SoK: Public Key Encryption with Openings

Carlo Brunetta[1], Hans Heum[2], and Martijn Stam[1]

[1] Simula UiB, Bergen, Norway.
carlob,martijn@simula.no
[2] NTNU - Norwegian University of Science and Technology, Trondheim, Norway.
hans.heum@ntnu.no*

Abstract. When modelling how public key encryption can enable secure communication, we should acknowledge that secret information, such as private keys or the randomness used for encryption, could become compromised. Intuitively, one would expect unrelated communication to remain secure, yet formalizing this intuition has proven challenging. Several security notions have appeared that aim to capture said scenario, ranging from the multi-user setting with corruptions, via selective opening attacks (SOA), to non-committing encryption (NCE). Remarkably, how the different approaches compare has not yet been systematically explored.

We provide a novel framework that maps each approach to an underlying philosophy of confidentiality: indistinguishability versus simulatability based, each with an a priori versus an a posteriori variant, leading to four distinct philosophies. In the absence of corruptions, these notions are largely equivalent; yet, in the presence of corruptions, they fall into a hierarchy of relative strengths, from IND-CPA and IND-CCA at the bottom, via indistinguishability SOA and simulatability SOA, to NCE at the top.

We provide a concrete treatment for the four notions, discuss subtleties in their definitions and asymptotic interpretations and identify limitations of each. Furthermore, we re-cast the main implications of the hierarchy in a concrete security framework, summarize and contextualize other known relations, identify open problems, and close a few gaps.

We end on a survey of constructions known to achieve the various notions. We identify and name a generic random-oracle construction that has appeared in various guises to prove security in seemingly different contexts. It hails back to Bellare and Rogaway's seminal work on random oracles (CCS'93) and, as previously shown, suffices to meet one of the strongest notions of our hierarchy (single-user NCE with bi-openings).

Keywords: Selective Opening Attacks · Multi-User Security · Non-Committing Encryption · Corruptions

* Work by Hans Heum partially performed as part of his PhD studies at Simula UiB.

Table of Contents

1 Introduction 3
2 Preliminaries 6
  2.1 Notation 6
  2.2 PKE Syntax 7
  2.3 Security Notions 7
3 Confidentiality with Openings 9
  3.1 Four Kinds of Opening 9
  3.2 Four Philosophies of Confidentiality 10
  3.3 A Priori Indistinguishability with Selective Openings (IND) 12
  3.4 A Posteriori Indistinguishability with Selective Opening (ISO) 17
  3.5 A Posteriori Simulatability with Selective Opening (SSO) 22
  3.6 A Priori Simulatability with Selective Opening (NCE) 27
4 Relations 31
  4.1 Hybridization 32
  4.2 Implications 33
  4.3 Separations 35
  4.4 Further Relations 36
  4.5 Selected Related Notions 37
5 Constructions 39
  5.1 Bellare–Rogaway Encryption 39
  5.2 Other Constructions 40
References 45
A ISO Tightly Implies IND with Bi-openings 50

1 Introduction

A group of crypto friends want to exchange their latest ideas with each other. Obviously, they want to do so confidentially and, being slightly old-fashioned, imagine they use public key encryption to secure their communication. Yet, an adversary, mistakenly reckoning our friends are all about the other crypto, sets out to break into the devices used by some of the cryptographers, recovering a number of private keys in the hope of a big score. Ideally, the communication that wasn't intended for the compromised bunch remains secure, so our somewhat disillusioned adversary cannot scoop their research ideas.

Scenarios similar to the one above, involving public key encryption (PKE) with multiple, say κ, users some of whom may be corrupted, have been used to motivate a range of different security notions for PKE above and beyond the now classical left-or-right indistinguishability under chosen ciphertext attacks (IND-CCA). These novel notions range from multi-user indistinguishability with corruptions, via various flavours of selective opening attacks (SOA), to (non-interactive) non-committing encryption (NCE). Given the plethora of theoretical notions based on very similar, practical motivations, the question arises what are the pros and cons of the various notions, and whether some notions should be preferred over others.

In addition to leaking keys, similar notions have also studied security when the randomness used for encryption leaks, or when both keys and randomness leak. Thus we end up with a host of notions, and a sizeable literature exploring how they relate and how each may be achieved. Navigating this literature can be a daunting task for the uninitiated, particularly given diverging formalisms, subtle variations in security definitions (which may or may not turn out consequential to any particular result), and even contradicting claims.

This is the situation that the present Systematization of Knowledge (SoK) aims to remedy. Our goal is *not* to provide a complete classification of all possible variations, but rather to provide a roadmap of the high level choices, and highlight possible pitfalls when designing notions of security with openings. Our generalized definitions, given in Sect. 3.3–Sect. 3.6, may serve as inspiration, yet we stress that for applications, utility of the definitions should be the guiding principle.

Our systematization identifies four philosophies underlying the notions of confidentiality of messages, using two orthogonal considerations: whether a notion is based on indistinguishability or simulation, and whether it is an a priori or a posteriori variant. Each notion then falls into one of the four categories: *a priori indistinguishability* for left-or-right indistinguishability-based (multi-user) notions (IND); *a posteriori indistinguishability* for indistinguishability SOA (ISO); *a posteriori simulatability* for simulation SOA (SSO); and a *priori simulatability* for non-committing encryption (NCE). Each notion furthermore comes in four variants depending on whether only keys leak (receiver opening, denoted $\star$), only messages and randomness leak (sender opening, denoted $\odot$), just messages (transmission opening, denoted $\diamond$), or all of the above (bi-opening, denoted $\circledast$); the exception is NCE, for which only sender, receiver, and bi-opening is defined. Finally, each notion comes in a CPA and a CCA variant, making for 30 notions of security in total—some of which are studied here for the first time. An overview of our notation is given in Table 1.

To begin with, we observe that the four philosophies of confidentiality, while polynomially equivalent sans openings, seem to fall into a strict hierarchy of strength whenever receiver and/or sender openings are accounted for, see Fig. 1. Transmission openings on the other hand are significantly weaker, and all but one of the main notions are known to be polynomially equivalent to IND-CPA/IND-CCA when only transmission openings are included (with the remaining equivalence being conjectured, see Open Problem 11).

Findings. We give generalized definitions of the four main philosophies that all include multiple users, challenges, full adaptivity, and bi-openings; message samplers and simulators are all stateful, which simplifies (and in the case of message samplers, strengthens) the formalizations. In addition, we provide novel a priori indistinguishability notions of security in the presence of transmission, sender and bi-openings (see Def. 3), which are notable for being equivalent to IND-CPA resp. IND-CCA, at a loss for sender openings (Thm. 1) and bi-openings (Thm. 3), but almost completely tightly in the case of transmission openings (Thm. 2).

While ideally we would give definite formalizations of each of the notions considered herein, there are a number of subtle definitional choices available with often no clear best one (for all contexts). We have leaned towards generality wherever possible, but some choices are less clear in terms of benefits and generality. For example, which inputs to a distinguisher should be provided by the game and which may be simulated (for SSO and NCE notions); specifically, whether the simulator may generate the

Table 1. Overview of the different formalizations of capturing confidentiality, the auxiliary adversarial power, and the adversarial opening prowess, together with their relevant associated oracle(s): here, $\mathcal{E}$ are (challenge) encryption oracles, $\mathcal{C}$ is a challenge oracle, $\mathcal{D}$ is the decryption oracle, and $\mathcal{T}$, $\mathcal{S}$, and $\mathcal{R}$ are the transmission, sender, and receiver opening oracles, respectively (see Sect. 3.1).

Shorthand	Associated oracle(s)	Name
κ		number of users (receivers)
β		number of challenge bits
IND	$\mathcal{E}$	indistinguishability
ISO	$\mathcal{E}$ $\mathcal{C}$	indistinguishability-based selective opening
SSO	$\mathcal{E}$	simulation-based selective opening
NCE	$\mathcal{E}$	non-committing encryption
CPA		chosen plaintext attack
CCA	$\mathcal{D}$	chosen ciphertext attack
$\diamond$	$\mathcal{T}$	transmission opening
$\odot$	$\mathcal{S}$	sender opening
$\star$	$\mathcal{R}$	receiver opening
$\circledast$	$\mathcal{S}$ $\mathcal{R}$	bi-opening (sender + receiver)

parameters and public keys itself. We opted for a notion where the simulator generates the public keys but not the parameters, see Def. 6.

As another contribution, we provide a concrete-security treatment of topics that have traditionally been studied asymptotically, recasting several known implications in a concrete light, with asymptotic interpretations that connect to the literature (see Thm. 4–Thm. 6).

Our systematization allowed us to uncover a number of open problems, which we highlight in Open Problem 1–Open Problem 21. Of particular interest are the many relations known for the CPA setting that remain open in the CCA setting (see Fig. 16). Surprisingly, it is unclear whether notions of SSO-CCA imply notions of ISO-CCA in the presence of openings, as a straightforward reduction runs into trouble with simulating the decryption oracle, see Open Problem 5. Slightly more technical, we expand an earlier CPA result by showing an equivalence between IND-CCA and ISO-CCA$\diamond$ (Thm. 7), but a similar expansion demonstrating equivalence of SSO-CCA$\diamond$ and IND-CCA runs into trouble simulating decryptions (Open Problem 11).

Message samplers play a central role in definitions of SOA (the “a posteriori” philosophies), and here again we find a number of open problems. For example, for which (restricted) classes of message samplers do κ-ISO-CPA$\star$ and κ-ISO-CCA$\star$, like their sender opening counterparts ($\odot$), become equivalent to IND-CPA resp. IND-CCA (Open Problem 10)?

Yet other open problems concern achievability: for example, can κ-ISO-CPA$\circledast$ be achieved in the standard model (Open Problem 19)? For achievable notions, a secondary question becomes how tightly, for instance for our novel notion of (κ, β)-IND-CPA$\circledast$ (Open Problem 21)?

We hope that highlighting these open problems can serve as inspiration.

Applications. With the hierarchy in hand, one may ask what notion is really needed for a given application: going higher in the hierarchy broadens applicability at the cost of achievability, with several impossibility results sitting towards the top. Take for instance NCE with receiver openings under chosen plaintext attacks (NCE-CPA$\star$): impossible to achieve in the plain model, yet easily achieved in the programmable random oracle model; and, once achieved, it allows one to show security of a plethora of multi-party computation protocols that were, before the appearance of the notion, only known to be secure against adaptive adversaries assuming information-theoretically secure channels between the parties [6, 100].

For some protocols, for instance involving secret sharing with encrypted shares, the more achievable ISO may suffice. For yet other applications, e.g. building authenticated key exchange (AKE), the multi-user with corruptions notion—the only notion of the hierarchy equivalent to IND-CCA—suffices [2].

As we present the various notions in Sect. 3, we also discuss their pros and cons in terms of usability and achievability. A summary follows (for references, see the relevant sections):

- The main challenge of a priori indistinguishability (IND) is that these notions traditionally do not allow an adversary to open challenge ciphertexts, making them unsuitable for modelling sender and

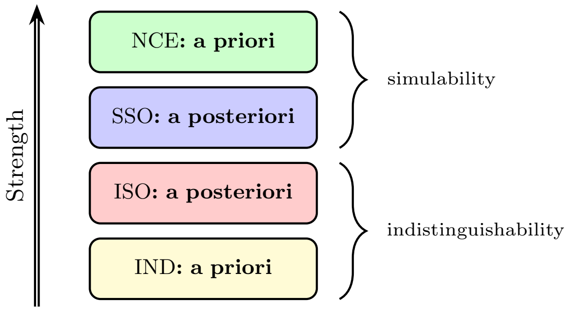

Fig. 1. The main hierarchy of philosophies, and their associated security notion with opening. The notions become stronger as we go up the hierarchy in the presence of sender or receiver opening (or both).

transmission openings, and for primitives such as secret sharing that rely on the adversary learning a subset of the challenge plaintexts. As we show, the use of multiple challenge bits allows also sender and transmission openings to be modelled in the a priori indistinguishability setting, though at the expense of composition challenges and tightness losses. On the positive side, notions of a priori indistinguishability are all implied by IND-CPA/IND-CCA, meaning many well-established constructions are known to achieve them, although not necessarily tightly so (see Sect. 3.3). Furthermore, if plaintexts are interrelated in only a limited way, it is possible that IND-CPA/ IND-CCA already suffices for ISO$\odot$-security, see Sect. 4.2.

- A posteriori indistinguishability (ISO) does allow for the opening of challenge ciphertexts, but at the cost of messages now being sampled rather than simply chosen. Thus, a heavier formalism involving message samplers becomes necessary, and, particular to a posteriori indistinguishability, involve conditional resampling, which may or may not be efficient for a particular application (see Sect. 3.4).
- A posteriori simulatability (SSO) does away with the resampling phase, leading to yet broader applicability, but at the cost of a significantly stronger notion: one that, in the presence of receiver openings (SSO$\star$), is even known to be unachievable in the standard model (see Sect. 3.5).
- Finally, a priori simulatability (NCE) does away with the need for message samplers, making for an arguably simpler notion, but at the cost of being the strongest and thus hardest to achieve. In particular, in the presence of receiver openings, it cannot be achieved in the standard model. On a more technical level, we disallow challenging a corrupted receiver in notions of a priori simulatability, and so if this restriction clashes with a particular use case then the lower three levels of the hierarchy might be more suitable (see Sect. 3.6).

One central question remains: when modelling openings, which notion of security is the right one? While there may not exist a definite answer, a few general recommendations come to mind:

- Firstly, assuming the goal is to model realistic settings, consider choosing a notion with bi-openings, as messages, encryption randomness, and keys all leak in the real world.
- Secondly, when proving your construction secure in the presence of openings, aim for the highest possible notion in the hierarchy, and (assuming it falls short of the top), consider providing some intuition as to why your construction does not reach higher (e.g. it is a standard model construction, and thus can not achieve simulatability notions with receiver openings; or it is binding, and thus can not achieve simulatability notions with sender openings; or simply that attempts at a proof ran into trouble).
- Thirdly, if you are applying PKE as a primitive in a larger protocol, stick to the lower notions unless you have a good reason to go higher, such as a need to open challenges while staying single-challenge-bit (ISO), or the need to support message distributions that are not efficiently conditionally resampleable (SSO). This ensures that the broadest possible set of PKE constructions can be used to instantiate your protocol.

Ultimately, we hope that this work can provide a helpful framework for comparing and contextualizing alternative notions of security with openings, and encourage their application in protocol design by untangling the surrounding literature. For example, as the main usefulness of notions of SOA lie in their ability to open a subset of challenge ciphertexts to unveil sampled plaintexts to the adversary—even when said plaintexts are arbitrarily interrelated—we expect that such notions have the potential to be particularly useful in the context of threshold cryptography. And yet we are aware of only a single work to date that applies a primitive achieving a SOA-like notion of security to prove a higher level protocol secure [17]. We hope the present systematization will help expand this list.

Future directions. We limit our investigation to perfect correctness, as most of the literature cited does not consider imperfectly correct schemes. However, with the rise of post-quantum cryptography, and lattice-based schemes in particular, the study of imperfect correctness is of increasing importance, and we view it as an important open problem to re-establish the relations studied herein in such a setting.

On a related note: several results considered herein rely heavily on (programmable) random oracles, or similar idealized models, which necessitate quantum upgrades such as the quantum random oracle model (QROM) [20] in order to be able to claim post-quantum security; thus re-establishing results in the post-quantum setting provides an intriguing double challenge.

Besides confidentiality, related goals have also been studied in the presence of openings, such as non-malleability [81] and anonymity (key privacy) [80]. While beyond our present scope, we find it would be interesting and worthwhile to re-establish our hierarchy in these settings, and investigate the notions of security the four philosophies give rise to when combined with the various openings.

Roadmap. After laying some groundwork in Sect. 2, we discuss the four kinds of openings in Sect. 3.1, and the four philosophies in Sect. 3.2. We then move our way up the hierarchy, from multi-user indistinguishability with openings (IND) at the bottom in Sect. 3.3, via indistinguishability-SOA (ISO) in Sect. 3.4 and simulation-SOA (SSO) in Sect. 3.5, to non-committing encryption (NCE) in Sect. 3.6. For each, we give a generalized definition stated in a unified formalism, survey other possible choices in the definitions, discuss the pros and cons in terms of applicability, highlight open problems, give a concrete-security proof that the notion in question implies the notion below it in the hierarchy, and provide some historical remarks. Additionally, in Sect. 3.3, we study for the first time notions of transmission, sender and bi-openings in the a priori indistinguishability setting, which can be modelled with security games employing more than one challenge bit, and show them to be poly-equivalent to the single- and multi-user notions of IND-CPA/IND-CCA at a loss linear in the number of challenge bits (tighter in the case of transmission openings).

In Sect. 4, we map out known relations, in terms of hybridization in the number of users (Sect. 4.1), implications between notions (Sect. 4.2), and separations between notions (Sect. 4.3). We identify a large number of open problems along the way, and we close a few. In Sect. 4.5, we discuss a selection of related notions that have appeared over the years, and where they sit in relation to our main hierarchy.

In Sect. 5 we move on to constructions. We identify in Sect. 5.1 a proof technique, which we name Bellare–Rogaway Encryption after its (first) inventors [12], which has appeared repeatedly in constructions aiming at the various notions, and with which one can in fact achieve the notion sitting at the very top of the hierarchy: non-committing encryption with bi-openings (the caveat being that the proof, by necessity, relies on programming a random oracle). Finally, in Sect. 5.2 we survey what other constructions are known to achieve the respective notions, both in idealized models like the programmable random oracle model (ROM), and in the standard model; Table 2 provides a helpful menu of options.

2 Preliminaries

2.1 Notation

For a positive integer n, we let $[n]$ denote the set $\{1, \ldots, n\}$. For a bit string $x \in \{0,1\}^*$, $|x|$ denotes its length. For a finite set $\mathcal{X}$, $|\mathcal{X}|$ is the cardinality of $\mathcal{X}$. For a list $\mathtt{X}$, $\mathtt{X}[i]$ retrieves the ith element of the list, and, by convention for an index set $\mathcal{I}$, $\mathtt{X}[\mathcal{I}]$ indicates the list $\mathtt{X}$ restricted to the elements whose indices are contained in $\mathcal{I}$. For $n \in \mathbb{Z}_{>0}$, we write $\mathtt{X}\llbracket n \rrbracket$ for $\mathtt{X}$ restricted to its first n elements, so it equals $\mathtt{X}[\mathcal{I}]$ with $\mathcal{I} = [n]$.

We use code-based experiments, where by convention all sets, lists, and lazily sampled functions are initialized empty. We use $\Pr[\mathsf{Code} : \mathsf{Event} \mid \mathsf{Condition}]$ to denote the conditional probability of Event occurring when Code is executed, conditioned on $\mathsf{Condition}$. We omit Code when it is clear from the context and $\mathsf{Condition}$ when it is not needed. In our experiments, we will assume that no oracle calls are made (by the adversary) with evidently out-of-bounds inputs (e.g. list indices or key handles).

In our (pseudo)code, we denote with $x \leftarrow y$ the deterministic assignment of y to the variable x and use the shorthand $\mathcal{X} \xleftarrow{\cup} x$ for the operation $\mathcal{X} \leftarrow \mathcal{X} \cup \{x\}$ and $\mathtt{X} \xleftarrow{\frown} x$ for appending the element x to a list $\mathtt{X}$. Furthermore, we use the shorthand $x \leftarrow\!\!\$\; \mathcal{X}$ to denote uniform sampling from the finite set $\mathcal{X}$ and $x \leftarrow\!\!\$\; Y(\cdot)$ for assigning the output of the probabilistic algorithm Y to x. We can make the randomness r of Y explicit by writing $x \leftarrow Y(\,\cdot\,; r)$, for previously sampled r.

We will also consider stateful algorithms, for instance when we write $m \leftarrow\!\!\$\; \mathsf{M}_{\langle s \rangle}(\alpha)$ the algorithm M takes as state s and as input the value α and outputs m, but it may change its state s as part of its processing (code-snippet taken from Fig. 8).

2.2 PKE Syntax

A public-key encryption scheme PKE consists of five algorithms. The probabilistic parameter generation algorithm $\mathsf{PKE.Pm}$ on input a security parameter λ outputs shared, public system parameters pm (these might for instance be the description of an elliptic curve group with generator for an ECDLP-based system); we use the shorthand $\mathsf{PKE}[\lambda]$ for a concrete instantiation of the scheme for the given security parameter. The probabilistic key generation algorithm $\mathsf{PKE.Kg}$ on input pm outputs a public/private key pair $(\mathsf{pk}, \mathsf{sk})$; without loss of generality, we assume that any algorithm that receives pk implicitly receives pm as well, and any algorithm receiving sk also receives pk. The probabilistic algorithm $\mathsf{PKE.Rnd}$ on input pm samples random coins r for encryption. Subsequently, the deterministic encryption algorithm $\mathsf{PKE.Enc}$ on input a public key pk, a message $m \in \mathcal{M}$ and randomness r outputs a ciphertext c. We can also treat $\mathsf{PKE.Enc}$ as a probabilistic algorithm by folding in the randomness generation by $\mathsf{PKE.Rnd}$, simply writing $c \leftarrow\!\!\$\; \mathsf{PKE.Enc}_{\mathsf{pk}}(m)$. Finally, the typically deterministic decryption algorithm $\mathsf{PKE.Dec}$ on input a private key sk and a ciphertext c outputs either a message $m \in \mathcal{M}$ or some failure symbol $\perp$. Henceforth, we assume that the message space $\mathcal{M}$ consists of (a subset of) arbitrary, finite length bit strings, i.e. $\mathcal{M} \subseteq \{0,1\}^*$ (which includes fixed-length binary representations of an algebraic structure). In addition, ciphertexts are typically bit strings whose length depends on, and hence leaks, the message length.

We typically assume the schemes to be perfectly correct, so

$$\Pr[r \leftarrow\!\!\$\; \mathsf{PKE.Rnd}(\mathsf{pm}) \;:\; \mathsf{PKE.Dec}_{\mathsf{sk}}(\mathsf{PKE.Enc}_{\mathsf{pk}}(m; r)) = m] = 1$$

for all parameters pm, all key pairs $(\mathsf{pk}, \mathsf{sk})$ generated by $\mathsf{PKE.Kg}(\mathsf{pm})$, and all messages $m \in \mathcal{M}$.

Remark 1. Some modern schemes, like LWE-based ones, allow a small probability of incorrectness, where decryption of an honestly generated ciphertext may occasionally return a wrong message or fail. Some classical schemes, such as hybrid KEM–DEM ones, may return distinct error messages (e.g. a KEM failure $\perp_{\mathsf{KEM}}$ versus a DEM failure $\perp_{\mathsf{DEM}}$). Other schemes might loosen the link between message length and ciphertext length to partially hide the former [51, 110]. To deal with these more general scenarios, some of the security definitions might require some subtle changes (beyond pure syntactical ones) to best capture the changed reality; moreover, even when the definitions remain the same, security proofs might rely on perfect correctness. While we occasionally highlight specific challenges when faced with a more general scenario, we stress that results shown to hold for single-error, length-regular, perfectly correct schemes do not automatically port to either of those more general scenarios. Specifically, whether known results on SOA and NCE still hold in the imperfect correctness scenario is unclear and we leave open the challenge of re-establishing relations in a setting with imperfect correctness.

Open Problem 1. *How are known security definitions and their relations affected when lifted to a setting with imperfect correctness, multiple error messages, or deviations from length regularity?*

2.3 Security Notions

Concrete advantages. Most security notions can be phrased in terms of an adversarial goal and the adversary's powers. This separation in goals and powers is reflected in the notation GOAL-POWER, such as IND-CPA for indistinguishability under chosen plaintext attacks.

We are primarily concerned with indistinguishability-style notions that task the adversary with guessing a single bit. These notions are modelled by a distinguishing experiment $\mathsf{Exp}^{\text{goal-power}}_{\mathsf{PKE}[\lambda],\ldots}(\mathbb{A})$ that, at the end of the game, either outputs true (if the adversary guesses correctly) or false (if not), leading to a distinguishing advantage (Def. 1). We will occasionally conflate the Boolean values true and false with the bits 1 and 0, respectively, and some games may manipulate the adversary's output prior to checking its guess, for instance to ensure that "bad" adversarial behaviour cannot be advantageous.

Definition 1 (Distinguishing advantage). *Let* $\mathsf{PKE}[\lambda]$ *be given and let* $\mathsf{Exp}^{\text{goal-power}}_{\mathsf{PKE}[\lambda],\ldots}(\mathbb{A})$ *be a distinguishing experiment for the notion* GOAL-POWER*, where the dots represent other possible dependencies of the notion (such as simulators or message samplers). Then an adversary* $\mathbb{A}$ *'s distinguishing advantage is defined as*

$$\mathsf{Adv}^{\text{goal-power}}_{\mathsf{PKE}[\lambda],\ldots}(\mathbb{A}) := 2 \cdot \Pr\Big[\mathsf{Exp}^{\text{goal-power}}_{\mathsf{PKE}[\lambda],\ldots}(\mathbb{A})\Big] - 1\,.$$

Unless otherwise stated, all our notions, including single-user notions, are multi-challenge, modelled by one or more challenge oracles (that depend on the bit to be guessed). Their precise formalizations depend on the underlying philosophy related to capturing "nothing leaks" (see Sect. 3).

Many powers are associated with a helper oracle, like the decryption oracle $\mathcal{D}$ used to model chosen ciphertext attacks (CCA), or the various opening oracles (see Sect. 3.1). As part of the goal, we prefix notions by the number of users κ and/or challenge bits β, and as part of the powers we postfix notions by the opening oracles available to them ($\diamond$, $\odot$, $\star$ or $\circledast$), see Table 1. If κ resp. β is set to one we typically omit the prefix, and likewise the postfix absent openings.

Although our focus will be on CPA and CCA security, there are many other powers sitting in between in strength. Without going into details, these include plaintext checking attacks (PCA) [101], ciphertext verification attacks (CVA) [39, 88], constrained chosen ciphertext attacks (CCCA) [74], and replayable chosen ciphertext attacks (RCCA) [32]. In the context of selective opening attacks, the CPA and CCA worlds display some remarkable divergence, raising the question where the various intermediate notions would sit.

Whenever possible, we state results using the concrete security approach. Thus, advantages are studied directly without a clear concept of what entails "security". Furthermore, advantages are well-defined without having to specify whether, say, a simulator may depend on the adversary or vice versa. Of course, when providing concrete reductions, the order of quantifiers does matter and speaks to the generality of a reduction rather than to the strength of a notion. A good, concrete reduction will result in a bound that, when combined with contemporary cryptanalytic hardness estimates of any underlying problem, results in concrete parameter choices for the scheme, targeting any desired security level. A good, concrete reduction should also be instrumental to derive asymptotic security statements as corollaries.

For simplicity, we will assume that adversaries will always terminate and do so with correctly formatted output; similarly, we assume adversaries respect the input–output interface of their oracles. These conventions are without loss of generality as the format checks could easily be added explicitly without yielding the adversary any additional advantage; the termination is implicitly captured by bounding the number of oracle calls an adversary may make and, if need be, modelling the adversary as a random system instead of an algorithm.

The asymptotic alternative. For asymptotic security, algorithms are modelled as uniform, probabilistic interactive Turing machines (ITM) that should work for all λ. An algorithm is deemed efficient if its ITM runs in worst-case time polynomial in λ; a scheme is defined secure for a given security notion if, for all efficient adversaries, its advantage is negligible in the security parameter λ [53]. For more complicated notions, we also need to resolve further dependencies (the "..." in Def. 1) and here, the order of quantifiers—whether a simulator may depend on the adversary or vice versa—does play a crucial role in defining the strength of a particular definition of security. Moreover, for security notions defined relative to message samplers, these samplers will be restricted to (efficient) classes for which security holds.

As a scheme can be secure or not in the asymptotic setting, the conditions under which a notion is achievable becomes a topic of study. Relatedly, we will also meet several impossibility results, saying that security under a notion is unachievable for all schemes (relative to some general conditions).

Comparing security notions. Both in the concrete and asymptotic settings, security notions can be compared with one another, but often with slightly different purposes. In the concrete setting, bounds are explicit and what they look like matters. For instance, if a reduction has a security loss factor 1 with little to no computational overhead, then the reduction is tight and, if there are tight reductions in both directions, then the notions are tightly equivalent. If a bound, tight or otherwise, cannot be improved upon (for instance there is a matching attack or metareduction), then the bound is sharp.

In the asymptotic setting, the emphasis is usually on implications, which is shown by the existence of a polynomial reduction, i.e. one whose security loss and computational overhead are polynomial in the security parameter. If there are polynomial reductions in both directions, then the notions are (polynomially) equivalent. If a security notion is not implied by another, then a separation may be shown, for instance by providing a (conditional) counterexample: a scheme that fulfils the definition of a PKE and can be proven secure under one notion, but for which there is an efficient adversary with a significant (typically overwhelming) advantage under the other notion.

If there is an implication in one direction and a separation in the other, then one notion is strictly stronger than the other; if there are separations in both directions, then the notions are incomparable. If there is neither an implication nor a separation between two notions, there is work to do!

Concrete security, specifically tightness and sharpness of bounds, only really becomes relevant when implications are known to exist. For instance, a priori indistinguishability with openings (Sect. 3.3) is polynomially equivalent to IND-CCA, but not tightly so, and is therefore only of interest in a concrete security setting. The remaining philosophies with openings have on the other hand been almost exclusively studied in the asymptotic setting, and the map of relations given in Fig. 16 is built on polynomial implications and separations.

Several of the security notions are furthermore defined relative to message samplers and/or simulators. In the concrete setting, reductions are then given relative to concrete samplers/simulators (see e.g. Thm. 4). In the asymptotic setting, they are instead quantified over as part of the definition of security, and implications between security notions are given relative to classes of samplers/simulators. Thus the asymptotic setting facilitates more general reductions, at the cost of concreteness.

We generally prefer the concrete approach over the asymptotic approach where possible. We state our results accordingly, but revert to the asymptotic mindset for the purpose of recalling many known results. We strive to make it clear from context and language usage which mindset we are working in at any time (e.g. tightness/lossiness in the concrete setting, implications/separations in the asymptotic setting).

Most of the reductions we consider are black-box, in the sense that they treat the adversary they build upon as a black-box. Furthermore, for generality, we often state that our reductions are type-preserving, which means that the type of queries the reduction makes, matches those of the underlying adversary. Type-preserving reductions are convenient to show simultaneously that, for instance, both CCA security of one flavour implies CCA security of another flavour and CPA security of that one flavour implies CPA security of that other flavour.

3 Confidentiality with Openings

3.1 Four Kinds of Opening

In a system with many users, one would like a guarantee that uncompromised traffic remains confidential even if a subset of the users are compromised. We will address different confidentiality guarantees in the next section and first explore user compromises themselves. Typical deployment of PKE involves two distinct user roles: that of sender and that of receiver. Their compromises need to be modelled separately, leading to four distinct flavours of openings.

Transmission openings. The weakest form of opening allows an adversary to open challenge ciphertexts to retrieve the underlying message. It differs from a chosen ciphertext attack as it specifically targets challenge ciphertexts, which are explicitly prohibited for a CCA-style decryption oracle. The notion is relatively rare, but we include it for completeness. We use the name transmission openings, let $\mathcal{T}$ denote the transmission opening oracle, and indicate its presence with the suffix $\diamond$.

Transmission openings can model partial compromises on both sender and receiver end: a sender might still have (a copy of) the message lying around, but have erased the ephemeral randomness used to encrypt; a receiver might take strong measures not to leak its long-term private key but might not care too much about the contents of a single message leaking.

The added power of the transmission opening oracle to an adversary appears minimal, in contrast to the stronger sender and receiver openings that we will discuss next. Indeed, Bellare and Yilek [16] (henceforth BY12) showed transmission-SOA equivalent to IND-CPA in the context of simulation-based selective opening attacks (which we will consider as a posteriori simulatability in the next section), see Sect. 4.2. Similarly, for a priori indistinguishability, we show that the added power is minimal (see Thm. 2), while for a priori simulatability, a transmission opening oracle would not add anything to the notion (as the adversary already knows the plaintext, see Def. 7).

Sender openings. Here an adversary can open any challenge ciphertext to receive both the message and underlying randomness; it models the compromise of a sender incapable of securely erasing said randomness. The study of SOA started out as a study of security in the presence of randomness reveals [41], as motivated by the setting of multi-party computation, where erasures are notoriously tricky [10].

Sender openings can be considered in a multiple-senders/single-receiver setting, so there is only one public–private key pair, yet the opening is per ciphertext. We let $\mathcal{S}$ denote the sender opening oracle, and indicate its presence with the suffix $\odot$.

Compared to transmission openings, and depending on the formalism, sender openings are much harder to deal with formally: the core technical difficulty sits with the committing property of most encryption schemes, as any adversary receiving the randomness and message corresponding to a challenge ciphertext can re-encrypt to verify that challenges and openings are consistent.

Receiver openings. An adversary who fully corrupts a receiver will obtain that user's private key. These kind of openings are found in both the multi-user literature [2,3,85] and the SOA literature [9,63,91].

For the latter, it is customary to reveal not just the private key, but also all the challenge messages that were encrypted under the corresponding public key. For us, receiver openings will only reveal the private key. For perfectly correct schemes, this choice is without loss of generality, as evidently, an adversary with access to both a ciphertext and the private key can simply run the decryption algorithm to obtain the originally encrypted message. When perfect correctness is not guaranteed, having one oracle that only reveals the private key and another that reveals the messages (cf. transmission openings) should result in a finer-grained notion.

We let $\mathcal{R}$ denote the receiver opening oracle, and indicate its presence with the suffix $\star$. For receiver openings to be meaningful, one should consider multiple receivers; we let κ indicate the number of receivers (used as prefix for security notions). As for sender openings, the core difficulty of receiver openings is the adversary's ability to verify that challenges and openings are consistent.

Bi-openings. Finally, an adversary might be granted access to both sender and receiver opening oracles (and thus also transmission openings), as indicated by the suffix $\circledast$. Bi-openings have been the standard in the NCE setting from the get-go [6,29], but only recently appeared in the SOA literature [91].

3.2 Four Philosophies of Confidentiality

To contrast and compare different confidentiality notions with openings, we first revisit four different philosophies to formalize confidentiality without. All four approaches aim to capture that an adversary "learns nothing" and hark back to Shannon's concept of perfect secrecy and its (near) equivalent notions [107]. The notions split in two, depending on whether they are indistinguishability versus simulatability based, and for both we consider an *a priori* and an *a posteriori* variant (see Fig. 1). We next describe each, going roughly in order of increasing strength.

A priori indistinguishability. In the information-theoretic, symmetric setting (where $\mathcal{M}$ consists of fixed-length bistrings) the idea is that, given any two messages, the same distribution over ciphertexts is induced, which can be formalized by stating that, for all m_0 and m_1 (and c),

$$\Pr[C = c \mid M = m_0] = \Pr[C = c \mid M = m_1],$$

where the probability is primarily over the choice of the secret key. In the computational, public key setting, the formalization differs, leading to classic left-or-right indistinguishability: an adversary, given access to a single public key pk, selects two (equal length) messages m_0 and m_1, receives the encryption of one of them under pk, and needs to figure out which one. Generalizations to allow multiple challenges and multiple users [7] form the basis for multi-user indistinguishability with corruptions (κ-IND-CPA$\star$ and κ-IND-CCA$\star$, see Sect. 3.3).

A posteriori indistinguishability. Here the concept is that, given a ciphertext, any two messages are equally likely. In the information-theoretic setting, it can be formalized as

$$\Pr[M = m_0 \mid C = c] = \Pr[M = m_1 \mid C = c],$$

which hides a dependency on the message distribution underlying M: the notion can only be satisfied iff the a priori distribution on the messages is uniform (in which case it is equivalent to a priori indistinguishability via Bayes's theorem).

For a computational PKE version, consider an experiment where an adversary specifies a message distribution M over $\mathcal{M} \subseteq \{0,1\}^*$. The game would sample a message m_0 according to M and encrypt m_0 to obtain ciphertext c. It would then sample a second message m_1 according to M, conditioned on $|m_0| = |m_1|$. Finally, the game returns (m_b, c) to the adversary, who has to guess b.

A posteriori indistinguishability is rather a rare notion for PKE, nonetheless it is the route taken for indistinguishability-based notions of selective opening attacks (ISO for short, see Sect. 3.4). In order for the notion to imply IND-CPA, however, the adversary should be allowed adaptive control over the distribution. In the concrete setting, we define the notion relative to a conditional resampler. In the asymptotic setting, this conditional resampler causes a bifurcation of the notion depending on whether said resampler can be efficiently implemented (effectively restricting the message samplers) or not.

Remark 2. While rare for PKE, a posteriori indistinguishability is reminiscent of standard notions of indistinguishability for KEMs, in which the adversary, given a ciphertext, must distinguish between the encapsulated symmetric key and a freshly drawn key [37]. With transmission openings, a posteriori indistinguishability closely matches Enhanced IND-CPA/IND-CCA for KEMs, in which the adversary receives challenge ciphertexts and is given the choice of whether to open it, revealing the encapsulated key, or to challenge it to receive either the encapsulated key, or a freshly drawn one [57].

A posteriori simulatability. Shannon's definition of perfect secrecy captured that, given a ciphertext c, each message m is as likely as it was before seeing c, or more formally

$$\Pr[M = m \mid C = c] = \Pr[M = m] \,;$$

again there is a dependency on the message distribution, but this time it can be arbitrary (so it is less troublesome than for a posteriori indistinguishability).

Perfect secrecy captures that, given a ciphertext, nothing should be leaked about the message. Goldwasser and Micali [54] famously captured this concept in a computational PKE setting as semantic security (SEM). Goldreich [53] later refined the notion by introducing a simulator: an adversary outputs a message distribution M, the game samples m according to M, encrypts m and returns the resulting ciphertext c to the adversary. Whatever the adversary subsequently computes, possibly using additional oracles, a simulator should be able to simulate. Some definitional variations are possible [111], based for instance on whether the adversary outputs only a single bit or an arbitrary output (to be simulated).

Like perfect secrecy, semantic security is arguably the most intuitive notion. Compared to a posteriori indistinguishability, there is no need to put any onerous restrictions on M, that is on how messages are sampled (although in the computational setting, it does need to be efficient). Simulation-based notions for selective opening attacks (SSO) follow this philosophy, see Sect. 3.5 for details.

A priori simulatability. Finally, we can require that the likelihood of observing a ciphertext c is independent of m, or

$$\Pr[C = c \mid M = m] = \Pr[C = c] \,.$$

In the information-theoretic setting, Shannon already showed this formalization equivalent to perfect secrecy.

The concept can be reinterpreted to say that, given a message, nothing is learned about the ciphertexts that might result from encrypting it, which can be captured by saying one can produce, or simulate, fully convincing ciphertexts without access to the message (before an adversary even gets involved). Hence, a priori simulatability.

Restricting to chosen-plaintext attacks for PKE, a priori simulatability can be defined by allowing an adversary to obtain the public key, select a message m, and then either receive an encryption of m or a simulated ciphertext, with the simulator only given access to the public key and the length of the message. The most common incarnation of a priori simulatability in the PKE setting fixes the simulator to simply select a random message of matching length and encrypting it, and the resulting notion is known as real-or-random indistinguishability (ROR) [8].

Upgrading to chosen-ciphertext attacks, the decryption oracle might have to be simulated as well; precise formalizations of a priori simulatability can be found in the universal composability (UC) framework [27, 90]. When allowing opening oracles, these need to be simulated as well; the resulting notion is known as (non-interactive) non-committing encryption (NCE), as discussed in Sect. 3.6.

```
Experiment $\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$
    $b \leftarrow\!\!\$\ \{0,1\}$
    $\mathsf{pm} \leftarrow\!\!\$\ \mathsf{PKE.Pm}(\lambda)$
    $\forall_{i\in[\kappa]}(\mathsf{pk}_i,\mathsf{sk}_i) \leftarrow\!\!\$\ \mathsf{PKE.Kg}(\mathsf{pm})$
    $\hat{b} \leftarrow\!\!\$\ \mathbb{A}^{\mathcal{E},\mathcal{D},\mathcal{R}}(\mathsf{pk}_1,\ldots,\mathsf{pk}_\kappa)$
    if $\mathcal{K}\cap\mathcal{I}\neq\emptyset : \hat{b} \leftarrow\!\!\$\ \{0,1\}$
    return $b = \hat{b}$

Oracle $\mathcal{E}(i,m_0,m_1)$
    if $|m_0| \neq |m_1|$ : return $\lightning$
    $\mathcal{K} \xleftarrow{\cup} i$
    $c \leftarrow\!\!\$\ \mathsf{PKE.Enc}_{\mathsf{pk}_i}(m_b)$
    $\mathcal{C}_i \xleftarrow{\cup} c$
    return $c$

Oracle $\mathcal{D}(i,c)$
    if $c \in \mathcal{C}_i$ : return $\lightning$
    $m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}_i}(c)$
    return $m$

Oracle $\mathcal{R}(i)$
    $\mathcal{I} \xleftarrow{\cup} i$
    return $\mathsf{sk}_i$
```

Fig. 2. A priori indistinguishability with receiver openings, also known as multi-user indistinguishability with corruptions.

Discussion. If we exclude a posteriori indistinguishability, then the remaining three notions are all equivalent in the information-theoretic setting, as already proven by Shannon. The same is true in the asymptotic computational setting: again excluding a posteriori indistinguishability, IND-CPA, SEM-CPA, and ROR-CPA are all polynomially equivalent [8, 54], as are IND-CCA, SEM-CCA, and UC-CCA [75, 111].

A posteriori indistinguishability appears not to have been studied for PKE without openings, although equivalence for message samplers satisfying conditional resamplability follows from Sect. 4.

3.3 A Priori Indistinguishability with Selective Openings (IND)

The multi-user setting [7] fits within the framework of a priori indistinguishability. Originally, openings were not considered, yet modelling multi-user security with receiver openings, also known as corruptions has seen an uptick [2, 58, 68, 92]. Several formalizations are possible, depending for instance on whether a single challenge bit is used across all κ public keys or whether each public key is allocated its own challenge bit. For receiver openings, we concentrate on the former, more standard approach (see Def. 2).

With only a single challenge bit, the opening of challenges would enable a trivial win and so must be disallowed. Thus, there is no meaningful notion of transmission or sender opening in such experiments. Generalizing to multiple challenge bits alleviates this issue: in the following, after recalling the canonical single-bit multi-user notion with receiver openings, we study a single-user multi-bit notion with sender openings, which to the best of our knowledge is studied here for the first time. We then collect the pieces and present the first notion of a priori indistinguishability with bi-openings, which, by allowing both multiple users and multiple challenge bits, strictly generalizes the prior notions.

Receiver openings. As explained, our formalization of a priori indistinguishability with receiver openings matches canonical notions of multi-user indistinguishability with corruptions, and we adopt a recent formalization [68].

Definition 2. *The* κ-IND-CCA$\star$ *advantage* $\mathsf{Adv}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$ *of an adversary* $\mathbb{A}$ *against public key encryption scheme* $\mathsf{PKE}[\lambda]$ *is the distinguishing advantage against the game* $\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$ *(see Fig. 2).*

Uses and Limitations. By a straightforward hybrid argument, the multi-user setting with receiver openings is implied by the single-user setting with a κ tightness loss, which entails polynomial equivalence [7].

For concrete instantiations, there are schemes known to be tightly secure in the multi-user setting with corruptions in the programmable random oracle model (see Sect. 5.2). The notion furthermore benefits from ease of composition, e.g. in constructing tightly secure hybrid encryption from a KEM and a DEM [92].

The main limitation of the notion is its segregation of opening versus challenging: adversaries cannot gain any advantage through both challenging and opening a user, as enforced by overwriting the adversary's output with a uniform value in the final stage of $\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$. This makes the notion inadequate for e.g. threshold security [105], for which security should hold as long as not too many challenges are opened.

Opening challenges. Employing multiple challenge bits, opening challenges becomes a viable strategy, provided the adversary outputs an uncompromised bit handle and guess at the end. For receiver openings, each user may for instance be associated a challenge bit; so that users can now be both challenged and

```
Experiment $\mathsf{Exp}^{\beta\text{-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$
    $\forall_{i\in[\beta]} b_i \leftarrow\$ \{0,1\}$
    $\mathsf{pm} \leftarrow\$ \mathsf{PKE.Pm}(\lambda)$
    $(\mathsf{pk},\mathsf{sk}) \leftarrow\$ \mathsf{PKE.Kg}(\mathsf{pm})$
    $(i,\hat{b}_i) \leftarrow\$ \mathbb{A}^{\mathcal{E},\mathcal{D},(\mathcal{T},)\mathcal{S}}(\mathsf{pk})$
    if $i \in \mathcal{I} : \hat{b} \leftarrow\$ \{0,1\}$
    return $b_i = \hat{b}_i$
```

```
Oracle $\mathcal{E}(i, m_0, m_1)$
    if $|m_0| \neq |m_1|$ : return $\lightning$
    $r \leftarrow\$ \mathsf{PKE.Rnd}(\mathsf{pm})$
    $c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}}(m_{b_i}; r)$
    $\mathtt{E} \xleftarrow{\frown} (i, m_{b_i}, r)$
    $\mathcal{C} \xleftarrow{\cup} c$
    return $c$

Oracle $\mathcal{D}(c)$
    if $c \in \mathcal{C}$ : return $\lightning$
    $m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}}(c)$
    return $m$
```

```
Oracle $\mathcal{T}(j)$
    $(i, m, r) \leftarrow \mathtt{E}[j]$
    $\mathcal{I} \xleftarrow{\cup} i$
    return $m$

Oracle $\mathcal{S}(j)$
    $(i, m, r) \leftarrow \mathtt{E}[j]$
    $\mathcal{I} \xleftarrow{\cup} i$
    return $(m, r)$
```

Fig. 3. A multiple-challenge-bit a priori indistinguishability game with transmission and sender openings.

corrupted, as long as at least one uncompromised user remains by the end. A KEM version of this multiple-challenge-bit security notion suffices for a construction achieving tightly (multi-challenge-bit) secure authenticated key exchange [2].

However, having multiple challenge bits comes with its own set of challenges: in particular, it typically leads to lossy composition theorems [68, 84]. The notion might also remain inadequate for e.g. threshold schemes, as having multiple challenge bits makes it hard to keep challenges consistent with each other.

Nonetheless, it facilitates the study of transmission and sender (and therefore also bi-) opening in the a priori indistinguishability setting.

Sender openings. We consider a single-receiver notion without receiver openings and define $\mathsf{Adv}^{\beta\text{-ind-cca}\odot}_{\mathsf{PKE}}(\mathbb{A})$ in the vein of Def. 2 using Fig. 3. At the start of the game, β challenge bits are drawn, representing β senders. The adversary can learn the value of a challenge bit by opening any challenge for which $m_0 \neq m_1$, which will no longer make for a valid guess. The notion is implied by the single-user notion (Thm. 1), essentially through a guessing argument, leading to a tightness loss linear in the number of challenge bits.

A comparable loss, namely the κ security loss from IND-CCA to κ-IND-CCA$\star$ is known to be sharp, both in the sense that there are schemes that meet the bound [7], and through meta-reduction showing that no black-box reduction can achieve a better bound [3]. One can wonder whether Thm. 1 is similarly sharp, as expressed in Open Problem 2.

Open Problem 2. *How sharp is the bound of Thm. 1?*

Theorem 1. *Let $\mathsf{PKE}[\lambda]$ be given. Then there is a type-preserving black-box reduction $\mathbb{B}_{\text{ind}}$ such that, for all $\mathbb{A}_\odot$,*

$$\mathsf{Adv}^{\beta\text{-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_\odot) \leq \beta \cdot \mathsf{Adv}^{\text{ind-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ind}})\,.$$

The runtime of $\mathbb{B}_{\text{ind}}$ is upper bounded by that of $\mathbb{A}_\odot$ plus q_e encryptions and q_d decryptions, where q_e and q_d are $\mathbb{A}_\odot$'s number of challenge oracle calls and decryption oracle calls respectively, and some small overhead.

Proof. The proof relies on two lemmas (stated and proved below), that together imply the stated result. The first one, Lemma 1, shows that 1-IND-CCA$\odot$ (Fig. 3 with $\beta = 1$) implies β-IND-CCA$\odot$ with a β tightness loss. The second one, Lemma 2, shows that when there is only one challenge bit, transmission/sender opening oracles do not help the adversary at all; in that case IND-CCA tightly implies 1-IND-CCA$\odot$. □

One could of course chain the reductions from the proofs of Lemma 1 and Lemma 2 in an attempt to create a more direct proof of Thm. 1. In essence, such a combined reduction would still make an initial guess i' on which sender i the adversary will end up attacking, forwarding challenges to its own challenge oracle whenever a challenge is asked on that handle and otherwise simulating the oracles off-line. The question though is what the reduction should do in the event that a challenge constructed under the guessed bit handle is requested to be opened. It cannot simulate either opening oracle honestly, and

```
Reduction $\mathbb{C}_{1\odot}(\mathsf{pk})$
$i' \leftarrow\!\!\$\ [\beta], k \leftarrow 0$
$\forall_{i\in[\beta]\setminus i'} b_i \leftarrow\!\!\$\ \{0,1\}$
$(i, \hat{b}_i) \leftarrow\!\!\$\ \mathbb{A}_{\odot}^{\mathcal{E},\mathcal{D},\mathcal{T},\mathcal{S}}(\mathsf{pk})$
if $i = i'$ : $\hat{b} \leftarrow \hat{b}_i$
else : $\hat{b} \leftarrow\!\!\$\ \{0,1\}$
return $\hat{b}$

If $\mathbb{A}_{\odot}$ calls $\mathcal{D}(c)$
if $c \in \mathcal{C}$ : return $\lightning$
$m \leftarrow \mathcal{D}_{\mathbb{C}}(c)$
return $m$

If $\mathbb{A}_{\odot}$ calls $\mathcal{E}_{\mathbb{A}}(i, m_0, m_1)$
if $i = i'$ :
    $k \leftarrow k + 1$
    $c \leftarrow \mathcal{E}_{\mathbb{C}}(m_0, m_1)$
    $\mathtt{E} \xleftarrow{\frown} (i, \perp, \perp, k)$
else :
    if $|m_0| \neq |m_1|$ : return $\lightning$
    $r \leftarrow\!\!\$\ \mathsf{PKE.Rnd}(\mathsf{pm})$
    $c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}}(m_{b_i}; r)$
    $\mathcal{C} \xleftarrow{\cup} c$
    $\mathtt{E} \xleftarrow{\frown} (i, m_{b_i}, r, \perp)$
return $c$

If $\mathbb{A}_{\odot}$ calls $\mathcal{T}(j)$
$(i, m, r, k) \leftarrow \mathtt{E}[j]$
if $i = i'$ :
    $m \leftarrow \mathcal{T}_{\mathbb{C}}(k)$
return $m$

If $\mathbb{A}_{\odot}$ calls $\mathcal{S}(j)$
$(i, m, r, k) \leftarrow \mathtt{E}[j]$
if $i = i'$ :
    $(m, r) \leftarrow \mathcal{S}_{\mathbb{C}}(k)$
return $(m, r)$
```

Fig. 4. The reduction $\mathbb{C}_{1\odot}$ simulating $\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{\beta\text{-ind-cca}\odot}$ for $\mathbb{A}_{\odot}$.

an inaccurate sender opening oracle is easily noticed by the adversary, so aborting the simulation (as the combined reduction would do) seems the natural option. Analysing the advantage of this guess-and-occassionally-abort reduction is doable with a sufficiently fine-grained case-analysis, but splitting the analysis in two Lemmas with their own reductions seemed the more modular and cleaner approach.

Lemma 1. *Let* $\mathsf{PKE}[\lambda]$ *be given. Then there is a type-preserving black-box reduction* $\mathbb{C}_{1\odot}$ *such that, for all* $\mathbb{A}_{\odot}$*,*

$$\mathsf{Adv}_{\mathsf{PKE}[\lambda]}^{\beta\text{-ind-cca}\odot}(\mathbb{A}_{\odot}) = \beta \cdot \mathsf{Adv}_{\mathsf{PKE}[\lambda]}^{1\text{-ind-cca}\odot}(\mathbb{C}_{1\odot})\,.$$

The runtime of $\mathbb{C}_{1\odot}$ *is upper bounded by that of* $\mathbb{A}_{\odot}$ *plus* q_e *encryptions and* q_d *decryptions, where* q_e *and* q_d *are* $\mathbb{A}_{\odot}$*'s number of challenge oracle calls and decryption oracle calls respectively, and some small overhead.*

Proof. The reduction $\mathbb{C}_{1\odot}$ (Fig. 4) makes a guess i' at the bit handle that $\mathbb{A}_{\odot}$ will end up attacking, forwarding challenge and opening oracle calls relating to that bit handle to its own oracles, and simulating the remaining challenge and opening oracle calls off-line. Decryption oracle calls are forwarded except if they involve a ciphertext issued as a challenge, in which case $\lightning$ is returned instead, as usual. Once $\mathbb{A}_{\odot}$ halts with a guess $(i, \hat{b}_i)$, $\mathbb{C}_{1\odot}$ checks whether the handle matches its guess, halting with $\hat{b}_i$ if yes and with a uniform guess otherwise.

As the simulation is perfect even if $\mathbb{A}_{\odot}$ asks to open a challenge produced under the guessed bit handle, $\mathbb{C}_{1\odot}$'s guess i' is information-theoretically hidden from $\mathbb{A}_{\odot}$, thus

$$\Pr[i = i'] = \frac{1}{\beta}\,.$$

Furthermore, if $\mathbb{C}_{1\odot}$ guessed incorrectly then its random bit $\hat{b}$ will be correct half of the time, so

$$\Pr\left[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{1\text{-ind-cca}\odot}(\mathbb{C}_{1\odot})\;\middle|\; i \neq i'\right] = \frac{1}{2}\,;$$

whereas if $\mathbb{C}_{1\odot}$ guessed correctly, its winning probability matches that of $\mathbb{A}_{\odot}$, regardless of whether or not $b_{i'}$ got compromised (if it did, the respective game mechanisms force the experiments of both $\mathbb{A}_{\odot}$ and $\mathbb{C}_{1\odot}$ to a comparison with a uniform random bit):

$$\Pr\left[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{1\text{-ind-cca}\odot}(\mathbb{C}_{1\odot})\;\middle|\; i = i'\right] = \Pr\left[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{\beta\text{-ind-cca}\odot}(\mathbb{A}_{\odot})\right].$$

Using these observations, we can express $\mathbb{C}_{1\odot}$'s winning probability as

$$\begin{aligned}
\Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\Big] &= \Pr[i = i']\Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\ \Big|\, i = i'\Big] \\
&\quad + \Pr[i \neq i']\Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\ \Big|\, i \neq i'\Big] \\
&= \frac{1}{\beta}\Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\ \Big|\, i = i'\Big] + (1 - \frac{1}{\beta})\frac{1}{2} \\
&= \frac{1}{\beta}\left(\Pr\Big[\mathsf{Exp}^{\beta\text{-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\odot})\Big] - \frac{1}{2}\right) + \frac{1}{2}\,;
\end{aligned}$$

and thus, after applying Def. 1, its advantage satisifies

$$\begin{aligned}
\mathsf{Adv}^{\text{1-ind-cca}\odot}(\mathbb{C}_{1\odot}) &= 2\left(\frac{1}{\beta}\left(\Pr\Big[\mathsf{Exp}^{\beta\text{-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\odot})\Big] - \frac{1}{2}\right) + \frac{1}{2}\right) - 1 \\
&= \frac{1}{\beta}\left(2\Pr\Big[\mathsf{Exp}^{\beta\text{-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\odot})\Big] - 1\right) \\
&= \frac{1}{\beta}\mathsf{Adv}^{\beta\text{-ind-cca}\odot}(\mathbb{A}_{\odot})\,.
\end{aligned}$$

□

Lemma 2. *Let* $\mathsf{PKE}[\lambda]$ *be given. Then there is a type-preserving black-box reduction* $\mathbb{B}_{\text{ind}}$ *such that, for all* $\mathbb{C}_{1\odot}$*,*

$$\mathsf{Adv}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot}) = \mathsf{Adv}^{\text{ind-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ind}})\,.$$

The runtime of $\mathbb{C}_{1\odot}$ *is upper bounded by that of* $\mathbb{B}_{\text{ind}}$*, plus some small overhead.*

Proof. The reduction $\mathbb{B}_{\text{ind}}$ simulates the game honestly by forwarding oracles up until the point that one of the opening oracles is called, at which point $\mathbb{B}_{\text{ind}}$ aborts the simulation and outputs a uniform guess. Let bad denote the event that an opening oracle is called. We have that

$$\Pr\Big[\mathsf{Exp}^{\text{ind-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ind}})\ \Big|\, \neg\mathsf{bad}\Big] = \Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\ \Big|\, \neg\mathsf{bad}\Big]$$

and

$$\Pr\Big[\mathsf{Exp}^{\text{ind-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ind}})\ \Big|\, \mathsf{bad}\Big] = \frac{1}{2}\,.$$

If $\mathbb{C}_{1\odot}$ calls one of its opening oracles, the real game will overwrite its guess with a uniform guess, and so

$$\Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\ \Big|\, \mathsf{bad}\Big] = \frac{1}{2}\,,$$

allowing us to conclude

$$\begin{aligned}
&\Pr\Big[\mathsf{Exp}^{\text{ind-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ind}})\Big] = \Pr\Big[\mathsf{Exp}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\Big] \\
\Longrightarrow\ &\mathsf{Adv}^{\text{ind-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ind}}) = \mathsf{Adv}^{\text{1-ind-cca}\odot}_{\mathsf{PKE}[\lambda]}(\mathbb{C}_{1\odot})\,.
\end{aligned}$$

□

Transmission openings. Fig. 3 simultaneously serves to define advantages for the β-IND-CCA$\diamond$ and β-IND-CCA notions (by omitting $\mathcal{S}$ for the former and both $\mathcal{S}$ and $\mathcal{T}$ for the latter). The notions β-IND-CCA and IND-CCA are tightly equivalent absent opening [68, Thm. 1]. Thus it is the reveal of messages and randomness, as opposed to the additional challenge bits, that gives the notions their strength, as the proof of tight equivalence fails in the presence of sender or transmission openings.

For transmission openings, we are able to show an almost tight equivalence to IND-CCA, losing only a factor 2 in one direction, while being trivially tight in the other direction. We conclude that transmission openings are of little interest in the a priori indistinguishability setting, adding at most a completely marginal strength.

Thm. 2 is inspired by BY12's proof of equivalence of IND-CPA and SSO-CPA$\diamond$ and shows that β-IND-CCA$\diamond$ is implied by ROR-CCA within a factor 2. Recall that ROR-CCA also implies IND-CCA within a factor 2 [8], leaving open the possibility that β-IND-CCA$\diamond$ and IND-CCA are in fact tightly equivalent (Open Problem 3, see also Fig. 5).

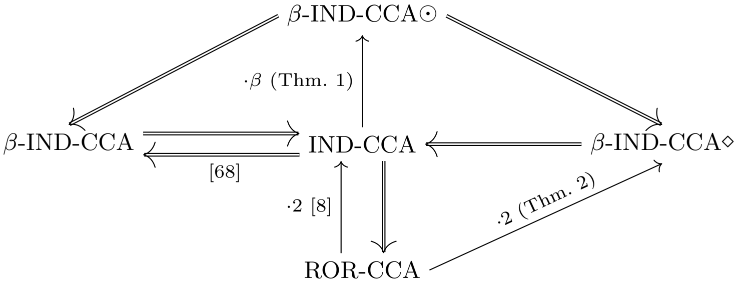

Fig. 5. Relations among the single-user single- and multi-bit notions of indistinguishability, with and without sender/transmission opening. (Double arrows = tight.)

```
Reduction $\mathbb{B}_{\mathrm{ror}}(\mathsf{pk})$
$\forall_{i\in[\beta]} d_i \leftarrow\!\!\$ \{0,1\}$
$(i,\hat{d}_i) \leftarrow\!\!\$ \mathbb{A}_\diamond^{\mathcal{E},\mathcal{D},\mathcal{T}}(\mathsf{pk})$
if $i \in \mathcal{I} : \hat{d}_i \leftarrow\!\!\$ \{0,1\}$
$\hat{b} \leftarrow \neg(d_i = \hat{d}_i)$
return $\hat{b}$
```

```
If $\mathbb{A}_\diamond$ calls $\mathcal{E}_{\mathbb{A}}(i, m_0, m_1)$
if $|m_0| \neq |m_1|$ : return $\lightning$
$c \leftarrow \mathcal{E}_{\mathbb{B}}(m_{d_i})$
$\mathtt{E} \xleftarrow{\frown} (i, m_{d_i})$
return $c$
```

```
If $\mathbb{A}_\diamond$ calls $\mathcal{T}(j)$
$(i,m) \leftarrow \mathtt{E}[j]$
$\mathcal{I} \xleftarrow{\cup} i$
return $m$
```

Fig. 6. The reduction $\mathbb{B}_{\mathrm{ror}}$ simulating $\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{\kappa\text{-ind-cca}\diamond}$ for $\mathbb{A}_\diamond$. (The decryption oracle is simply forwarded.)

Open Problem 3. *Are* IND-CCA *and* β-IND-CCA◇ *tightly equivalent?*

Theorem 2. *Let* $\mathsf{PKE}[\lambda]$ *be given. Then there is a type-preserving black-box reduction* $\mathbb{B}_{\mathrm{ror}}$ *such that, for all* $\mathbb{A}_\diamond$,

$$\mathsf{Adv}_{\mathsf{PKE}[\lambda]}^{\beta\text{-ind-cca}\diamond}(\mathbb{A}_\diamond) \leq 2 \cdot \mathsf{Adv}_{\mathsf{PKE}[\lambda]}^{\text{ror-cca}}(\mathbb{B}_{\mathrm{ror}})\,.$$

The runtime of $\mathbb{B}_{\mathrm{ror}}$ *is upper bounded by that of* $\mathbb{A}_\diamond$ *plus some small overhead.*

Proof. Reduction $\mathbb{B}_{\mathrm{ror}}$ is given in Fig. 6 (the $\mathcal{D}$ oracle is simply forwarded). Denote by b the challenge bit that $\mathbb{B}_{\mathrm{ror}}$ is tasked with guessing, with $b = 0$ corresponding to "real" and $b = 1$ to "random". We will bound its advantage next.

Intuitively, if $\mathbb{A}_\diamond$ returns a compromised bit handle, then its β-IND-CCA◇ advantage becomes 0; in this case $\mathbb{B}_{\mathrm{ror}}$ overwrites $\mathbb{A}_\diamond$'s guess with a uniform value, thus also gaining advantage 0 and equality holds. Otherwise there are two cases: either $b = 0$, in which case the simulation is completely faithful, or $b = 1$, in which case all uncompromised d_i are information-theoretically hidden from $\mathbb{A}_\diamond$. The reduction makes the guess $\hat{b} = 0$ (real) if $\mathbb{A}_\diamond$ makes a correct guess and $\hat{b} = 1$ (random) if not.

Let us look at each case separately: for $b = 0$,

$$\begin{aligned}
\Pr\left[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{\text{ror-cca}}(\mathbb{B}_{\mathrm{ror}}) \;\middle|\; b = 0\right] &= \Pr\left[\hat{b} = 0 \;\middle|\; b = 0\right] \\
&= \Pr\left[\hat{d}_i = d_i \;\middle|\; b = 0\right] \\
&= \Pr\left[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{\kappa\text{-ind-cca}\diamond}(\mathbb{A}_\diamond)\right],
\end{aligned}$$

where the final equality holds due to the simulation being completely faithful. For $b = 1$,

$$\begin{aligned}
\Pr\left[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{\text{ror-cca}}(\mathbb{B}_{\mathrm{ror}}) \;\middle|\; b = 1\right] &= \Pr\left[\hat{b} = 1 \;\middle|\; b = 1\right] \\
&= \Pr\left[\hat{d}_i \neq d_i \;\middle|\; b = 1\right] = \frac{1}{2},
\end{aligned}$$

where the final equality follows from unopened d_i being information-theoretically hidden from $\mathbb{A}_\diamond$ and opened $\hat{d}_i$ being supplanted by a uniform bit by the reduction.

Inserting the distinguishing advantage (Def. 1) and reordering yields the statement. □

```
Experiment $\mathsf{Exp}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$
$\forall_{j\in[\beta]} b_j \leftarrow\$ \{0,1\}$
$\mathsf{pm} \leftarrow\$ \mathsf{PKE.Pm}(\lambda)$
$\forall_{i\in[\kappa]} (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow\$ \mathsf{PKE.Kg}(\mathsf{pm})$
$(j, \hat{b}_j) \leftarrow\$ \mathbb{A}^{\mathcal{E},\mathcal{D},(\mathcal{T},)\mathcal{S},\mathcal{R}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
if $\exists i \in \mathcal{K}_j \cap \mathcal{I} : \hat{b}_j \leftarrow\$ \{0,1\}$
if $j \in \mathcal{J} : \hat{b}_j \leftarrow\$ \{0,1\}$
return $b_j = \hat{b}_j$
```

```
Oracle $\mathcal{E}(i, j, m_0, m_1)$
if $|m_0| \neq |m_1|$ : return $\lightning$
$\mathcal{K}_j \xleftarrow{\cup} i$
$r \leftarrow\$ \mathsf{PKE.Rnd}(\mathsf{pm})$
$c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}_i}(m_{b_j}; r)$
$\mathtt{E} \xleftarrow{\frown} (j, m_{b_j}, r)$
$\mathcal{C}_i \xleftarrow{\cup} c$
return $c$
```

```
Oracle $\mathcal{D}(i, c)$
if $c \in \mathcal{C}_i$ : return $\lightning$
$m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}_i}(c)$
return $m$
```

```
Oracle $\mathcal{T}(k)$
$(j, m, r) \leftarrow \mathtt{E}[k]$
$\mathcal{J} \xleftarrow{\cup} j$
return $m$
```

```
Oracle $\mathcal{S}(k)$
$(j, m, r) \leftarrow \mathtt{E}[k]$
$\mathcal{J} \xleftarrow{\cup} j$
return $(m, r)$
```

```
Oracle $\mathcal{R}(i)$
$\mathcal{I} \xleftarrow{\cup} i$
return $\mathsf{sk}_i$
```

Fig. 7. A priori indistinguishability with multiple challenge bits and bi-opening.

Bi-openings. With multiple users and multiple challenge bits, it is possible to model bi-opening in the a priori indistinguishability setting. For full generality, users and challenge bits should be decoupled; otherwise, if each user is allocated a separate challenge bit, it is unclear if the resulting notion is really stronger than single-bit indistinguishability with receiver opening [68]. Thus, the adversary is now free to choose which user to challenge and under which challenge bit the challenge should be constructed. What we end up with is a so-called "free-bit" notion, matching recent formalizations [68, 85] except for the addition of a sender and a transmission opening oracle.

Definition 3. *The* (κ, β)-IND-CCA$\circledast$ *advantage* $\mathsf{Adv}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$ *of an adversary* $\mathbb{A}$ *against public key encryption scheme* $\mathsf{PKE}[\lambda]$ *is the distinguishing advantage against the game* $\mathsf{Exp}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A})$ *(see Fig. 7).*

The notion is implied by IND-CCA with a $\kappa \cdot \beta$ tightness loss, which follows from it being implied by κ-IND-CCA$\star$ with a β tightness loss as we show next.

Theorem 3. *Let* $\mathsf{PKE}[\lambda]$ *be given. Then there is a type-preserving black-box reduction* $\mathbb{B}_\star$ *such that, for all* $\mathbb{A}_\circledast$*,*

$$\mathsf{Adv}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_\circledast) \leq \beta \cdot \mathsf{Adv}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_\star)\,.$$

The runtime of $\mathbb{B}_\star$ *is upper bounded by that of* $\mathbb{A}_\circledast$ *plus* q_e *encryptions and* q_d *decryptions, where* q_e *and* q_d *are* $\mathbb{A}_\circledast$*'s number of challenge oracle calls and decryption calls respectively, and some small overhead.*

Proof (sketch). The proof is essentially the same as that of Thm. 1, except with multiple users, and it is again useful to split it up in the equivalent of Lemma 1 and Lemma 2: at the outset the reduction, playing $(\kappa, 1)$-IND-CCA$\circledast$, guesses a bit handle j' and simulates (κ, β)-IND-CCA$\circledast$ by forwarding challenge calls using that bit handle to its own challenge oracle; it simulates non-matching oracle calls using the public keys and encrypting honestly. Once again, the simulation is perfect, even in the event that a call to an opening oracle would compromise $b_{j'}$ and, as before, the probability that the guess matches the returned handle is $\Pr[j = j'] = {}^{1}\!/\!_{\beta}$, leading to the stated tightness loss. Finally, we can show that a reduction playing κ-IND-CCA$\star$ gains the advantage of any adversary playing $(\kappa, 1)$-IND-CCA$\circledast$ by simply aborting the simulation and outputting a uniform guess in the case that the challenge bit is compromised through opening. □

3.4 A Posteriori Indistinguishability with Selective Opening (ISO)

A posteriori indistinguishability, as described in Sect. 3.2, is defined relative to a message sampler M: rather than choosing challenge messages m_0 and m_1, the adversary is allowed to affect the sampling through input α. During a subsequent challenge phase, the adversary receives either the originally sampled message(s) or a resampled version thereof. The concept lends itself well to modelling opening attacks: the resampling can be refined to conditional resampling, thus ensuring consistency with the opening and

```
Experiment $\mathsf{Exp}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A})$
$b \leftarrow\$ \{0,1\}$
$\mathsf{challenged} \leftarrow \mathsf{false}$
$q \leftarrow 0, s \leftarrow \epsilon$
$\mathsf{pm} \leftarrow\$ \mathsf{PKE.Pm}(\lambda)$
$\forall_{i\in[\kappa]}(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow\$ \mathsf{PKE.Kg}(\mathsf{pm})$
$\hat{b} \leftarrow\$ \mathbb{A}^{\mathcal{E},\mathcal{C},\mathcal{D},(\mathcal{T},)\mathcal{S},\mathcal{R}}(\mathsf{pk}_1,\ldots,\mathsf{pk}_\kappa)$
return $b = \hat{b}$
```

```
Oracle $\mathcal{D}(i,c)$
if $c \in \mathcal{C}_i$ : return $\lightning$
$m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}_i}(c)$
return $m$
```

```
Oracle $\mathcal{E}(i,\alpha)$
if challenged : return $\lightning$
$q \leftarrow q+1$
$\mathtt{K} \xleftarrow{\frown} i$
$\mathtt{A} \xleftarrow{\frown} \alpha$
$m \leftarrow\$ \mathsf{M}_{\langle s\rangle}(\alpha)$
$\mathtt{L} \xleftarrow{\frown} |m|, \mathtt{M}^0 \xleftarrow{\frown} m$
$r \leftarrow\$ \mathsf{PKE.Rnd}(\mathsf{pm})$
$\mathtt{R} \xleftarrow{\frown} r$
$c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}_i}(m;r)$
$\mathcal{C}_i \xleftarrow{\cup} c$
return $c$
```

```
Oracle $\mathcal{C}$
if challenged : return $\lightning$
$\mathsf{challenged} \leftarrow \mathsf{true}$
for $j \in [q]$
    if $\mathtt{K}[j] \in \mathcal{I} : \mathcal{J} \xleftarrow{\cup} j$
$\mathtt{M}^1 \leftarrow\$ \mathsf{S}(\mathtt{A},\mathtt{L},\mathcal{J},\mathtt{M}^0[\mathcal{J}])$
return $\mathtt{M}^b$
```

```
Oracle $\mathcal{T}(j)$
if challenged : return $\lightning$
$\mathcal{J} \xleftarrow{\cup} j$
return $\mathtt{M}^0[j]$
```

```
Oracle $\mathcal{S}(j)$
if challenged : return $\lightning$
$\mathcal{J} \xleftarrow{\cup} j$
return $(\mathtt{M}^0[j], \mathtt{R}[j])$
```

```
Oracle $\mathcal{R}(i)$
if challenged : return $\lightning$
$\mathcal{I} \xleftarrow{\cup} i$
return $\mathsf{sk}_i$
```

Fig. 8. A posteriori indistinguishability, also known as indistinguishability SOA, with bi-opening.

avoiding that the challenge bit leaks trivially. Moreover, unlike a priori indistinguishability, the experiment poses no limitations on which ciphertexts may be opened, while simultaneously retaining the preferred single-challenge-bit structure.

When formalizing an ISO notion, the way in which messages get sampled (and resampled) plays an important role and, as we will survey shortly, different abstractions are possible. Our notion of ISO-CCA⊛ uses BY12's idea of a fixed, stateful sampling algorithm M which, on adversarial input α, outputs a single message. We generalize BY12's notion by, in addition to sender and transmission opening, also allowing receiver opening and chosen ciphertext attacks. Furthermore, we make a syntactical distinction between the message sampler M and its resampler S. An adversary $\mathbb{A}$'s advantage (against a given PKE) will be relative to both this message sampler M and resampler S, as made explicit in Def. 4 below. This definition simultaneously serves to define weaker notions such as κ-ISO-CPA⋄, κ-ISO-CCA⋆, etc., by changing which oracles the adversary has access to.

Definition 4. *The κ-ISO-CCA⊛ advantage $\mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A})$ of an adversary $\mathbb{A}$ against public key encryption scheme $\mathsf{PKE}[\lambda]$, relative to message sampler M and resampler S, is the distinguishing advantage against the game $\mathsf{Exp}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A})$ (see Fig. 8).*

A run of the game proceeds in two stages. In the first stage, the adversary has access to its encryption oracle $\mathcal{E}$, as well as to any of its auxiliary oracles (to open and decrypt). Each encryption query will result in a single challenge ciphertext and the game keeps track of the corresponding encrypted messages (across queries) in the list $\mathtt{M}^0$. The opening oracle(s) will reveal some of $\mathtt{M}^0$, either directly through $\mathcal{T}$ or $\mathcal{S}$ or indirectly through $\mathcal{R}$; the shorthand $\mathtt{M}^0[\mathcal{J}]$, for $(\mathtt{M}^0[j])_{j\in\mathcal{J}}$, indicates the opened messages.

The second stage commences once the adversary calls its challenge oracle $\mathcal{C}$, which blocks access to all oracles apart from $\mathcal{D}$; the flag $\mathsf{challenged}$ enforces the access control. The challenge oracle itself creates a full list of resampled messages $\mathtt{M}^1$ using resampler S and returns either the real or resampled list, depending on the challenge bit.

To avoid trivial wins, $\mathtt{M}^1$ needs to be consistent with $\mathtt{M}^0$ relative to what an adversary trivially learns about the latter: the opening oracles reveal $\mathtt{M}^0[\mathcal{J}]$ (and $\mathcal{J}$); the queries α, collected in $\mathtt{A}$, carry information about the distribution; and we usually assume that ciphertext lengths leak message lengths, collected in the list $\mathtt{L}$. Hence, the resampler S is given the input $(\mathtt{A},\mathtt{L},\mathcal{J},\mathtt{M}^0[\mathcal{J}])$ to facilitate conditional resampling.

Ideally, the resampler S samples exactly from the same distribution as M, conditioned on S's input. Let $\mathring{\mathsf{S}}$ be this ideal, not necessarily efficient, resampler.

Definition 5 (Resampling error). *Let* M *be a stateful sampling algorithm with ideal resampler* $\mathring{\mathsf{S}}$*, and let* S *be a resampling algorithm. Let* $\ell \in \mathbb{Z}_{>0}$ *correspond to the number of* M *calls, and define the support* $\mathrm{Supp}_{\ell,\lambda}(\mathsf{M})$ *as the set of all tuples* $(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$ *subject to* $|\mathtt{A}| = |\mathtt{L}| = \ell$ *that may occur, i.e. for which there exists an adversary* $\mathbb{A}$ *and PKE scheme* PKE *such that the probability that* $\mathsf{Exp}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A})$ *results in* $(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$ *being input to* S *is non-zero. For* $(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}]) \in \mathrm{Supp}_{\ell,\lambda}(\mathsf{M})$*, let* $\delta_{\mathring{\mathsf{S}},\mathsf{S}}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$ *be the statistical distance between* $\mathring{\mathsf{S}}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$ *and* $\mathsf{S}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$*. Then the* resampling error *of* S *is*

$$\epsilon^{\ell}_{\mathring{\mathsf{S}},\mathsf{S}}(\lambda) = \max_{(\mathtt{A},\mathtt{L},\mathcal{J},\mathtt{M}^0[\mathcal{J}])\in\mathrm{Supp}_{\ell,\lambda}(\mathsf{M})} \delta_{\mathring{\mathsf{S}},\mathsf{S}}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])\,.$$

Remark 3. Although BY12 mention the requirement that resampling should result in a distribution statistically close to the ideal conditional resampler, their formalization differs in a number of aspects. Firstly, they define ideal resampling algorithmically with respect to the coins of the original samplers; our approach is more abstract and simply accepts that the relevant conditional distribution is well-defined. Secondly, they define ideal resampling on inputs outside $\mathrm{Supp}_{\ell,\lambda}(\mathsf{M})$ to yield $\perp$ and expect a resampler to behave the same; we do not pose any demands on the resampler in that case (consequently, their resampler must be able to distinguish $\mathrm{Supp}_{\ell,\lambda}(\mathsf{M})$ whereas ours does not). Thirdly, they define the resampling error in terms of a game played by an unbounded adversary, thus hiding the dependency on ℓ. As any statistical distance can be realized as distinguishing advantage by an unbounded adversary—and no unbounded adversary can do any better—their advantage would be more akin to taking the supremum over $\ell \in \mathbb{Z}_{>0}$ of $\epsilon^{\ell}_{\mathring{\mathsf{S}},\mathsf{S}}(\lambda)$. Finally, their distinguishing game randomly samples parameters according to the public key encryption scheme at hand, making their advantage essentially an expectation of the resampling error over the choice of said parameters.

Compared to BY12, our definition of resampling error Def. 5 has the benefit of being easier to work with, by avoiding unbounded adversaries that interact in a resampling experiment, and connecting instead to the intuitive language of statistical distance. Due to the slightly different choices made, somewhat surprisingly the definitions appear technically incomparable; we leave open the task of convincing consolidation of the concept of conditional resamplability, especially in relation to the notion of efficient conditional resamplability (see the "Asymptotic interpretation" paragraph later in this section).

Lemma 3. *Let* $\mathsf{PKE}[\lambda]$ *and* κ*-ISO-CCA*$\circledast$ *adversary* $\mathbb{A}$*, making at most* ℓ $\mathcal{E}$*-queries, be given. Let* M *be a message sampler with ideal resampler* $\mathring{\mathsf{S}}$ *and let* S *be an arbitrary resampler. Then*

$$\left|\mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A}) - \mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A})\right| \leq \epsilon^{\ell}_{\mathring{\mathsf{S}},\mathsf{S}}(\lambda)$$

Proof. The games $\mathsf{Exp}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A})$ (using S) and $\mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A})$ (using $\mathring{\mathsf{S}}$) are identical if $b = 0$ for both, and, if $b = 1$, they are identical until $\mathbb{A}$ makes its single call to the $\mathcal{C}$ oracle. Moreover, by definition of the resampling error, the statistical distance between $\mathcal{C}$'s output in the two experiments (using S versus $\mathring{\mathsf{S}}$) is at most $\epsilon^{\ell}_{\mathring{\mathsf{S}},\mathsf{S}}(\lambda)$, which therefore bounds the computational distinguishing advantage of any $\mathbb{A}$. □

Alternative samplers. BY12's stateful sampler differs from more common formulations of selective opening attacks, where the adversary can directly specify a stateless, multi-message sampler (or distribution), so $\mathcal{E}$ would return a vector of challenge ciphertexts. For convenience, we include some works from an a posteriori simulation-based perspective in the discussion below, see also Sect. 3.5.

Historically, sender openings and receiver openings had been studied separately, which was reflected in the samplers as well. For sender openings, typically only a single public key is created (so $\kappa = 1$), a vector of ℓ messages is sampled once and subsequently encrypted under this lone public key [10]; in contrast, for receiver openings, there are $\kappa > 1$ public keys and a vector of exactly κ messages is sampled and encrypted one-each under the κ public keys (receivers) of the system [63].

Definitional choices on how messages may be sampled, affect the strength of the resulting notion. There are two main choices on how sampling is modelled, leading to four possible styles of sampler.

The first choice is between stateless and stateful samplers, where stateful sampling allows message dependencies across calls to the $\mathcal{E}$-oracle. Intuitively, stateless sampling allows a hybrid argument to reduce (non-tightly) the number of $\mathcal{E}$-queries to 1 [87], whereas for stateful sampling such a hybrid

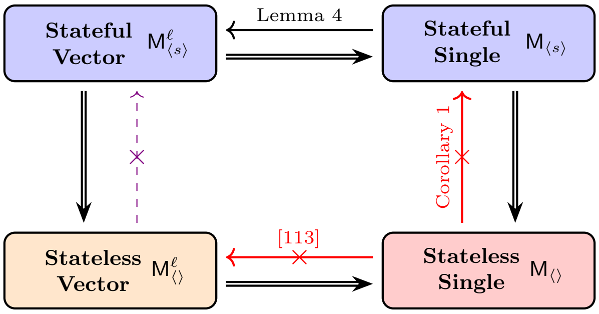
Stateful Vector $\mathsf{M}^{\ell}_{\langle s\rangle}$
Lemma 4
Stateful Single $\mathsf{M}_{\langle s\rangle}$
Corollary 1
[113]
Stateless Vector $\mathsf{M}^{\ell}_{\langle\rangle}$
Stateless Single $\mathsf{M}_{\langle\rangle}$

Fig. 9. Relations between notions of SOA employing different styles of samplers. (Double arrows = trivial; purple-dashed = conjectured).

argument looks challenging (and is in fact not possible, see below). We can make the distinction between stateless and stateful samplers explicit by their subscript, writing M respectively $\mathsf{M}_{\langle s\rangle}$.

The second choice asks whether, for each sampling, a single message should be returned, or a message vector of length ℓ. Returning a vector allows for a joint distribution over the messages even when sampling is stateless. We can make the distinction between single-message and vector samplers explicit by their superscript, writing $\mathsf{M}_{\langle s\rangle}$ respectively $\mathsf{M}^{\ell}_{\langle s\rangle}$. For vector samplers, Fig. 8's challenge oracle $\mathcal{E}$ takes as its first input a list $\mathtt{I}$ of key handles that the adversary wishes to challenge ($\mathcal{E}$ then samples $\ell = |\mathtt{I}|$ messages and encrypts each sampled message under the corresponding public key). Existing vector sample notions often put restrictions on which $\mathtt{I}$ are allowed, for instance each key handle exactly once [63] or an ℓ-fold repetition of a single key handle [113].

The four resulting notions and how they relate are presented in Fig. 9. We will argue that stateful samplers yield potentially stronger notions of security than the more common (stateless) vector sampling.

Our first observation is that for stateful samplers it is irrelevant whether one message is sampled per oracle call or a fixed number ℓ: the two resulting notions are tightly equivalent. Showing that stateful vector sampling is at least as general as stateful single-message sampling is of course trivial when considering vectors of length $\ell = 1$ (and for $\ell > 1$ one can define $\mathsf{M}^{\ell}_{\langle s\rangle}$ using its first element to encode $\mathsf{M}_{\langle s\rangle}$ and use some fixed, known message for the remaining $\ell - 1$ elements). We show the converse, that stateful single-message sampling is as general as stateful vector sampling, next.

Lemma 4. *Let $\mathsf{PKE}[\lambda]$ be given, let $\mathsf{M}^{\ell}_{\langle s\rangle}$ be a stateful message sampler (with ideal resampler $\mathring{\mathsf{S}}$) returning ℓ messages per call, and let $\mathsf{M}_{\langle s'\rangle}$ be the stateful message sampler sampling from the same distribution as $\mathsf{M}^{\ell}_{\langle s\rangle}$ by calling $\mathsf{M}^{\ell}_{\langle s\rangle}$ initially and then after every ℓth call, yet returning only a single message per call (so for $\ell - 1$ of its calls, $\mathsf{M}_{\langle s'\rangle}$ ignores its input). Let $\mathring{\mathsf{S}}'$ be $\mathsf{M}_{\langle s'\rangle}$'s ideal resampler. Then there is a type-preserving black-box reduction $\mathbb{B}$ such that, for all $\mathbb{A}$,*

$$\mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M}^{\ell}_{\langle s\rangle},\mathring{\mathsf{S}}}(\mathbb{A}) \leq \mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M}_{\langle s'\rangle},\mathring{\mathsf{S}}'}(\mathbb{B})\,.$$

The runtime of $\mathbb{B}$ is upper bounded by that of $\mathbb{A}$, except that if $\mathbb{A}$ makes q encryption oracle calls, then $\mathbb{B}$ makes $\ell \cdot q$ encryption oracle calls.

Proof (sketch). Whenever $\mathbb{A}$ calls $\mathcal{E}(\mathtt{I}, \alpha)$, $\mathbb{B}$ calls its own encryption oracle ℓ times sequencing through $\mathtt{I}$ for the i inputs and using the same input α throughout, returning the resulting ℓ-length ciphertext vector. This simulation is perfect as any relations between the sampled messages will be preserved using the sampler's internal state. □

As stateful vector-sampling $\mathsf{M}^{\ell}_{\langle s\rangle}$ trivially implies stateless vector-sampling $\mathsf{M}^{\ell}_{\langle\rangle}$ (by ignoring the state), as a corollary stateful single-message sampling $\mathsf{M}_{\langle s\rangle}$ implies stateless vector sampling $\mathsf{M}^{\ell}_{\langle\rangle}$ (top-right to bottom-left in Fig. 9).

In contrast, stateless single-message sampling does not imply vector-sampling, which intuitively follows from a hybrid argument (see Sect. 4.1).

ISO implies IND. We next present a concrete variation of BY12's result that ISO security implies IND security [16, Thm. 4.3]. Specifically, we show that, for a suitably chosen message sampler (see Lemma 5),

κ-ISO-CCA$\star$ implies κ-IND-CCA$\star$ with only a factor 2 loss (the CPA case follows from the reduction's type-preservation). In contrast, BY12 only showed that single-sample κ-IND-CPA$\diamond$ implies single challenge IND-CPA (with a factor 2 loss), thus our result is both tighter for multi-challenge situations and more general by allowing additional oracles.

Conversely, general separations in the form of counterexamples are known (see Sect. 4 for details), indicating that a posteriori indistinguishability is a strictly stronger notion than a priori indistinguishability in the presence of receiver openings.

Lemma 5. *Let $\mathsf{M}_{\langle s \rangle}$ be the sampler that as input α only takes message pairs (m_0, m_1) subject to both $|m_0| = |m_1|$ and $m_0 \neq m_1$. On first invocation (when $s = \varepsilon$), it draws $s \leftarrow\$ \{0,1\}$ and, on all invocations, on input $\alpha = (m_0, m_1)$, it returns m_s. Let $\mathring{\mathsf{S}}$ be its ideal resampler.*

Consider S that on input $(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$, first checks whether $\mathtt{M}^0[\mathcal{J}]$ is non-empty. If so, it contains at least one opened m_s drawn from (m_0, m_1) satisfying $m_0 \neq m_1$ and S sets $s' \leftarrow s$; otherwise it draws $s' \leftarrow\$ \{0,1\}$. Finally, S sets and returns $\mathtt{M}^1 \leftarrow \mathtt{A}_{s'}$, where $\mathtt{A}$ is interpreted as $(\mathtt{A}_0, \mathtt{A}_1)$ based on the special form of the α.

Then $\mathsf{S} = \mathring{\mathsf{S}}$.

Proof. By inspection.

Theorem 4. *Let $\mathsf{PKE}[\lambda]$ be given and let $\mathsf{M}_{\langle s \rangle}$ be as given in Lemma 5. Then there is a type-preserving black-box reduction $\mathbb{B}_{\text{iso}}$ such that, for all $\mathbb{A}_{\text{ind}}$,*

$$\mathsf{Adv}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}}) \leq 2 \cdot \mathsf{Adv}^{\kappa\text{-iso-cca}\star}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{B}_{\text{iso}}) \,.$$

The runtime of $\mathbb{B}_{\text{iso}}$ is upper bounded by that of $\mathbb{A}_{\text{ind}}$.

Proof. Without loss of generality, we may assume that $\mathbb{A}_{\text{ind}}$ does not call $\mathcal{E}_{\mathbb{A}}(i, m_0, m_1)$ with either $|m_0| \neq |m_1|$ or $m_0 = m_1$, nor does it corrupt and challenge on the same key. Technically, one can create an intermediate reduction $\mathbb{B}_{\text{ind}}$ that runs $\mathbb{A}_{\text{ind}}$ and, playing the same game, forwards everything to its own oracles but those pointless $\mathcal{E}$ calls, which it can easily simulate by responding with $\lightning$ and $\mathsf{PKE.Enc}(m_0)$, respectively; if $\mathbb{A}_{\text{ind}}$ makes a call that would trigger $\mathcal{K} \cap \mathcal{I} \neq \emptyset$ then $\mathbb{B}_{\text{ind}}$ samples $\hat{b} \leftarrow\$ \{0,1\}$ and terminates with that bit. By inspection, $\mathbb{B}_{\text{ind}}$'s advantage is at least $\mathbb{A}_{\text{ind}}$'s (it could be larger for instance when $\mathbb{A}_{\text{ind}}$ correctly guesses the bit, while having both corrupted a key and challenged it with $m_0 = m_1$).

Let $\mathbb{B}_{\text{iso}}$ be such that if $\mathbb{A}_{\text{ind}}$ calls $\mathcal{E}_{\mathbb{A}}(i, m_0, m_1)$ subject to both $|m_0| = |m_1|$ and $m_0 \neq m_1$, it sets $\alpha = (m_0, m_1)$ and calls $\mathcal{E}_{\mathbb{B}}(i, \alpha)$, returning the resulting c; if $\mathbb{A}_{\text{ind}}$ calls any other oracles, $\mathbb{B}_{\text{iso}}$ forwards the call and returns the result. When $\mathbb{A}_{\text{ind}}$ halts with a guess $\hat{s}$, $\mathbb{B}_{\text{iso}}$ calls $\mathcal{C}$ and receives $\mathtt{M}^b$. Since $\mathbb{B}_{\text{iso}}$ can maintain its own perfect copy of $\mathtt{A}$, it can check whether $\mathtt{M}^b = \mathtt{A}_{\hat{s}}$. If so, $\mathbb{B}_{\text{iso}}$ halts with output $\hat{b} = 0$ (indicating a guess that the returned messages were the real ones), otherwise $\mathbb{B}_{\text{iso}}$ halts with output $\hat{b} = 1$.

We can rephrase $\mathbb{B}_{\text{iso}}$'s distinguishing advantage

$$\mathsf{Adv}^{\kappa\text{-iso-cca}\star}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{B}_{\text{iso}}) = \Pr\left[\hat{b} = 0 \,\middle|\, b = 0\right] - \Pr\left[\hat{b} = 0 \,\middle|\, b = 1\right]$$

and analyse each term individually. Based on $\mathbb{B}_{\text{iso}}$'s description, the event $\hat{b} = 0$ is equivalent to the event $\mathtt{M}^b = \mathtt{A}_{\hat{s}}$.

If $b = 0$, then $\mathtt{M}^0 = \mathtt{A}_s$, so the first term is equivalent to $\Pr[\mathtt{A}_{\hat{s}} = \mathtt{A}_s \mid b = 0]$. Given that $\mathtt{A}_0 \neq \mathtt{A}_1$ (by our assumption on $\mathbb{A}_{\text{ind}}$ not making queries with $m_0 = m_1$) we can simplify further to $\Pr[\hat{s} = s \mid b = 0]$. At this point, the conditional $b = 0$ becomes irrelevant as it is independent of both $\hat{s}$ and s (jointly). Finally, $\Pr[\hat{s} = s]$ equals $\Pr\left[\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}})\right]$ as, by design of $\mathbb{B}_{\text{iso}}$ and M, $\mathbb{A}_{\text{ind}}$ is provided with an environment that perfectly matches that of $\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\cdot)$, where M's bit s plays the role of the bit $\mathbb{A}_{\text{ind}}$ has to guess. Thus,

$$\Pr\left[\hat{b} = 0 \,\middle|\, b = 0\right] = \Pr\left[\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}})\right] .$$

If $b = 1$, then $\mathtt{M}^1 = \mathtt{A}_{s'}$, so the second term is equivalent to $\Pr[\mathtt{A}_{\hat{s}} = \mathtt{A}_{s'} \mid b = 1]$, which simplifies to $\Pr[\hat{s} = s' \mid b = 1]$. As we assumed $\mathbb{A}_{\text{ind}}$ maintained $\mathcal{K} \cap \mathcal{I} = \emptyset$ as invariant (for its game), it follows that, for $\mathbb{B}_{\text{iso}}$'s game, $\mathcal{J} = \emptyset$ and hence $\mathring{\mathsf{S}}$ will have drawns s' uniformly at random, independently of $\hat{s}$. Thus,

$$\Pr\left[\hat{b} = 0 \,\middle|\, b = 1\right] = \frac{1}{2} \,.$$

Putting the pieces together, we obtain

$$2 \cdot \mathsf{Adv}^{\kappa\text{-iso-cca}\star}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{B}_{\text{iso}}) = 2 \cdot \Pr\Big[\mathsf{Exp}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}})\Big] - 1 \geq \mathsf{Adv}^{\kappa\text{-ind-cca}\star}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}}),$$

where the final inequality takes into account the intermediate $\mathbb{B}_{\text{ind}}$ reduction to justify our assumption on $\mathbb{A}_{\text{ind}}$'s behaviour (from the beginning of the proof).

□

Asymptotic interpretation. The advantage specified in Def. 4 leads to two potential asymptotic interpretations: weak ISO and "full" ISO. Full ISO security is achieved for a scheme PKE if, for all PPT $\mathbb{A}$ and all PPT M, the advantage $\mathsf{Adv}^{\text{iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A})$ is negligible in λ (where of course security for CPA and other openings are defined analogously). Importantly, the quantification over the message samplers M is irrespective of how efficient $\mathring{\mathsf{S}}$ might be implemented. For weak ISO security, the class of samplers is restricted to those for which there exists an efficient resampler S, that is a PPT S such that for all polynomials poly the maximum resampling error $\max_{\ell \leq \mathsf{poly}(\lambda)} \epsilon^{\ell}_{\mathring{\mathsf{S}},\mathsf{S}}(\lambda)$ is negligible (in λ).

Clearly, restricting the class of samplers M as for weak ISO results in a significantly weaker notion, as further explored by Böhl et al. [19]. There are in fact further gradations in the security notion depending on how "efficient resampling" is formalized. We already explained the subtle differences between our formalization of the sampling error and BY12's. Other works [19, 63] require the efficient resampler to have zero statistical distance; such a stricter requirement on the resampler leads to a potentially weaker notion of security.

How ISO is formalized asymptotically (full versus weak) affects how the alternative types of sampling relate to each other, specifically when interpreting Lemma 4. For Full ISO, an almost immediate consequence of Lemma 4 is that stateful vector sampling and stateful single-message sampling are equivalent: if $\mathsf{M}^{\ell}_{\langle s \rangle}$ is efficient (PPT), then so is the derived $\mathsf{M}_{\langle s' \rangle}$ (and vice versa), without having to worry about the efficiency of the ideal resampler. However, for weak ISO, resampling efficiency comes into play and even if $\mathsf{M}^{\ell}_{\langle s \rangle}$'s ideal resampler $\mathring{\mathsf{S}}$ has an efficient resampler S, there is no guaranteee that the same holds for $\mathsf{M}_{\langle s' \rangle}$'s ideal resampler $\mathring{\mathsf{S}}'$. The main technical difficulty here is that $\mathring{\mathsf{S}}'$ also needs to deal with message vectors that are not a multiple of ℓ messages and it is possible that "partial" vectors are harder to resample efficiently (for instance, if the lengths of the missing messages might assist the resampling).

We conclude that, in an asymptotic setting, additional definitional choices increase the number of possible notions and, lacking clear use cases, determining which notion makes most sense is tricky. For instance, Full-ISO has mostly fallen out of favour, as allowing inefficient resampling seems to yield an artificially strong notion of security, with currently no known schemes achieving it (see also Sect. 4.5).

Luckily, even the weakest version of weak ISO allows us to conclude from Thm. 4 that ISO security implies IND security. Both the message sampler and its ideal resampler used in that theorem (as specified in Lemma 5) are clearly efficient and hence included in any reasonably class of message samplers (used to define a flavour of weak ISO security).

Further remarks. ISO-CPA was initially defined relative to message distributions independent of the public key [10]. Böhl et al. [19] noted that as a result, the notion did not imply IND-CPA; with the converse being open, the notions seemed incomparable. Allowing message samplers an input α resolves this issue [16].

Given that (κ, β)-IND-CCA$\circledast$ is implied by κ-IND-CCA$\star$ with a β loss (Thm. 3), we may conclude that κ-ISO-CCA$\star$ implies (κ, β)-IND-CCA$\circledast$ with a $2 \cdot \beta$ security loss. In fact, a tighter reduction that loses only a factor 2 is possible by starting from κ-ISO-CCA$\circledast$ (instead of κ-ISO-CCA$\star$ as in Thm. 4), see Thm. 8 (App. A) for details.

3.5 A Posteriori Simulatability with Selective Opening (SSO)

As explained in Sect. 3.2, the main idea of a posteriori simulatability is to have a simulator Sim simulate the computations of $\mathbb{A}$ without seeing the ciphertexts, thus capturing the idea that the ciphertexts leak nothing about the plaintexts. Unlike the ISO notion from the previous section, for SSOthe relevant security game does not contain a conditional resampling phase, making the notion suitable when such resampling is problematic. A potential downside is that the presence of simulators and distinguishers can complicate reductions compared to indistinguishability-based alternatives.

```
Experiment $\mathsf{Exp}^{\kappa\text{-sso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{A},\mathbb{D})$
$b \leftarrow\!\$\ \{0,1\}$
$s \leftarrow \varepsilon$
$\mathsf{pm} \leftarrow\!\$\ \mathsf{PKE.Pm}(\lambda)$
$\forall_{i\in[\kappa]}(\mathsf{pk}_i,\mathsf{sk}_i) \leftarrow\!\$\ \mathsf{PKE.Kg}(\mathsf{pm})$
if $b = 0$ :
    $\mathsf{out} \leftarrow\!\$\ \mathbb{A}^{\mathcal{E},\mathcal{D},\mathcal{S},\mathcal{R},\mathcal{T}}(\mathsf{pk}_1,\ldots,\mathsf{pk}_\kappa)$
    for $j \in [|\mathtt{K}|]$
        if $\mathtt{K}[j] \in \mathcal{I} : \mathcal{J} \xleftarrow{\cup} j$
else :
    $\mathsf{out} \leftarrow\!\$\ \mathsf{Sim}^{\mathcal{E},\mathcal{T}}(\mathsf{pm})$
$\hat{b} \leftarrow\!\$\ \mathbb{D}(\mathsf{pm},\mathtt{A},\mathtt{M},\mathcal{J},\mathsf{out})$
return $b = \hat{b}$
```

```
Oracle $\mathcal{E}(i,\alpha)$
$\mathtt{K} \xleftarrow{\frown} i, \mathtt{A} \xleftarrow{\frown} \alpha$
$m \leftarrow\!\$\ \mathsf{M}_{\langle s\rangle}(\alpha)$
$\mathtt{M} \xleftarrow{\frown} m$
$r \leftarrow\!\$\ \mathsf{PKE.Rnd}(\mathsf{pm})$
$c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}_i}(m;r)$
$\mathtt{R} \xleftarrow{\frown} r, \mathcal{C}_i \xleftarrow{\cup} c$
if $b = 0$ : return $c$
else : return $|m|$
```

```
Oracle $\mathcal{D}(i,c)$
if $c \in \mathcal{C}_i$ : return $\lightning$
$m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}_i}(c)$
return $m$
```

```
Oracle $\mathcal{T}(j)$
$\mathcal{J} \xleftarrow{\cup} j$
return $\mathtt{M}[j]$
```

```
Oracle $\mathcal{S}(j)$
$\mathcal{J} \xleftarrow{\cup} j$
return $(\mathtt{M}[j],\mathtt{R}[j])$
```

```
Oracle $\mathcal{R}(i)$
$\mathcal{I} \xleftarrow{\cup} i$
return $\mathsf{sk}_i$
```

Fig. 10. κ-SSO-CCA⊛ security game, for which a distinguisher $\mathbb{D}$ is tasked with guessing whether it received the view of adversary $\mathbb{A}$ playing the real game or a simulated view (by Sim).

In the CPA setting, with either sender or receiver opening SSO-CPA is known to be strictly stronger (and therefore harder to achieve) than ISO-CPA [63], and so we place a posteriori simulatability above a priori indistinguishability in our hierarchy (Fig. 1). With only transmission openings present, the notion is tightly implied by a priori indistinguishability [16], see Sect. 4.2.

On the other hand, in the CCA setting the relationship between a posteriori simulatability and a posteriori indistinguishability remains largely open: in particular, while we conjecture that SSO-CCA implies ISO-CCA with any (matching) openings (see Fig. 16), a proof thereof has to the best of our knowledge yet to appear, see Open Problem 5.

Our formalization (Fig. 10) is based on BY12's SSO-CPA⊙ notion, where we added multiple users, receiver and thus bi-openings, as well as a CCA oracle. The joint advantage of adversary $\mathbb{A}$ and distinguisher $\mathbb{D}$ is therefore relative to the stateful, single-message sampler M (see Sect. 3.4) as well as the simulator Sim. (The equivalent of Lemma 4, extending to vector samplers, holds also in the current SSO setting.)

Definition 6. *The* κ-SSO-CCA⊛ *advantage* $\mathsf{Adv}^{\kappa\text{-sso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{A},\mathbb{D})$ *of an adversary* $\mathbb{A}$ *and distinguisher* $\mathbb{D}$ *against public key encryption scheme* $\mathsf{PKE}[\lambda]$*, relative to message sampler* M *and simulator* Sim*, is the distinguishing advantage against the game* $\mathsf{Exp}^{\kappa\text{-sso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{A},\mathbb{D})$ *(see Fig. 10).*

And so our formalization of SSO comes with three players: the adversary $\mathbb{A}$, the simulator Sim, and the distinguisher $\mathbb{D}$. In the real game ($b = 0$), $\mathbb{A}$ gets access to the public keys and all oracles, and the encryption oracle $\mathcal{E}$ returns encryptions of the sampled messages. Once the adversary is content, it halts with some output out. Intuitively, $\mathbb{A}$'s goal is to make it as easy as possible for $\mathbb{D}$ to guess correctly, and without loss of generality $\mathbb{A}$ will simply output its view, i.e. the transcript of its interactions with the game as well as its internal randomness.

In the ideal game ($b = 1$), the game instead calls Sim who does not get access to the ciphertexts, and whose goal is to fool the distinguisher, and so Sim will want to do everything in its power to make out look like it originated from an adversary who did have access to the real ciphertexts. Since we usually assume that ciphertexts leak message lengths, the simulator does receive message lengths (in place of ciphertexts) to facilitate its job; additionally, it can open individual messages through the transmission opening oracle. The sender and receiver oracles are not present as, in the ideal game, there are no keys to be opened, nor is there any randomness sampled by the encryption oracle.

Eventually, $\mathbb{D}$ makes a decision on whether out was produced by someone with access to the real ciphertexts and opening oracles or not, halting with a guess $\hat{b} = 0$ for "real", or $\hat{b} = 1$ for "ideal". Crucially, the distinguisher receives all the sampled messages $\mathtt{M}$ directly from the experiment, in addition to the real parameters pm, sampler inputs $\mathtt{A}$ and list of opened challenges $\mathcal{J}$.

Remark 4. The strength of the notion is governed by which of these additional inputs $\mathbb{D}$ receives directly from the game, as those inputs effectively bind the simulator to be honest. For instance, denying the

```
Experiment $\mathsf{Exp}^{\kappa\text{-sso}'\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{A}, \mathbb{D})$
$b \leftarrow\$ \{0,1\}$
$s \leftarrow \varepsilon$
$\mathsf{pm} \leftarrow\$ \mathsf{PKE.Pm}(\lambda)$
$\forall_{i \in [\kappa]} (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow\$ \mathsf{PKE.Kg}(\mathsf{pm})$
$\mathbf{if}\ b = 0 :$
    $\mathsf{out} \leftarrow\$ \mathbb{A}^{\mathcal{E},\mathcal{D},\mathcal{S},\mathcal{R},\mathcal{T}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
$\mathbf{else}\ :$
    $\mathsf{out} \leftarrow\$ \mathsf{Sim}^{\mathcal{E},\mathcal{R}',\mathcal{T}}(\mathsf{pm})$
$\hat{b} \leftarrow\$ \mathbb{D}(\mathsf{pm}, \mathtt{K}, \mathtt{A}, \mathtt{M}, \mathcal{I}, \mathcal{J}, \mathsf{out})$
$\mathbf{return}\ b = \hat{b}$
```

```
Oracle $\mathcal{E}(i, \alpha)$
$\mathtt{K} \overset{\frown}{\longleftarrow} i, \mathtt{A} \overset{\frown}{\longleftarrow} \alpha$
$m \leftarrow\$ \mathsf{M}_{\langle s \rangle}(\alpha)$
$\mathtt{M} \overset{\frown}{\longleftarrow} m$
$r \leftarrow\$ \mathsf{PKE.Rnd}(\mathsf{pm})$
$c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}_i}(m; r)$
$\mathtt{R} \overset{\frown}{\longleftarrow} r, \mathcal{C}_i \overset{\cup}{\longleftarrow} c$
$\mathbf{if}\ b = 0 : \mathbf{return}\ c$
$\mathbf{else\ if}\ i \in \mathcal{I} : \mathbf{return}\ m$
$\mathbf{else}\ : \mathbf{return}\ |m|$
```

```
Oracle $\mathcal{D}(i, c)$
$\mathbf{if}\ c \in \mathcal{C}_i : \mathbf{return}\ \lightning$
$m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}_i}(c)$
$\mathbf{return}\ m$
```

```
Oracle $\mathcal{T}(j)$
$\mathcal{J} \overset{\cup}{\longleftarrow} j$
$\mathbf{return}\ \mathtt{M}[j]$
```

```
Oracle $\mathcal{S}(j)$
$\mathcal{J} \overset{\cup}{\longleftarrow} j$
$\mathbf{return}\ (\mathtt{M}[j], \mathtt{R}[j])$
```

```
Oracle $\mathcal{R}(i)$
$\mathcal{I} \overset{\cup}{\longleftarrow} i$
$\mathbf{return}\ \mathsf{sk}_i$
```

```
Oracle $\mathcal{R}'(i)$
$\mathcal{I} \overset{\cup}{\longleftarrow} i$
$\mathbf{for}\ j \in [|\mathtt{K}|]$
    $\mathbf{if}\ \mathtt{K}[j] = i : \mathtt{L} \overset{\frown}{\longleftarrow} \mathtt{M}[j]$
$\mathbf{return}\ \mathtt{L}$
```

Fig. 11. An alternative κ-SSO$'$-CCA$\circledast$ security game, where the simulator Sim's behaviour is bound by additional direct inputs (by the game) to the distinguisher $\mathbb{D}$. The simulator has access to its $\mathcal{T}$ oracle whenever the adversary has access to $\mathcal{S}$ or $\mathcal{T}$ (so the $\diamond$, $\odot$, and $\circledast$ notions) and to $\mathcal{R}'$ whenever $\mathbb{A}$ has access to $\mathcal{R}$ (so the $\star$ and $\circledast$ notions).

distinguisher access to the list $\mathcal{J}$ of opened challenges yields a vacuous notion: a simulator could run a copy of the experiment with $\mathbb{A}$ and, whenever $\mathbb{A}$ makes an $\mathcal{E}$-query, Sim would call its own $\mathcal{E}$-oracle, immediately open that challenge using $\mathcal{T}$ to receive the underlying message, and then simply encrypt that message to obtain a ciphertext to return to $\mathbb{A}$ (and eventually Sim uses $\mathbb{A}$'s out as output). Conditioned on only $\mathsf{pm}, \mathtt{A}$, and $\mathtt{M}$, this simulator's out will be identically distributed to that of a real adversary.

For receiver openings, our mechanism only provides the distinguisher with the indices of the messages that were opened as a logical consequence of revealing private keys. Instead, one could provide the distinguisher directly with the list $\mathcal{I}$ of the opened keys [63], plus the information ($\mathtt{K}$) needed to identify which ciphertexts were encrypted under which key (for past single-shot, stateless vector sampling formalizations, restrictions on $\mathcal{E}$ typically made $\mathtt{K}$ superfluous). In the case of bi-openings, one would then provide both the list $\mathcal{I}$ and $\mathcal{J}$ to an adversary [91]. For completeness, we have included the alternative notion κ-SSO$'$-CCA$\circledast$ in Fig. 11, that captures this finer-grained mechanism in the context of adaptive, stateful sampling.

Considering sender openings only, Bellare, Hofheinz and Yilek [10] opt for another mechanism instead: their advantage statement is parameterized by the number of openings allowed and both the adversary and simulator are restricted to making at most that many openings (the restriction on the number of openings made by the simulator is not made explicit in the published versions [10, 69], but follows from one of the full versions [14]). Inspired by this mechanism, one could in our formalism replace the distinguisher's input $\mathcal{J}$ by only the cardinality $|\mathcal{J}|$ of said list; effectively, it allows the simulator a bit more freedom to deviate from what an adversary is doing, but not enough to render the notion vacuous as above.

Not providing $\mathbb{D}$ with the parameters pm gives an alternative, weaker notion of SSO [113]: since the simulator is now free to produce the parameters itself, it opens for strategies that e.g. involve inserting trapdoors in pm. Conversely, providing the public keys pk_i to the distinguisher takes away the ability for a simulator to use 'fake' keys in its output out; in that case, for the notion to make sense, Sim would have to be provided the pk_i as input, as well as oracle access to $\mathcal{R}$, furthermore the list $\mathcal{I}$ of corrupted parties would then be a secondary additional input to the distinguisher. We let our notion of a posteriori simulatability be one in which $\mathbb{D}$, and thus also Sim, are given pm (matching BY12), but not the pk_i.

Remark 5. A further weakening of SSO$\circledast$ (and SSO$\star$) is possible by restricting adversarial access to the challenge oracle $\mathcal{E}$ on already corrupted key handles, by adding a line to the top of $\mathcal{E}$ that, whenever $i \in \mathcal{I}$, immediately return $\lightning$; we denote this notion by SSO*. Past definitions of SSO with receiver openings often implicitly included this restriction by virtue of being staged and using stateless vector sampling: for

```
Distinguisher $\mathbb{D}_{\mathrm{sso}}(\mathtt{A}, \mathtt{M}, \mathcal{J}, \mathsf{out})$
$\mathtt{M}^0 \leftarrow \mathtt{M},\ \forall_{i \in |\mathtt{M}|} \mathtt{L}[i] \leftarrow |\mathtt{M}[i]|$
$\mathtt{M}^1 \leftarrow\!\!\$\ \mathring{\mathsf{S}}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$
$d \leftarrow\!\!\$\ \{0, 1\}$
$\mathsf{state} \leftarrow \mathsf{out}$
Run $\mathbb{A}^2_{\mathrm{iso}}(\mathtt{M}^d)$ using $\mathsf{state}$
if $\mathbb{A}^2_{\mathrm{iso}}$ terminates within time $T$ with output $\hat{d}$
    $\hat{b} \leftarrow \neg(d = \hat{d})$
else
    $\hat{b} \leftarrow 1$
return $\hat{b}$
```

Fig. 12. The distinguisher $\mathbb{D}_{\mathrm{sso}}$ of Thm. 5.

instance, an adversary would have a single shot to receive a vector of challenge ciphertexts after which it could non-adaptively corrupt a set of keys [63].

There is no obvious reason for such a restriction, although intuitively challenging on an already corrupted key handle seems of little benefit to an adversary as the messages involved are not considered confidential: the call's corresponding index j is guaranteed to end up in $\mathcal{J}$, so a simulator Sim can access the message as well. Yet, the newly sampled message can depend on the sampler's state, and corrupting a private key possibly allows an adversary to trigger a subsequent call to the sampler that reveals information about its state relating to past, unopened messages.

Curiously, for NCE we will soon see (Def. 7) that the restriction is inevitable and, as a consequence, we can only show that NCE implies this weaker, restricted version of SSO.

Open Problem 4. *How do notions of* SSO, SSO$'$, *and* SSO* *relate?*

SSO implies ISO. A posteriori simulatability with selective opening implies a posteriori indistinguishability with selective opening in the CPA-setting, as shown by BY12 for transmission and sender openings [16, Theorem 3.3]. We next provide a concrete, updated statement to include receiver and bi-openings. Our proof corrects a subtle mistake in BY12's original, as explained inline.

Theorem 5. *Let* $\mathsf{PKE}[\lambda]$ *be given, and let* M *be a sampler with ideal resampler* $\mathring{\mathsf{S}}$. *Then, for any adversary* $\mathbb{A}_{\mathrm{iso}}$ *playing* $\mathsf{Exp}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\cdot)$ *and making at most q challenge queries, there exist (non black-box) type-preserving reduction* $\mathbb{B}_{\mathrm{sso}}$ *and distinguisher* $\mathbb{D}_{\mathrm{sso}}$ *such that, for all simulators* Sim,

$$\mathsf{Adv}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\mathrm{iso}}) \leq 2 \cdot \mathsf{Adv}^{\kappa\text{-sso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{B}_{\mathrm{sso}}, \mathbb{D}_{\mathrm{sso}}) \,,$$

The combined runtime of $\mathbb{B}_{\mathrm{sso}}$ *and* $\mathbb{D}_{\mathrm{sso}}$ *is upper bounded by twice that of* $\mathbb{A}_{\mathrm{iso}}$ *plus that of one call to* $\mathring{\mathsf{S}}$ *and some small overhead.*

Proof. Let T be an upper bound on the time $\mathbb{A}_{\mathrm{iso}}$ can take in the experiment $\mathsf{Exp}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\cdot)$. Without loss of generality, we assume $\mathbb{A}_{\mathrm{iso}}$ always calls $\mathcal{C}$ exactly once, so we can split $\mathbb{A}_{\mathrm{iso}}$ in two: let $\mathbb{A}^1_{\mathrm{iso}}$ be $\mathbb{A}_{\mathrm{iso}}$ up until the point it calls $\mathcal{C}$ and denote with state its state at that point; furthermore, let $\mathbb{A}^2_{\mathrm{iso}}$ be $\mathbb{A}_{\mathrm{iso}}$ after it receives the response $\mathtt{M}^b$ of calling $\mathcal{C}$, continuing from state state.

Define $\mathbb{B}_{\mathrm{sso}}$ as running $\mathbb{A}^1_{\mathrm{iso}}$, forwarding all oracles appropriately and, once $\mathbb{A}^1_{\mathrm{iso}}$ terminates (by making its $\mathcal{C}$ call) with state state, $\mathbb{B}_{\mathrm{sso}}$ sets $\mathsf{out} \leftarrow \mathsf{state}$ and terminates with output out.

The distinguisher (Fig. 12) receives the real message vector $\mathtt{M}^0$ as its input and uses the ideal resampling algorithm $\mathring{\mathsf{S}}$ to produce a resampled message vector $\mathtt{M}^1$, draws a bit d, and, depending on its value, runs $\mathbb{A}^2_{\mathrm{iso}}$ on either the real or resampled message vector, using its own input out as $\mathbb{A}^2_{\mathrm{iso}}$'s initial state. As there is no a priori guarantee that the distinguisher's input out is a valid state (namely one that could have been generated by $\mathbb{A}^1_{\mathrm{iso}}$), the runtime of $\mathbb{A}^2_{\mathrm{iso}}$ on $\mathsf{state} = \mathsf{out}$ is not bound by that of $\mathbb{A}_{\mathrm{iso}}$ itself, which is why the distinguisher checks $\mathbb{A}^2_{\mathrm{iso}}$'s runtime explicitly. (BY12, using an asymptotic framework, implicitly and erroneously assume that $\mathbb{A}^2_{\mathrm{iso}}$ inherits $\mathbb{A}_{\mathrm{iso}}$'s polynomial runtime, even when run on simulated states.) If $\mathbb{A}^2_{\text{iso}}$ does finish in time then $\mathbb{D}_{\text{sso}}$'s output depends on whether $\mathbb{A}^2_{\text{iso}}$ guessed d correctly or not.

Let b denote the challenge bit of the $\mathsf{Exp}^{\kappa\text{-sso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\cdot,\cdot)$ game, then $\mathbb{B}_{\text{sso}}$ and $\mathbb{D}_{\text{sso}}$ win with the following probability.

$$\Pr\left[\mathsf{Exp}^{\kappa\text{-sso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{B}_{\text{sso}},\mathbb{D}_{\text{sso}})\right] = \frac{1}{2}\left(\Pr\left[b=\hat{b}\;\middle|\;b=0\right] + \Pr\left[b=\hat{b}\;\middle|\;b=1\right]\right) .$$

For the first term, $\mathbb{B}_{\text{sso}}$ and $\mathbb{D}_{\text{sso}}$ essentially run $\mathbb{A}_{\text{iso}}$ start to finish with $\mathbb{A}^1_{\text{iso}}$'s finishing state equaling $\mathbb{A}^2_{\text{iso}}$ starting state. Thus, $\mathbb{D}_{\text{sso}}$ will never have to time-cap $\mathbb{A}^2_{\text{iso}}$ and (with some logic deciphering) $\Pr\left[b=\hat{b}\;\middle|\;b=0\right] = \Pr\left[d=\hat{d}\;\middle|\;b=0\right]$. Moreover, $\mathbb{A}_{\text{iso}}$ is given a faithful simulation of $\mathsf{Exp}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\cdot)$, thus $\Pr\left[d=\hat{d}\;\middle|\;b=0\right] = \Pr\left[\mathsf{Exp}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A}_{\text{iso}})\right]$.

For the second term, let bad_T denote the event that the distinguisher time-caps $\mathbb{A}^2_{\text{iso}}$ and hence sets $\hat{b} \leftarrow 1$. Then we can rewrite

$$\Pr\left[b=\hat{b}\;\middle|\;b=1\right] = \Pr[\mathsf{bad}_T] + (1-\Pr[\mathsf{bad}_T])\Pr\left[d\neq\hat{d}\;\middle|\;b=1\wedge\neg\mathsf{bad}_T\right] ,$$

where we also used that $\Pr[\mathsf{bad}_T \mid b=0] = 0$. For the final conditional probability, out and hence state was produced by Sim without access to unopened messages, so that the challenge bit d is information-theoretically hidden from $\mathbb{A}^2_{\text{iso}}$ and the probability that $d \neq \hat{d}$ equals a half.

Finally, collecting the pieces and some algebraic manipulation yields the theorem statement:

$$\begin{aligned}
2\cdot\mathsf{Adv}^{\kappa\text{-sso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}}(\mathbb{B}_{\text{sso}},\mathbb{D}_{\text{sso}}) &= 2\cdot\Pr\left[b=\hat{b}\;\middle|\;b=0\right] + 2\cdot\Pr\left[b=\hat{b}\;\middle|\;b=1\right] - 2\\
&= \left(2\cdot\Pr\left[\mathsf{Exp}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{S}}(\mathbb{A}_{\text{iso}})\right] - 1\right)\\
&\quad + 2\cdot\Pr[\mathsf{bad}_T] + 2\cdot(1-\Pr[\mathsf{bad}_T])\cdot\frac{1}{2} - 1\\
&= \mathsf{Adv}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\text{iso}}) + \Pr[\mathsf{bad}_T]\\
&\geq \mathsf{Adv}^{\kappa\text{-iso-cpa}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\text{iso}}) .
\end{aligned}$$

□

The proof of Thm. 5 does not carry over to the CCA setting, as the distinguisher $\mathbb{D}_{\text{sso}}$ would have to provide a simulation of the decryption oracle $\mathcal{D}$ to $\mathbb{A}^2_{\text{iso}}$ without access to the private keys or any oracles itself (indeed, already for weaker notions like PCA, CVA, etc., the proof breaks down). Whether SSO-CCA implies ISO-CCA with either kind of opening remains open, see Open Problem 5.

Open Problem 5. *Does* SSO-CCA *imply* ISO-CCA *in the presence of sender, receiver, or transmission opening?*

Asymptotic interpretation. As we are primarily concerned with concrete security, our security definition is indifferent as to whether the distinguisher may depend on the simulator or vice versa. When moving to asymptotic security, the choice of quantifiers leads to a weak and strong version of SSO security. A given scheme PKE is deemed weak SSO secure (for either opening, CPA or CCA) if, for all PPT $\mathbb{A}$, PPT M, and PPT $\mathbb{D}$ there exists a PPT simulator Sim such that the relevant advantage is negligible in λ. Alternatively, PKE is deemed to be strong SSO secure (again, for either opening, CPA or CCA) if for every PPT $\mathbb{A}$ and PPT M there is a PPT simulator Sim such that all PPT distinguishers $\mathbb{D}$ have negligible distinguishing advantage. Both options have appeared in the literature, for instance BY12 opted for the weaker variant, whereas more recent work went for the stronger one [91].

Similar to ISO, the mechanism on how to sample can affect the definition of SSO security (whether weak or strong). The equivalent of Lemma 4 holds for SSO, but without any need for ideal resamplers. Consequently, for both weak and strong SSO security, stateful (fixed-length) vector samplers and stateful single-message samplers are equivalent.

For an asymptotic interpretation of Thm. 5, we observe that the runtime of the derived distinguisher $\mathbb{D}$ includes that of the ideal sampler $\mathring{\mathsf{S}}$, thus this $\mathbb{D}$ might not be efficient. To ensure $\mathbb{D}$ is efficient, we could replace its call to $\mathring{\mathsf{S}}$ with that of an efficient resampler S and rely on Lemma 3 to argue that this change only affects negligible change in the distinguisher's advantage.

SSO⋆ is unachievable (without programming). Our formalization of SSO is multi-challenge, allowing non-trivial relations between sampled messages (through stateful or vector sampling, see Sect. 3.4). Yang et al. [113] showed κ-SSO-CPA⋆ to be unachievable in the non-programmable random oracle model, in the sense that private keys would have to be at least as long as the total number of plaintext bits to be encrypted [113, Thm. 3.1]. Thus, κ-SSO-CPA⊛, κ-SSO-CCA⋆, and κ-SSO-CCA⊛ must all be similarly unachievable. This mirrors Nielsen's earlier impossibility of non-interactive NCE in the non-programmable oracle model [100] (see Sect. 3.6).

Intuitively, the impossibility works as follows: the adversary $\mathbb{A}$, who wants to help distinguisher $\mathbb{D}$, commits to the public keys and challenge ciphertexts using the (non-programmable) random oracle. Then, it communicates the commitment, i.e. the digest of the random oracle, to $\mathbb{D}$, through the set of opened key handles, $\mathcal{I}$ (as $\mathcal{I}$ is given to $\mathbb{D}$ in their experiment, cf. Fig. 11). This is done by opening key number i iff the digest at position i has bit value 1. Then, $\mathbb{D}$ can recompute the commitment and check that it matches the form of $\mathcal{I}$. This strategy is hard to simulate, as it would require a simulator to commit to the ciphertexts before opening the corresponding messages. With the ability to program the random oracle, the commitment can be delayed until after opening.

The analysis relies on the uniformity of the random oracle and combinatorial bounds on the possible ways to choose the various cryptographic objects: for private keys larger than the number of bits encrypted, the analysis fails, leading to the stated requirement. Also note how the strategy puts a lower bound on the number of users in the system: concretely, if the random oracle output length is h, Yang et al. set $\kappa = h + 1$.

Yang et al.'s impossibility involved a notion of SSO in which the simulator was allowed to produce the parameters itself, cf. Remark 4; it implies impossibility for stronger notions where the simulator has less freedom.

Historical remarks. When Dwork et al. [41] first formalized selective opening attacks (for commitment schemes), they provided three definitions based on the framework of semantic security: one simulator based one, one based on a relation-predicate and one based on a function-predicate. The first two are equivalent, whereas the final one appeared weaker. Their simulation-based definition was then adapted to the public key setting with sender openings [10]. As with ISO, message sampling initially happened independently of the public key, meaning the notion did not imply IND-CPA; as detailed in Sect. 3.4, this shortcoming was later patched by BY12, who simultaneously introduced the notion of stateful samplers.

An SSO notion with receiver openings was subsequently developed [9]. Recently bi-openings were considered, including a "weak" version, for which an adversary must choose between access to either $\mathcal{S}$ or $\mathcal{R}$, but not both; it already has to make this choice after seeing the public keys, so before seeing any of the challenge ciphertexts. This weak version already turned out strictly stronger than either SSO⊙ and SSO⋆ individually [91].

As a final observation, the seeming difficulty in showing an implication from a posteriori simulatability to a posteriori indistinguishability with opening in the CCA setting as compared to in the CPA setting (see Open Problem 5) echoes the great time gap between establishing the equivalence of a posteriori simulatability (semantic security) and a priori indistinguishability without openings in the CPA setting (1986) [98] versus the CCA setting (2003) [111].

3.6 A Priori Simulatability with Selective Opening (NCE)

As explained in Sect. 3.2, a priori simulatability captures the idea that knowledge of a message should not be a prerequisite for producing ciphertexts that can pass off as encryptions of it, or in other words: irrespective of m, a simulator (without access to m) should be able to create a ciphertext c that cannot be distinguished from a real encryption of m.

Our formalization tasks a simulator Sim with simulating the view of adversary $\mathbb{A}$, taking on the role of game rather than player. This has the benefit of involving no message samplers: single messages are simply chosen by the adversary and given to the encryption oracle, who receives a (real or simulated) encryption in return. This mechanism makes a transmission oracle superfluous, as $\mathbb{A}$ always knows what message a challenge was supposed to encrypt; we therefore exclude the $\mathcal{T}$ oracle from Fig. 13.

Definition 7. *The* κ-NCE-CCA⊛ *advantage* $\mathsf{Adv}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}}(\mathbb{A})$ *of an adversary* $\mathbb{A}$ *against public key encryption scheme* $\mathsf{PKE}[\lambda]$*, relative to simulator* Sim*, is the distinguishing advantage against the game* $\mathsf{Exp}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}}(\mathbb{A})$ *(see Fig. 13).*

```
Experiment $\mathsf{Exp}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}}(\mathbb{A})$
    $b \leftarrow\$ \{0,1\}$
    $\mathcal{O}_b \leftarrow (\mathcal{E}_b, \mathcal{D}_b, \mathcal{S}_b, \mathcal{R}_b)$
    $s \leftarrow \varepsilon, q \leftarrow 0$
    $\mathsf{pm} \leftarrow\$ \mathsf{PKE.Pm}(\lambda)$
    if $b = 0$ :
        $\forall_{i\in[\kappa]}(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow\$ \mathsf{PKE.Kg}(\mathsf{pm})$
    else :
        $\forall_{i\in[\kappa]}\mathsf{pk}_i \leftarrow\$ \mathsf{Sim}_{\langle s\rangle}(\mathsf{Ini}, \mathsf{pm})$
    $\hat{b} \leftarrow\$ \mathbb{A}^{\mathcal{O}_b}(\mathsf{pm}, \mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
    return $b = \hat{b}$
```

```
Oracle $\mathcal{E}_0(i, m)$
    if $i \in \mathcal{I}$ : return $\lightning$
    $r \leftarrow\$ \mathsf{PKE.Rnd}(\mathsf{pm})$
    $c \leftarrow \mathsf{PKE.Enc}_{\mathsf{pk}_i}(m; r)$
    $\mathtt{R} \xleftarrow{\frown} r$
    $\mathcal{C}_i \xleftarrow{\cup} c$
    return $c$

Oracle $\mathcal{D}_0(i, c)$
    if $c \in \mathcal{C}_i$ : return $\lightning$
    $m \leftarrow \mathsf{PKE.Dec}_{\mathsf{sk}_i}(c)$
    return $m$

Oracle $\mathcal{S}_0(j)$
    $r \leftarrow \mathtt{R}[j]$
    return $r$

Oracle $\mathcal{R}_0(i)$
    $\mathcal{I} \xleftarrow{\cup} i$
    return $\mathsf{sk}_i$
```

```
Oracle $\mathcal{E}_1(i, m)$
    if $i \in \mathcal{I}$ : return $\lightning$
    $q \leftarrow q + 1$
    $c \leftarrow\$ \mathsf{Sim}_{\langle s\rangle}(\mathsf{Enc}, i, |m|)$
    $\mathtt{M} \xleftarrow{\frown} m, \mathtt{M}_i \xleftarrow{\frown} (q, m)$
    $\mathcal{C}_i \xleftarrow{\cup} c$
    return $c$

Oracle $\mathcal{D}_1(i, c)$
    if $c \in \mathcal{C}_i$ : return $\lightning$
    $m \leftarrow\$ \mathsf{Sim}_{\langle s\rangle}(\mathsf{Dec}, i, c)$
    return $m$

Oracle $\mathcal{S}_1(j)$
    $m \leftarrow \mathtt{M}[j]$
    $r \leftarrow\$ \mathsf{Sim}_{\langle s\rangle}(\mathsf{Sen}, j, m)$
    return $r$

Oracle $\mathcal{R}_1(i)$
    $\mathcal{I} \xleftarrow{\cup} i$
    $\mathsf{sk}_i \leftarrow\$ \mathsf{Sim}_{\langle s\rangle}(\mathsf{Rec}, i, \mathtt{M}_i)$
    return $\mathsf{sk}_i$
```

Fig. 13. The NCE experiment. Middle column = real, right column = ideal.

The oracles of Fig. 13 each come in two variants: one "real" (for $b = 0$), and one simulated, or "ideal" ($b = 1$). The real encryption, decryption, and opening oracles behave as expected: in particular, the real challenge encryption oracle simply encrypts the chosen message and returns the ciphertext.

Ideal oracles, meanwhile, call the corresponding subroutine of Sim. For the encryption oracle, Sim is asked to produce a ciphertext (under the relevant key handle) seeing only the length of the message. The message (or messages) is later revealed to Sim in the event that an opening oracle is called. To avoid trivial wins by the adversary, in the case of an $\mathcal{S}$ call Sim must come up with randomness such that re-encrypting the message produces the same ciphertext; similarly, in the case of $\mathcal{R}$ it should provide a private key such that decrypting *any* of the ciphertexts previously provided as a challenge under the corresponding key handle yields the correct message.

In our previous notions (Sect. 3.3–3.5), $\mathbb{A}$ was free to continue challenging a key handle i after the key had been opened. For NCE, such behaviour must be restricted, lest the notion becomes trivially unachievable. To see how, consider an adversary that calls $\mathcal{R}_b(i)$, receiving private key $\widetilde{\mathsf{sk}}_i$, followed by $\mathcal{E}_b(i, m)$ for a uniformly at random chosen message m of pre-determined length, receiving ciphertext c, and subsequently the adversary outputs $\hat{b} = 0$ iff $m = \mathsf{PKE.Dec}_{\widetilde{\mathsf{sk}}_i}(c)$. In the $b = 0$ world, the decryption check is guaranteed to succeed (assuming perfect correctness) so $\hat{b} = 0$ is guaranteed; yet, in the $b = 1$ world, m is information-theoretically hidden from Sim beyond its length $|m|$, thus the decryption check will hold with probability at most $2^{-|m|}$, yielding a significant distinguishing advantage of $1 - 2^{-|m|}$.

Even when a simulator would be allowed to program a random oracle, the adversary above is troublesome. Realistically, Sim's only hope to fool $\mathbb{A}$ is to program the random oracle for the calls that $\mathbb{A}$ makes to perform the check $m = \mathsf{PKE.Dec}_{\widetilde{\mathsf{sk}}_i}(c)$. In order to do so successfully, Sim would somehow have to learn m based on the queries $\mathbb{A}$ makes, but herein lies the rub: consider the sequence of oracle calls that honest decryption would make, and suppose there is a first oracle call whose input allows Sim to extract non-trivial information about m. Then that non-trivial information is already extractable from the answers Sim itself has provided so far, creating a complication (as m's contents would still have been information-theoretically hidden up to then based on prior calls). Essentially, even when Sim is allowed to program, it still ends up having to bootstrap its own knowledge of m (which is impossible).

```
Reduction $\mathbb{B}_{\text{nce}}(\mathsf{pm}, \mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
$s \leftarrow \varepsilon$
$\mathsf{out} \leftarrow\$ \mathbb{A}_{\text{sso}}^{\mathcal{E},\mathcal{D},(\mathcal{T},)\mathcal{S},\mathcal{R}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
for $j \in [|\mathtt{K}|]$ :
    if $\mathtt{K}[j] \in \mathcal{I} : \mathcal{J} \xleftarrow{\cup} j$
$\hat{b} \leftarrow\$ \mathbb{D}_{\text{sso}}(\mathsf{pm}, \mathtt{A}, \mathtt{M}, \mathcal{J}, \mathsf{out})$
return $\hat{b}$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{E}(i, \alpha)$
$\mathtt{K} \xleftarrow{\frown} i, \mathtt{A} \xleftarrow{\frown} \alpha$
$m \leftarrow\$ \mathsf{M}_{\langle s \rangle}(\alpha)$
$\mathtt{M} \xleftarrow{\frown} m$
$c \leftarrow \mathcal{E}_{\mathbb{B}}(i, m)$
return $c$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{T}(j)$
$\mathcal{J} \xleftarrow{\cup} j$
return $\mathtt{M}[j]$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{S}(j)$
$\mathcal{J} \xleftarrow{\cup} j$
$r \leftarrow \mathcal{S}_{\mathbb{B}}(j)$
return $(\mathtt{M}[j], r)$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{R}(i)$
$\mathcal{I} \xleftarrow{\cup} i$
$\mathsf{sk}_i \leftarrow \mathcal{R}_{\mathbb{B}}(i)$
return $\mathsf{sk}_i$
```

Fig. 14. The reduction $\mathbb{B}_{\text{nce}}$, simulating κ-SSO*-CCA$\circledast$ for $\mathbb{A}_{\text{sso}}$ and $\mathbb{D}_{\text{sso}}$ (the decryption oracle is simply forwarded).

As in previous sections, Sim is stateful, and we again allow it (in the ideal game) to produce the public keys but not the parameters: as with SSO (cf. Remark 4), the notion is meaningful even when the parameters are generated by Sim. Restricting the simulator by disallowing the generation of its own pm yields a potentially stronger notion and, as it matches the formalism of SSO better, is our preferred option.

Remark 6. Our formalization employs stateful simulators; earlier formalizations of NCE often consider NCE schemes to be tuples of algorithms extended to include a faking and an opening algorithm. For example, Hazay et al. [63] define the algorithms $\mathsf{PKE.Enc}^*$ and $\mathsf{PKE.Open}$ such that any ciphertext produced by $\mathsf{PKE.Enc}^*$ (on input the public key and message length) may be opened using $\mathsf{PKE.Open}$ (on input the message, the public and private keys, a trapdoor produced by $\mathsf{PKE.Enc}^*$, and the ciphertext to be opened). In our nomenclature, this corresponds to a simulator with precisely prescribed state and behaviour, potentially strengthening the notion without clear benefits. (Hazay et al. further strengthen their notion by insisting the simulator uses an externally, honestly generated public key of which it only learns the private key when running $\mathsf{PKE.Open}$, but not yet when running $\mathsf{PKE.Enc}^*$.)

NCE implies SSO. We next show (Thm. 6) how a priori simulatability (NCE) tightly implies a posteriori simulatability (SSO) in the presence of bi-openings. The reduction comes with a crucial caveat though: due to our NCE notion's restriction that opened keys cannot subsequently be challenged, the notion of SSO implied by NCE is weakened to one that makes the same restriction.

Theorem 6. *Let* SSO* *be defined as* SSO*(Fig. 10), except that a query $\mathcal{E}(i, \alpha)$ with $i \in \mathcal{I}$ leads to the immediate return of $\lightning$. Let $\mathsf{PKE}[\lambda]$ be given, then there exists a type-preserving black-box reduction $\mathbb{B}_{\text{nce}}$ (with black-box access to $\mathbb{A}_{\text{sso}}, \mathbb{D}_{\text{sso}}$, and M to be quantified later) such that for all* NCE *simulators $\mathsf{Sim}_{\text{nce}}$, there is a (black-box)* SSO *simulator $\mathsf{Sim}_{\text{sso}}$ such that for all adversaries $\mathbb{A}_{\text{sso}}$, message samplers M, and distinguishers $\mathbb{D}_{\text{sso}}$,*

$$\mathsf{Adv}^{\kappa\text{-sso}^*\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}_{\text{sso}}}(\mathbb{A}_{\text{sso}}, \mathbb{D}_{\text{sso}}) = \mathsf{Adv}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}_{\text{nce}}}(\mathbb{B}_{\text{nce}})\,.$$

The runtime of $\mathbb{B}_{\text{nce}}$ is upper bounded by that of $\mathbb{A}_{\text{sso}}$ and $\mathbb{D}_{\text{sso}}$ combined, plus that of running M q_e times, where q_e is the number of oracle calls made by $\mathbb{A}_{\text{sso}}$ to $\mathcal{E}$. The runtime of $\mathsf{Sim}_{\text{sso}}$ is upper bounded by q times that of $\mathsf{Sim}_{\text{nce}}$, where q is the total number of oracle calls made by $\mathbb{A}_{\text{sso}}$ to $\mathcal{E}$, $\mathcal{D}$, $\mathcal{T}$, $\mathcal{S}$, and $\mathcal{R}$, and $\mathsf{Sim}_{\text{sso}}$ additionally calls its $\mathcal{T}$ oracle $|\mathcal{J}|$ times, where $|\mathcal{J}|$ is the total number of opened challenge ciphertexts.

Proof. Let $\mathbb{A}_{\text{sso}}$ be an adversary playing the SSO* game; without loss of generality, we may assume that $\mathbb{A}_{\text{sso}}$ does not make any (pointless) $\mathcal{E}(i, \alpha)$ queries for which $i \in \mathcal{I}$. Let $\mathbb{D}_{\text{sso}}$ be a distinguisher. As the games $\mathsf{Exp}^{\kappa\text{-sso}^*\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}_{\text{sso}}}(\cdot)$ and $\mathsf{Exp}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}_{\text{nce}}}(\cdot)$ are both independent of their simulators $\mathsf{Sim}_{\ldots}$ conditioned on their respective bits $b = 0$, we can create a reduction $\mathbb{B}_{\text{nce}}$ (Fig. 14) that calls $\mathbb{A}_{\text{sso}}$ and $\mathbb{D}_{\text{sso}}$, simulating their $\mathsf{Exp}^{\kappa\text{-sso}^*\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}_{\text{sso}}}(\cdot)$ environment such that, if $\mathbb{B}_{\text{nce}}$'s oracles and inputs are real ($b = 0$), then so are those of $\mathbb{A}_{\text{sso}}$ and $\mathbb{D}_{\text{sso}}$. Thus

$$\Pr\left[\mathsf{Exp}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}_{\text{nce}}}(\mathbb{B}_{\text{nce}}) \,\middle|\, b = 0\right] = \Pr\left[\mathsf{Exp}^{\kappa\text{-sso}^*\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}_{\text{sso}}}(\mathbb{A}_{\text{sso}}, \mathbb{D}_{\text{sso}}) \,\middle|\, b = 0\right]$$

```
Simulator $\mathsf{Sim}_{\text{sso}}(\mathsf{pm})$
$q \leftarrow 0, s \leftarrow \varepsilon$
$\forall_{i \in [\kappa]} \mathsf{pk}_i \leftarrow \mathsf{Sim}_{\text{nce}\langle s \rangle}(\mathsf{Ini}, \mathsf{pm})$
$\mathsf{out} \leftarrow\!\!\$\ \mathbb{A}_{\text{sso}}^{\mathcal{E},\mathcal{D},(\mathcal{T},)\mathcal{S},\mathcal{R}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
return $\mathsf{out}$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{D}(i, c)$
if $c \in \mathtt{C}_i$ : return $\lightning$
$m \leftarrow\!\!\$\ \mathsf{Sim}_{\text{nce}\langle s \rangle}(\mathsf{Dec}, i, c)$
return $m$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{E}(i, \alpha)$
$q \leftarrow q + 1$
$\mathtt{M}_i \xleftarrow{\frown} (q, \bot)$
$|m| \leftarrow \mathcal{E}_{\mathsf{Sim}}(i, \alpha)$
$c \leftarrow\!\!\$\ \mathsf{Sim}_{\text{nce}\langle s \rangle}(\mathsf{Enc}, i, |m|)$
$\mathtt{C}_i \xleftarrow{\cup} c$
return $c$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{T}(j)$
$m \leftarrow \mathcal{T}_{\mathsf{Sim}}(j)$
return $m$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{S}(j)$
$m \leftarrow \mathcal{T}_{\mathsf{Sim}}(j)$
$r \leftarrow \mathsf{Sim}_{\text{nce}\langle s \rangle}(\mathsf{Sen}, j, m)$
return $(m, r)$
```

```
If $\mathbb{A}_{\text{sso}}$ calls $\mathcal{R}(i)$
for $j \in |\mathtt{M}_i|$ :
    $(q, \bot) \leftarrow \mathtt{M}_i[j]$
    $m \leftarrow \mathcal{T}_{\mathsf{Sim}}(q)$
    $\mathtt{M}_i[j] \leftarrow (q, m)$
$\mathsf{sk}_i \leftarrow\!\!\$\ \mathsf{Sim}_{\text{nce}\langle s \rangle}(\mathsf{Rec}, i, \mathtt{M}_i)$
return $\mathsf{sk}_i$
```

Fig. 15. An SSO simulator mimicking the behaviour of $\mathbb{A}_{\text{sso}}$ in the $b = 1$ case of $\mathbb{B}_{\text{nce}}$'s simulation (Fig. 14). Implicit in the figure is $\mathsf{Sim}_{\text{sso}}$'s state s, which it uses to keep track of global variables.

(where we slightly abused notation by conflating the b bits from the two distinct games).

If $b = 1$, both $\mathsf{Exp}^{\kappa\text{-sso}^*\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}_{\text{sso}}}(\cdot)$ and $\mathsf{Exp}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}_{\text{nce}}}(\cdot)$ do depend on their simulator. We claim that for any simulator $\mathsf{Sim}_{\text{nce}}$, the simulator $\mathsf{Sim}_{\text{sso}}$ (Fig. 15) perfectly mimics $\mathbb{A}_{\text{sso}}$'s behaviour in the ideal ($b = 1$) world. We may therefore conclude that

$$\Pr\left[\mathsf{Exp}^{\kappa\text{-sso}^*\text{-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathsf{Sim}_{\text{sso}}}(\mathbb{A}_{\text{sso}}, \mathbb{D}_{\text{sso}}) \;\middle|\; b = 1\right] = \Pr\left[\mathsf{Exp}^{\kappa\text{-nce-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{Sim}_{\text{nce}}}(\mathbb{B}_{\text{nce}}) \;\middle|\; b = 1\right].$$

Combining the $b = 0$ and $b = 1$ cases with the definition of a distinguishing advantage yields the theorem statement. □

Remark 7. The restriction on $i \notin \mathcal{I}$ for $\mathcal{E}(i, \alpha)$-calls appears unfortunate and is largely an artefact of our modelling choice to allow an adversary adaptive oracle access to challenge encryptions. Clearly, our theorem suffices to show that (our notion of) NCE tightly implies the more customary staged, single-shot version of SSO (with stateless vector-sampling), as there the game only allows corruptions after the challenge.

For instance, implications similar to ours were previously shown for sender openings [45, Lemma 1] and receiver openings [63, Theorem 4.4] (based on variants of NCE⊙ and NCE⋆, respectively). However, both those results use a staged SSO notion; moreover, both results use a single-key single-shot NCE notion, resulting in a reduction loss linear in the number of challenge ciphertexts (i.e. equal to the length of the vector of messages being sampled). In comparison, our result is tighter and more general and seems to resolve an open problem identified by BY12, who imagine that the general NCE⊛ (referred to as "MPC definition" by BY12) implies both SSO⊙ and SSO⋆.

The proof of Thm. 6 can readily be adapted to use the ($i \notin \mathcal{I}$)-restricted variant of SSO′ (Fig. 11) as target notion, if so desired. However, whether NCE suffices to imply (unrestricted) SSO we leave open (Open Problem 6).

Open Problem 6. *Does* NCE *imply* SSO *(without restriction) and if so, how tightly?*

Asymptotic interpretation. Our concrete definition of an NCE advantage (Def. 7) leads to two possible asymptotic definitions of security, again depending on the order of quantifiers (cf. SSO's asymptotic interpretation). For the "weak" option, a scheme PKE is deemed secure if, for all PPT $\mathbb{A}$ there exists a PPT simulator Sim resulting in a negligible advantage, whereas for the "strong" option, a scheme is only deemed secure if there exists a universal PPT simulator Sim that works for all PPT $\mathbb{A}$ (i.e. will result in a negligible advantage for all PPT $\mathbb{A}$).

The way our concrete implication is stated (Thm. 6) allows us to conclude that weak NCE implies the corresponding weak SSO* and strong NCE implies the corresponding strong SSO*.

NCE⋆ is unachievable (without programming). Just like κ-SSO-CPA⋆, κ-NCE-CPA⋆ is unachievable in the non-programmable random oracle model (in the sense that private keys would have to be at least as long as the total number of plaintext bits to be encrypted), as first shown by Nielsen [100]. Since κ-NCE-CPA$\star$ implies κ-SSO-CPA$\star$, this impossibility also follows from the (later) impossibility of κ-SSO-CPA$\star$ [113] (see Sect. 3.5).

Intuitively, Nielsen's proof is combinatorial in nature. There are at most $2^{|\mathsf{sk}|}$ possible decryptions of each ciphertext (one for each key); since chosen messages are independent of the private keys, if the messages are chosen from a set significantly larger than $2^{|\mathsf{sk}|}$, then it becomes likely that the chosen message is not in the set of possible decryptions for the simulated ciphertext.

The unachievability of κ-NCE-CPA$\circledast$, κ-NCE-CCA$\star$, and κ-NCE-CCA$\circledast$ in the non-programmable random oracle model follows.

Historical remarks. The formulation of NCE arose from the study of multi-party computation (MPC), whose protocols typically rely on suitably secure channels between any pair of parties. Before NCE, many protocols were only shown secure against adaptive adversaries in the information-theoretic setting; in the computational setting, it was only known how to achieve security against static adversaries, i.e. lacking the ability to adaptively choose which parties to open (after seeing the challenges). NCE arose naturally as a cryptographic security notion for secure peer-to-peer channels, permitting MPC schemes secure against adaptive adversaries.

Beaver and Haber [6] were the first to achieve such adaptively secure channels, but their solution crucially relied on secure erasures, meaning their scheme was not secure in the presence of (receiver) openings. Later schemes achieved security against adaptive adversaries in the presence of openings, at the cost of being interactive (three-round) protocols [5,29,38]; these are all proven secure in the standard model. Nielsen then completed the picture by constructing non-interactive NCE (sometimes known as NINCE) in the programmable random oracle model, and proving that programming is necessary to achieve NCE-CPA$\circledast$ non-interactively [100].

Camenisch et al. [25] gave an updated definition of 1-NCE-CCA$\circledast$ ("FULL–SIM" [25, Def. 6]) that allowed the adversary adaptive access to a sender and a receiver opening oracle, and featuring a stateful simulator. Thus their formalization closely resembles ours (Def. 7), with the exception that $\kappa = 1$, and that they allowed for public key encryption with labels [26]. Moreover, they allow for challenges to be issued even after corrupting the (one) user by altering the mechanism so that, post-compromise, the full message is fed to the simulator, enabling simulation (and bypassing the counterexample we used to argue to not allow post-compromise challenges). One can update Def. 7 to include such a mechanism.

Jaeger [82] recently introduced a generalized definition of NCE, referred to as SIM*, in which access to a programmable random oracle is part of the security notion (rather than the construction). Several composition results not known to hold in the a priori simulatability setting, such as the composition of a KEM and a DEM to form a PKE, were shown to hold for notions of SIM*. Intuitively, as simulator, adversary, and reduction alike are all allowed to program the same random oracle, programming can be made consistent across several game hops. He furthermore showed that κ-SIM*-CCA$\circledast$ hybridizes in the number of users κ, a result not known to hold for notions of NCE. (To see that SIM* is a strict generalization of NCE, observe that the adversary is strengthened compared to NCE adversaries due to its ability to program the random oracle, while the simulator is restricted compared to NCE simulators due to only being given the ability to program one specific oracle; thus, notions of SIM* trivially imply the corresponding notion of NCE.)

4 Relations

The road to understand the relations between the various notions presented in Sect. 3 has been long and winding, and as we have seen with the various open problems, the journey is still ongoing. We provide overviews of known relations in the CPA and CCA settings in Fig. 16. Here, bold arrows highlight results new to the current work, while purple dashed arrows represent relations that have, to the best of our knowledge, yet to be formally established, but that mirror other known results.

We stress that Fig. 16 reflects asymptotic interpretations of the various notions as, historically, these notions and their relations have primarily been considered asymptotically. As we mentioned in Sect. 2.3, the relevant implications and separations are given relative to classes of samplers/simulators in that case. For instance, an implication SSO-CPA$\odot$ $\Rightarrow$ ISO-CPA$\odot$ should be interpreted that, for all efficient PKE, if for the class of all efficient samplers SSO-CPA$\odot$ security holds, then ISO-CPA$\odot$ security also holds for the class of all samplers with efficient resamplability. For this particular implication, one can instead show the more concrete statement that for any efficient PKE and any sampler with efficient resamplability, SSO-CPA$\odot$ security with respect to that particular sampler implies ISO-CPA$\odot$ security with respect to that same sampler. In addition to the trivial, 'drop-an-oracle' concrete implications, we provided concrete counterparts for the main "downwards" implications in Sect. 3.

In contrast, we leave the remaining (non-trivial) results using their original asymptotic phrasing, unless explicitly stated otherwise. In addition to (full) implications as explained above, these results include partial implications and various kinds of separations. A partial implication IND-CPA $\Rightarrow$ SSO-CPA$\odot$ states that for all efficient PKE, if IND-CPA security holds, then for all samplers in some class, SSO-CPA$\odot$ security holds (the partiality of the implication may also restrict the class of PKE).

Separations show that implications cannot be proven, for instance IND-CPA $\not\Rightarrow$ SSO-CPA$\odot$ would indicate that there exists an efficient IND-CPA-secure scheme and some efficient sampler for which SSO-CPA$\odot$ security does not hold. From a strict, logical perspective, such a separation has to be conditional on the existence of an efficient IND-CPA-secure scheme to begin with. Often separations only have a partial scope, for instance by showing that they only hold for efficient PKE with certain properties, or when considering black-box reductions only. Although we will always indicate the rough scope, we routinely omit the precise details and conditions under which a separation has been shown and refer to the relevant source for details instead.

Finally, we will encounter some semi-separations, in the sense that partial security of one notion is insufficient to establish an impliciation. For instance, Cor. 1's semi-separation ISO-CCA$\odot$ $\not\Rightarrow$ SSO-CPA$\odot$ indicates that there plausibly exists an efficient PKE for which ISO-CCA$\odot$ security holds with respect to some strict subclass of message samplers, yet SSO-CPA$\odot$ security with respect to some relevant class of message samplers does not.

Remark 8. Formalizations often differ between works in subtle but important ways (for instance the order of quantifiers or the exact interfaces of adversaries, simulators and distinguishers), and the implications and separations given in Fig. 16 should therefore be interpreted as going between families of notions, in the sense that there is for instance a formalization of SSO-CPA$\odot$ known to imply a formalization of ISO-CPA$\odot$; or the other way around, that some formalization of ISO-CPA$\odot$ cannot imply some formalization of SSO-CPA$\odot$.Consequently, one should take some care when interpreting the figures, as arrows may not trivially compose. To alleviate some of these complications, we provided concrete, more systematic versions of several known implications in Sect. 3 with updated proofs, and our asymptotic interpretations clarify which implications do hold and between which formalizations subtleties arise. One striking example where composition is not straightforward is our reduction from NCE to SSO (Thm. 6), as it requires challenging compromised key handles to be disallowed also in the SSO experiment, as it is in the NCE experiment. Thus, implications from NCE to notions further down the hierarchy also do not immediately follow without putting similar restrictions on each notion.

4.1 Hybridization

A natural question when presented with any fully adaptive, multi-user security notion is whether the notion hybridizes in the number of users and challenges (i.e. the number and type of calls to the challenge encryption oracle $\mathcal{E}$), as is the case for multi-user indistinguishability without openings [7]. Indeed, in the closely related a priori indistinguishability setting the answer is yes as we noted in Sect. 3.3 already: the original hybrid proof works equally well with receiver openings (i.e. for κ-IND-CPA$\star$/κ-IND-CCA$\star$) [68], and is unaffected by the presence of sender or transmission openings. Of more interest are the other three settings, especially as historically they were often presented in a single challenge setting and, in the case of sender openings only, with only a single user. Our first observation is that in all three settings the availability of openings does seem to, at the very least, complicate any hybrid argument.

Users. On the positive side, ISO$\diamond$ and SSO-CPA$\diamond$ should hybridize due to their equivalence with IND (see Sect. 4.2 below); we conjecture that so does SSO-CCA$\diamond$.

On the other hand, for receiver openings, corrupting a single user does not benefit an adversary, thus IND $\Rightarrow$ 1-SSO$\star$ $\Rightarrow$ 1-ISO$\star$. As hybridization in the number of users would mean that 1-ISO$\star$ $\Rightarrow$ κ-ISO$\star$, resp. 1-SSO$\star$ $\Rightarrow$ κ-SSO$\star$ and neither κ-ISO$\star$ nor κ-SSO$\star$ can be equivalent to IND-CPA (see Sect. 4.3 below), hybridization is not possible.

Jaeger [82] noted that providing a proof of hybridization for NCE$\star$ appears difficult, however a proof of impossibility remains elusive.

Open Problem 7. *Determine whether* NCE$\star$ *hybridizes in the number of users.*

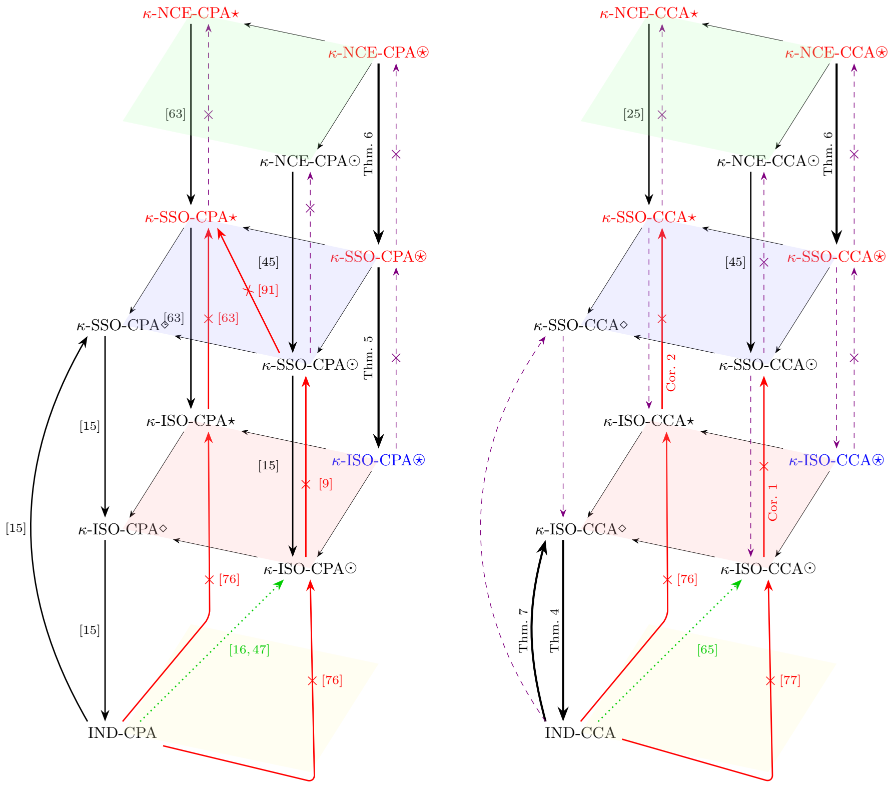

Fig. 16. Known and conjectured relations in the CPA and CCA settings. Highlighted in red are notions known to be unachievable in the standard model, and in blue notions for which standard model achievability remains open. Slim arrows (without references) are trivial, bold arrows highlight implications shown for the first time here, and violet dashed arrows are open problems (shown as the relation one would expect based on known results). Finally, green dotted arrows represent conditional implications.

In the presence of sender openings only, the question of hybridization appears mostly unexplored.

Open Problem 8. *Do notions of* ISO$\odot$, SSO$\odot$, *and/or* NCE$\odot$ *hybridize in the number of users?*

Challenges. NCE does not hybridize in the number of challenges in general; this follows from the impossibility result: there are standard model schemes secure against single-challenge NCE (as long as $|\mathsf{sk}| \geq |m|$) that will be insecure against q-challenge NCE (for insufficient $|\mathsf{sk}|$).

In the ISO and SSO settings, hybridization in the number of challenges depends on the sampler (see "Alternative samplers" in Sect. 3.4). For stateless samplers, hybridization is possible [16, 87]. However, in the stateless setting, hybridization in the vector length ℓ is not always possible: specifically, for κ-SSO-CCA$\star$ vector message sampling $\mathsf{M}^{\ell}_{\langle\rangle}$ where multiple messages per key are sampled simultaneously (so ℓ is a multiple of κ) is strictly stronger than the setting where for each challenge only a single message is sampled per key (so $\ell = \kappa$), at least without the ability to program oracles, such as a random oracle [113]. Consequently, for stateful samplers in that same κ-SSO-CCA$\star$ setting, no generic hybrid argument is possible for the number of challenges. On the other hand, for (stateless) product samplers a hybrid argument is possible in the length of the messages vectors which was initially shown for commitment schemes [41] and which is similar to the hybrid argument used for the number of challenges; in general, the lack of hybridization in one setting (κ-SSO-CCA$\star$) need not imply the lack thereof in other ones.

Open Problem 9. *When is hybridization in the number of challenges/length of the sampling vector (im)possible in the* ISO *and* SSO *settings?*

4.2 Implications

Trivial implications from ignoring oracles. Almost any time a security notion comes with a helper oracle, such as the decryption oracle $\mathcal{D}$ of CCA or the various opening oracles (see Table 1), there is a trivial reduction to a security notion without the oracle that simply ignores the oracle in question. Possible, non-trivial complications may arise for simulator-based notions where the simulator also loses oracle access and hence becomes less powerful; for our two simulator-based notions no such complications can arise as NCE's simulator has no oracle access whatsoever and SSO's simulator access to $\diamond$ is effectively restricted (technically, security with some opening includes universal quantification over all adversaries that do not open and, for those adversaries, the simulator cannot call its $\diamond$ oracle without being immediately noticed by a distinguisher checking its $\mathscr{J}$ input). So for instance, any CCA notion implies the CPA equivalent, and any notion with bi-openings implies any notion with transmission, sender, or receiver openings only. In Fig. 16, these trivial implications are represented by slim arrows that outline each plane of the hierarchy, going from more powerful openings to weaker ones. Likewise, any notion in the rightmost column (CCA) trivially implies the corresponding notion on the left (CPA).

Downward implications. The hierarchical structure between the different philosophies presented in Fig. 1 (from a priori indistinguishability up to a priori simulatability) is well-established in the CPA setting, as follows from the asymptotic interpretations of Thm. 6 (NCE implies SSO [45,60,63]), Thm. 5 (SSO implies ISO [15]), and Thm. 4 (ISO implies IND [15]), respectively. The main caveat that we uncovered is that NCE only implies a slightly restricted version of our more general SSO notion, although it still suffices to imply the older SSO version (with a stateless vector sampler). Various separation results, to be discussed in Sect. 4.3, reinforce the hierarchy by ruling out that 'lower' notions are in fact equivalent to the corresponding 'higher' one (with the noticeable exception of CPA$\diamond$, for which IND, ISO and SSO are all equivalent). By contrast, the CCA hierarchy currently contains a glaring hole as far as implications go, namely whether SSO-CCA implies ISO-CCA in the presence of openings, see Open Problem 5.

A priori indistinguishability and tightness. As mentioned in Sect. 4.1, the multi-user-with-corruptions notions of the a priori indistinguishability setting hybridize, meaning they are all implied by the single-user IND-CPA/IND-CCA notions with a tightness loss linear in the number of users [7], and the number of challenge bits in the case of β-IND-CCA$\odot$ and (κ, β)-IND-CCA$\circledast$ (Thm. 3), and so these notions are all asymptotically equivalent. Since Fig. 16 maps asymptotic implications and separations, it therefore does not distinguish between single- and multi-user notions of a priori indistinguishability. (We will revisit tightness in Sect. 5.2.)

IND partially implies ISO$\odot$. There are classes of message samplers for which IND-CPA does imply κ-ISO-CPA$\odot$: Fuchsbauer et al. [47] showed that these include samplers inducing product distributions (i.e. independent message sampling), Markov distributions, and more generally any graph-induced message distribution for which the underlying directed graph can be traversed in polynomial time (for a certain definition of "traversed"). Heuer later showed that the result transfers to the CCA setting [65]; extending to receiver openings remains open. These partial implications appear as green dotted arrows in Fig. 16.

Open Problem 10. *For which classes of message samplers does* IND-CPA *security imply* κ-ISO-CPA$\star$ *security, or even* κ-ISO-CPA$\circledast$ *security? How about in the* CCA *setting?*

IND implies notions of SOA with transmission openings. As shown by BY12, IND-CPA implies κ-SSO-CPA$\diamond$, and when starting from κ-IND-CPA, the reduction is furthermore tight [16, Thm. 4.1]. As we will show momentarily, for ISO a similar implication holds also in the CCA setting, which follows from a reduction to multi-user real-or-random indistinguishability (Thm. 7).

For SSO, it is unclear whether a similar implication holds in the CCA setting; a straightforward upgrade of BY12's proof runs into technical difficulties with simulating access to a decryption oracle (intuitively, the simulator would not have access to a decryption oracle, and we cannot rule out the possibility that an adversary could somehow convince the distinguisher that it does have access to one).

Open Problem 11. *Does* IND-CCA *security imply* SSO-CCA$\diamond$ *security?*

Theorem 7. *Let* $\mathsf{PKE}[\lambda]$ *be given, let* $q \in \mathbb{Z}_{>0}$*, and let* M *be a sampler with ideal resampler* $\mathring{\mathsf{S}}$*. Then there is a reduction* $\mathbb{B}_{\text{ror}}$ *such that for any adversary* $\mathbb{A}_{\text{iso}}$ *making at most q challenge oracle calls,*

$$\mathsf{Adv}^{\kappa\text{-iso-cca}\diamond}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\text{iso}}) = 2 \cdot \mathsf{Adv}^{\kappa\text{-ror-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\text{ror}}) .$$

The runtime of $\mathbb{B}_{\text{ror}}$ *is upper bounded by that of* $\mathbb{A}_{\text{iso}}$*, plus that of running* M *q times and* $\mathring{\mathsf{S}}$ *once, and some small overhead.*

```
Reduction $\mathbb{B}_{\mathrm{ror}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
  $d \leftarrow\!\!\$\ \{0,1\}$
  $\mathsf{challenged} \leftarrow \mathsf{false}$
  $\hat{d} \leftarrow\!\!\$\ \mathbb{A}_{\mathrm{iso}}^{\mathcal{E},\mathcal{C},\mathcal{D},\mathcal{T}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
  $\hat{b} \leftarrow \neg(d = \hat{d})$
  return $\hat{b}$

If $\mathbb{A}_{\mathrm{iso}}$ calls $\mathcal{D}(i, c)$
  $m \leftarrow \mathcal{D}_{\mathbb{B}}(i, c)$
  return $m$

If $\mathbb{A}_{\mathrm{iso}}$ calls $\mathcal{E}(i, \alpha)$
  if challenged : return $\lightning$
  $\mathtt{A} \overset{\frown}{\longleftarrow} \alpha$
  $m \leftarrow\!\!\$\ \mathsf{M}_{\langle s \rangle}(\alpha)$
  $\mathtt{L} \overset{\frown}{\longleftarrow} |m|$
  $\mathtt{M}^0 \overset{\frown}{\longleftarrow} m$
  $c \leftarrow \mathcal{E}_{\mathbb{B}}(i, m)$
  return $c$

If $\mathbb{A}_{\mathrm{iso}}$ calls $\mathcal{C}$
  if challenged : return $\lightning$
  $\mathsf{challenged} \leftarrow \mathsf{true}$
  $\mathtt{M}^1 \leftarrow\!\!\$\ \mathring{\mathsf{S}}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$
  return $\mathtt{M}^d$

If $\mathbb{A}_{\mathrm{iso}}$ calls $\mathcal{T}(j)$
  if challenged : return $\lightning$
  $\mathcal{J} \overset{\cup}{\longleftarrow} j$
  return $\mathtt{M}[j]$
```

Fig. 17. The reduction $\mathbb{B}_{\mathrm{ror}}$ of Thm. 7 simulating κ-ISO-CCA$\diamond$ for $\mathbb{A}_{\mathrm{iso}}$.

Proof. The reduction $\mathbb{B}_{\mathrm{ror}}$ (Fig. 17) simulates κ-ISO-CCA$\diamond$ for $\mathbb{A}_{\mathrm{iso}}$. Let b denote the challenge bit of the $\mathsf{Exp}^{\kappa\text{-ror-cca}}_{\mathsf{PKE}[\lambda]}(\cdot,\cdot)$ game, then $\mathbb{B}_{\mathrm{ror}}$ wins with the following probability.

$$\begin{aligned}
\Pr\Big[\mathsf{Exp}^{\kappa\text{-ror-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\mathrm{ror}})\Big] &= \frac{1}{2}\left(\Pr\Big[\hat{b} = 0 \,\Big|\, b = 0\Big] + \Pr\Big[\hat{b} = 1 \,\Big|\, b = 1\Big]\right) \\
&= \frac{1}{2}\left(\Pr\Big[d = \hat{d} \,\Big|\, b = 0\Big] + \Pr\Big[d \neq \hat{d} \,\Big|\, b = 1\Big]\right) ,
\end{aligned}$$

as follows from $\mathbb{B}_{\mathrm{ror}}$'s description (Fig. 17). The first term represents the probability that $\mathbb{A}_{\mathrm{iso}}$ wins a faithful simulation of the ISO-CCA$\diamond$ game, and so

$$\Pr\Big[d = \hat{d} \,\Big|\, b = 0\Big] = \Pr\Big[\mathsf{Exp}^{\kappa\text{-iso-cca}\diamond}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\mathrm{iso}})\Big] .$$

For the second term, the challenge bit is information-theoretically hidden from $\mathbb{A}_{\mathrm{iso}}$ (as resampling is ideal), so that

$$\Pr\Big[d \neq \hat{d} \,\Big|\, b = 1\Big] = \frac{1}{2} .$$

Recombining the two cases gives the theorem statement:

$$\begin{aligned}
2 \cdot \mathsf{Adv}^{\kappa\text{-ror-cca}}_{\mathsf{PKE}[\lambda]}(\mathbb{B}_{\mathrm{ror}}) &= 2 \cdot \Pr\Big[d = \hat{d} \,\Big|\, b = 0\Big] + 2 \cdot \Pr\Big[d \neq \hat{d} \,\Big|\, b = 1\Big] - 2 \\
&= 2 \cdot \Pr\Big[\mathsf{Exp}^{\kappa\text{-iso-cca}\diamond}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\mathrm{iso}})\Big] + 2 \cdot \frac{1}{2} - 2 \\
&= \mathsf{Adv}^{\kappa\text{-iso-cca}\diamond}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{A}_{\mathrm{iso}}) .
\end{aligned}$$

□

Asymptotic interpretation. For an asymptotic interpretation of Thm. 7, we observe that the runtime of the reduction $\mathbb{B}_{\mathrm{ror}}$ includes that of the ideal sampler $\mathring{\mathsf{S}}$, and thus the reduction might not be efficient. Like for Thm. 5, we can ensure $\mathbb{B}_{\mathrm{ror}}$ is efficient by replacing its call to $\mathring{\mathsf{S}}$ with that of an efficient resampler S, and rely on Lemma 3 to argue that this change only effects a negligible change in the reduction's advantage.

4.3 Separations

IND does not imply ISO (with sender or receiver openings). Standard IND-CCA security implies neither ISO-CPA$\odot$ nor κ-ISO-CPA$\star$ in general: Hofheinz et al. [76, 77] showed that IND-CCA does not imply a closely related notion of threshold ISO-CPA$\odot$ security, for which the adversary gets a number of encrypted shares of a secret and must uncover the secret. Meanwhile, ISO-CPA$\odot$ does imply threshold ISO-CPA$\odot$, establishing a separation. Adapting the strategy to the receiver opening setting via an analogous notion of receiver-threshold security (in which each share is encrypted under a random public key rather than a single, fixed public key), they were able to likewise establish a separation from IND-CCA to κ-ISO-CPA$\star$ [76].

As, as mentioned in the introduction, actually developed use cases of SOA-notions are rare, the separations above simultaneously support our belief in the notions' usefulness in the context of threshold security, where IND-CCA is insufficient.

ISO does not imply SSO (with sender or receiver openings). Bellare et al. [9], generalizing an earlier result due to Hofheinz [10, 69], showed that no *committing* PKE scheme (i.e. one that implies a computationally binding non-interactive commitment scheme) can achieve SSO-CPA$\odot$, even when message samplers are restricted to independently sampled uniform bitstrings. (Hofheinz only showed that there exists a sampler relative to which such schemes cannot be proven SSO-CPA$\odot$-secure via black-box reductions to standard cryptographic assumptions.) Since there are IND-CCA schemes that are committing, for instance the Cramer–Shoup encryption scheme [35], they concluded that IND-CCA cannot imply SSO-CPA$\odot$.

As explained in Sect. 4.2, Heuer [65] lifted an earlier result in the CPA setting [47] to the CCA setting, and thereby showed that for many message samplers, including those that sample independently (i.e. that infer product distributions), IND-CCA does imply ISO-CCA$\odot$.

Combining the two results, we see that for a large class of message samplers, IND-CCA implies ISO-CCA$\odot$ while being separated from SSO-CPA$\odot$ (even when restricted to that same class of message samplers). Thus we reach a semi-separation (previously a full separation was observed in the CPA setting only [9]):

Corollary 1. *Class-restricted* ISO-CCA$\odot$ *security does not imply (class-restricted)* SSO-CPA$\odot$ *security.*

Turning to receiver openings, Bellare et al. [9] also identified a feature called decryption verifiability, which captures that anyone given a tuple consisting of public key, private key, ciphertext and message, should be able to verify that the ciphertext is an honest encryption of the message (for some definition of "honest"). They showed that no scheme that is decryption verifiable can be SSO-CPA$\star$ secure, even for independent, uniform sampling. Since many natural IND-CPA schemes are decryption verifiable, for instance ElGamal [43], a separation is established.

Both their results (committing schemes do not achieve SSO-CPA$\odot$; decryption verifiable schemes do not achieve κ-SSO-CPA$\star$) build on the same underlying proof idea as the impossibility of SSO$\star$, already discussed in Sect. 3.5: the adversary computes the hash of the challenge ciphertexts, and communicates the hash value to the distinguisher using the set of opened indices ($\mathcal{J}$ for $\odot$; $\mathcal{I}$ for $\star$); for any scheme that satisfies the respective property, no efficient simulator can provide a convincing simulation of this strategy.

However, with the standard model impossibility of κ-SSO-CPA$\star$ in hand, combined with the existence of schemes achieving κ-ISO-CCA$\star$ in the standard model (under reasonable computational assumptions, see Sect. 5.2), we can again reach a stronger conclusion (previously only shown for the CPA setting [63]):

Corollary 2. κ-ISO-CCA$\star$ *security does not imply* κ-SSO-CPA$\star$ *security.*

The separation above leaves room for partial implications, where κ-SSO-CPA$\star$ is implied for restricted classes of message samplers (that are not captured by the impossibility). Furthermore, that impossibility result itself relies on the ability of an adversary to use the index set of corrupted challenges ($\mathcal{J}$) or keys ($\mathcal{I}$) to 'communicate' with the distinguisher. Using a formalization of κ-SSO-CPA$\star$ where only the cardinality of those sets is passed directly to the distinguisher would require re-examination of those impossibility results (to determine whether they can be ported to the new formalization).

Technically, the situation is still partly open for bi-opening, as no scheme has yet been shown to achieve κ-ISO-CPA$\circledast$ in the standard model (see Open Problem 19); any such scheme would immediately lead to a separation, namely that κ-ISO-CPA$\circledast$ security does not imply κ-SSO-CPA$\circledast$ security, either in part (if the scheme only achieves CPA) or in full (if it also achieves CCA).

4.4 Further Relations

SSO and NCE. Remarkably, whether SSO and NCE are separated or equivalent appears largely open, beyond the logical consequences of various unachievability results for both SSO$\star$ and NCE$\star$ in the standard model.

Camenisch et al. [25] claim that NCE$\star$ is strictly stronger than notions of SSO$\star$, however this claim was based on the former being unachievable in the standard model, and the belief that the latter was achievable in the standard model. With neither achievable in the standard model, their relation becomes less relevant from a logical perspective (unless perhaps in an idealized model, such as the programmable random oracle).

Hazay et al. [63] gave a separation from κ-SSO-CPA$\star$ to κ-NCE-CPA$\star$, however their algorithmic formalization of NCE was stronger than our Def. 7, cf. Remark 6: in particular, and unlike our formalization, their NCE simulator was not allowed to produce the public keys (while their SSO simulator was). This technicality turns out to be central to their separation result.

Open Problem 12. *Which, plausibly achievable notions of* SSO *imply notions of* NCE *in general and which are separated? For settings where neither notion is achievable in the standard model, which relations can be drawn in (specific) idealized models?*

Sender vs. receiver opening. Although sender openings and receiver openings model quite different scenarios, their formalizations have a lot in common; moreover, ideas for one scenario have subsequently often been adapted to the other. Yet, the question when one type of opening implies the other, or conversely, when no such implication can exist, has not achieved much direct attention (Open Problem 13). Of course, juxtaposing feasibility results for one with impossibility results for the other, does provide some answers.

Together, the unachievability of κ-SSO-CPA$\star$ in the standard model and combined with the plausible existence of a standard model construction achieving NCE-CCA$\odot$ [45] (see Sect. 5.2) seem to indicate that a priori simulatability with sender opening (NCE$\odot$) cannot imply a posteriori simulatability with receiver opening (SSO$\star$). However, the construction was only shown 1-NCE-CCA$\odot$ secure in the standard model. As it is unclear whether NCE$\odot$ hybridizes in the number of users (see Sect. 4.1), and SSO$\star$, unlike NCE$\star$, is only relevant for $\kappa > 1$ (for $\kappa = 1$ it becomes equivalent to a notion without openings), we can only formally conclude that 1-NCE-CCA$\odot$ does not imply 1-NCE-CPA$\star$. Of course, the achievability of 1-NCE-CCA$\odot$ feeds the impression that κ-NCE-CCA$\odot$ is achievable as well, which would restore a full separation.

BY12 gave a construction that achieves κ-SSO-CPA$\odot$ in the standard model for arbitrary κ; we may therefore conclude that κ-SSO-CPA$\odot$ does not imply κ-SSO-CPA$\star$.

The reverse direction is uninteresting for unrestricted message spaces in the standard model (as SSO$\star$ and NCE$\star$ are both unachievable then). Lai et al. [91] observed that the Cramer–Shoup scheme [35] achieves a single-message-bit variant of κ-SSO-CCA$\star$ in the standard model [78], but as the scheme is committing, it cannot achieve SSO-CPA$\odot$ [69]; we therefore conjecture that notions of SSO$\star$ and SSO$\odot$ are separated in both directions, and thus incomparable.

Implications between notions with various openings of course do hold for a priori indistinguishability (at a loss, see Sect. 3.3), while the situation appears open for ISO.

Open Problem 13. *How do notions with sender openings (only) and notions with receiver openings (only) relate to each other?*

4.5 Selected Related Notions

So far, we have concentrated on the main notions that can arise when formalizing openings in the context of public key encryption. Below we will briefly address a few additional notions (we already touched upon these in Sect. 3). Yet, many more related security notions have been proposed in the past, that we will not expand upon. Some of these further notions are subtle variants of what we have already seen, for instance weak bi-opening [91] or tweaked NCE [63], while other notions go beyond message confidentiality, for instance by considering non-malleability [81] or anonymity/key-privacy [80].

SIM*. As NCE$\star$ cannot be realized in the random oracle model without programming (let alone the standard model), studying the notion only makes sense in idealized models. One difficulty that surfaces when including programming directly in the security model is the need to keep said programming consistent, especially when considering the composition of various reductions that might all vie with the simulator in their programming.

When Jaeger [82] recently introduced a strengthened notion of a priori simulatability, named SIM*-AC (for "Adaptive Compromise"), hereafter SIM*, he baked the ideal object into the notion, allowing simulators, reductions and, crucially, adversaries alike to program the ideal object through the same oracle

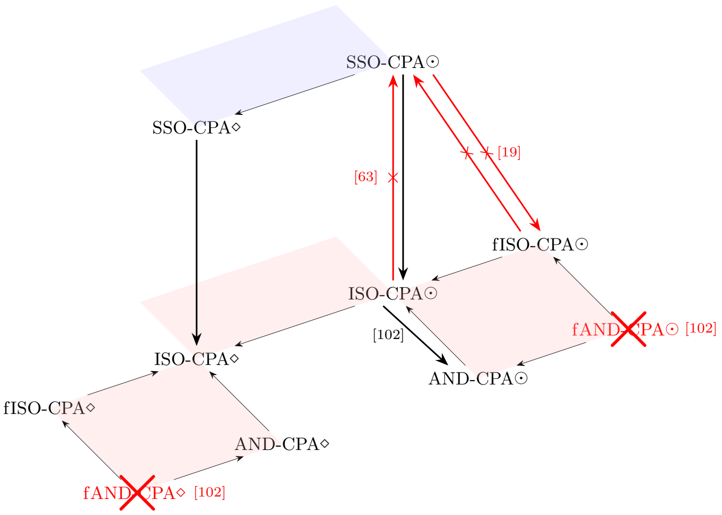

Fig. 18. Full ISO, full AND, and how they relate to the middle section of Fig. 16 (ISO-CPA⊙ and SSO-CPA⊙). A red cross means the notion is unachievable even in idealized models (under mild conditions).

interface. He showed that SIM* does hybridize in the number of users and that the security of several common constructions, such as the KEM/DEM hybrid PKE and the Fujisaki–Okamoto transform, carries over to the a priori simulatability setting. However, even for SIM* hybridization in the number of challenges appears tricky [83].

As SIM* security also covers adversaries that do not program, any SIM* notion tightly implies its NCE counterpart (in the programmable oracle model, where only the simulator and reduction may program).

Deniable encryption. Deniable public key encryption, as introduced by Cannetti et al. [28], shares many features with NCE: in particular, like NCE, one should be able to open a ciphertext to plaintexts other than the one originally encrypted. The main difference is that, for deniable encryption, this capability should be available to the end user, while for NCE we only require it to be accessible to the simulator (i.e. by programming ideal objects). Thus, non-interactive deniable PKE with bi-opening is strictly stronger than NCE, and thus it must also unachievable in the standard model. Unlike NCE however, non-interactive deniable PKE with receiver opening is unachievable *in general*: to see this, recall that programming random oracles is necessary to achieve NCE with receiver openings [100]; we have no way of granting the end user such powers.

On the other hand, non-interactive deniable encryption with sender opening *is* achievable, although the only known such scheme requires encryption to be a quantum algorithm [34]. Additionally there is a generic transform that takes any sender-deniable encryption scheme to a receiver-deniable encryption scheme by adding extra interaction (so it is no longer non-interactive, bypassing the impossibility result) [28], and there are interactive schemes achieving deniability with bi-opening [33].

If the requirement that a given ciphertext can be opened to any message is weakened to, for instance, only two predetermined messages (under two different keys), one can construct non-interactive schemes with this limited deniability functionality. Several common blockcipher modes-of-operation, including AES-GCM, sport such deniability-like properties [56], and prior works have rather focused on how to rid these schemes of that deniability, and thus achieve fully *committing* authenticated encryption [40, 56].

Full ISO-CPA and friends. Böhl et al. [19] noted how allowing resampling to be inefficient leads to a significantly stronger notion of ISO. They named the resulting notion “Full” ISO (fISO), and showed that fISO-CPA⊙ and SSO-CPA⊙ are incomparable. No scheme to date has been shown to achieve fISO-CPA (not even for transmission openings, though it has mainly been studied for sender openings, see Open Problem 14) and, as allowing inefficient resampling seems to yield an artificially strong notion of security, fISO has mostly fallen out of favour. fISO therefore sits outside of our main hierarchy (Fig. 16).

Open Problem 14. *Is* fISO-CPA *achievable in the presence of sender, receiver, or transmission opening?*

A closely related notion, which we will refer to as Full AND (fAND), works exactly as Full ISO, except that in the challenge phase (see oracle $\mathcal{C}$ in Fig. 8), the adversary gets to see both the original and the resampled messages [102]. Although (non-full) AND-CPA⊙ and ISO-CPA⊙ are equivalent, fAND turns out to be unachievable (for public key spaces supporting Σ-protocols) [102]. Furthermore, the proof that fAND-CPA is unachievable arlready holds in the presence of transmission opening (by inspection), in stark contrast to AND-CPA⋄, which is equivalent to IND-CPA (Thm. 7). These results are summarized in Fig. 18, with the large red crosses representing unachievability.

The unachievability of fAND-CPA⋄ lends support to our belief that fISO-CPA⊙ is an unachievable notion and that its abandonment is warranted.

5 Constructions

The literature is rich with constructions achieving the various notions of the hierarchy, from well-known constructions achieving IND-CPA at the bottom, to the unachievable-without-programming κ-NCE-CCA⊛ at the top.

In this section, we survey the state of the art, and group families of constructions based on their main underlying assumption. In addition, we identify a ROM construction technique due to Bellare and Rogaway [12], henceforth Bellare–Rogaway Encryption (BRE).

5.1 Bellare–Rogaway Encryption

Bellare–Rogaway Encryption combines a random oracle with one-time pad encryption, and the technique is as simple as it is elegant. It gives highly efficient schemes with very strong security guarantees, with two important caveats: first, the security proof relies on programming random oracles (which as we have seen is necessary to achieve simulatability with receiver opening, see Sect. 3.5–Sect. 3.6); and second, that the security reduces to one-wayness of a Key Encapsulation Mechanism (KEM) in the presence of plaintext checking attacks (OW-PCA), rather than just CPA. Such KEMs can be instantiated with the GapCDH assumption, or other such "gap" assumptions that exploit (conjectured) differences between computational and decisional hardness; there are also generic transformations taking OW-CPA to OW-PCA in the (Q)ROM [71]. OW-PCA has furthermore been shown to be necessary to prove CCA security of common BRE schemes [109] like REACT [101].

It goes as follows: let encryption, on input a public key pk and message m, encapsulate a seed K_{kem} as c_1, and let the seed be fed to an extendable output function (XOF) $\mathcal{F}$ to retrieve a pseudorandom bitstring K_{otp} of length m; then, use K_{otp} as key for one-time padding (OTP) the message, yielding c_2; the ciphertext c is the tuple (c_1, c_2). Decryption essentially repeats the process: decapsulate c_1 to recover seed K_{kem}, compute K_{otp}, and one-time pad it with c_2 to recover m. As indicated by the notation, key encapsulation mechanisms (KEMs) are optimal primitives for instantiating such schemes, provided they achieve OW-PCA. (Bellare and Rogaway originally employed trapdoor permutations in place of KEMs, for which plaintext checking comes for free [12].)

When modelling $\mathcal{F}$ as a random oracle, reprogramming enables a simulator: upon encryption, it simply chooses the message encryptions c_2 uniformly at random, and stores it for later. The properties of the OTP then ensure that, for any message m and ciphertext c_2, there exists a key K_{otp} such that $m \oplus K_{\text{otp}} = c_2$. Thus, when this message is later revealed to the simulator, it can simply reprogram $\mathcal{F}$ such that on input seed K_{kem}, the key $K_{\text{otp}} = m \oplus c_2$ is output. (If CCA-security is desired, the seed can be expanded to include a MAC key, with which the message can be authenticated.)

An adversary can notice the simulation by querying the random oracle on the correct seed before the message is revealed to the simulator, disrupting the planned programming. Thus, the security of the construction reduces to the one-wayness of the KEM: if a reduction can recognize such a query (through the use of the plaintext-checking oracle), then it can win the OW-PCA game.

To summarize: BRE is the combination of a KEM, a XOF modelled as an RO, the OTP, and possibly a MAC, with the proof relying on the ability of a simulator to make an OTP ciphertext decrypt to any message by programming the random oracle.

The BRE technique has appeared numerous times, including:

- to achieve IND-CPA and IND-CCA security (as originally conceived) [12];

- generalized as the REACT transformation, which takes any OW-PCA PKE to an IND-CCA PKE [101] and which also allows for other symmetric schemes in place of the OTP;
- modularized as a TagKEM hybrid PKE achieving tight multi-instance IND-CCA security [23];
- to achieve κ-SSO-CCA with sender [66], receiver [93], and bi-opening [91];
- to achieve (single-user) NCE with bi-opening in both CPA [100] and CCA [24,25] settings.

While terminology and implementation details often differ between the constructions (particularly in how exactly they achieve CCA security, for example by use of encrypt-then-MAC [23,66,79,91,101], or by use of MAC-and-encrypt [25]), these schemes all share the same underlying idea. As we see, the technique is very powerful, as it comes close to achieving the strongest notion of our hierarchy: Camenisch et al. [25] gave a construction achieving 1-NCE-CCA⊛, although generalizing to multiple users remains open.

Open Problem 15. *How can the BRE technique be used to achieve* κ-NCE-CCA⊛ *for* $\kappa > 1$*?*

On another note, as pointed out by Jaeger [82], the issue with modularity in the context of NCE lies with challenges in keeping random oracle programming consistent through multiple reductions. (And to paraphrase Krawczyk [89], modularity is simplicity's best friend.) As we have seen (Sect. 3.6), Jaeger resolved this issue by upgrading to notions of SIM*, i.e. notions of a priori simulatability in which all players—simulators, adversaries and reductions alike—are given the ability to reprogram the random oracle, facilitating consistent programming across several game hops. Such an approach could allow for a modularized BRE technique, for instance by first transforming the OW-PCA KEM into a suitably κ-SIM*-CCA⊛ secure TagKEM, and then show that combining it with the OTP yields a hybrid PKE that also achieves κ-SIM*-CCA⊛ (taking inspiration from Brunetta et al. [23]).

Open Problem 16. *To what extent can the BRE technique be used to achieve* κ-SIM*-CCA⊛ *in a modular manner?*

5.2 Other Constructions

Although the BRE technique gives simple and efficient constructions achieving one of the strongest notions of the hierarchy, modelling a primitive like the XOF as a programmable random oracle is, admittedly, a leap of faith; preferably, results are stated in the standard model, reducing to standard (falsifiable) assumptions, such as the pseudorandomness of the XOF's output or the intractability of finding colliding inputs (for a sufficiently long XOF output).

While neither NCE$\star$ nor SSO$\star$ is achievable in the standard model, there is a rich literature of constructions that do in ideal models, as well as those that achieve the various other notions in the standard model. Roughly speaking, moving down the hierarchy and/or from bi-openings to receiver or sender openings usually makes achievability easier in the sense of allowing simpler constructions based on weaker assumptions.

We provide an overview of the various constructions next: Table 2 groups constructions together in families based on their model, underlying assumption, or general approach, and places them in the hierarchy based on the strongest notion that a member of a family has been shown to achieve.

For the ideal models, we primarily consider the aforementioned programmable random oracle model [12], where specifically the simulator is allowed to program (cf. programming by a reduction [46]). An alternative to the programmable ROM is the programmable Ideal Cipher Model (ICM). In the ideal cipher model, a symmetric cipher is modelled as a family of random permutations; the ideal cipher model is primarily known from its non-programming incarnation [18,107,112], with programming relatively rare (whereas for random oracles, programming is fairly established). In Table 2, we concentrate on the ideal model being used by the various constructions. Obviously, there will still be a standard computational assumption needed; these will be explained in the accompanying text.

For constructions in the standard model, a wide variety of different primitives and hardness assumptions have been used to construct PKE schemes secure against various types and notions of opening. These underpinning principles are listed below.

- Lossy Encryption (LE) [10], a PKE scheme with an additional "lossy key generation" algorithm, such that 1) lossy keypairs are indistinguishable from normal ones, and 2) ciphertexts produced under lossy keys statistically hide the message. It follows that correctness cannot hold for lossy keypairs; informally speaking, encrypting with a lossy key should produce only garbage. It also follows that for any lossy keypair $(\mathsf{pk}, \mathsf{sk})$, ciphertext c, and message m, there exists (with high probability) randomness r such that encrypting m under lossy key pk and randomness r results in c, (thus revealing r opens c to m).

Table 2. An overview of the state of the art of constructions achieving the various notions (arrows point in the direction of superseding results). Highlighted in red are impossibility results, and in blue settings concerned with tightness (a priori indistinguishability); constructions that additionally achieve CCA security are highlighted in **bold**.

		Simulability		Indistinguishability	
		A Priori	A Posteriori	A Posteriori	A Priori
Ideal Model	⋆	↓	**ICM** [78]	←	**ROM** [92]
	⊛	**ROM** [25, 100]	←	←	↓
	⊙	↑	**ICM** [67]	←	Open Problem 21
Standard Model	⋆	× [100]	× [113]	**HPS** [87]	Open Problem 20, **MDDH** [58]
	⊛	↑	↑	Open Problem 19	Open Problem 21
	⊙	**HPS** [45]	LEEO [10], **TailKEM** [95, 97], **ABM-LTF** [70, 94]	LE [10]	Cor. 3

- Lossy Encryption with Efficient Opening (LEEO) [10], an LE scheme for which opening, as described above, can be done efficiently.
- All-But-Many Lossy Trapdoor Functions (ABM-LTF) [70], generalizations of Lossy Trapdoor Functions (LTFs) [104], themselves generalizations of trapdoor functions enhanced to two modes of operation: either, given access to a trapdoor, the entire pre-image can be recovered from the image, or a statistically significant amount of information about the pre-image is unrecoverable from the image and trapdoor alone; the two functionalities should furthermore be indistinguishable without access to the trapdoor. Thus they resemble lossy encryption as described above, and can indeed be used to instantiate LE schemes [10]. All-but-many LTFs are LTFs that are additionally parametrized by tags, which can be either injective or lossy, yielding injective or lossy trapdoor functions, respectively. A master trapdoor allows for the generation of an arbitrary number of lossy tags, while the production of lossy tags should be infeasible without access to the master trapdoor.
- Hash Proof Systems (HPS) [36], a primitive hailing back to the Cramer–Shoup encryption scheme [35] that resembles a non-interactive, designated verifier zero-knowledge proof system. HPS come in universal and weak [61] variants, with the latter resembling a KEM variant of lossy encryption [87].
- Tailored KEM (TailKEM) [95], KEMs with two additional properties: firstly, the support of the encapsulation algorithm should cover only a small subset of the full ciphertext space, and secondly, an ephemeral key and its encapsulation, as output by the encapsulation algorithm, should be indistinguishable from elements chosen uniformly at random from the respective spaces; additionally, the TailKEM should satisfy a "tailored" variant of CCCA security [74].
- the Matrix Decisional Diffie–Hellman assumption (MDDH) [44], which is a generalization of DDH and the n-linear assumption (nLin) [74].

As shown in Sect. 3, each layer of the hierarchy implies the layers below (conjectured in the case of SSO-CCA ⇒ ISO-CCA, see Open Problem 5). Therefore, any construction that achieves a notion in the hierarchy will also achieve the notions below it; and likewise, any construction achieving security under bi-opening trivially achieves security under receiver and sender (and transmission) openings. For brevity, we restrict Table 2 to show the most general placement(s) for each category only, with arrows pointing to stronger results. We make an exception for a priori indistinguishability, as asymptotically speaking these notions are all achieved by any IND-CPA/IND-CCA PKE scheme, and hence we focus on works explicitly targeting tightness under the notion in question (highlighted in blue in Table 2).

We will work our way through Table 2 next, first covering idealized settings, then achievability in the standard model for each of the notions going down the hierarchy (left to right in the table), before finishing with the quest for tightness in the a priori indistinguishability setting. Formal treatments of the various construction families are beyond our scope, as is comparing schemes in terms of efficiency and useability;

we will restrict ourselves to brief summaries of each, highlighting open questions of achievability as we go along.

On achievability. When we speak of a notion being achievable, we mean that, under some common computational assumption, there exists an efficient construction achieving (asymptotic) security under the notion, for arbitrary message spaces and for any number of challenges. As neither NCE$\star$ nor SSO$\star$ can be achieved in the standard model so generally (as private key lengths would have to exceed the total number of bits to be encrypted), several standard model constructions only support polynomially sized message spaces [10, 60, 63, 73, 78, 86, 91, 96]. These constructions, while potentially useful for limited applications, do not appear in Table 2.

On tightness. So far, we have used the term "tight" to mean that the multiplicative loss of a concrete bound equals 1. In the following we will instead adopt the asymptotic convention of Han et al. [58] and say that a bound is "tight" if the multiplicative loss is independent of the number of oracle queries q, the number of users κ, and the security parameter λ, and "almost-tight" if the loss is independent of q and κ and at most linear in λ.

Ideal model constructions. We consider various ideal models that allow programming, such as the Random Oracle Model (ROM) [12] and the Ideal Cipher Model (ICM).

Further ROM constructions. Hofheinz et al. [77] showed that for a large class of group-based PKE constructions, IND-CPA implies ISO-CPA$\odot$ when some hash function in the PKE construction is modelled as a non-programmable random oracle and the underlying group is treated generically using a programmable interpretation of Shoup's Generic Group Model (GGM) [108].

Heuer et al. [66] showed that RSA-OAEP [13] achieves SSO-CCA$\odot$, and that the DHIES transformation [1], originally shown to take a OW-PCA KEM to a IND-CCA PKE in the ROM, actually yields PKE schemes achieving SSO-CCA$\odot$. Whether these constructions achieve even stronger notions remains open.

Open Problem 17. *Do the OAEP and DHIES transformations yield PKE schemes achieving* κ-NCE-CCA$\odot$*? What about* κ-NCE-CCA$\circledast$*?*

Heuer [65] showed that the Fujisaki–Okamoto (FO) transformation [48] turns a OW-CPA PKE satisfying a certain condition (γ-spreadness) into an ISO-CCA$\odot$ secure PKE. Pan et al. [103] gave three constructions that tightly (and compactly) achieve SSO-CCA$\odot$ in the ROM, based on the Computational Diffie–Hellman (CDH) assumption, the Strong CDH assumption, and the DDH assumption, respectively. They furthermore showed that the FO transformation tightly turns any LE scheme into an ISO-CCA$\odot$ secure scheme, and any LEEO scheme tightly into an SSO-CCA$\odot$ secure scheme. More recently, Jaeger [82] showed that the FO transformation turns any OW-PCA KEM into a κ-SIM*-CCA$\circledast$ secure PKE, which tightly implies κ-NCE-CCA$\circledast$ security (see Sect. 3.6).

ICM constructions. The programmable ICM allows for strong security guarantees for schemes that already see widespread use, such as KEM/DEM hybrid encryption (with the DEM based on the ideal cipher). The constructions below achieve SSO$\odot$ and SSO$\star$, respectively (neither construction seems to achieve SSO$\circledast$).

Heuer et al. [67] show that any hybrid PKE combining an IND-CCA secure KEM with an NCE-like data encapsulation mechanism ("simulatable" DEM) satisfies SSO-CCA$\odot$. They go on to show that common blockcipher modes of operation such as CTR, CBC, CCM, and GCM achieve this non-committing property in the programmable ideal cipher model.

Huang et al. [78] show that the Canetti–Halevi–Katz (CHK) transformation [30], originally taking any CPA-secure identity-based encryption (IBE) [106] to an IND-CCA secure PKE in the standard model, gives PKE schemes achieving κ-SSO-CCA$\star$ when instantiated with an IBE satisfying the IBE equivalent of SSO-CPA$\star$; such IBE schemes are then constructed in the programmable ideal cipher model. However, the CHK transformation does not seem to apply to the sender opening setting [64], making it a less promising route towards bi-opening resilience.

A priori simulatability (NCE) in the standard model. As already covered (Sect. 3.6), a priori simulatability is not achieveable in the standard model in the presence of receiver openings, nor, as a consequence, for bi-openings (see Sect. 3.6). For sender openings only however, the situation is different: Fehr et al. [45] gave constructions that achieve NCE-CPA$\odot$ and NCE-CCA$\odot$, respectively, in the standard model, from HPS for the former, and from HPS combined with a new primitive called a Cross-Authentication Code (XAC) for the latter. (Interestingly, neither scheme satisfies perfect correctness, cf. Remark 1.)

A variant of NCE has been studied under a restriction wherein each ciphertext need only be decryptable to one of ℓ pre-determined messages [49,62]. For ℓ sufficiently small, the standard model impossibility of NCE$\star$ is thus bypassed. We view the possibilities and limitations of such notions in the standard model as an interesting avenue for further exploration.

Canetti et al. [31] provided an alternative path to overcoming the standard-model impossibility of NCE$\circledast$, namely by letting the private key evolve over time (while keeping the public key fixed). However, their solution does rely on the secure erasure of older private key material, making it somewhat unsatisfactory.

A posteriori simulatability (SSO) in the standard model. Similarly to NCE, we have seen in Sect. 3.5 that a posteriori simulatability is unachievable in the standard model in the presence of receeiver openings, and so we will only discuss constructions for SSO$\odot$ below.

SSO$\odot$ *from LEEO.* Bellare et al. [10] defined lossy encryption, and showed that lossy encryption suffices to achieve ISO-CPA$\odot$ (without being necessary [102]); if opening is furthermore efficient, then the resulting LEEO scheme achieves SSO-CPA$\odot$. They go on to show that such schemes can be constructed in the standard model; indeed, they observe that the classical Goldwasser–Micali encryption scheme [55], based on the quadratic residuosity assumption, is a lossy encryption scheme with efficient opening.

BY12 later updated their definition of κ-SSO-CPA$\odot$ to include multiple challenges, users, and stateful samplers (as detailed in Sect. 3.4), and showed that LEEO achieves also this updated notion (at an additional security loss in the number of users and challenges) [16, Thm. 6.2].

Although LEEO on its own does not suffice to argue SSO-CCA$\odot$ security [45], there are specific LEEO schemes achieving SSO-CCA$\odot$ [94], see below.

SSO$\odot$ *from ABM-LTF.* Hofheinz [70] showed that SSO-CCA$\odot$ can be achieved from ABM-LTFs, and that ABM-LTFs can be constructed from the Decisional Composite Residuosity (DCR) assumption. (They do so by first showing that the scheme satisfies ISO-CCA$\odot$, and then by constructing an efficient opener, making the scheme a LEEO.)

Libert et al. [94] later achieved an almost-tight reduction from ABM-LTFs to SSO-CCA$\odot$, by the construction of a CCA-secure LEEO, while simultaneously constructing ABM-LTFs from Learning With Errors (LWE); however their construction relied on a non-standard PRF property, achieved through reduction to a stronger-than-usual LWE assumption (Non-uniform LWE [21]).

SSO$\odot$ *from TailKEM.* Generalizing the approach of Fehr et al. (see NCE in the standard model paragraph above), Liu et al. [95] introduced TailKEMs and showed how they could be used to construct PKE schemes that achieve SSO-CCA$\odot$ in the standard model. They furthermore showed that TailKEMs can be instantiated from HPS, indistinguishability obfuscation (i$\mathcal{O}$ [4], obfuscating programs so that an adversary with white-box access to the program learns nothing beyond its input–output behaviour, for some definition of "learning nothing"), or nLin.

Subsequently, Lyu et al. [97] showed that further refining the definition to cover multiple instances allows for a tight proof of security, and that such refined TailKEMs can be almost-tightly instantiated from MDDH.

A posteriori indistinguishability (ISO) in the standard model. While SSO$\star$ is unachievable in the standard model, ISO$\star$ has been achieved in the standard model, as detailed below; as has ISO$\odot$, though not ISO$\circledast$ (see Open Problem 19).

On a similar note, for sender opening, the impossibility for committing (binding) schemes to achieve SSO-CPA$\odot$ does not hold for ISO-CPA$\odot$ [9,10].

Generic transformation from IND-CPA *to* κ-ISO-CCA$\star$*.* Jia et al. [87] gave a generic way to lift any κ-ISO-CPA$\star$ scheme to κ-ISO-CCA$\star$ in the standard model based on the Naor–Yung paradigm [99], combining the CPA scheme with an IND-CCA scheme and a non-interactive zero-knowledge proof. Combined with a κ-ISO-CPA$\star$ construction from any weak HPS, and Hazay et al.'s construction of weak HPS from any IND-CPA PKE [61], we get a transformation taking any IND-CPA secure PKE to a κ-ISO-CCA$\star$ secure PKE in the standard model.

Interestingly, the κ-ISO-CPA$\star$ construction [61] was originally aimed at achieving bounded leakage resilience [42], begging the question whether it hints to deeper connections between the two fields of study.

Open Problem 18. *How do notions of leakage resilience relate to notions of selective opening attacks?*

ISO$\star$ *from HPS.* In addition to the above-mentioned transformation, Jia et al. [87] showed that the HPS-based schemes due to Cramer and Shoup [36] achieve κ-ISO-CCA$\star$.

ISO$\odot$ *from LE.* As mentioned, Bellare et al. [10] showed that lossy encryption achieves ISO-CPA$\odot$ even if opening is not efficient. They furthermore showed that, aside from the Goldwasser–Micali encryption scheme (see LEEO paragraph above), LE schemes can be instantiated in the standard model from the DDH assumption, as well as from lossy trapdoor functions. As mentioned in Sect. 3.4, BY12 later updated their security notions to include multiple users and challenges, and showed that the same constructions achieve multi-user, multi-challenge ISO-CPA$\odot$.

ISO$\odot$ *from ABM-LTF.* Hofheinz [70] constructed an ABM-LTF not only from the DCR assumption, but also from pairings,which then yielded a PKE scheme achieving ISO-CCA$\odot$. As the pairing-based ABM-LTF is not known to support efficient opening, its SSO-CCA$\star$ security is left open (unlike the aforementioned DCR-based scheme).

Boyen et al. [22] constructed ABM-LTFs from LWE, yielding a PKE construction achieving ISO-CCA$\odot$. While their scheme achieves a potentially weaker notion of security from LWE compared to the concurrent work Libert et al. [94] (ISO-CCA$\odot$ vs. SSO-CCA$\odot$, see SSO section above), it does have a tight security reduction to standard hardness assumptions.

On the possibility of ISO$\circledast$*.* As both ISO$\odot$ and ISO$\star$ are achievable in the standard model, it stands to reason that a posteriori simulatability with bi-opening is as well, though such a scheme has yet to appear. Given that there are respective HPS schemes achieving ISO$\odot$ [87] and ISO$\star$ [45] in the standard model, possibly ISO$\circledast$ can be achieved from the same primitive in the standard model.

Indeed, Lai et al. [91] gave a standard model HPS construction targeting SSO-CCA under a variant of bi-opening called weak bi-opening (see Sect. 4.5); as their security notion implies κ-SSO-CCA$\star$, their construction is bound by that notion's impossibility result, and indeed their scheme requires private keys whose size is lower bounded by the total number of bits to be encrypted.

Open Problem 19. *How can a posteriori indistinguishability with bi-openings (*κ-ISO-CPA$\circledast$*/*κ-ISO-CCA$\circledast$*) be achieved in the standard model?*

A priori indistinguishability (IND) and tightness. There are of course many schemes that achieve a priori indistinguishability with openings: being polynomially equivalent to single-user, single-challenge IND-CPA (resp. IND-CCA) make the notions mainly interesting in the context of concrete security and tightness.

Early constructions of tightly secure PKE in the multi-user setting did not consider openings [50, 52, 72], yet results establishing that blackbox reductions cannot be tight do rely on receiver opening [3].

Nonetheless, only a few works have to date targeted tight a priori indistinguishability with receiver openings. The first one is a Naor–Yung [99] inspired hybrid PKE that tightly achieves κ-IND-CCA$\star$ in the programmable ROM [92]. The second is a PKE employing a generalization of HPS known as Quasi-Adaptive HPS [59], that achieves κ-IND-CCA$\star$ with an almost-tight reduction to MDDH in the standard model [58].

No scheme to date has tightly achieved IND$\star$, for either CPA or CCA, in the standard model.

Open Problem 20. *How can* κ-IND-CPA$\star$*/*κ-IND-CCA$\star$ *be achieved tightly in the standard model?*

(Note that any scheme tightly achieving ISO$\star$ would also achieve tightness under IND$\star$ via Thm. 4.)

Regarding IND$\odot$ *and* IND$\circledast$. We introduced a priori indistinguishability notions with sender and receiver opening in Sect. 3.3, and so, naturally, the degree to which tightness under these multiple-challenge-bit security notions can be achieved has not been much studied. However, given that there are schemes that tightly achieve ISO-CCA$\odot$ in the standard model [22], and given that ISO-CCA$\odot$ tightly implies β-IND-CCA$\odot$ as follows from Thm. 8 (for $\kappa = 1$, ignoring the receiver opening oracle), we conclude that the latter notion is also tightly achievable.

Corollary 3. β-IND-CCA$\odot$ *is tightly achievable in the standard model.*

Security under multi-challenge-bit indistinguishability notions has also been studied in the context of multi-instance security [11], where tightness was recently achieved in the programmable ROM by Brunetta et al. [23] from a BRE-like hybrid PKE (see Sect. 5.1). Multi-instance security notions fix one challenge bit per key (as opposed to the more general free-bit approach of our notions) and do not consider sender openings. We nonetheless conjecture that there are BRE schemes that tightly achieve (κ, β)-IND-CPA$\circledast$.

Open Problem 21. *Can we (almost-)tightly achieve* (κ, β)-IND-CPA$\circledast$ *in 1) the ROM? 2) the standard model? What about for* CCA*?*

Acknowledgement

The authors would like to thank Joseph Jaeger for many helpful comments and discussions.

References

1. Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (Apr 2001). https://doi.org/10.1007/3-540-45353-9_12
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). 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, Part II. LNCS, vol. 9666, pp. 273–304. Springer, Heidelberg (May 2016). https://doi.org/10.1007/978-3-662-49896-5_10
4. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer, Heidelberg (Aug 2001). https://doi.org/10.1007/3-540-44647-8_1
5. Beaver, D.: Plug and play encryption. In: Kaliski Jr., B.S. (ed.) CRYPTO'97. LNCS, vol. 1294, pp. 75–89. Springer, Heidelberg (Aug 1997). https://doi.org/10.1007/BFb0052228
6. Beaver, D., Haber, S.: Cryptographic protocols provably secure against dynamic adversaries. In: Rueppel, R.A. (ed.) EUROCRYPT'92. LNCS, vol. 658, pp. 307–323. Springer, Heidelberg (May 1993). https://doi.org/10.1007/3-540-47555-9_26
7. 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). https://doi.org/10.1007/3-540-45539-6_18
8. 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). https://doi.org/10.1007/BFb0055718
9. Bellare, M., Dowsley, R., Waters, B., Yilek, S.: Standard security does not imply security against selective-opening. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 645–662. Springer, Heidelberg (Apr 2012). https://doi.org/10.1007/978-3-642-29011-4_38
10. 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 (Apr 2009). https://doi.org/10.1007/978-3-642-01001-9_1
11. 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 (Aug 2012). https://doi.org/10.1007/978-3-642-32009-5_19
12. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) ACM CCS 93. pp. 62–73. ACM Press (Nov 1993). https://doi.org/10.1145/168588.168596
13. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: Santis, A.D. (ed.) EUROCRYPT'94. LNCS, vol. 950, pp. 92–111. Springer, Heidelberg (May 1995). https://doi.org/10.1007/BFb0053428

14. Bellare, M., Yilek, S.: Encryption schemes secure under selective opening attack. Cryptology ePrint Archive, Report 2009/101 (original full version) (2009), https://eprint.iacr.org/2009/101, version 20090302:083605
15. Bellare, M., Yilek, S.: Encryption schemes secure under selective opening attack. Cryptology ePrint Archive, Report 2009/101 (2009), https://eprint.iacr.org/2009/101
16. Bellare, M., Yilek, S.: Encryption schemes secure under selective opening attack. Cryptology ePrint Archive, Report 2009/101 (updated full version) (2012), https://eprint.iacr.org/2009/101, version 20120923:212424
17. Benhamouda, F., Gentry, C., Gorbunov, S., Halevi, S., Krawczyk, H., Lin, C., Rabin, T., Reyzin, L.: Can a public blockchain keep a secret? In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part I. LNCS, vol. 12550, pp. 260–290. Springer, Heidelberg (Nov 2020). https://doi.org/10.1007/978-3-030-64375-1_10
18. Black, J., Rogaway, P., Shrimpton, T., Stam, M.: An analysis of the blockcipher-based hash functions from PGV. Journal of Cryptology **23**(4), 519–545 (Oct 2010). https://doi.org/10.1007/s00145-010-9071-0
19. Böhl, F., Hofheinz, D., Kraschewski, D.: On definitions of selective opening security. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 522–539. Springer, Heidelberg (May 2012). https://doi.org/10.1007/978-3-642-30057-8_31
20. Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 41–69. Springer, Heidelberg (Dec 2011). https://doi.org/10.1007/978-3-642-25385-0_3
21. Boneh, D., Lewi, K., Montgomery, H.W., Raghunathan, A.: Key homomorphic PRFs and their applications. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 410–428. Springer, Heidelberg (Aug 2013). https://doi.org/10.1007/978-3-642-40041-4_23
22. Boyen, X., Li, Q.: All-but-many lossy trapdoor functions from lattices and applications. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 298–331. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_11
23. Brunetta, C., Heum, H., Stam, M.: Multi-instance secure public-key encryption. In: Boldyreva, A., Kolesnikov, V. (eds.) PKC 2023, Part II. LNCS, vol. 13941, pp. 336–367. Springer, Heidelberg (May 2023). https://doi.org/10.1007/978-3-031-31371-4_12
24. Camenisch, J., Lehmann, A., Neven, G., Samelin, K.: Virtual smart cards: How to sign with a password and a server. In: Zikas, V., De Prisco, R. (eds.) SCN 16. LNCS, vol. 9841, pp. 353–371. Springer, Heidelberg (Aug / Sep 2016). https://doi.org/10.1007/978-3-319-44618-9_19
25. Camenisch, J., Lehmann, A., Neven, G., Samelin, K.: UC-secure non-interactive public-key encryption. In: Köpf, B., Chong, S. (eds.) CSF 2017 Computer Security Foundations Symposium. pp. 217–233. IEEE Computer Society Press (2017). https://doi.org/10.1109/CSF.2017.14
26. Camenisch, J., Shoup, V.: Practical verifiable encryption and decryption of discrete logarithms. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 126–144. Springer, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_8
27. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. Cryptology ePrint Archive, Report 2000/067 (2000), https://eprint.iacr.org/2000/067
28. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable encryption. In: Kaliski Jr., B.S. (ed.) CRYPTO'97. LNCS, vol. 1294, pp. 90–104. Springer, Heidelberg (Aug 1997). https://doi.org/10.1007/BFb0052229
29. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: 28th ACM STOC. pp. 639–648. ACM Press (May 1996). https://doi.org/10.1145/237814.238015
30. Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222. Springer, Heidelberg (May 2004). https://doi.org/10.1007/978-3-540-24676-3_13
31. 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 (Feb 2005). https://doi.org/10.1007/978-3-540-30576-7_9
32. Canetti, R., Krawczyk, H., Nielsen, J.B.: Relaxing chosen-ciphertext security. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 565–582. Springer, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_33
33. Canetti, R., Park, S., Poburinnaya, O.: Fully deniable interactive encryption. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part I. LNCS, vol. 12170, pp. 807–835. Springer, Heidelberg (Aug 2020). https://doi.org/10.1007/978-3-030-56784-2_27
34. Coladangelo, A., Goldwasser, S., Vazirani, U.V.: Deniable encryption in a quantum world. In: Leonardi, S., Gupta, A. (eds.) STOC'22. pp. 1378–1391. ACM (2022). https://doi.org/10.1145/3519935.3520019
35. Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO'98. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (Aug 1998). https://doi.org/10.1007/BFb0055717
36. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (Apr / May 2002). https://doi.org/10.1007/3-540-46035-7_4

37. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM Journal on Computing **33**(1), 167–226 (2003). https://doi.org/10.1137/S0097539702403773
38. Damgård, I., Nielsen, J.B.: Improved non-committing encryption schemes based on a general complexity assumption. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 432–450. Springer, Heidelberg (Aug 2000). https://doi.org/10.1007/3-540-44598-6_27
39. Das, A., Dutta, S., Adhikari, A.: Indistinguishability against chosen ciphertext verification attack revisited: The complete picture. In: Susilo, W., Reyhanitabar, R. (eds.) ProvSec 2013. LNCS, vol. 8209, pp. 104–120. Springer, Heidelberg (Oct 2013). https://doi.org/10.1007/978-3-642-41227-1_6
40. Dodis, Y., Grubbs, P., Ristenpart, T., Woodage, J.: Fast message franking: From invisible salamanders to encryptment. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 155–186. Springer, Heidelberg (Aug 2018). https://doi.org/10.1007/978-3-319-96884-1_6
41. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. In: 40th FOCS. pp. 523–534. IEEE Computer Society Press (Oct 1999). https://doi.org/10.1109/SFFCS.1999.814626
42. Dziembowski, S.: Intrusion-resilience via the bounded-storage model. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 207–224. Springer, Heidelberg (Mar 2006). https://doi.org/10.1007/11681878_11
43. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. In: Blakley, G.R., Chaum, D. (eds.) CRYPTO'84. LNCS, vol. 196, pp. 10–18. Springer, Heidelberg (Aug 1984)
44. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (Aug 2013). https://doi.org/10.1007/978-3-642-40084-1_8
45. Fehr, S., Hofheinz, D., Kiltz, E., Wee, H.: Encryption schemes secure against chosen-ciphertext selective opening attacks. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 381–402. Springer, Heidelberg (May / Jun 2010). https://doi.org/10.1007/978-3-642-13190-5_20
46. 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). https://doi.org/10.1007/978-3-642-17373-8_18
47. Fuchsbauer, G., Heuer, F., Kiltz, E., Pietrzak, K.: Standard security does imply security against selective opening for Markov distributions. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016-A, Part I. LNCS, vol. 9562, pp. 282–305. Springer, Heidelberg (Jan 2016). https://doi.org/10.1007/978-3-662-49096-9_12
48. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg (Aug 1999). https://doi.org/10.1007/3-540-48405-1_34
49. Garay, J.A., Wichs, D., Zhou, H.S.: Somewhat non-committing encryption and efficient adaptively secure oblivious transfer. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 505–523. Springer, Heidelberg (Aug 2009). https://doi.org/10.1007/978-3-642-03356-8_30
50. Gay, R., Hofheinz, D., Kohl, L.: Kurosawa-desmedt meets tight security. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 133–160. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_5
51. Gellert, K., Jager, T., Lyu, L., Neuschulten, T.: On fingerprinting attacks and length-hiding encryption. In: Galbraith, S.D. (ed.) CT-RSA 2022. LNCS, vol. 13161, pp. 345–369. Springer, Heidelberg (Mar 2022). https://doi.org/10.1007/978-3-030-95312-6_15
52. Giacon, F., Kiltz, E., Poettering, B.: Hybrid encryption in a multi-user setting, revisited. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part I. LNCS, vol. 10769, pp. 159–189. Springer, Heidelberg (Mar 2018). https://doi.org/10.1007/978-3-319-76578-5_6
53. Goldreich, O.: Foundations of Cryptography: Basic Tools, vol. 1. Cambridge University Press, Cambridge, UK (2001)
54. Goldwasser, S., Micali, S.: Probabilistic encryption and how to play mental poker keeping secret all partial information. In: 14th ACM STOC. pp. 365–377. ACM Press (May 1982). https://doi.org/10.1145/800070.802212
55. Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences **28**(2), 270–299 (1984)
56. Grubbs, P., Lu, J., Ristenpart, T.: Message franking via committing authenticated encryption. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 66–97. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_3
57. Han, S., Liu, S., Gu, D.: Key encapsulation mechanism with tight enhanced security in the multi-user setting: Impossibility result and optimal tightness. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part II. LNCS, vol. 13091, pp. 483–513. Springer, Heidelberg (Dec 2021). https://doi.org/10.1007/978-3-030-92075-3_17
58. Han, S., Liu, S., Gu, D.: Almost tight multi-user security under adaptive corruptions & leakages in the standard model. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part III. LNCS, vol. 14006, pp. 132–162. Springer, Heidelberg (Apr 2023). https://doi.org/10.1007/978-3-031-30620-4_5

59. Han, S., Liu, S., Lyu, L., Gu, D.: Tight leakage-resilient CCA-security from quasi-adaptive hash proof system. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 417–447. Springer, Heidelberg (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_15
60. Hara, K., Kitagawa, F., Matsuda, T., Hanaoka, G., Tanaka, K.: Simulation-based receiver selective opening CCA secure PKE from standard computational assumptions. In: Catalano, D., De Prisco, R. (eds.) SCN 18. LNCS, vol. 11035, pp. 140–159. Springer, Heidelberg (Sep 2018). https://doi.org/10.1007/978-3-319-98113-0_8
61. Hazay, C., López-Alt, A., Wee, H., Wichs, D.: Leakage-resilient cryptography from minimal assumptions. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 160–176. Springer, Heidelberg (May 2013). https://doi.org/10.1007/978-3-642-38348-9_10
62. Hazay, C., Patra, A.: Efficient one-sided adaptively secure computation. Journal of Cryptology **30**(1), 321–371 (Jan 2017). https://doi.org/10.1007/s00145-015-9222-4
63. Hazay, C., Patra, A., Warinschi, B.: Selective opening security for receivers. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 443–469. Springer, Heidelberg (Nov / Dec 2015). https://doi.org/10.1007/978-3-662-48797-6_19
64. Hemenway, B., Libert, B., Ostrovsky, R., Vergnaud, D.: Lossy encryption: Constructions from general assumptions and efficient selective opening chosen ciphertext security. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 70–88. Springer, Heidelberg (Dec 2011). https://doi.org/10.1007/978-3-642-25385-0_4
65. Heuer, F.: On the selective opening security of public-key encryption. Doctoral thesis, Ruhr-Universität Bochum, Universitätsbibliothek (2017)
66. Heuer, F., Jager, T., Kiltz, E., Schäge, S.: On the selective opening security of practical public-key encryption schemes. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 27–51. Springer, Heidelberg (Mar / Apr 2015). https://doi.org/10.1007/978-3-662-46447-2_2
67. Heuer, F., Poettering, B.: Selective opening security from simulatable data encapsulation. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part II. LNCS, vol. 10032, pp. 248–277. Springer, Heidelberg (Dec 2016). https://doi.org/10.1007/978-3-662-53890-6_9
68. Heum, H., Stam, M.: Tightness subtleties for multi-user pke notions. In: Paterson, M.B. (ed.) Cryptography and Coding. pp. 75–104. Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-92641-0_5
69. Hofheinz, D.: Possibility and impossibility results for selective decommitments. Journal of Cryptology **24**(3), 470–516 (Jul 2011). https://doi.org/10.1007/s00145-010-9066-x
70. Hofheinz, D.: All-but-many lossy trapdoor functions. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 209–227. Springer, Heidelberg (Apr 2012). https://doi.org/10.1007/978-3-642-29011-4_14
71. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 341–371. Springer, Heidelberg (Nov 2017). https://doi.org/10.1007/978-3-319-70500-2_12
72. Hofheinz, D., Jager, T.: Tightly secure signatures and public-key encryption. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 590–607. Springer, Heidelberg (Aug 2012). https://doi.org/10.1007/978-3-642-32009-5_35
73. Hofheinz, D., Jager, T., Rupp, A.: Public-key encryption with simulation-based selective-opening security and compact ciphertexts. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 146–168. Springer, Heidelberg (Oct / Nov 2016). https://doi.org/10.1007/978-3-662-53644-5_6
74. Hofheinz, D., Kiltz, E.: Secure hybrid encryption from weakened key encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (Aug 2007). https://doi.org/10.1007/978-3-540-74143-5_31
75. Hofheinz, D., Müller-Quade, J., Steinwandt, R.: On modeling IND-CCA security in cryptographic protocols. Cryptology ePrint Archive, Report 2003/024 (2003), https://eprint.iacr.org/2003/024
76. Hofheinz, D., Rao, V., Wichs, D.: Standard security does not imply indistinguishability under selective opening. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 121–145. Springer, Heidelberg (Oct / Nov 2016). https://doi.org/10.1007/978-3-662-53644-5_5
77. Hofheinz, D., Rupp, A.: Standard versus selective opening security: Separation and equivalence results. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 591–615. Springer, Heidelberg (Feb 2014). https://doi.org/10.1007/978-3-642-54242-8_25
78. Huang, Z., Lai, J., Chen, W., Au, M.H., Peng, Z., Li, J.: Simulation-based selective opening security for receivers under chosen-ciphertext attacks. Des. Codes Cryptogr. **87**(6), 1345–1371 (2019)
79. Huang, Z., Lai, J., Chen, W., ul Haq, M.R., Jiang, L.: Practical public key encryption with selective opening security for receivers. Information Sciences **478**, 15–27 (2019)
80. Huang, Z., Lai, J., Han, S., Lyu, L., Weng, J.: Anonymous public key encryption under corruptions. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part III. LNCS, vol. 13793, pp. 423–453. Springer, Heidelberg (Dec 2022). https://doi.org/10.1007/978-3-031-22969-5_15
81. Huang, Z., Liu, S., Mao, X., Chen, K.: Non-malleability under selective opening attacks: Implication and separation. In: Malkin, T., Kolesnikov, V., Lewko, A.B., Polychronakis, M. (eds.) ACNS 15. LNCS, vol. 9092, pp. 87–104. Springer, Heidelberg (Jun 2015). https://doi.org/10.1007/978-3-319-28166-7_5

82. Jaeger, J.: Let attackers program ideal models: Modularity and composability for adaptive compromise. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part III. LNCS, vol. 14006, pp. 101–131. Springer, Heidelberg (Apr 2023). https://doi.org/10.1007/978-3-031-30620-4_4
83. Jaeger, J.: Personal communication (2023)
84. Jager, T., Kiltz, E., Riepel, D., Schäge, S.: Tightly-secure authenticated key exchange, revisited. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 117–146. Springer, Heidelberg (Oct 2021). https://doi.org/10.1007/978-3-030-77870-5_5
85. 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, Part I. LNCS, vol. 10677, pp. 409–441. Springer, Heidelberg (Nov 2017). https://doi.org/10.1007/978-3-319-70500-2_14
86. Jia, D., Lu, X., Li, B.: Receiver selective opening security from indistinguishability obfuscation. In: Dunkelman, O., Sanadhya, S.K. (eds.) INDOCRYPT 2016. LNCS, vol. 10095, pp. 393–410. Springer, Heidelberg (Dec 2016). https://doi.org/10.1007/978-3-319-49890-4_22
87. Jia, D., Lu, X., Li, B.: Constructions secure against receiver selective opening and chosen ciphertext attacks. In: Handschuh, H. (ed.) CT-RSA 2017. LNCS, vol. 10159, pp. 417–431. Springer, Heidelberg (Feb 2017). https://doi.org/10.1007/978-3-319-52153-4_24
88. Joye, M., Quisquater, J.J., Yung, M.: On the power of misbehaving adversaries and security analysis of the original EPOC. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 208–222. Springer, Heidelberg (Apr 2001). https://doi.org/10.1007/3-540-45353-9_16
89. Krawczyk, H.: The joy of cryptography: A personal journey (2023), https://crypto.iacr.org/2023/files/667slides.pdf, IACR Distinguished Lecture, presented at Crypto'23
90. Küsters, R., Tuengerthal, M.: Joint state theorems for public-key encryption and digital signature functionalities with local computation. In: Sabelfeld, A. (ed.) CSF 2008 Computer Security Foundations Symposium. pp. 270–284. IEEE Computer Society Press (2008). https://doi.org/10.1109/CSF.2008.18
91. Lai, J., Yang, R., Huang, Z., Weng, J.: Simulation-based bi-selective opening security for public key encryption. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part II. LNCS, vol. 13091, pp. 456–482. Springer, Heidelberg (Dec 2021). https://doi.org/10.1007/978-3-030-92075-3_16
92. 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)
93. Lei, F., Chen, W., Chen, K.: A non-committing encryption scheme based on quadratic residue. In: International Symposium on Computer and Information Sciences. pp. 972–980. Springer (2006)
94. Libert, B., Sakzad, A., Stehlé, D., Steinfeld, R.: All-but-many lossy trapdoor functions and selective opening chosen-ciphertext security from LWE. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 332–364. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_12
95. Liu, S., Paterson, K.G.: Simulation-based selective opening CCA security for PKE from key encapsulation mechanisms. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 3–26. Springer, Heidelberg (Mar / Apr 2015). https://doi.org/10.1007/978-3-662-46447-2_1
96. Lu, Y., Hara, K., Tanaka, K.: Receiver selective opening CCA secure public key encryption from various assumptions. In: Nguyen, K., Wu, W., Lam, K.Y., Wang, H. (eds.) ProvSec 2020. LNCS, vol. 12505, pp. 213–233. Springer, Heidelberg (Nov / Dec 2020). https://doi.org/10.1007/978-3-030-62576-4_11
97. Lyu, L., Liu, S., Han, S., Gu, D.: Tightly SIM-SO-CCA secure public key encryption from standard assumptions. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part I. LNCS, vol. 10769, pp. 62–92. Springer, Heidelberg (Mar 2018). https://doi.org/10.1007/978-3-319-76578-5_3
98. Micali, S., Rackoff, C., Sloan, B.: The notion of security for probabilistic cryptosystems. In: Odlyzko, A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 381–392. Springer, Heidelberg (Aug 1987). https://doi.org/10.1007/3-540-47721-7_27
99. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC. pp. 427–437. ACM Press (May 1990). https://doi.org/10.1145/100216.100273
100. Nielsen, J.B.: Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 111–126. Springer, Heidelberg (Aug 2002). https://doi.org/10.1007/3-540-45708-9_8
101. Okamoto, T., Pointcheval, D.: REACT: Rapid Enhanced-security Asymmetric Cryptosystem Transform. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 159–175. Springer, Heidelberg (Apr 2001). https://doi.org/10.1007/3-540-45353-9_13
102. Ostrovsky, R., Rao, V., Visconti, I.: On selective-opening attacks against encryption schemes. In: Abdalla, M., Prisco, R.D. (eds.) SCN 14. LNCS, vol. 8642, pp. 578–597. Springer, Heidelberg (Sep 2014). https://doi.org/10.1007/978-3-319-10879-7_33
103. Pan, J., Zeng, R.: Compact and tightly selective-opening secure public-key encryption schemes. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part III. LNCS, vol. 13793, pp. 363–393. Springer, Heidelberg (Dec 2022). https://doi.org/10.1007/978-3-031-22969-5_13
104. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 187–196. ACM Press (May 2008). https://doi.org/10.1145/1374376.1374406
105. Shamir, A.: How to share a secret. Communications of the Association for Computing Machinery **22**(11), 612–613 (Nov 1979)

106. Shamir, A.: Identity-based cryptosystems and signature schemes. In: Blakley, G.R., Chaum, D. (eds.) CRYPTO'84. LNCS, vol. 196, pp. 47–53. Springer, Heidelberg (Aug 1984)
107. Shannon, C.E.: Communication theory of secrecy systems. Bell Systems Technical Journal **28**(4), 656–715 (1949)
108. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT'97. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (May 1997). https://doi.org/10.1007/3-540-69053-0_18
109. Steinfeld, R., Baek, J., Zheng, Y.: On the necessity of strong assumptions for the security of a class of asymmetric encryption schemes. In: Batten, L.M., Seberry, J. (eds.) ACISP 02. LNCS, vol. 2384, pp. 241–256. Springer, Heidelberg (Jul 2002). https://doi.org/10.1007/3-540-45450-0_20
110. Tezcan, C., Vaudenay, S.: On hiding a plaintext length by preencryption. In: Lopez, J., Tsudik, G. (eds.) ACNS 11. LNCS, vol. 6715, pp. 345–358. Springer, Heidelberg (Jun 2011). https://doi.org/10.1007/978-3-642-21554-4_20
111. Watanabe, Y., Shikata, J., Imai, H.: Equivalence between semantic security and indistinguishability against chosen ciphertext attacks. In: Desmedt, Y. (ed.) PKC 2003. LNCS, vol. 2567, pp. 71–84. Springer, Heidelberg (Jan 2003). https://doi.org/10.1007/3-540-36288-6_6
112. Winternitz, R.S.: A secure one-way hash function built from DES. In: IEEE Symposium on Security and Privacy. pp. 88–90. IEEE Computer Society (1984). https://doi.org/10.1109/SP.1984.10027
113. Yang, R., Lai, J., Huang, Z., Au, M.H., Xu, Q., Susilo, W.: Possibility and impossibility results for receiver selective opening secure PKE in the multi-challenge setting. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part I. LNCS, vol. 12491, pp. 191–220. Springer, Heidelberg (Dec 2020). https://doi.org/10.1007/978-3-030-64837-4_7

A ISO Tightly Implies IND with Bi-openings

From the fact that (κ, β)-IND-CCA⊛ is implied by κ-IND-CCA⋆ with a β loss (Thm. 3), we concluded in Sect. 3.4 that κ-ISO-CCA⋆ implies (κ, β)-IND-CCA⊛ with a $2 \cdot \beta$ security loss. We next show that there is a message sampler such that the reduction only loses a factor 2.

Lemma 6. *Let $\mathsf{M}_{\langle s \rangle}$ be as given in Fig. 19 (left): its input α is interpeted as (j, m_0, m_1), where the message pair m_0, m_1 is subject to $|m_0| = |m_1|$, and index j is subject to $j \in [\beta]$. On first invocation (when $s = \varepsilon$), it draws $s \leftarrow\!\!\$ \{0,1\}^\beta$ and, on all invocations, on input $\alpha = (j, m_0, m_1)$, it returns $m_{s[j]}$. Let $\mathring{\mathsf{S}}$ be its ideal resampler.*

Consider S (Fig. 19, right) that on input $(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$, for each compromised call sets $s'[j] \leftarrow s[j]$ for the relevant j, and for uncompromised j draws $s'[j] \leftarrow\!\!\$ \{0,1\}$, and appends $m_{s'[j]}$ to the resampled message vector $\mathtt{M}^1$; and, once this process is completed for all challenge calls, returns $\mathtt{M}^1$.

Then $\mathsf{S} = \mathring{\mathsf{S}}$.

Proof. By inspection.

Theorem 8. *Let $\mathsf{PKE}[\lambda]$ be given, and let $\mathsf{M}_{\langle s \rangle}$ be as given in Lemma 6. Then there is a type-preserving black-box reduction $\mathbb{B}_{\text{iso}}$ such that, for all $\mathbb{A}_{\text{ind}}$,*

$$\mathsf{Adv}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}}) \leq 2 \cdot \mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}(\mathbb{B}_{\text{iso}}) \,,$$

where $\mathbb{B}_{\text{iso}}$ calls $\mathbb{A}_{\text{ind}}$ once.

The proof follows in the vein of that of Thm. 4, only updated to accomodate the message sampler given in Lemma 6 and the more general experiments.

Proof. Without loss of generality, we may assume that the adversary $\mathbb{A}_{\text{ind}}$ does not call $\mathcal{E}_{\mathbb{A}}(i, j, m_0, m_1)$ with either $|m_0| \neq |m_1|$ or $m_0 = m_1$, and that if $\mathbb{A}_{\text{ind}}$ outputs guess $(j, \hat{b}_j)$ then $\mathbb{A}_{\text{ind}}$ made a call to $\mathcal{E}$ with bit handle j at least once, and did not compromise b_j through calls to its opening oracles. Technically, one can create an intermediate reduction $\mathbb{B}_{\text{ind}}$ that runs $\mathbb{A}_{\text{ind}}$ and, playing the same game, forwards everything to its own oracles but those pointless calls to $\mathcal{E}$ (which it can easily simulate), and in the event that $\mathbb{A}_{\text{ind}}$ either halts with an index j that has not been challenged, or halts with an index j that has been compromised, or tries to compromise the last uncompromised challenge bit, $\mathbb{B}_{\text{ind}}$ chooses a bit handle j' that has been challenged at least once (making an extra call to $\mathcal{E}$ in the case that $\mathbb{A}_{\text{ind}}$ halts without ever calling $\mathcal{E}$), samples $\hat{b}_{j'} \leftarrow\!\!\$ \{0,1\}$, and terminates with the output $(j', \hat{b}_{j'})$. By inspection, $\mathbb{B}_{\text{ind}}$'s advantage is at least $\mathbb{A}_{\text{ind}}$'s.

```
Sampler $\mathsf{M}_{\langle s\rangle}(j, m_0, m_1)$
if $s = \varepsilon : s \leftarrow\$ \{0,1\}^{\beta}$
return $m_{s[j]}$
```

```
Resampler $\mathsf{S}(\mathtt{A}, \mathtt{L}, \mathcal{J}, \mathtt{M}^0[\mathcal{J}])$
for $k \in \mathcal{J}$ :
    $(j, m_0, m_1) \leftarrow \mathtt{A}[k]$
    if $s'[j] = \varepsilon \wedge m_0 \neq m_1$ :
        if $\exists b \in \{0,1\}$ s.t. $\mathtt{M}^0[k] = m_b : s'[j] \leftarrow b$
        else return $\lightning$
for $k \in |\mathtt{A}|$ :
    $(j, m_0, m_1) \leftarrow \mathtt{A}[k]$
    if $s'[j] = \varepsilon : s'[j] \leftarrow\$ \{0,1\}$
    $\mathtt{M}^1 \overset{\frown}{\longleftarrow} m_{s'[j]}$
return $\mathtt{M}^1$
```

Fig. 19. The M and S of Lemma 6.

```
Reduction $\mathbb{B}_{\text{iso}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
$k \leftarrow 0$
$(j, \hat{s}[j]) \leftarrow\$ \mathbb{A}_{\text{ind}}^{\mathcal{E},\mathcal{D},(\mathcal{T},)\mathcal{S},\mathcal{R}}(\mathsf{pk}_1, \ldots, \mathsf{pk}_\kappa)$
if $k_j = \varepsilon : \hat{b} \leftarrow\$ \{0,1\}$
else : $\mathtt{M}^b \leftarrow \mathcal{C}$
    $(j, m_0, m_1) \leftarrow \mathtt{A}[k_j]$
    if $\mathtt{M}^b[k_j] = m_{\hat{s}[j]} : \hat{b} \leftarrow 0$
    else : $\hat{b} \leftarrow 1$
return $\hat{b}$
```

```
If $\mathbb{A}_{\text{ind}}$ calls $\mathcal{E}(i, j, m_0, m_1)$
if $|m_0| \neq |m_1|$ : return $\lightning$
$k \leftarrow k + 1, k_j \leftarrow k$
$\alpha \leftarrow (j, m_0, m_1), \mathtt{A} \overset{\frown}{\longleftarrow} \alpha$
$c \leftarrow \mathcal{E}_{\mathbb{B}}(i, \alpha)$
return $c$
```

Fig. 20. The reduction $\mathbb{B}_{\text{iso}}$ simulating $\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{(\kappa,\beta)\text{-ind-cca}\circledast}$ for $\mathbb{A}_{\text{ind}}$. The decryption oracle and opening oracles are simply forwarded.

Let $\mathbb{B}_{\text{iso}}$ be as given in Fig. 20: if $\mathbb{A}_{\text{ind}}$ calls $\mathcal{E}_{\mathbb{A}}(i, j, m_0, m_1)$, subject to both $|m_0| = |m_1|$ and $m_0 \neq m_1$, it sets $\alpha = (j, m_0, m_1)$ and calls $\mathcal{E}_{\mathbb{B}}(i, \alpha)$, returning the resulting c; if $\mathbb{A}_{\text{ind}}$ calls any other oracles, $\mathbb{B}_{\text{iso}}$ forwards the call and returns the result. When $\mathbb{A}_{\text{ind}}$ halts with a guess $\hat{s}[j]$, $\mathbb{B}_{\text{iso}}$ calls $\mathcal{C}$ and receives $\mathtt{M}^b$. Since $\mathbb{B}_{\text{iso}}$ can maintain its own perfect copy of $\mathtt{A}$, it can check whether calls made to bit handle j are consistent with the guess $\hat{s}[j]$ (checking one such call suffices, as the rest are guaranteed to be consistent with it). If so, $\mathbb{B}_{\text{iso}}$ halts with output $\hat{b} = 0$ (indicating a guess that the returned messages were the real ones), otherwise $\mathbb{B}_{\text{iso}}$ halts with output $\hat{b} = 1$.

We can rephrase $\mathbb{B}_{\text{iso}}$'s distinguishing advantage

$$\mathsf{Adv}_{\mathsf{PKE}[\lambda],\mathsf{M},\mathring{\mathsf{S}}}^{\kappa\text{-iso-cca}\circledast}(\mathbb{B}_{\text{iso}}) = \Pr\Big[\hat{b} = 0 \;\Big|\; b = 0\Big] - \Pr\Big[\hat{b} = 0 \;\Big|\; b = 1\Big]$$

and analyse each term individually. Based on $\mathbb{B}_{\text{iso}}$'s description, and given the assumptions that j was used as part of at least one challenge oracle call, the event $\hat{b} = 0$ is equivalent to the event $\mathtt{M}^b[k_j] = m_{\hat{s}[j]}$, where k_j represents the last call to $\mathcal{E}$ using bit handle j, and m_0, m_1 are the messages that were provided for that call.

If $b = 0$, then $\mathtt{M}^0[k_j] = m_{s[j]}$, and so (given our assumption that calls are not made with $m_0 = m_1$) the first term is equivalent to $\Pr[\hat{s}[j] = s[j] \mid b = 0]$. At this point the conditional becomes irrelevant, as it is independent of both $\hat{s}$ and s (jointly). Finally, $\Pr[\hat{s}[j] = s[j]]$ equals $\Pr\Big[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{(\kappa,\beta)\text{-ind-cca}\circledast}(\mathbb{A}_{\text{ind}})\Big]$ as, by design of $\mathbb{B}_{\text{iso}}$ and M, $\mathbb{A}_{\text{iso}}$ is provided with an environment that perfectly matches that of $\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{(\kappa,\beta)\text{-ind-cca}\circledast}(\cdot)$, where the individual bits of M's bit string s play the role of the challenge bits, one free choice of which $\mathbb{A}_{\text{ind}}$ has to guess. Thus,

$$\Pr\Big[\hat{b} = 0 \;\Big|\; b = 0\Big] = \Pr\Big[\mathsf{Exp}_{\mathsf{PKE}[\lambda]}^{(\kappa,\beta)\text{-ind-cca}\circledast}(\mathbb{A}_{\text{ind}})\Big]\,.$$

If $b = 1$, then $\mathtt{M}^1[k_j] = m_{s'[j]}$, so the second term is equivalent to $\Pr[\hat{s}[j] = s'[j] \mid b = 1]$. As we assumed $\mathbb{A}_{\text{ind}}$ did not make opening oracle calls that would compromise $s[j]$, $\mathring{\mathsf{S}}$ will have drawn $s'[j]$

uniformly at random, independently of $\hat{s}[j]$. Thus,

$$\Pr\left[\hat{b} = 0 \,\middle|\, b = 1\right] = \frac{1}{2}\,.$$

Putting the pieces together, we obtain

$$\begin{aligned} 2 \cdot \mathsf{Adv}^{\kappa\text{-iso-cca}\circledast}_{\mathsf{PKE}[\lambda],\mathsf{M},\hat{\mathsf{S}}}(\mathbb{B}_{\text{iso}}) &= 2 \cdot \Pr\left[\mathsf{Exp}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}})\right] - 1 \\ &\geq \mathsf{Adv}^{(\kappa,\beta)\text{-ind-cca}\circledast}_{\mathsf{PKE}[\lambda]}(\mathbb{A}_{\text{ind}})\,, \end{aligned}$$

where the final inequality takes into account the intermediate $\mathbb{B}_{\text{ind}}$ reduction to justify our assumption on $\mathbb{A}_{\text{ind}}$'s behaviour (from the beginning of the proof). □