{B}$ play a single-bit game.

Corollary 5 ($\text{ROR}^{\not{k},1} \overset{\leq 2\beta}{\Longrightarrow} \text{LOR}^{\not{k},\beta}$)**.** *For every adversary $\mathbb{A}$, there is an adversary $\mathbb{B}$, such that*

$$\text{LOR-CXA}_{\text{PKE}}^{\not{k},\beta}(\mathbb{A}) \leq 2 \cdot \beta \cdot \text{ROR-CXA}_{\text{PKE}}^{\not{k},1}(\mathbb{B}).$$

$\mathbb{B}$ *calls $\mathbb{A}$ once, and additionally uses the resources needed to draw β uniformly random bits.*

Proof (Sketch). Theorem 5 states that $\text{ROR}^{\not k,1} \stackrel{\leq 2}{\Longrightarrow} \text{LOR}^{\not k,1}$, while Theorem 2 states that $\text{LOR}^{\not k,1} \stackrel{\leq \beta}{\Longrightarrow} \text{LOR}^{\not k,\beta}$, allowing us to conclude that $\text{ROR}^{\not k,1} \stackrel{\leq 2\beta}{\Longrightarrow} \text{LOR}^{\not k,\beta}$.

Interestingly, the tightest known relation from the diagonal-bit $\text{ROR}^{\kappa,\boxtimes}$ to $\text{LOR}^{\kappa,\boxtimes}$ is one that loses a factor 2κ, even in the absense of corruptions. This is once again achieved going through the single-user notion.

Corollary 6 $\left(\text{ROR} \stackrel{\leq 2\kappa}{\Longrightarrow} \text{LOR}^{u,\boxtimes}\right)$**.** *There is a reduction* $\mathbb{B}$ *such that, for every adversary* $\mathbb{A}$,

$$\text{LOR-CXA}^{u,\boxtimes}_{\text{PKE}}(\mathbb{A}) \leq 2 \cdot \kappa \cdot \text{ROR-CXA}_{\text{PKE}}(\mathbb{B})\,.$$

$\mathbb{B}$ *calls* $\mathbb{A}$ *once, and additionally uses the resources needed to draw* κ *fresh key-pairs and* κ *uniformly random bits.*

Proof (Sketch). It is well established [6] that $\text{ROR} \stackrel{\leq 2}{\Longrightarrow} \text{LOR}$, and we know from Theorem 4 that $\text{LOR} \stackrel{\leq \kappa}{\Longrightarrow} \text{LOR}^{\not k,\boxtimes}$, allowing us to conclude that $\text{ROR} \stackrel{\leq 2\kappa}{\Longrightarrow} \text{LOR}^{\not k,\boxtimes}$.

E   Deferred Proof of Theorem 2

Theorem 2 $\left(\text{IND}^{\not k,1} \stackrel{\leq \beta}{\Longrightarrow} \text{IND}^{\not k,\beta}\right)$**.** *There is an SFBB reduction* $\mathbb{B}$ *such that, for any adversary* $\mathbb{A}$,

$$\text{IND-CXA}^{\not k,\beta}_{\text{PKE}}(\mathbb{A}) \leq \beta \cdot \text{IND-CXA}^{\not k,1}_{\text{PKE}}(\mathbb{B})\,,$$

where $\mathbb{B}$*'s overhead consists of drawing* β *uniformly random bits.*

We will show the result for $\text{IND} = \text{LOR}$ and $\text{CXA} = \text{CCA}$; the proof transfers directly to the ROR and CPA cases.

Proof. We will prove the statement by constructing an adversary $\mathbb{B}$ that achieves the claimed advantage by leveraging any advantage an adversary $\mathbb{A}$ has against the free-bit game, and making a guess on the bit that $\mathbb{A}$ is going to attack. $\mathbb{B}$ will guess correctly with probability $1/\beta$, leading to the β security loss. The proof is very similar to that of Theorem 4, the main complication being that we now need to keep track of compromised challenge bits, instead of just which keys are corrupted.

$\mathbb{B}$ is given in Fig. 9. In the following, let b be the challenge bit of $\mathbb{B}$'s game $\text{Exp}^{\text{lor-cca},\not k,1}_{\text{PKE}}$ (see Fig. 1, with $\beta = 1$), let the set of compromised bits (i.e., bits used by $\mathbb{A}$ to challenge a corrupted key) be denoted by $\mathsf{J}^{\mathcal{K}}$, and assume that $\mathbb{A}$ returns the guess $(j, \hat{b}_j)$. Finally, note that the value of j^* is

$\mathbb{B}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$	if $\mathbb{A}$ calls $\mathcal{E}(i, j, m_0, m_1)$	if $\mathbb{A}$ calls $\mathcal{D}(i, c)$
$j^* \leftarrow\!\!\$\, [\beta]$	if $m_0 \not\sim m_1$ **then return** ξ	if $c \in \mathsf{C}_i$ **then return** ξ
for $j \in [\beta], j \neq j^*$ **do** :	if $j = j^*$	$m \leftarrow \mathcal{D}(i, c)$
$\quad b_j \leftarrow\!\!\$\, \{0, 1\}$	$\quad c^* \leftarrow \mathcal{E}(i, m_0, m_1)$	**return** m
$\mathsf{C}_1, \ldots, \mathsf{C}_\kappa \leftarrow \emptyset$	**else** $c^* \leftarrow\!\!\$\, \mathsf{Pk.Enc}_{pk_i}(m_{b_j})$	
$(j, \hat{b}_j) \leftarrow\!\!\$\, \mathbb{A}^{\mathcal{E},\mathcal{D},\mathcal{K}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$	$\mathsf{C}_i \leftarrow c^*$	if $\mathbb{A}$ calls $\mathcal{K}(i)$
if $j \neq j^*$ **then return** 0	**return** c^*	$\mathsf{sk}_i \leftarrow \mathcal{K}(i)$
return $\hat{b}_j$		**return** sk_i

Fig. 9. The adversary $\mathbb{B}$, playing $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca},\not{k},1}$ while simulating $\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca},\not{k},\beta}$ for $\mathbb{A}$.

information-theoretically hidden from $\mathbb{A}$. Then, $\mathbb{B}$ achieves the following advantage, $\Pr\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca},\not{k},1}(\mathbb{B}) = 1\right]$

$$= \Pr\left[j = j^* \wedge j^* \notin \mathsf{J}^{\mathcal{K}} \wedge b_{j^*} = \hat{b}_{j^*}\right] + \Pr\left[j = j^* \wedge j^* \in \mathsf{J}^{\mathcal{K}} \wedge \delta = 1\right]$$

$$\quad + \Pr\left[j \neq j^* \wedge j^* \notin \mathsf{J}^{\mathcal{K}} \wedge b = 0\right] + \Pr\left[j \neq j^* \wedge j^* \in \mathsf{J}^{\mathcal{K}} \wedge \delta = 1\right]$$

$$= \Pr[j = j^*]\left(\Pr\left[j^* \notin \mathsf{J}^{\mathcal{K}} \wedge b_{j^*} = \hat{b}_{j^*} \,\middle|\, j = j^*\right] + \Pr\left[j^* \in \mathsf{J}^{\mathcal{K}} \wedge \delta = 1 \,\middle|\, j = j^*\right]\right)$$

$$\quad + \Pr[j \neq j^*]\left(\Pr[b = 0] \cdot \Pr\left[j^* \notin \mathsf{J}^{\mathcal{K}} \,\middle|\, j \neq j^*\right] + \Pr[\delta = 1] \cdot \Pr\left[j^* \in \mathsf{J}^{\mathcal{K}} \,\middle|\, j \neq j^*\right]\right)$$

$$= \frac{1}{\beta}\left(\Pr\left[j \notin \mathsf{J}^{\mathcal{K}} \wedge b_j = \hat{b}_j\right] + \Pr\left[j \in \mathsf{I}^{\mathcal{K}} \wedge \delta = 1\right]\right)$$

$$\quad + \frac{1}{2}\left(1 - \frac{1}{\kappa}\right)\left(\Pr\left[j^* \notin \mathsf{J}^{\mathcal{K}} \,\middle|\, j \neq j^*\right] + \Pr\left[j^* \in \mathsf{J}^{\mathcal{K}} \,\middle|\, j \neq j^*\right]\right)$$

$$= \frac{1}{\beta}\Pr\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca},\not{k},\beta}(\mathbb{A}) = 1\right] + \frac{1}{2}\left(1 - \frac{1}{\kappa}\right)$$

$$= \frac{1}{2\beta}\left(2 \cdot \Pr\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca},\not{k},\beta}(\mathbb{A}) = 1\right] - 1\right) + \frac{1}{2}$$

which implies that $\mathrm{LOR\text{-}CCA}_{\mathrm{PKE}}^{\not{k},1}(\mathbb{B})$

$$= 2 \cdot \Pr\left[\mathsf{Exp}_{\mathrm{PKE}}^{\mathrm{lor\text{-}cca},\not{k},1}(\mathbb{B}) = 1\right] - 1$$

$$= 2 \cdot \left(\frac{1}{2\beta}\mathrm{LOR\text{-}CCA}_{\mathrm{PKE}}^{\not{k},\beta}(\mathbb{A}) + \frac{1}{2}\right) - 1$$

$$= \frac{1}{\beta} \cdot \mathrm{LOR\text{-}CCA}_{\mathrm{PKE}}^{\not{k},\beta}(\mathbb{A}),$$

which is what we set out to show. $\square$

F   Multi-bit Composability of Hybrid Encryption

As shown by Cramer and Shoup [14], one can combine the practicality of asymmetric encryption with the efficiency of symmetric encryption into a highly efficient public key encryption system. The idea is to encrypt the message under an ephemeral symmetric key, which is itself encapsulated under a public key. This paradigm, which already saw widespread use at the time, has become known as the KEM/DEM paradigm, after its constituent Key Encapsulation Mechanism and Data Encapsulation Mechanism; it is also known as *hybrid* encryption.

Recently, Lee et al. [28] built on earlier work by Giacon et al. [17] and showed that a KEM and a DEM tightly compose to a PKE in a single-bit multi-user setting with corruptions. We paraphrase their result in Theorem 3.

Theorem 3 (Lee, Lee, Park, DCC'20). *There are SFBB reductions* $\mathbb{B}$ *and* $\mathbb{C}$ *such that, for every adversary* $\mathbb{A}$,

$$\text{LOR-CXA}_{\text{PKE}}^{\not{k},1}(\mathbb{A}) \leq 2 \cdot \text{ROR-CCA}_{\text{KEM}}^{\not{k},1}(\mathbb{B}) + 1\text{LOR-CCA}_{\text{DEM}}(\mathbb{C}).$$

Here, 1LOR means "one-time left-or-right"; see their paper for definitions and proof. Combining their result with Theorem 2 yields the following, more general, corollary.

Corollary 7 (Free-bit composability). *There are SFBB reductions* $\mathbb{B}$ *and* $\mathbb{C}$ *such that, for every adversary* $\mathbb{A}$,

$$\text{LOR-CXA}_{\text{PKE}}^{\not{k},\beta}(\mathbb{A}) \leq 2 \cdot \beta \cdot \text{ROR-CCA}_{\text{KEM}}^{\not{k},1}(\mathbb{B}) + \beta \cdot 1\text{LOR-CCA}_{\text{DEM}}(\mathbb{C}).$$

Proof. Immediately follows from Theorems 2 and 3.

While lossy in the number of challenge bits, it matches Lee et al.'s Theorem for $\beta = 1$. However, the implication to the diagonal-bit notion, with $\beta = \kappa$, results in a rather lossy composition, as made explicit below.

Corollary 8 (Diagonal-bit composability). *There are SFBB reductions* $\mathbb{B}$ *and* $\mathbb{C}$ *such that, for every adversary* $\mathbb{A}$,

$$\text{LOR-CXA}_{\text{PKE}}^{\not{k},\boxed{}}(\mathbb{A}) \leq 2 \cdot \kappa \cdot \text{ROR-CCA}_{\text{KEM}}(\mathbb{B}) + \kappa \cdot 1\text{LOR-CCA}_{\text{DEM}}(\mathbb{C}).$$

Proof. Follows from Theorems 4 and 3.

No tighter composition is known for multi-bit security notions, for much the same reason that no tight composition is known for AKE: as pointed out by Jager et al. [24], the multi-bit KEM notion does not easily allow for a game hop in which real keys are exchanged for fake ones, making the resulting game something in between the 'real' and 'random' worlds. Any attempt to circumvent this issue (without specialising to specific constructions) seems to lead to hybrid or guessing arguments, yielding similar linear losses.

References

1. Auerbach, B., Giacon, F., Kiltz, E.: Everybody's a target: scalability in public-key encryption. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12107, pp. 475–506. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45727-3_16

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. LNCS, vol. 9014, pp. 629–658. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46494-6_26

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. LNCS, vol. 9666, pp. 273–304. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_10

4. Baecher, P., Brzuska, C., Fischlin, M.: Notions of black-box reductions, revisited. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013. LNCS, vol. 8269, pp. 296–315. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-42033-7_16

5. 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 (2000). https://doi.org/10.1007/3-540-45539-6_18

6. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: 38th FOCS, pp. 394–403. IEEE Computer Society Press, October 1997. https://doi.org/10.1109/SFCS.1997.646128

7. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among notions of security for public-key encryption schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055718

8. Bellare, M., Hofheinz, D., Kiltz, E.: Subtleties in the definition of IND-CCA: when and how should challenge decryption be disallowed? J. Cryptol. **28**(1), 29–48 (2015). https://doi.org/10.1007/s00145-013-9167-4

9. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01001-9_1

10. Bellare, M., Ristenpart, T., Tessaro, S.: Multi-instance security and its application to password-based cryptography. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 312–329. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_19

11. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055716

12. Canetti, R., Halevi, S., Katz, J.: Adaptively-secure, non-interactive public-key encryption. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 150–168. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-30576-7_9

13. Cohn-Gordon, K., Cremers, C., Gjøsteen, K., Jacobsen, H., Jager, T.: Highly efficient key exchange protocols with optimal tightness. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11694, pp. 767–797. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_25

14. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055717

15. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. In: 40th FOCS, pp. 523–534. IEEE Computer Society Press, October 1999. https://doi.org/10.1109/SFFCS.1999.814626

16. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016. LNCS, vol. 9665, pp. 1–27. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_1

17. Giacon, F., Kiltz, E., Poettering, B.: Hybrid encryption in a multi-user setting, revisited. In: Abdalla, M., Dahab, R. (eds.) PKC 2018. LNCS, vol. 10769, pp. 159–189. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76578-5_6

18. Han, S., Liu, S., Gu, D.: Key encapsulation mechanism with tight enhanced security in the multi-user setting: impossibility result and optimal tightness. In: ASIACRYPT 2021. LNCS. Springer, Heidelberg (2021, to appear)

19. Hazay, C., Patra, A., Warinschi, B.: Selective opening security for receivers. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9452, pp. 443–469. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48797-6_19

20. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017. LNCS, vol. 10677, pp. 341–371. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70500-2_12

21. Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. Des. Codes Cryptogr. **80**(1), 29–61 (2016)

22. Hofheinz, D., Nguyen, N.K.: On tightly secure primitives in the multi-instance setting. In: Lin, D., Sako, K. (eds.) PKC 2019. LNCS, vol. 11442, pp. 581–611. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17253-4_20

23. Hofheinz, D., Rao, V., Wichs, D.: Standard security does not imply indistinguishability under selective opening. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9986, pp. 121–145. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53644-5_5

24. Jager, T., Kiltz, E., Riepel, D., Schäge, S.: Tightly-secure authenticated key exchange, revisited. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12696, pp. 117–146. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77870-5_5

25. Jager, T., Stam, M., Stanley-Oakes, R., Warinschi, B.: Multi-key authenticated encryption with corruptions: reductions are lossy. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017. LNCS, vol. 10677, pp. 409–441. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70500-2_14

26. Jager, T., Stam, M., Stanley-Oakes, R., Warinschi, B.: Multi-key authenticated encryption with corruptions: reductions are lossy. Cryptology ePrint Archive, Report 2017/495 (2017). https://eprint.iacr.org/2017/495

27. Katz, J., Lindell, Y.: Introduction to Modern Cryptography, 2nd edn. Chapman & Hall/CRC, London (2015)

28. Lee, Y., Lee, D.H., Park, J.H.: Tightly CCA-secure encryption scheme in a multi-user setting with corruptions. Des. Codes Cryptogr. **88**(11), 2433–2452 (2020)

29. Lewko, A., Waters, B.: Why proving HIBE systems secure is difficult. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 58–76. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_4

30. Luykx, A., Mennink, B., Paterson, K.G.: Analyzing multi-key security degradation. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10625, pp. 575–605. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70697-9_20
31. Reingold, O., Trevisan, L., Vadhan, S.: Notions of reducibility between cryptographic primitives. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 1–20. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24638-1_1