Selective Opening Security for Receivers

Carmit Hazay[1(✉)], Arpita Patra[2], and Bogdan Warinschi[3]

[1] Faculty of Engineering, Bar-Ilan University, Ramat Gan, Israel
carmit.hazay@biu.ac.il
[2] Department of Computer Science & Automation,
Indian Institute of Science, Bengaluru, India
arpita@csa.iisc.ernet.in
[3] Department of Computer Science, University of Bristol, Bristol, UK
csxbw@bristol.ac.uk

Abstract. In a selective opening (SO) attack an adversary breaks into a subset of honestly created ciphertexts and tries to learn information on the plaintexts of some untouched (but potentially related) ciphertexts. Contrary to intuition, standard security notions do not always imply security against this type of adversary, making SO security an important standalone goal. In this paper we study *receiver security*, where the attacker is allowed to obtain the decryption keys corresponding to some of the ciphertexts.

First we study the relation between two existing security definitions, one based on simulation and the other based on indistinguishability, and show that the former is strictly stronger. We continue with feasibility results for both notions which we show can be achieved from (variants of) non-committing encryption schemes. In particular, we show that indistinguishability-based SO security can be achieved from a tweaked variant of non-committing encryption which, in turn, can be instantiated from a variety of basic, well-established, assumptions. We conclude our study by showing that SO security is however strictly weaker than all variants of non-committing encryption that we consider, leaving potentially more efficient constructions as an interesting open problem.

Keywords: Selective opening attacks · Encryption schemes · Non-committing encryption

1 Introduction

Security notions for encryption come in many forms that reflect different attacker goals (e.g. one-wayness, indistinguishability for plaintexts or non-malleability of ciphertexts), variations in possible attack scenarios (e.g. chosen plaintext or ciphertext attacks) and definitional paradigms (e.g. through games or simulation). A class of attacks motivated by practical considerations are those where the adversary may perform *selective openings* (SO). Here, an adversary is allowed to break into a subset of honestly created ciphertexts leaving untouched other (potentially related) ciphertexts.

T. Iwata and J.H. Cheon (Eds.): ASIACRYPT 2015, Part I, LNCS 9452, pp. 443–469, 2015.
DOI: 10.1007/978-3-662-48797-6_19

This attack scenario was first identified in the context of adaptively secure multi-party computation (MPC) where communication is over encrypted channels visible to the adversary. The standard trust model for MPC considers an adversary who, based on the information that he sees, can decide to corrupt parties and learn their internal state. In turn, this may allow the attacker to determine the parties' long term secret keys and/or the randomness used to create the ciphertexts. The broader context of Internet communication also naturally gives rise to SO attacks. Attackers that access and store large amount of encrypted Internet traffic are a reality, and getting access to the internal states of honest parties can be done by leveraging design or implementation weaknesses of deployed systems. For example the Heartbleed attack allowed a remote party to extract (among other things) the encryption keys used to protect OpenSSL connections.

Security against SO attacks comes in several distinct flavors. Depending on the attack scenario, we distinguish two settings that fall under the general idea of SO attacks. In *sender security*, we have n senders and one receiver. The receiver holds a secret key relative to a public key known to all the senders. The senders encrypt messages for the receiver and the adversary is allowed to corrupt some of the senders (and learn the messages and randomness underlying some of the ciphertexts). The concern is that the messages sent by uncorrupted senders stay secret. The second scenario deals with *receiver security*. Here we consider one sender and n receivers who hold independently generated public and secret keys. The attacker is allowed to corrupt some of the receivers (and learn the secret keys that decrypt some of the observed ciphertexts). Security in this setting is concerned with the messages received by uncorrupted receivers. For each of these settings, security can be defined using either the standard indistinguishability paradigm or simulation-based definitions. Importantly, both scenarios capture realistic attacks in secure computation where usually every party acts as either a sender or a receiver at some point of time during a protocol execution.

Since most of the existent encryption schemes have been analyzed w.r.t. traditional notions of security (e.g. indistinguishability under chosen plaintext or chosen ciphertext attacks (**ind-cpa**, **ind-cca**)), a central question in this area is to understand how security against SO attacks relates to the established definitions. Despite compelling intuition that the only information that an adversary obtains is what it can glean from the opened plaintexts, progress towards confirming or disproving this conjecture has been rather slow. Perhaps the most interesting and surprising results are due to Bellare et al. [1,2] who showed that selective sender security as captured via *simulation* based definitions is strictly stronger than indistinguishability under chosen plaintext attacks [15] (denoted by **ind-cpa** security). The gap between standard notions of security and SO security is uncomfortable: while SO attacks may naturally occur we do not have a clear understanding of the level of security that existing constructions offer nor do we have many ideas on how to achieve security against such attacks.

In this paper we study receiver security. This setting is less studied than sender security yet it corresponds to more plausible attacks (e.g. the Heartbleed

attack). In a nutshell, we clarify the relation between various security notions for receiver security and propose novel constructions. Before we describe our contributions in detail we overview existing work in the area and take this opportunity to introduce more carefully the different security notions of SO security.

1.1 Related Work

Selective opening attacks were first introduced in [12] in the context of commitment schemes. In the context of encryption schemes, the first rigorous definitions were proposed by Bellare, Hofheinz and Yilek [2]. They studied SO security for public key encryption (PKE), for both the receiver and the sender settings and for each setting proposed two types of definitions, indistinguishability-based and simulation-based ones.

Very roughly, the indistinguishability-based definition (denoted by **ind-so**) requires that an adversary that sees a vector of ciphertexts cannot distinguish the true plaintexts of the unopened ciphertexts from independently sampled plaintexts. This is required even with access to the randomness used for generating the opened ciphertexts (in the sender corruption setting), or with access to the secret keys that decrypt the opened ciphertexts (in the receiver corruption setting). This definition requires messages to come from a distribution that is *efficiently resamplable*. A stronger security variant that does not restrict the message distribution called *full* **ind-so** has been introduced later by Böhl, Hofheinz and Kraschewski [5]. The simulation based notion (denoted by **sim-so**) is reminiscent of the definitional approach of Dwork et al. [12] and requires computational indistinguishability between an idealized execution and the real one.

The first feasibility results for security against SO attacks are for the sender setting and leverage an interesting relation with lossy encryption: a lossy PKE implies **ind-so** for sender security [2]. Furthermore, if the PKE scheme has an efficient opening algorithm of ciphertexts, then the scheme also satisfies **sim-so** security. The work of Hemenway et al. [18] shows that lossy (and therefore **ind-so**) PKE can be constructed based on several generic cryptographic primitives.

For primitives that benefit from multiple security notions, a central question is to understand how these notions relate to each other. This type of results are important as they clarify the limitations of some of the notions and enable tradeoffs between security and efficiency (to gain efficiency, a scheme with weaker guarantees may be employed, if the setting allows it). The relation between traditional security notions of encryption and security against SO attacks was a long-standing open problem that was solved by Bellare et al. [1]. Their result is that standard **ind-cpa** security *does not* imply **sim-so** (neither in the sender nor in the receiver setting). There is no fully satisfactory result concerning the relation between **ind-cpa** and **ind-so**. Here, the best result is that these two notions imply each other in the generic group model [19] and that for the chosen-chiphertext attacks variant (CCA) the two notions are distinct.

Relations between the different notions for selective opening have mainly been studied in the sender setting. Böhl et al. establish that full **ind-so** and **sim-so** are incomparable. Recently, [23] introduced an even stronger variant of the full **ind-so**

definition, and showed that many **ind-cpa**, **ind-so** and **sim-so** secure encryption schemes are insecure according to their new notion. They further showed that **sim-so** definition does not imply lossy encryption even without efficient openability. Finally, SO security has been considered for CCA attacks [13,20] and in the identity-based encryption scenario [3].

1.2 Our Contribution

With only two exceptions [1,2] prior work on SO security has addressed mainly the sender setting. We concentrate on the receiver setting. Though theoretically the feasibility for SO security for the receiver is implied by the existence of non-committing encryption schemes [6,8,9,22], the state of the art constructions still leave many interesting open problems in terms of relations between notions and feasibility results. This is the focus of this work.

For relation between notions, similarly to prior separating results in the SO setting [5,19,23], we demonstrate the existence of a separating scheme that is based on generic assumptions and can be instantiated under various concrete assumptions. For constructions, we find it useful to leverage the close relation between (variants of) non-committing encryption and security under SO attacks. For example, we show that **ind-so** security follows from a tweaked variant of non-committing encryption which, in turn, we show how to instantiate from a variety of standard assumptions. Interestingly, we also show a separation between SO security and non-committing encryption (which leaves open the question of potentially more efficient constructions that meet the former notion but not the latter). Below, we elaborate on our results in details.

Notation-wise, we denote the indistinguishability and simulation-based definitions in the receiver setting by **rind-so** and **rsim-so**, respectively. For the corresponding notions in the sender setting we write **sind-so** and **ssim-so**, respectively. That is, we prepend "s" or "r" to indicate if the definition is for sender security or receiver security.

The relation between **rind-so** *and* **rsim-so**. First, we study the relation between the indistinguishability and simulation-based security notions in the receiver setting. We establish that the **rind-so** notion is strictly weaker (and therefore easier to realize) than the notion of **rsim-so**, by presenting a concrete pubic key scheme that meets the former but not the latter level of security. Loosely speaking, a ciphertext includes a commitment to the plaintext together with encryptions of the opening information of this commitment (namely, the plaintext and the corresponding randomness). We then prove that when switching to an alternative fake mode the hiding properties of our building blocks (commitment and encryption schemes) imply that the ciphertext does not contain any information about the plaintext. Nevertheless, simulation always fails since it would require breaking the binding property of the commitment. Applying the

observation that **rsim-so** implies **rind-so** security,[1] we obtain the result that **rind-so** is strictly weaker.

In more details, our separating scheme is built from a commitment scheme and a primitive called non-committing encryption for the receiver (NCER) [7] that operates in two indistinguishable ciphertexts modes: valid and fake, where a fake ciphertext can be decrypted into any plaintext using a proper secret key. This property is referred to as *secret key equivocation* and is implied by the fact that fake ciphertexts are lossy which, in turn, implies **rind-so** security. Specifically, the security of our scheme implies that:

Theorem 1.1 (Informal). *There exists a PKE that is* **rind-so** *secure but is not* **rsim-so** *secure.*

Somewhat related to our work, [1] proved that the standard **ind-cpa** security does not imply **rsim-so** security via the notion of *decryption verifiability* – the idea that it is hard to decrypt a ciphertext into two distinct messages (even using two different secret keys). Specifically, [1] showed that any **ind-cpa** secure PKE that is decryption verifiable cannot be **rsim-so** secure. Compared with their result, our result implies that **rsim-so** security is strictly stronger than **rind-so** security (which may turn out to be stronger than **ind-cpa** security).

The feasibility of **rind-so** *and* **rsim-so**. We recall that in the sender setting, the notions **sind-so** and **ssim-so** are achievable from lossy encryption and lossy encryption with efficient openability.[2] We identify a security notion (and a variant) which plays for receiver security the role that lossy encryption plays in sender security. Specifically, we prove that NCER implies **rsim-so** and that a variant of NCER, which we refer as *tweaked NCER* (formally defined in Tweaked NCER subsection of Sect. 3), implies **rind-so**. Loosely speaking, the security of tweaked NCER is formalized as follows. Similarly to NCER, tweaked NCER has the ability to create fake ciphertexts that are computationally indistinguishable from real ciphertexts. Nevertheless, while in NCER a fake ciphertext can be efficiently decrypted to any plaintext (by producing a matching secret key), in tweaked NCER a fake ciphertext can only be efficiently decrypted to a concrete predetermined plaintext. Informally, our results are captured by the following theorem:

Theorem 1.2 (Informal). *Assume the existence of tweaked NCER and NCER, then there exist PKE schemes that are* **rind-so** *and* **rsim-so** *secure, respectively.*

[1] This can be derived from the fact that the adversary's view is identical for any two simulated executions with different sets of unopened messages, as the simulator never gets to see these messages.

[2] Recall that a lossy encryption scheme is a public key encryption with the additional ability to generate fake indistinguishable public keys so that a fake ciphertext (that is generated using a fake public key) is lossy and is a non-committing ciphertext with respect to the plaintext. A lossy encryption implies the existence of an opening algorithm (possibly inefficient) that can compute a randomness for a given fake ciphertext and a message.

Interestingly, we show that the converse implications do not hold. That is, a **rsim-so** secure PKE is not necessarily a tweaked NCER or a NCER. This further implies that a **rind-so** secure PKE is not necessarily a tweaked NCER or NCER. This result is reminiscent of the previous result that **sim-so** and **rind-so** secure PKE do not imply lossy encryption even without efficient openability [23].

Our separating scheme is based on an arbitrary key-simulatable PKE scheme. Intuitively, in such schemes, it is possible to produce a public key without sampling the corresponding secret key. The set of obliviously sampled public keys may be larger than the the set of public keys sampled together with their associated secret key, yet it is possible to explain a public key sampled along with a secret key as one sampled without. In these schemes we also require that the two type of keys are also computationally indistinguishable. Our proof holds for the case that the set of obliviously sampled keys is indeed larger, so that not every obliviously sampled public key can be explained to possess a secret key. In summary, we prove that:

Theorem 1.3 (Informal). *Assume the existence of key-simulatable PKE, then there exists a PKE scheme that is* **rsim-so** *secure but is neither tweaked NCER nor NCER.*

Our constructions show that **rsim-so** (and **rind-so**) security can be achieved under the same assumptions as key-simulatable PKE – there are results that show that the latter can be constructed from a variety of hardness assumptions such as Decisional Diffie-Hellman (DDH) and Decisional Composite Residuosity (DCR). They also show that we can construct schemes from any hardness assumption that implies simulatable PKE [9] (where both public keys and ciphertexts can be obliviously sampled).

Realizing tweaked NCER. Finally, we demonstrate the broad applicability of this primitive and show how to construct it from various important primitives: key-simultable PKE, two-round honest-receiver statistically-hiding $\binom{2}{1}$ oblivious transfer (OT) and hash proof systems (HPS). We stress that it is not known how to build NCER under these assumptions (or any other generic assumption), which implies that tweaked NCER is much easier to realize. In addition, we prove that the two existing NCER schemes [7] with security under the DDH and DCR hardness assumptions imply the tweaked NCER notion, where surprisingly, the former construction that is a secure NCER for only polynomial-size message spaces, is a tweaked NCER for exponential-size message spaces (this further hints that tweaked NCER may be constructed more efficiently than NCER). These results imply that tweaked NCER (and thus **rind-so**) can be realized based on DDH, DCR, RSA, factoring and learning with errors (LWE) hardness assumptions.

Our results are summarized in Fig. 1.

The relation between **sind-so** *and* **ssim-so**. As a side result, we study the relation between the indistinguishability and simulation based security definitions in the sender setting. We show that **sind-so** is strictly weaker than the notion of **ssim-so**

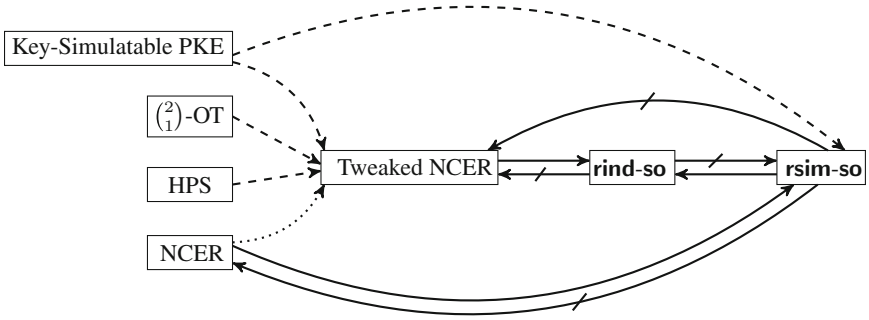

Fig. 1. The arrows can be read as follows: *solid arrows* denote implication, *crossed arrows* denote counterexamples, *dashed arrows* denote assumption-wise implication and *dotted arrows* denote implication with respect to concrete instances (where the implication may not hold in general). The implication of receiver indistinguishability security by simulation security is a known result.

by presenting a concrete public key scheme that meets the former but not the latter level of security. Our separating scheme is built using the two primitives lossy public key encryption and commitment scheme. We exploit the hiding properties of these building blocks to prove that our scheme implies **sind-so** security. On the other hand, simulation always fails since it implies breaking the binding property of the commitment scheme. Informally, we prove the following theorem:

Theorem 1.4 (Informal). *There exists a PKE that is* **sind-so** *secure but is not* **ssim-so** *secure.*

We stress that this was already demonstrated indirectly in [4] (by combining two separation results). Here we design a concrete counter example to demonstrate the same in a simpler manner. A similar result has been shown for *full* **ind-so** and **sim-so** in [5], demonstrating that these definitions do *not* imply each other in the sender setting.

To sum up, we study the different levels for receiver security in the presence of SO attacks. We clarify the relation between these notions and provide constructions that meet them using the close conceptual relation between SO security and non-committing encryption. From a broader perspective, our results position more precisely SO security for the receiver in the spectrum of security notions for encryption.

2 Preliminaries

Basic notations. For $x, y \in \mathbb{N}$ with $x < y$, let $[x] := \{1, \ldots, x\}$ and $[x, y] := \{x, x+1, \ldots, y\}$. We denote the computational security parameter by k and statistical security parameter by s. A function $\mu(\cdot)$ is *negligible* in security parameter κ if for every polynomial $p(\cdot)$ there exists a value N such that for all

$\kappa > N$ it holds that $\mu(k) < \frac{1}{p(\kappa)}$, where κ is either k or s. For a finite set S, we denote by $s \leftarrow S$ the process of sampling s uniformly. For a distribution X, we denote by $x \leftarrow X$ the process of sampling x from X. For a deterministic algorithm A, we write $a \leftarrow A(x)$ the process of running A on input x and assigning y the result. For a randomized algorithm A, we write $a \leftarrow A(x; r)$ the process of running A on input x and randomness r and assigning a the result. At times we skip r in the parenthesis to avoid mentioning it explicitly. We write PPT for probabilistic polynomial-time. For a PKE (or commitment) scheme C, we use the notation $\mathcal{M}_C$ and respectively $\mathcal{R}_C$ to denote the input and the randomness space of the encryption (or commitment) algorithm of C. We use bold fonts to denote vectors. If $\mathbf{m}$ is an n dimensional vector, we write $\mathbf{m}_i$ for the i-th entry in $\mathbf{m}$; if $\mathcal{I} \subseteq [n]$ is a set of indices we write $\mathbf{m}_{\mathcal{I}}$ for the vector of dimension $|\mathcal{I}|$ obtained by projecting $\mathbf{m}$ on the coordinates in $\mathcal{I}$.

2.1 Public Key Encryption

A public key encryption (PKE) scheme PKE with message space $\mathcal{M}$ consists of three PPT algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. The key generation algorithm $\mathsf{Gen}(1^k)$ outputs a public key pk and a secret key sk. The encryption algorithm $\mathsf{Enc}_{pk}(m; r)$ takes pk and a message $m \in \mathcal{M}$ and randomness $r \in \mathcal{R}$, and outputs a ciphertext c. The decryption algorithm $\mathsf{Dec}_{sk}(c)$ takes sk and a ciphertext c and outputs a message m. For correctness, we require that $m = \mathsf{Dec}_{sk}(c)$ for all $m \in \mathcal{M}$ and all $(pk, sk) \leftarrow \mathsf{Gen}(1^k)$ and all $c \leftarrow \mathsf{Enc}_{pk}(m)$. The standard notion of security for PKE is indistinguishability under chosen plaintext attacks, denoted by **ind-cpa** [15] (and the corresponding experiment is denoted as $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathsf{cpa}_{\mathsf{PKE}}$). *As a general remark, we note that whenever we refer to a secret key, we refer to the randomness used to generate it by the key generation algorithm.*

2.2 Selective Opening Security

Depending on the attack scenario, we distinguish two settings that fall under the general idea of SO attacks. In *sender security*, we have n senders and one receiver. The receiver holds a secret key relative to a public key known to all senders. The senders send messages to the receiver and the adversary is allowed to corrupt some of the senders (and learn the messages and randomness underlying some of the ciphertexts). The concern is that the messages sent by uncorrupted users stay secret. The second scenario deals with *receiver security*. Here we consider one sender and n receivers who hold independently generated public and secret keys. The attacker is allowed to learn the secret keys of some of the receivers. Security is concerned with the messages received by uncorrupted receivers.

For each of these settings we consider two types of definitions from the literature [2]: (1) an indistinguishability based definition and (2) a simulation based definition. Indistinguishability-based definitions require that an adversary that sees a vector of ciphertexts cannot distinguish the true plaintexts of the ciphertexts from independently sampled plaintexts, even in the presence of the randomness used for generating the opened ciphertexts (in the sender corruption

setting), or the secret keys that decrypt the opened ciphertexts (in the receiver corruption setting). The indistinguishability based definitions use the notion of *efficiently resamplable* message distributions which we recall next following [5].

Definition 2.1 (Efficiently Resamplable Distribution). *Let $n = n(k) > 0$ and let* Dist *be a joint distribution over $\left(\{0,1\}^k\right)^n$. We say that* Dist *is efficiently resamplable if there is a PPT algorithm* $\mathsf{Resamp}_{\mathsf{Dist}}$ *such that for any $\mathcal{I} \subseteq [n]$ and any partial vector $\mathbf{m}'_{\mathcal{I}} \in \left(\{0,1\}^k\right)^{|\mathcal{I}|}$, $\mathsf{Resamp}_{\mathsf{Dist}}(\mathbf{m}'_{\mathcal{I}})$ returns a vector $\mathbf{m}$ sampled from $\mathsf{Dist}|\mathbf{m}'_{\mathcal{I}}$, i.e. $\mathbf{m}'$ is sampled from* Dist *conditioned on $\mathbf{m}_{\mathcal{I}} = \mathbf{m}'_{\mathcal{I}}$.*

Below, we recall indistinguishability and simulation based definitions for security in the presence of selective opening attacks[3]. We present the definitions for sender and receiver security. To avoid heavy notation we follow the following conventions when naming the security notions: we use "ind" or "sim" to indicate if the definition is indistinguishability-based or simulation-based, and prepend "s" or "r" to indicate if the definition is for sender security or receiver security; we keep "so" in the name of the notion to indicate that we deal with selective opening attacks. We also note that we consider chosen plaintext attacks only, but avoid showing this explicitly in the names of the security notions.

Experiment 1 $\mathbf{Exp}^{\mathsf{sind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A}, k)$

$b \leftarrow \{0,1\}$
$(pk, sk) \leftarrow \mathsf{Gen}(1^k)$
$(\mathsf{Dist}, \mathsf{Resamp}_{\mathsf{Dist}}, \mathrm{state}_1) \leftarrow \mathsf{A}(pk)$
$\mathbf{m} := (m_i)_{i\in[n]} \leftarrow \mathsf{Dist}$
$\mathbf{r} := (r_i)_{i\in[n]} \leftarrow \mathcal{R}^n_{\mathsf{PKE}}$
$\mathbf{e} := (e_i)_{i\in[n]} \leftarrow (\mathsf{Enc}_{pk}(m_i; r_i))_{i\in[n]}$
$(\mathcal{I}, \mathrm{state}_2) \leftarrow \mathsf{A}(\mathbf{e}, \mathrm{state}_1)$
$\mathbf{m}' \leftarrow \mathsf{Resamp}(\mathbf{m}_{\mathcal{I}})$
$\mathbf{m}^* = \mathbf{m}$ *if* $b = 0$, *else* $\mathbf{m}^* = \mathbf{m}'$
$b' \leftarrow \mathsf{A}(\mathbf{r}_{\mathcal{I}}, \mathbf{m}^*, \mathrm{state}_2)$,
Return 1 if $b = b'$, *and 0 otherwise.*

Experiment 2 $\mathbf{Exp}^{\mathsf{rind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A}, k)$

$b \leftarrow \{0,1\}$
$(\mathbf{pk}, \mathbf{sk}) := (pk_i, sk_i) \leftarrow (\mathsf{Gen}(1^k))_{i\in[n]}$
$(\mathsf{Dist}, \mathsf{Resamp}_{\mathsf{Dist}}, \mathrm{state}_1) \leftarrow \mathsf{A}(\mathbf{pk})$
$\mathbf{m} := (m_i)_{i\in[n]} \leftarrow \mathsf{Dist}$
$\mathbf{r} := (r_i)_{i\in[n]} \leftarrow \mathcal{R}^n_{\mathsf{PKE}}$
$\mathbf{e} := (e_i)_{i\in[n]} \leftarrow (\mathsf{Enc}_{pk_i}(m_i; r_i))_{i\in[n]}$
$(\mathcal{I}, \mathrm{state}_2) \leftarrow \mathsf{A}(\mathbf{e}, \mathrm{state}_1)$
$\mathbf{m}' \leftarrow \mathsf{Resamp}(\mathbf{m}_{\mathcal{I}})$
$\mathbf{m}^* = \mathbf{m}$ *if* $b = 0$, *else* $\mathbf{m}^* = \mathbf{m}'$
$b' \leftarrow \mathsf{A}(\mathbf{sk}_{\mathcal{I}}, \mathbf{m}^*, \mathrm{state}_2)$
Return 1 if $b = b'$, *and 0 otherwise.*

Definition 2.2 (Indistinguishability Based SO Security). *For a PKE scheme* $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, *a polynomially bounded function $n = n(k) > 0$ and a stateful PPT adversary* A, *consider the following two experiments; the left experiment corresponds to sender corruptions, whereas, the right experiment corresponds to receiver corruptions.*

In the above experiments we only assume adversaries that are well-behaved in that they always output efficiently resamplable distributions together with resampling algorithms.

[3] We remark that a stronger security notion that does not does require efficient resemplability is possible, but no constructions that satisfy this stronger notion are known.

We say that PKE *is* **sind-so** *secure if for a well-behaved PPT* A *there exists a negligible function* $\mu = \mu(k)$ *such that*

$$\mathbf{Adv}^{\mathbf{sind}}\text{-}\mathbf{so}_{\mathsf{PKE}}(\mathsf{A}, k) := 2\left|\Pr[\mathbf{Exp}^{\mathbf{sind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A}, k) = 1] - \frac{1}{2}\right| \leq \mu.$$

We say that PKE *is* **rind-so** *secure if for a well-behaved PPT* A *there exists a negligible function* $\mu = \mu(k)$ *such that*

$$\mathbf{Adv}^{\mathbf{rind}}\text{-}\mathbf{so}_{\mathsf{PKE}}(\mathsf{A}, k) := 2\left|\Pr[\mathbf{Exp}^{\mathbf{rind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A}, k) = 1] - \frac{1}{2}\right| \leq \mu.$$

$\Pr\left[\mathbf{Exp}^{\mathbf{sind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A}, k) = 1\right]$ *and* $\Pr\left[\mathbf{Exp}^{\mathbf{rind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A}, k) = 1\right]$ *denote the* **winning probability** *of* A *in the respective experiments.*

Simulation based security is defined, as usual, by comparing an idealized execution with the real one. Again, we consider both sender and receiver security.

Experiment 3 $\mathbf{Exp}^{\mathbf{SSIM\text{-}SO}\text{-}\mathbf{real}}_{\mathsf{PKE}}(\mathsf{A}, k)$
$(pk, sk) \leftarrow \mathsf{Gen}(1^k)$
$(\mathsf{Dist}, \mathrm{state}_1) \leftarrow \mathsf{A}(pk)$
$\mathbf{m} := (m_i)_{i\in[n]} \leftarrow \mathsf{Dist}$
$\mathbf{r} := (r_i)_{i\in[n]} \leftarrow \mathcal{R}^n_{\mathsf{PKE}}$
$\mathbf{e} := (e_i)_{i\in[n]} \leftarrow (\mathsf{Enc}_{pk}(m_i; r_i))_{i\in[n]}$
$(\mathcal{I}, \mathrm{state}_2) \leftarrow \mathsf{A}(\mathbf{e}, \mathrm{state}_1)$
$\mathrm{output} \leftarrow \mathsf{A}(\mathbf{r}_{\mathcal{I}}, \mathbf{m}_{\mathcal{I}}, \mathrm{state}_2)$
Return $(\mathbf{m}, \mathsf{Dist}, \mathcal{I}, \mathrm{output})$.

Experiment 4 $\mathbf{Exp}^{\mathbf{SSIM\text{-}SO}\text{-}\mathbf{ideal}}_{\mathsf{PKE}}(\mathsf{S}, k)$
$(\mathsf{Dist}, \mathrm{state}_1) \leftarrow \mathsf{S}(\cdot)$
$\mathbf{m} := (m_i)_{i\in[n]} \leftarrow \mathsf{Dist}$
$(\mathcal{I}, \mathrm{state}_2) \leftarrow \mathsf{S}(\mathrm{state}_1)$
$\mathrm{output} \leftarrow \mathsf{S}(\mathbf{m}_{\mathcal{I}}, \mathrm{state}_2)$
Return $(\mathbf{m}, \mathsf{Dist}, \mathcal{I}, \mathrm{output})$.

Experiment 5 $\mathbf{Exp}^{\mathbf{rsim\text{-}so}\text{-}\mathbf{real}}_{\mathsf{PKE}}(\mathsf{A}, k)$
$(\mathbf{pk}, \mathbf{sk}) := (pk_i, sk_i) \leftarrow (\mathsf{Gen}(1^k))_{i\in[n]}$
$(\mathsf{Dist}, \mathrm{state}_1) \leftarrow \mathsf{A}(\mathbf{pk})$
$\mathbf{m} := (m_i)_{i\in[n]} \leftarrow \mathsf{Dist}$
$\mathbf{r} := (r_i)_{i\in[n]} \leftarrow \mathcal{R}^n_{\mathsf{PKE}}$
$\mathbf{e} := (e_i)_{i\in[n]} \leftarrow (\mathsf{Enc}_{pk_i}(m_i; r_i))_{i\in[n]}$
$(\mathcal{I}, \mathrm{state}_2) \leftarrow \mathsf{A}(\mathbf{e}, \mathrm{state}_1)$
$\mathrm{output} \leftarrow \mathsf{A}(\mathbf{sk}_{\mathcal{I}}, \mathbf{m}_{\mathcal{I}}, \mathrm{state}_2)$
Return $(\mathbf{m}, \mathsf{Dist}, \mathcal{I}, \mathrm{output})$.

Experiment 6 $\mathbf{Exp}^{\mathbf{rsim\text{-}so}\text{-}\mathbf{ideal}}_{\mathsf{PKE}}(\mathsf{S}, k)$
$(\mathsf{Dist}, \mathrm{state}_1) \leftarrow \mathsf{S}(\cdot)$
$\mathbf{m} := (m_i)_{i\in[n]} \leftarrow \mathsf{Dist}$
$(\mathcal{I}, \mathrm{state}_2) \leftarrow \mathsf{S}(\mathrm{state}_1)$
$\mathrm{output} \leftarrow \mathsf{S}(\mathbf{m}_{\mathcal{I}}, \mathrm{state}_2)$
Return $(\mathbf{m}, \mathsf{Dist}, \mathcal{I}, \mathrm{output})$.

Definition 2.3 (Simulation Based SO Security). *For a PKE scheme* $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$*, a polynomially bounded function* $n = n(k) > 0$*, a PPT adversary* A *and a PPT algorithm* S*, we define the following pairs of experiments.*

We say that PKE *is* **ssim-so** *secure iff for every PPT* A *there is a PPT algorithm* S*, a PPT distinguisher* D *with binary output and a negligible function* $\mu = \mu(k)$ *such that*

$$\mathbf{Adv}^{\mathbf{ssim}}\text{-}\mathbf{so}_{\mathsf{PKE}}(\mathsf{D}, k) :=$$
$$\left|\Pr[1 \leftarrow \mathsf{D}(\mathbf{Exp}^{\mathbf{SSIM\text{-}SO}}_{\mathsf{PKE}}\text{-}\mathbf{real}(\mathsf{A}, k))] - \Pr[1 \leftarrow \mathsf{D}(\mathbf{Exp}^{\mathbf{SSIM\text{-}SO}}_{\mathsf{PKE}}\text{-}\mathbf{ideal}(\mathsf{S}, k))]\right| \leq \mu.$$

We say that PKE *is* **rsim-so** *secure iff for every PPT* A *there is a PPT algorithm* S*, a PPT distinguisher* D *with binary output and a negligible function* $\mu = \mu(k)$ *such that*

$$\mathbf{Adv}^{\mathbf{rsim}}\text{-}\mathbf{so}_{\mathsf{PKE}}(\mathsf{D}, k) :=$$
$$\left|\Pr[1 \leftarrow \mathsf{D}(\mathbf{Exp}^{\mathbf{rsim\text{-}so}}_{\mathsf{PKE}}\text{-}\mathbf{real}(\mathsf{A}, k))] - \Pr[1 \leftarrow \mathsf{D}(\mathbf{Exp}^{\mathbf{rsim\text{-}so}}_{\mathsf{PKE}}\text{-}\mathbf{ideal}(\mathsf{S}, k))]\right| \leq \mu.$$

Our definitions consider non-adaptive attacks, where the adversary corrupts the parties in one go. Our results remain unaffected even in the face of an adaptive adversary [5].

3 Building Blocks

Our constructions employ a number of fundamental cryptographic building blocks as well as a new primitive which we call tweaked NCER.

Commitment Schemes. We define a non-interactive statistically hiding commitment scheme (NISHCOM).

Definition 3.1 (NISHCOM). *A non-interactive commitment scheme* nisCom *consists of two algorithms* $(\mathsf{nisCommit}, \mathsf{nisOpen})$ *defined as follows. Given a security parameter* k*, message* $m \in \mathcal{M}_{\mathsf{nisCom}}$ *and random coins* $r \in \mathcal{R}_{\mathsf{nisCom}}$*, PPT algorithm* $\mathsf{nisCommit}$ *outputs commitment* c*. Given* k*, commitment* c *and message* m*, (possibly inefficient) algorithm* $\mathsf{nisOpen}$ *outputs* r*. We require the following properties:*

- Correctness. *We require that* $c = \mathsf{nisCommit}(m; r)$ *for all* $m \in \mathcal{M}_{\mathsf{nisCom}}$ *and* $r \leftarrow \mathsf{nisOpen}(c, m)$.
- Security. *A NISHCOM* nisCom *is* **stat-hide** *secure if commitments of two distinct messages are statistically indistinguishable. Specifically, for any* unbounded powerful *adversary* A*, there exists a negligible function* $\mu = \mu(s)$ *such that*
  $\mathbf{Adv}^{\mathbf{stat}}_{\mathsf{nisCom}}\text{-}\mathbf{hide}(\mathsf{A}, k) := |\Pr[1 \leftarrow \mathsf{A}(c_0)] - \Pr[1 \leftarrow \mathsf{A}(c_1)]| \leq \mu$
  for $c_i \leftarrow \mathsf{nisCommit}(m_i)$*,* $i \in \{0, 1\}$ *and* $m_0, m_1 \in \mathcal{M}_{\mathsf{nisCom}}$.
  A NISHCOM nisCom *is* **comp-bind** *secure if no commitment can be opened to two different messages in polynomial time. Specifically, the advantage* $\mathbf{Adv}^{\mathbf{comp}}_{\mathsf{nisCom}}\text{-}\mathbf{bind}(\mathsf{A}, k)$ *of* A *is defined by* $\Pr[(m_0, r_0, m_1, r_1) \leftarrow \mathsf{A}(k) : \mathsf{nisCommit}(prm, m_0; r_0) = \mathsf{nisCommit}(prm, m_1; r_1)]$ *(with the probability over the choice of the coins of* A*) is smaller than some negligible function* $\mu = \mu(k)$*. A NISHCOM* nisCom *is called* **secure** *it is* $\{\mathbf{stat\text{-}hide}, \mathbf{comp\text{-}bind}\}$ *secure.*

Non-committing Encryption for Receiver (NCER). A non-committing encryption for receiver [7,21] is a PKE scheme with the property that there is a way to generate fake ciphertexts which can then be decrypted (with the help of a trapdoor) to any plaintext. Intuitively, fake ciphertexts are generated in a lossy way so that the plaintext is no longer well defined given the ciphertext and the public key. This leaves enough entropy for the secret key to be sampled in a way that determines the desired plaintext. We continue with a formal definition of NCER and its security notion referred as **ind-ncer** security.

Definition 3.2 (NCER). *An NCER* nPKE *consists of five PPT algorithms* $(\mathsf{nGen}, \mathsf{nEnc}, \mathsf{nEnc}^*, \mathsf{nDec}, \mathsf{nOpen})$ *defined as follows. Algorithms* $(\mathsf{nGen}, \mathsf{nEnc}, \mathsf{nDec})$ *form a PKE. Given the public key* pk*, the fake encryption algorithm* nEnc^* *outputs a ciphertext* e^* *and a trapdoor* t*. Given the secret key* sk*, the public key* pk*, fake ciphertext* e^**, trapdoor* t *and plaintext* m*, algorithm* nOpen *outputs* sk^*.

- Correctness. *We require that* $m = \mathsf{nDec}_{sk}(c)$ *for all* $m \in \mathcal{M}$*, all* $(pk, sk) \leftarrow \mathsf{nGen}(1^k)$ *and all* $c \leftarrow \mathsf{nEnc}_{pk}(m)$.

– Security. *An NCER scheme* nPKE *is* **ind-ncer** *secure if the real and fake ciphertexts are indistinguishable. Specifically, for a PPT adversary* A*, consider the experiment* $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{nPKE}}$**-ncer** *defined as follows.*

Experiment 7 $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}(\mathsf{A}, k)$
$b \leftarrow \{0,1\}$
$(pk, sk_0) \leftarrow \mathsf{nGen}(1^k)$
$m \leftarrow \mathsf{A}(pk)$
$e_0 \leftarrow \mathsf{nEnc}_{pk}(m)$
$(e_1, t) \leftarrow \mathsf{nEnc}^*_{pk}(1^k)$, $sk_1 \leftarrow \mathsf{nOpen}(sk_0, pk, e_1, t, m)$
$b' \leftarrow \mathsf{A}(sk_b, e_b)$
Return 1 if $b = b'$*, and 0 otherwise.*

We say that nPKE *is* **ind-ncer***-secure if for a PPT adversary* A*, there exists a negligible function* $\mu = \mu(k)$ *such that*

$$\mathbf{Adv}^{\mathbf{ind\text{-}ncer}}_{\mathsf{nPKE}}(\mathsf{A}, k) := 2\left|\Pr[\mathbf{Exp}^{\mathbf{ind}}\text{-}\mathbf{ncer}_{\mathsf{nPKE}}(\mathsf{A}, k) = 1] - \frac{1}{2}\right| \leq \mu.$$

An NCER nPKE *is* **secure** *if it is* **ind-ncer** *secure.*

Tweaked NCER. We introduce a variant of NCER which modifies the definition of NCER in the following two ways. First, the opening algorithm nOpen may be *inefficient*. In addition, the fake encryption algorithm is required to output a fake ciphertext e^* given the secret key sk and a plaintext m, so that decryption is "correct" with respect to e^* and m. We call the resulting notion, which we formalize below, *tweaked NCER.*

Definition 3.3 (Tweaked NCER). *A tweaked NCER scheme* tPKE *is a PKE that consists of five algorithms* $(\mathsf{tGen}, \mathsf{tEnc}, \mathsf{tEnc}^*, \mathsf{tDec}, \mathsf{tOpen})$ *defined as follows. Algorithms* $(\mathsf{tGen}, \mathsf{tEnc}, \mathsf{tDec})$ *form a PKE. Given the secret key* sk *and the public key* pk*, and a plaintext* m*, the PPT fake encryption algorithm* tEnc^* *outputs a ciphertext* e^**. Given the secret key* sk *and the public key* pk*, fake ciphertext* e^* *such that* $e^* \leftarrow \mathsf{tEnc}^*_{pk}(sk, m')$ *for some* $m' \in \mathcal{M}_{\mathsf{tPKE}}$ *and a plaintext* m*, the inefficient algorithm* tOpen *outputs* sk^* *such that* $m = \mathsf{tDec}_{sk^*}(e^*)$*.*

– Correctness. *We require that* $m = \mathsf{tDec}_{sk}(c)$ *for all* $m \in \mathcal{M}$*, all* $(pk, sk) \leftarrow \mathsf{tGen}(1^k)$ *and all* $c \leftarrow \mathsf{tEnc}_{pk}(m)$*.*
– Security. *A tweaked NCER scheme* tPKE *is* **ind-tcipher** *secure if real and fake ciphertexts are indistinguishable. Specifically, for a PPT adversary* A*, consider the experiment* $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{tPKE}}$**-tcipher** *defined as follows.*

Experiment 8 $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}(\mathsf{A}, k)$
$b \leftarrow \{0,1\}$
$(pk, sk) \leftarrow \mathsf{tGen}(1^k)$
$m \leftarrow \mathsf{A}(pk)$
$e_0 \leftarrow \mathsf{tEnc}_{pk}(m)$
$e_1 \leftarrow \mathsf{tEnc}^*_{pk}(sk, m)$
$b' \leftarrow \mathsf{A}(sk, e_b)$
Return 1 if $b = b'$*, and 0 otherwise.*

Experiment 9 $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tncer}(\mathsf{A}, k)$
$b \leftarrow \{0,1\}$
$(pk, sk_0) \leftarrow \mathsf{tGen}(1^k)$
$m \leftarrow \mathsf{A}(pk)$
$e_0 \leftarrow \mathsf{tEnc}^*_{pk}(sk_0, m)$
$e_1 \leftarrow \mathsf{tEnc}^*_{pk}(sk_0, m')$ *for* $m' \in \mathcal{M}_{\mathsf{tPKE}}$
$sk_1 \leftarrow \mathsf{tOpen}(e_1, m)$
$b' \leftarrow \mathsf{A}(sk_b, e_b)$
Return 1 if $b = b'$*, and 0 otherwise.*

We say that tPKE *is* **ind-tcipher** *secure if for a PPT adversary* A*, there exists a negligible function* $\mu = \mu(k)$ *such that*

$$\mathbf{Adv}^{\textsf{ind-tcipher}}_{\textsf{tPKE}}(\mathsf{A}, k) := 2\big|\Pr[\mathbf{Exp}^{\textsf{ind}}\text{-}\textbf{tcipher}_{\textsf{tPKE}}(\mathsf{A}, k) = 1] - \frac{1}{2}\big| \leq \mu.$$

We say that tPKE *is* **ind-tncer** *secure if for an unbounded adversary* A*, there exists a negligible function* $\mu = \mu(s)$ *such that*

$$\mathbf{Adv}^{\textsf{ind-tncer}}_{\textsf{tPKE}}(\mathsf{A}, k) := 2\big|\Pr[\mathbf{Exp}^{\textsf{ind}}\text{-}\textbf{tncer}_{\textsf{tPKE}}(\mathsf{A}, k) = 1] - \frac{1}{2}\big| \leq \mu.$$

A tweaked NCER tPKE *is* **secure** *if it is* {**ind-tcipher**, **ind-tncer**} *secure.*

Key-Simulatable PKE. A key-simulatable public key encryption scheme is a PKE in which the public keys can be generated in two modes. In the first mode a public key is picked together with a secret key, whereas the second mode implies an oblivious public key generation without the secret key. Let $\mathcal{V}$ denote the set of public keys generated in the first mode and $\mathcal{K}$ denote the set of public keys generated in the second mode. Then it is possible that $\mathcal{K}$ contains $\mathcal{V}$ (i.e., $\mathcal{V} \subseteq \mathcal{K}$). Moreover, in case $\mathcal{V} \subset \mathcal{K}$ the set of public keys from $\mathcal{K} \backslash \mathcal{V}$ is not associated with any secret key. We respectively denote the keys in $\mathcal{V}$ and $\mathcal{K} \backslash \mathcal{V}$ as *valid* and *invalid* public keys. In addition to the key generation algorithms, key-simulatable PKE also consists of an efficient key faking algorithm that explains a public key from $\mathcal{V}$, that was generated in the first mode, as an obliviously generated public key from $\mathcal{K}$ that was generated without the corresponding secret key. The security requirement asserts that it is hard to distinguish a random element from $\mathcal{K}$ from a random element from $\mathcal{V}$. The formal definition follows. We note that the notion of key-simulatable PKE is very similar to the simulatable PKE [9] notion with the differences that the latter notion assumes that $\mathcal{K} = \mathcal{V}$ and further supports oblivious ciphertext generation and ciphertext faking.

Definition 3.4 (Key-simulatable PKE). *A key-simulatable public key encryption* sPKE *consists of five PPT algorithms* $(\mathsf{sGen}, \mathsf{sEnc}, \mathsf{sDec}, \widetilde{\mathsf{sGen}}, \widetilde{\mathsf{sGen}}^{-1})$ *defined as follows. Algorithms* (sGen, sEnc, sDec) *form a PKE. Given the security parameter* k*, the oblivious public key generator* $\widetilde{\mathsf{sGen}}$ *returns a public key* pk' *from* $\mathcal{K}$ *and the random coins* r' *used to sample* pk'*. Given a public key* $pk \in \mathcal{V}$*, the key faking algorithm returns some random coins* r*.*

- Correctness. *We require that* $m = \mathsf{sDec}_{sk}(c)$ *for all* $m \in \mathcal{M}$*, all* $(pk, sk) \leftarrow \mathsf{sGen}(1^k)$ *and all* $c \leftarrow \mathsf{sEnc}_{pk}(m)$*.*
- Security. *A key-simulatable scheme* sPKE *is* **ind-cpa** *secure if* (sGen, sEnc, sDec) *is* **ind-cpa** *secure. It is called* **ksim** *secure if it is hard to distinguish an obliviously generated key from a legitimately generated key. Specifically, for a PPT adversary* A*, there exists a negligible function* $\mu = \mu(k)$ *such that* $\mathbf{Adv}^{\textsf{ksim}}_{\textsf{sPKE}}(\mathsf{A}, k) := \big|\Pr\left[1 \leftarrow \mathsf{A}(r, pk)\right] - \Pr\left[1 \leftarrow \mathsf{A}(r', pk')\right]\big| \leq \mu$ *where* $(pk, sk) \leftarrow \mathsf{sGen}(1^k)$*,* $r \leftarrow \widetilde{\mathsf{sGen}}^{-1}(pk)$ *and* $(pk', r') \leftarrow \widetilde{\mathsf{sGen}}(1^k)$*.*
  A key-simulatable scheme sPKE *is* **secure** *if it is* {**ind-cpa**, **ksim**} *secure.*

An *extended* key-simulatable PKE is a secure key-simulatable where in addition $\mathcal{V} \subset \mathcal{K}$ and it holds that $\Pr\left[pk \in \mathcal{K}\backslash\mathcal{V} \mid (pk, r) \leftarrow \widetilde{\mathsf{sGen}}(1^k)\right]$ is non-negligible.

4 Selective Opening Security for the Receiver

In this section we provide negative and positive results regarding security for the receiver in the presence of selective opening attacks. First, we show that **rind-so** is strictly weaker than **rsim-so** security by constructing a scheme that meets the former but not the latter level of security. We then relate the different forms of security under SO attacks with non-committing encryption (for the receiver). Specifically, we show that secure NCER implies **rsim-so** and that secure tweaked NCER implies **rind-so**. Interestingly, we show that the converse implications do not hold. In terms of constructions, we show that tweaked NCER can be constructed from various primitives such as key-simulatable PKE, two-round honest-receiver statistically-hiding $\binom{2}{1}$-OT protocol, secure HPS and NCER. The DDH based secure NCER scheme of [7] that works for polynomial message space turns out to be secure tweaked NCER for exponential message space.

4.1 rind-so Secure PKE $\not\Rightarrow$ rind-so Secure PKE

Our construction is built from an **ind-ncer** secure scheme nPKE and a {**stat-hide**, **comp-bind**} secure NISHCOM nisCom that satisfy a compatibility condition. Specifically, we require that the message and randomness spaces of nisCom, denoted by $\mathcal{M}_{\mathsf{nisCom}}$ and $\mathcal{R}_{\mathsf{nisCom}}$, are compatible with the message space $\mathcal{M}_{\mathsf{nPKE}}$ of nPKE.

Definition 4.1. *An* **ind-ncer** *secure NCER* nPKE *and a* {**stat-hide**, **comp-bind**} *secure NISHCOM* nisCom *are said to be* compatible *if* $\mathcal{M}_{\mathsf{nPKE}} = \mathcal{M}_{\mathsf{nisCom}} = \mathcal{R}_{\mathsf{nisCom}}$.

Theorem 4.2. *Assume there exist an* **ind-ncer** *secure NCER and a* {**stat-hide**, **comp-bind**} *secure NISHCOM that are compatible. Then, there exists a PKE that is* **rind-so** *secure but is not* **rsim-so** *secure.*

Proof: We describe our separating encryption scheme first. Consider a scheme $\mathsf{nPKE} = (\mathsf{nGen}, \mathsf{nEnc}, \mathsf{nEnc}^*, \mathsf{nDec}, \mathsf{nOpen})$ that is secure NCER (cf. Definition 3.2) and an NISHCOM $\mathsf{nisCom} = (\mathsf{nisCommit}, \mathsf{nisOpen})$ (cf. Definition 3.1) that are compatible. We define the encryption scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ as follows.

$\mathsf{Gen}(1^k)$
$(pk_0, sk_0) \leftarrow \mathsf{nGen}(1^k)$
$(pk_1, sk_1) \leftarrow \mathsf{nGen}(1^k)$
$pk = (pk_0, pk_1)$
$sk = (sk_0, sk_1)$
Return (pk, sk)

$\mathsf{Enc}_{pk}(m)$
$c \leftarrow \mathsf{nisCommit}(m, r)$
$e_0 \leftarrow \mathsf{nEnc}_{pk_0}(m)$
$e_1 \leftarrow \mathsf{nEnc}_{pk_1}(r)$
Return $e = (e_0, e_1, c)$

$\mathsf{Dec}_{sk}(e)$
$e := (e_0, e_1, c)$
$m = \mathsf{nDec}_{sk_0}(e_0)$
$r = \mathsf{nDec}_{sk_1}(e_1)$
if $c = \mathsf{nisCommit}(m, r)$
  Return m
else Return $\perp$

The proof follows from Lemmas 4.3 and 4.7 below which formalize that PKE is **rind-so** secure but not **rsim-so** secure. ■

Lemma 4.3. *Assume that* nPKE *is* **ind-ncer** *secure and* nisCom *is* {**stat-hide**, **comp-bind**} *secure, then* PKE *is* **rind-so** *secure.*

Proof: More precisely we show that for any PPT adversary A attacking PKE there exist a PPT adversary B and an unbounded powerful adversary C such that

$$\mathbf{Adv}^{\mathbf{rind}}_{\mathsf{PKE}}\text{-}\mathbf{so}(\mathsf{A},k) \leq n\left(4\cdot\mathbf{Adv}^{\mathbf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}(\mathsf{B},k)+\mathbf{Adv}^{\mathbf{stat}}_{\mathsf{nisCom}}\text{-}\mathbf{hide}(\mathsf{C},k)\right).$$

We prove this lemma using the following sequence of experiments.

- $\mathbf{Exp}_0 = \mathbf{Exp}^{\mathbf{rind}}_{\mathsf{PKE}}\text{-}\mathbf{so}$.
- $\mathbf{Exp}_1$ is identical to $\mathbf{Exp}_0$ except that the first component of each ciphertext in the vector $\mathbf{e}$ is computed using nEnc^* of nPKE. That is, for all $i \in [n]$ ciphertext e_i is defined by (e^*_{i0}, e_{i1}, c_i) such that $(e^*_{i0}, t_{i0}) \leftarrow \mathsf{nEnc}^*_{pk_{i0}}(1^k)$. Furthermore, if $i \in \mathcal{I}$ (i.e., A asks to open the ith ciphertext), then $\mathbf{Exp}_1$ computes $sk^*_{i0} \leftarrow \mathsf{nOpen}(sk_{i0}, e^*_{i0}, t_{i0}, m_i)$ and hands (sk^*_{i0}, sk_{i1}) to A.
- $\mathbf{Exp}_2$ is identical to $\mathbf{Exp}_1$ except that the second component of each ciphertext in the vector $\mathbf{e}$ is computed using nEnc^* of nPKE, That is, for all $i \in [n]$ ciphertext e_i is defined by $(e^*_{i0}, e^*_{i1}, c_i)$ such that $(e^*_{i1}, t_{i1}) \leftarrow \mathsf{nEnc}^*_{pk_{i1}}(1^k)$. Furthermore, if $i \in \mathcal{I}$ (i.e., A asks to open the ith ciphertext), then $\mathbf{Exp}_2$ computes $sk^*_{i1} \leftarrow \mathsf{nOpen}(sk_{i1}, e^*_{i1}, t_{i1}, r_i)$ and hands (sk^*_{i0}, sk^*_{i1}) to A, where r_i is the randomness used to compute c_i.
- $\mathbf{Exp}_3$ is identical to $\mathbf{Exp}_2$ except that the third component of each ciphertext in the vector $\mathbf{e}$ is a commitment of a dummy message. That is, for all $i \in [n]$ ciphertext e_i is defined by $(e^*_{i0}, e^*_{i1}, c^*_i)$ such that $c^*_i \leftarrow \mathsf{nisCommit}(m^*_i; r^*_i)$, where m^*_i is a dummy message from $\mathcal{M}_{\mathsf{nisCom}}$ and $r^*_i \leftarrow \mathcal{R}_{\mathsf{nisCom}}$. Furthermore, if $i \in \mathcal{I}$ then $\mathbf{Exp}_3$ first computes $r_i \leftarrow \mathsf{nisOpen}(c^*_i, m_i)$. Then it computes $sk^*_{i1} \leftarrow \mathsf{nOpen}(sk_{i1}, e^*_{i1}, t_{i1}, r_i)$ and hands (sk^*_{i0}, sk^*_{i1}) to A, where r_i is the randomness returned by nisOpen.

We note that although the third experiment is not efficient (the experiment needs to equivocate the commitment without a trapdoor), it does not introduce a problem in our proof: an adversary that distinguishes between $\mathbf{Exp}_2$ and $\mathbf{Exp}_3$ gives rise to an unbounded adversary that breaks the statistical hiding property of the commitment scheme used by our construction.

Let ϵ_j be the advantage of A in $\mathbf{Exp}_j$, i.e. $\epsilon_j := 2\left|\Pr[\mathbf{Exp}_j(\mathsf{A},k)=1]-\frac{1}{2}\right|$. We first note that $\epsilon_3 = 0$ since in experiment $\mathbf{Exp}_3$ the adversary receives a vector of ciphertexts that are statistically independent of the encrypted plaintexts, implying that the adversary (even with unbounded computing power) outputs the correct bit b with probability $1/2$. Next we show that $|\epsilon_0 - \epsilon_1| \leq 2n\Delta_{\mathbf{ind}}\text{-}\mathbf{ncer}$ and $|\epsilon_1 - \epsilon_2| \leq 2n\Delta_{\mathbf{ind}}\text{-}\mathbf{ncer}$, where $\Delta_{\mathbf{ind}}\text{-}\mathbf{ncer} = \mathbf{Adv}^{\mathbf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}(\mathsf{B},k)$ for a PPT adversary B. Finally, we argue that $|\epsilon_2 - \epsilon_3| \leq n\Delta_{\mathbf{stat}}\text{-}\mathbf{hide}$ where $\Delta_{\mathbf{stat}}\text{-}\mathbf{hide} =$

$\mathbf{Adv}^{\mathbf{stat}}_{\mathsf{nisCom}}$-$\mathbf{hide}(\mathsf{C}, k)$ for an unbounded powerful adversary C. All together this implies that $|\epsilon_0 - \epsilon_3| \leq 4n\Delta_{\mathbf{ind}}$-$\mathbf{ncer} + n\Delta_{\mathbf{stat}}$-$\mathbf{hide}$ and that $\epsilon_0 \leq 4n\Delta_{\mathbf{ind}}$-$\mathbf{ncer} + n\Delta_{\mathbf{stat}}$-$\mathbf{hide}$, which proves the lemma.

Claim 4.4. *$|\epsilon_0 - \epsilon_1| \leq 2n\Delta_{\mathbf{ind}}$-$\mathbf{ncer}$, where $\Delta_{\mathbf{ind}}$-$\mathbf{ncer} = \mathbf{Adv}^{\mathbf{ind}}_{\mathsf{nPKE}}$-$\mathbf{ncer}(\mathsf{B}, k)$.*

Proof: We prove the claim by introducing n intermediate hybrids experiments between $\mathbf{Exp}_0$ and $\mathbf{Exp}_1$; the difference between two consequent hybrids is bounded by a reduction to **ind**-**ncer** security of nPKE. More specifically, we introduce $n-1$ intermediate hybrid experiments so that $E_0 = \mathbf{Exp}_0$, $E_n = \mathbf{Exp}_1$ and the ith hybrid experiment E_i is defined recursively. That is,

- $E_0 = \mathbf{Exp}_0$.
- For $i = [n]$, E_i is identical to E_{i-1} except that the ith ciphertext e_i is computed by (e^*_{i0}, e_{i1}, c_i) where $(e^*_{i0}, t_{i0}) \leftarrow \mathsf{nEnc}^*_{pk_{i0}}(1^k)$. Furthermore, if $i \in \mathcal{I}$ (i.e., if A asks to open the ith ciphertext), then E_i computes $sk^*_{i0} \leftarrow \mathsf{nOpen}(sk_{i0}, e^*_{i0}, t_{i0}, m_i)$ and hands (sk^*_{i0}, sk_{i1}) to A.

Clearly $E_n = \mathbf{Exp}_1$ where the first component of all ciphertext is computed using nEnc^*. Let γ_i define the advantage of A in E_i, i.e. $\gamma_i := 2\left|\Pr[E_i(\mathsf{A}, k) = 1] - \frac{1}{2}\right|$. Next we show that $|\gamma_{i-1} - \gamma_i| \leq 2\Delta_{\mathbf{ind}}$-$\mathbf{ncer}$ for all $i \in [n]$. This implies that $|\gamma_0 - \gamma_n| \leq 2n\Delta_{\mathbf{ind}}$-$\mathbf{ncer}$. Now, since $\gamma_0 = \epsilon_0$ and $\gamma_n = \epsilon_1$ we get $|\epsilon_0 - \epsilon_1| \leq 2n\Delta_{\mathbf{ind}}$-$\mathbf{ncer}$, thus proving the claim.

We fix $i \in [n]$ and prove that $|\gamma_{i-1} - \gamma_i| \leq 2\Delta_{\mathbf{ind}}$-$\mathbf{ncer}$. Specifically, we show that any adversary B that wishes to distinguish a real ciphertext from a fake one relative to nPKE can utilize the power of adversary A. Upon receiving pk from experiment $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{nPKE}}$-$\mathbf{ncer}$and i, B interacts with A as follows.

1. B samples first a bit $b \leftarrow \{0,1\}$ and sets $pk_{i0} = pk$. It then uses nGen to generate the rest of the public keys to obtain $\mathbf{pk}$ (and all but $(i0)$th secret key).[4] Finally, it hands $\mathbf{pk}$ to A that returns Dist and $\mathsf{Resamp}_{\mathsf{Dist}}$.
2. B samples $\mathbf{m} \leftarrow \mathsf{Dist}(1^k)$ and outputs m_i to $\mathbf{Exp}^{\mathbf{ind}}_{\mathsf{nPKE}}$-$\mathbf{ncer}$ that returns (sk, e). B then sets $sk_{i0} = sk$. (Note that this completes vector $\mathbf{sk}$ since B generated the rest of the secret keys in the previous step).
   - For $j \in [i-1]$, B computes the first component of ciphertext e_j by $(e_{j0}, t_{j0}) \leftarrow \mathsf{nEnc}^*_{pk_{j0}}(1^k)$. B completes e_j honestly (i.e., exactly as specified in Enc).
   - For $j = i$, B sets the first component of e_j to be e. B completes e_j honestly.
   - For $j \in [i+1, n]$, B computes ciphertext e_j honestly.

   Let $\mathbf{e} = (e_j)_{j \in [n]}$. B hands $\mathbf{e}$ to A that returns $\mathcal{I}$.
3. B resamples $\mathbf{m}' \leftarrow \mathsf{Resamp}_{\mathsf{Dist}}(\mathbf{m}_{\mathcal{I}})$. Subsequently it hands $\mathbf{m}^*$ to A as well as secret keys for all the indices that are specified in $\mathcal{I}$, where $\mathbf{m}^* = \mathbf{m}$ if $b = 0$, $\mathbf{m}^* = \mathbf{m}'$ otherwise. That is,
   - If $j \in \mathcal{I}$ lies in $[i-1]$, then B computes $sk^*_{j0} \leftarrow \mathsf{nOpen}(sk_{j0}, e_{j0}, t_{j0}, m_j)$ and hands (sk^*_{j0}, sk_{j1}).

[4] Recall that each public key within $\mathbf{pk}$ includes two public keys relative to nPKE.

– If $j \in \mathcal{I}$ equals i, then B hands (sk_{j0}, sk_{j1}) where sk_{j0} is same as sk that B had received from $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}$.
– If $j \in \mathcal{I}$ lies in $[i+1, n]$, then B returns (sk_{j0}, sk_{j1}).

4. B outputs 1 in experiment $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}$ if A wins.

Next, note that B perfectly simulates E_{i-1} if it received a real ciphertext e within (sk, e). Otherwise, B perfectly simulates E_i. This ensures that the probability that B outputs 1 in $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}$ given a real ciphertext is at least as good as the probability that A wins in E_{i-1}. On the other hand, the probability that B outputs 1 in $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}$ given a fake ciphertext is at least as good as the probability that A wins in E_i. Since the advantage of A in E_i is γ_i, its winning probability (cf. Definition 2.2) $\Pr[E_i(\mathsf{A}, k) = 1]$ in the experiment is $\frac{\gamma_i}{2} + \frac{1}{2}$. Similarly, the winning probability of A in experiment E_{i-1} is $\frac{\gamma_{i-1}}{2} + \frac{1}{2}$. Denoting the bit picked in $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{ncer}_{\mathsf{nPKE}}$ by c we get,

$$\underbrace{\Pr\left[1 \leftarrow \mathsf{B}(sk, e) \mid (pk, sk) \leftarrow \mathsf{nGen}(1^k) \wedge e \leftarrow \mathsf{nEnc}_{pk}(m_i)\right]}_{=\Pr[1 \leftarrow \mathsf{B} \mid c=0]} \geq \frac{\gamma_{i-1}}{2} + \frac{1}{2} \quad \text{and}$$

$$\underbrace{\Pr\left[1 \leftarrow \mathsf{B}(sk, e) \mid (pk, sk) \leftarrow \mathsf{nGen} \wedge (e, t_e) \leftarrow \mathsf{nEnc}^*_{pk} \wedge sk \leftarrow \mathsf{nOpen}(sk, e, t_e, m_i)\right]}_{=\Pr[1 \leftarrow \mathsf{B} \mid c=1]} \geq \frac{\gamma_i}{2} + \frac{1}{2}.$$

This implies that

$$\begin{aligned}
\Delta_{\mathsf{ind}}\text{-}\mathbf{ncer} = \mathbf{Adv}^{\mathsf{ind\text{-}ncer}}_{\mathsf{nPKE}}(\mathsf{B}, k) &= 2\left|\Pr[\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{ncer}_{\mathsf{nPKE}}(\mathsf{B}, k) = 1] - \frac{1}{2}\right| \\
&= 2\left|\Pr[0 \leftarrow \mathsf{B} \mid c = 0]\underbrace{\Pr(c=0)}_{=1/2} + \Pr[1 \leftarrow \mathsf{B} \mid c = 1]\underbrace{\Pr(c=1)}_{=1/2} - \frac{1}{2}\right| \\
&= |\Pr[0 \leftarrow \mathsf{B} \mid c = 0] + \Pr[1 \leftarrow \mathsf{B} \mid c = 1] - 1| \\
&= |\Pr[1 \leftarrow \mathsf{B} \mid c = 0] - \Pr[1 \leftarrow \mathsf{B} \mid c = 1]| \geq \frac{|\gamma_{i-1} - \gamma_i|}{2}.
\end{aligned}$$

□

The following claim follows by a similar hybrid argument as described above.

Claim 4.5. $|\epsilon_1 - \epsilon_2| \leq 2n\Delta_{\mathsf{ind}}\text{-}\mathbf{ncer}$, *where* $\Delta_{\mathsf{ind}}\text{-}\mathbf{ncer} = \mathbf{Adv}^{\mathsf{ind}}_{\mathsf{nPKE}}\text{-}\mathbf{ncer}(\mathsf{B}, k)$.

Finally, we prove the following claim.

Claim 4.6. $|\epsilon_2 - \epsilon_3| \leq n\Delta_{\mathsf{stat}}\text{-}\mathbf{hide}$, *where* $\Delta_{\mathsf{stat}}\text{-}\mathbf{hide} = \mathbf{Adv}^{\mathsf{stat}}_{\mathsf{nisCom}}\text{-}\mathbf{hide}(\mathsf{C}, k)$.

Proof: We prove the claim by introducing n intermediate hybrids experiments between $\mathbf{Exp}_2$ and $\mathbf{Exp}_3$; we show that each pair of consecutive experiments is statistically indistinguishable based on **stat-hide** security of the NISHCOM. These hybrid experiments are defined as follows:

– $H_0 = \mathbf{Exp}_2$.
– For $i = [n]$, H_i is identical to H_{i-1} except that the ith ciphertext e_i in $\mathbf{e}$ is computed as $(e^*_{i0}, e^*_{i1}, c^*_i)$ where $c^*_i \leftarrow \mathsf{nisCommit}(m^*_i; r^*_i)$, where m^*_i is a dummy message from $\mathcal{M}_{\mathsf{nisCom}}$ and $r^* \leftarrow \mathcal{R}_{\mathsf{nisCom}}$. Furthermore, if $i \in \mathcal{I}$, then H_i computes $r_i \leftarrow \mathsf{nisOpen}(c^*_i, m_i)$ and hands (sk^*_{i0}, sk^*_{i1}) to A.

We remark again that the hybrid experiments defined above are not efficient, but this is not an issue as we rely on the statistical security of the underlying NISHCOM.

Clearly, $H_n = \mathbf{Exp}_3$ where the third component of each ciphertext within $\mathbf{e}$ is computed using dummy messages. Let ν_i be the advantage of A in H_i, i.e., $\nu_i := 2\left|\Pr[H_i(\mathsf{A}, k) = 1] - \frac{1}{2}\right|$. Next, we show that $|\nu_{i-1} - \nu_i| \leq \Delta_{\mathsf{stat}}\text{-}\mathbf{hide}$ for all $i \in [n]$, where $\Delta_{\mathsf{stat}}\text{-}\mathbf{hide} = \mathbf{Adv}^{\mathsf{stat}}_{\mathsf{nisCom}}\text{-}\mathbf{hide}(\mathsf{C}, k)$. All together, this implies that $|\nu_0 - \nu_n| \leq n\Delta_{\mathsf{stat}}\text{-}\mathbf{hide}$. Since $\nu_0 = \epsilon_2$ and $\nu_n = \epsilon_3$ we get that $|\epsilon_2 - \epsilon_3| \leq n\Delta_{\mathsf{stat}}\text{-}\mathbf{hide}$ which proves the claim.

Fix $i \in [n]$. The only difference between experiments H_{i-1} and H_i is relative to the third component of ciphertext e_i. Namely, in H_{i-1}, the third component in e_i is a commitment to m_i where m_i is the ith element in $\mathbf{m}$. On the other hand, in H_i it is a commitment to a dummy message from $\mathcal{M}_{\mathsf{nisCom}}$. As the underlying NISHCOM satisfies statistical hiding property, even an unbounded adversary C cannot distinct H_{i-1} and H_i with probability better than $\Delta_{\mathsf{stat}}\text{-}\mathbf{hide}$, so $|\nu_{i-1} - \nu_i| \leq \Delta_{\mathsf{stat}}\text{-}\mathbf{hide}$ as desired. □

We conclude with the proof of the following lemma.

Lemma 4.7. PKE *is not* **rsim-so** *secure.*

Proof: We then rely on a result of [1] which establishes that no decryption verifiable **ind-cpa** secure is **rsim-so**. Informally, decryption verifiability implies the existence of an algorithm W (that either outputs accept or reject), such that it is hard to find pk, sk_0, sk_1, distinct m_0, m_1 and a ciphertext e where both $W(pk, sk_0, e, m_0)$ and $W(pk, sk_1, e, m_1)$ accept. Note that it is hard to find two valid secret keys and plaintexts as required since decryption follows successfully only if the commitment that is part of the ciphertext is also correctly opened. In particular, an adversary that produces a ciphertext that can be successfully decrypted into two distinct plaintexts (under two different keys) must break the **comp-bind** security of the underlying commitment scheme.[5] This implies that PKE is not **rsim-so** secure. □

Compatible Secure NCER and Secure NISHCOM. We instantiate the commitment scheme with the Paillier based scheme of Damgård and Nielsen [10, 11], which is comprised of the following algorithms that use public parameters (N, g) where N is a k-bit RSA composite and $g = x^N \bmod N^2$ for an uniformly random $x \leftarrow \mathbb{Z}^*_N$.

[5] Recall that the decryption algorithm verifies first whether the commitment within the ciphertext is consistent with the decrypted ciphertexts (that encrypt the committed message and its corresponding randomness for commitment).

– $\mathsf{nisCommit}$, given N, g and message $m \in \mathbb{Z}_N$, pick $r \leftarrow \mathbb{Z}_N^*$ and compute $g^m \cdot r^N \bmod N^2$.
– $\mathsf{nisOpen}$, given commitment c and message m, compute randomness r such that $c = g^m \cdot r^N \bmod N^2$. Namely, find first $\tilde{r}$ such that $c = \tilde{r}^N \bmod N^2$. This implies that $\tilde{r}^N = (x^N)^m \cdot r^N \bmod N^2$ for some $r \in \mathbb{Z}_N^*$, since we can fix $r = \tilde{r}/x^m$.

This scheme is computationally binding, as a commitment is simply a random Paillier encryption of zero. Furthermore, opening to two different values implies finding the Nth root of g (which breaks the underlying assumption of Paillier, i.e., DCR). Finally, the NCER can be instantiated with the scheme from [7] that is also based on the DCR assumption. The message space of these two primitives is $\mathbb{Z}_N$. In addition, the randomness of the commitment scheme is $\mathbb{Z}_N^*$ and thus can be made consistent with the plaintext spaces, as it is infeasible to find an element in $\mathbb{Z}_N/\mathbb{Z}_N^*$.

4.2 Secure Tweaked NCER $\Longrightarrow$ rind-so Secure PKE

In this section we prove that every secure tweaked NCER is a **rind-so** secure PKE. Intuitively, this holds since real ciphertexts are indistinguishable from fake ones, and fake ciphertexts do not commit to any fixed plaintext. This implies that the probability of distinguishing an encryption of one message from another is exactly half, even for an unbounded adversary.

Theorem 4.8. *Assume there exists an* $\{$**ind-tcipher**, **ind-tncer**$\}$ *secure tweaked NCER, then there exists a PKE that is* **rind-so** *secure.*

Proof: More precisely, let $\mathsf{tPKE} = (\mathsf{tGen}, \mathsf{tEnc}, \mathsf{tEnc}^*, \mathsf{tDec}, \mathsf{tOpen})$ denote a secure tweaked NCER. Then we prove that tPKE is **rind-so** secure, by proving that for any PPT adversary A attacking tPKE in the **rind-so** experiment there exist a PPT adversary B and an unbounded powerful adversary C such that

$$\mathbf{Adv}^{\mathsf{rind}}_{\mathsf{tPKE}}\text{-}\mathbf{so}(\mathsf{A}, k) \leq 2n \left(\mathbf{Adv}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}(\mathsf{B}, k) + \mathbf{Adv}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tncer}(\mathsf{C}, k) \right).$$

We modify experiment **rind-so** step by step, defining a sequence of $2n + 1$ experiments and bound the advantage of A in the last experiment. The proof is then concluded by proving that any two intermediate consecutive experiments are indistinguishable due to either **ind-tcipher** security or **ind-tncer** security of tPKE. Specifically, we define a sequence of hybrid experiments $\{\mathbf{Exp}_i\}_{i=0}^{2n}$ as follows.

– $\mathbf{Exp}_0 = \mathbf{Exp}^{\mathsf{rind}}_{\mathsf{tPKE}}\text{-}\mathbf{so}$.
– For all $i \in [n]$, $\mathbf{Exp}_i$ is identical to $\mathbf{Exp}_{i-1}$ except that the ith ciphertext in vector $\mathbf{e}$ is computed by $e_i^* \leftarrow \mathsf{tEnc}^*_{pk_i}(sk_i, m_i)$, so that if $i \in \mathcal{I}$ then $\mathbf{Exp}_i$ outputs the secret key sk_i computed by tGen and hands sk_i to adversary A (here we rely on the additional property of tEnc^*).

– For all $i \in [n]$, $\mathbf{Exp}_{n+i}$ is identical to $\mathbf{Exp}_{n+i-1}$ except that the ith ciphertext in vector $\mathbf{e}$ is computed by sampling a random message $m_i^* \in \mathcal{M}_{\mathsf{tPKE}}$ first and then computing $e_i^* \leftarrow \mathsf{tEnc}^*_{pk_i}(sk_i, m_i^*)$. Next, if $i \in \mathcal{I}$ then $\mathbf{Exp}_{n+i}$ computes a secret key $sk_i^* \leftarrow \mathsf{tOpen}(e_i^*, m_i)$ and hands sk_i^* to A.

Let ϵ_i denote the advantage of A in experiment $\mathbf{Exp}_i$ i.e., $\epsilon_i := |\Pr[\mathbf{Exp}_i(\mathsf{A}, k) = 1] - \frac{1}{2}|$. We first note that $\epsilon_{2n} = 0$ since in experiment $\mathbf{Exp}_{2n}$ the adversary receives a vector of ciphertexts that are statistically independent of the encrypted plaintexts, implying that the adversary outputs the correct bit b with probability 1/2. We next show that $|\epsilon_{i-1} - \epsilon_i| \leq 2\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher}$ for any $i \in [n]$, where $\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher} = \mathbf{Adv}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}(\mathsf{B}, k)$ for a PPT adversary B. Finally, we prove that $|\epsilon_{n+i-1} - \epsilon_{n+i}| \leq 2\Delta_{\mathsf{ind}}\text{-}\mathbf{tncer}$ for any $i \in [n]$, where $\Delta_{\mathsf{ind}}\text{-}\mathbf{tncer} = \mathbf{Adv}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tncer}(\mathsf{C}, k)$ for an unbounded powerful adversary C. Together this implies that $|\epsilon_0 - \epsilon_{2n}| \leq 2n(\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher} + \Delta_{\mathsf{ind}}\text{-}\mathbf{tncer})$. So we conclude that $\epsilon_0 \leq n(\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher} + \Delta_{\mathsf{ind}}\text{-}\mathbf{tncer}) + \epsilon_{2n} = 2n(\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher} + \Delta_{\mathsf{ind}}\text{-}\mathbf{tncer})$ which concludes the proof of the theorem for all $i \in [n]$. □

Claim 4.9. $|\epsilon_{i-1} - \epsilon_i| \leq 2n\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher}$, *where* $\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher} = \mathbf{Adv}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}(\mathsf{B}, k)$.

Proof: In the following, we prove that one can design an adversary B that distinguishes a real ciphertext from a fake one in $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}$, using adversary A. B interacts with A as follows:

1. Upon receiving pk from $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}$ and an integer i, B sets $pk_i = pk$. It picks a bit b randomly. It then generates the rest of the public and secret key pairs using tGen for all $j \in [n] \backslash i$, obtaining $\mathbf{pk}$. It hands $\mathbf{pk}$ to A who returns Dist and $\mathsf{Resamp}_{\mathsf{Dist}}$.
2. B samples $\mathbf{m} \leftarrow \mathsf{Dist}(1^k)$ and hands m_i to $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}$ which returns (sk, e). B fixes $e_i = e$ and completes $\mathbf{sk}$ by setting $sk_i = sk$. Next, for $j \in [i-1]$ it computes $e_j \leftarrow \mathsf{tEnc}^*_{pk_j}(sk_j, m_j)$, whereas for $j \in [i+1, n]$ it samples randomness $r_j \leftarrow \mathcal{R}_{\mathsf{tPKE}}$ and computes $e_j \leftarrow \mathsf{tEnc}_{pk_j}(m_j; r_j)$. Let $\mathbf{e} = (e_i)_{i \in [n]}$. B hands $\mathbf{e}$ to A who returns $\mathcal{I}$.
3. B samples $\mathbf{m}' \leftarrow \mathsf{Resamp}(\mathbf{m}_{\mathcal{I}})$ and hands A $\mathbf{m}^*$ and the following secret keys for all the indices that are specified in $\mathcal{I}$. Here $\mathbf{m}^*$ is $\mathbf{m}$ if $b = 0$ and $\mathbf{m}'$ otherwise. That is,
   – If $j \in \mathcal{I}$ lies in $[i-1]$ or in $[i+1, n]$, then B returns sk_j.
   – If $j \in \mathcal{I}$ equals i, then B returns sk.
4. B outputs 1 in $\mathbf{Exp}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tcipher}$ if A wins.

Next, note that B perfectly simulates $\mathbf{Exp}_{i-1}$ if it receives a real ciphertext e within (sk, e). On the other hand, B perfectly simulates $\mathbf{Exp}_i$ if e is a fake ciphertext. This ensures that the probability that B outputs 1 given a real ciphertext is at least as good as the probability that A wins in $\mathbf{Exp}_{i-1}$. On the other hand, the probability that B outputs 1 given a fake ciphertext is at least as good as the probability that A wins in $\mathbf{Exp}_i$. Since the advantage of A in $\mathbf{Exp}_i$ is ϵ_i, its winning probability (cf. Definition 2.2) $\Pr[\mathbf{Exp}_i(\mathsf{A}, k) = 1]$ in the experiment is

$\frac{\epsilon_i}{2}+\frac{1}{2}$. Similarly, the winning probability of A in experiment $\mathbf{Exp}_{i-1}$ is $\frac{\epsilon_{i-1}}{2}+\frac{1}{2}$. Denoting the bit picked in $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{tcipher}_{\mathsf{tPKE}}$ by c,

$$\underbrace{\Pr\left[1 \leftarrow \mathsf{B}(pk, sk, e, m_i) \mid (pk, sk) \leftarrow \mathsf{tGen}(1^k) \wedge e \leftarrow \mathsf{tEnc}_{pk}(m_i)\right]}_{=\Pr[1 \leftarrow \mathsf{B} \mid c=0]} \geq \frac{\epsilon_{i-1}}{2}+\frac{1}{2} \text{ and}$$

$$\underbrace{\Pr\left[1 \leftarrow \mathsf{B}(pk, sk, e^*, m_i) \mid (pk, sk) \leftarrow \mathsf{tGen}(1^k) \wedge e^* \leftarrow \mathsf{tEnc}^*_{pk}(sk, m_i)\right]}_{=\Pr[1 \leftarrow \mathsf{B} \mid c=1]} \geq \frac{\epsilon_i}{2}+\frac{1}{2}.$$

This implies that

$$\begin{aligned}
\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher} &= \mathbf{Adv}^{\mathsf{ind\text{-}tcipher}}_{\mathsf{tPKE}}(\mathsf{B}, k) = 2\left|\Pr[\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{tcipher}_{\mathsf{tPKE}}(\mathsf{B}, k) = 1] - \frac{1}{2}\right| \\
&= 2\left|\Pr[0 \leftarrow \mathsf{B} \mid c = 0]\underbrace{\Pr(c = 0)}_{=1/2} + \Pr[1 \leftarrow \mathsf{B} \mid c = 1]\underbrace{\Pr(c = 1)}_{=1/2} - \frac{1}{2}\right| \\
&= |\Pr[0 \leftarrow \mathsf{B} \mid c = 0] + \Pr[1 \leftarrow \mathsf{B} \mid c = 1] - 1| \\
&= |\Pr[1 \leftarrow \mathsf{B} \mid c = 0] - \Pr[1 \leftarrow \mathsf{B} \mid c = 1]| \geq \frac{|\epsilon_{i-1} - \epsilon_i|}{2}
\end{aligned}$$

□

Claim 4.10. $|\epsilon_{n+i-1} - \epsilon_{n+i}| \leq 2n\Delta_{\mathsf{ind}}\text{-}\mathbf{tcipher}$ *for all* $i \in [n]$, *where* $\Delta_{\mathsf{ind}}\text{-}\mathbf{tncer} = \mathbf{Adv}^{\mathsf{ind}}_{\mathsf{tPKE}}\text{-}\mathbf{tncer}(\mathsf{C}, k)$.

Proof: We prove that one can design an unbounded adversary C that distinguishes the two views generated in experiment **ind-tncer**, using adversary A. C interacts with A:

1. Upon receiving pk from $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{tncer}_{\mathsf{tPKE}}$ and an integer i, C sets $pk_i = pk$ and picks a bit b. It then generates the rest of the public and secret key pairs using tGen for all $j \in [n]\backslash\{i\}$, obtaining $\mathbf{pk}$. It hands $\mathbf{pk}$ to A who returns Dist and $\mathsf{Resamp}_{\mathsf{Dist}}$.
2. C samples $\mathbf{m} \leftarrow \mathsf{Dist}(1^k)$ and hands m_i to $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{tncer}_{\mathsf{tPKE}}$ which returns (sk, e). C fixes $e_i = e$ and completes $\mathbf{sk}$ by setting $sk_i = sk$. Next, for $j \in [i-1]$ it samples $\mathbf{m}^*_j \leftarrow \mathcal{M}_{\mathsf{tPKE}}$ and computes $e_j \leftarrow \mathsf{tEnc}^*_{pk_j}(sk_j, m^*_j)$, whereas for $j \in [i+1, n]$ it computes $e_j \leftarrow \mathsf{tEnc}^*_{pk_j}(sk_j, m_j)$. Let $\mathbf{e} = (e_j)_{j\in[n]}$. C hands $\mathbf{e}$ to A receiving $\mathcal{I}$.
3. C samples $\mathbf{m}' \leftarrow \mathsf{Resamp}(\mathbf{m}_{\mathcal{I}})$ and hands $\mathbf{m}^*$ to A and the following secret keys for all the indices that are specified in $\mathcal{I}$. Here $\mathbf{m}^*$ is $\mathbf{m}$ if $b = 0$ and $\mathbf{m}'$ otherwise. That is,
   - If $j \in \mathcal{I}$ lies in $[i-1]$, then C returns sk_j such that $sk_j = \mathsf{tOpen}(e_j, m_j)$.
   - If $j \in \mathcal{I}$ equals i, then C returns sk.
   - If $j \in \mathcal{I}$ lies in $[i+1, n]$, then C returns sk_j.

4. C outputs 1 in $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{tncer}_{\mathsf{tPKE}}$ if A wins.

Next, note that B perfectly simulates $\mathbf{Exp}_{n+i-1}$ if it receives a real ciphertext e within (sk, e). On the other hand, B perfectly simulates $\mathbf{Exp}_{n+i}$ if e is a fake ciphertext and sk is a secret key returned by tOpen. This ensures that the probability that B outputs 1 given a real ciphertext is at least as good as the probability that A wins in $\mathbf{Exp}_{n+i-1}$. On the other hand, the probability that B outputs 1 given a fake ciphertext is at least as good as the probability that A wins in $\mathbf{Exp}_{n+i}$. Since the advantage of A in $\mathbf{Exp}_i$ is ϵ_{n+i}, its winning probability (c.f Definition 2.2) $\Pr[\mathbf{Exp}_i(\mathsf{A}, k) = 1]$ in the experiment is $\frac{\epsilon_{n+i}}{2} + \frac{1}{2}$. Similarly, the winning probability of A in experiment $\mathbf{Exp}_{n+i-1}$ is $\frac{\epsilon_{n+i-1}}{2} + \frac{1}{2}$. Denoting the bit picked in $\mathbf{Exp}^{\mathsf{ind}}\text{-}\mathbf{tncer}_{\mathsf{tPKE}}$ by c we get,

$$\underbrace{\Pr\left[1 \leftarrow \mathsf{C}(sk, e) \mid (pk, sk) \leftarrow \mathsf{tGen}(1^k) \wedge e \leftarrow \mathsf{tEnc}^*_{pk}(sk, m_i)\right]}_{=\Pr[1 \leftarrow \mathsf{C} \mid c=0]} \geq \frac{\epsilon_{n+i-1}}{2} + \frac{1}{2} \text{ and}$$

$$\underbrace{\Pr\left[1 \leftarrow \mathsf{C}(sk^*, e^*) \mid (pk, sk) \leftarrow \mathsf{tGen} \wedge e^* \leftarrow \mathsf{tEnc}^*_{pk}(sk, m^*) \wedge sk^* \leftarrow \mathsf{tOpen}(e^*, sk, m_i)\right]}_{=\Pr[1 \leftarrow \mathsf{C} \mid c=1]} \geq \frac{\epsilon_{n+i}}{2} + \frac{1}{2}.$$

Following a similar argument as in the previous claim, we conclude that $2\Delta_{\mathsf{ind}}\text{-}\mathbf{tncer} \geq |\epsilon_{n+i-1} - \epsilon_{n+i}|$. ■

4.3 Secure NCER $\Longrightarrow$ rsim-so Secure PKE

In this section we claim that secure NCER implies selective opening security in the presence of receiver corruption. Our theorem is stated for the stronger simulation based security definition but holds for the indistinguishability definition as well. The proof is given in the full version [17].

Theorem 4.11. *Assume there exists an* **ind-ncer** *secure PKE, then there exists a PKE that is* **rsim-so** *secure.*

4.4 rsim-so Secure PKE $\not\Rightarrow$ Secure NCER and Tweaked NCER

In this section we prove that **rsim-so** does not imply both tweaked NCER and NCER by providing a concrete counter example based on an extended key-simulatable PKE (cf. see Key-Simulatable PKE subsubsection of Sect. 3). The key point in our proof is that in some cases simulatable public keys cannot be explained as valid public keys. Formally,

Theorem 4.12. *Assume there exists an* $\{$**ind-cpa**, **ksim**$\}$ *secure extended key-simulatable PKE, then there exists a PKE that is* **rsim-so** *secure but is not a* $\{$**ind-tcipher**, **ind-tncer**$\}$ *secure tweaked NCER nor a* **ind-ncer** *secure NCER.*

Proof: We describe our separating encryption scheme first; the complete proof is given in the full version [17]. Given an extended key-simulatable PKE $\mathsf{sPKE} = (\mathsf{sGen}, \mathsf{sEnc}, \mathsf{sDec}, \widetilde{\mathsf{sGen}}, \widetilde{\mathsf{sGen}}^{-1})$ for a plaintext space $\mathcal{M}_{\mathsf{sPKE}}$, we construct a new scheme $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with a binary plaintext space that is **rsim-so** secure, and thus also **rind-so** secure, yet it does not imply tweaked NCER. For simplicity, we assume that $\mathcal{M}_{\mathsf{sPKE}}$ is the binary space $\{0,1\}$. The DDH based instantiation of sPKE with $\mathcal{V} \subset \mathcal{K}$ from see Realizing Key-Simulatable and Extended Key-Simulatable PKE subsection of Sect. 4.4 is defined with respect to this space. ■

$\mathsf{Gen}(1^k)$
$\alpha \leftarrow \{0,1\}$
$(pk_\alpha, sk_\alpha) \leftarrow \mathsf{sGen}(1^k)$
$(pk_{1-\alpha}, r_{1-\alpha}) \leftarrow \widetilde{\mathsf{sGen}}(1^k)$
$pk = (pk_0, pk_1)$
$sk = (\alpha, sk_\alpha, r_{1-\alpha})$
Return (pk, sk)

$\mathsf{Enc}_{pk}(b)$
$e_0 \leftarrow \mathsf{Enc}_{pk_0}(b)$
$e_1 \leftarrow \mathsf{Enc}_{pk_1}(b)$
Return $e = (e_0, e_1)$

$\mathsf{Dec}_{sk}(e)$
$sk = (\alpha, sk_\alpha, r_{1-\alpha})$
$e := (e_0, e_1)$
$b = \mathsf{Dec}_{sk_\alpha}(e_\alpha)$
Return b

Realizing Key-Simulatable and Extended Key-Simulatable PKE. An example of a $\{$**ind-cpa**, **ksim**$\}$ secure key-simulatable PKE is the ElGamal PKE [14] where we set $\mathcal{K}$ to be equal to the set of *valid* public keys, i.e. $\mathcal{K} = \mathcal{V}$. In addition, note that any simulatable PKE as defined in [9] is also $\{$**ind-cpa**, **ksim**$\}$ secure key-simulatable PKE.

Below we provide an example of extended key-simulatable PKE with security under the DDH assumption. For simplicity we consider a binary plaintext space. Let $(g_0, g_1, p) \leftarrow \mathcal{G}(1^k)$ be an algorithm that given a security parameter k returns a group description $\mathbb{G} = \mathbb{G}_{g_0,g_1,p}$ specified by its generators g_0, g_1 and its order p. Furthermore, we set $\mathcal{K} = \mathbb{G}^2$ and $\mathcal{V} = \{(g_0^x, g_1^x) \in \mathbb{G}^2 \mid x \in \mathbb{Z}_p\}$. Then define the following extended key-simulatable PKE,

- sGen, given the security parameter k, set $(g_0, g_1, p) \leftarrow \mathcal{G}(1^k)$. Choose uniformly random $x \leftarrow \mathbb{Z}_p$ and compute $h_i = g_i^x$ for all $i \in \{0,1\}$. Output the secret key $sk = x$ and the public key $pk = (h_0, h_1)$.
- sEnc, given the public key pk and plaintext $m \in \{0,1\}$, choose a uniformly random $s, t \leftarrow \mathbb{Z}_p$. Output the ciphertext $(g_0^s g_1^t, g_0^m \cdot (h_0^s h_1^t))$.
- sDec, given the secret key x and ciphertext (g_c, h_c), output $h_c \cdot (g_c^x)^{-1}$.
- $\widetilde{\mathsf{sGen}}$, given 1^k, output two random elements from $\mathbb{G}$ and their bit sequence as the randomness.
- $\widetilde{\mathsf{sGen}}^{-1}$, given a legitimate public key h_0, h_1, simply returns the bit strings of h_0, h_1 as the randomness used to sample them from $\mathbb{G}^2$ by $\widetilde{\mathsf{sGen}}$.

We remark that a public key chosen randomly from $\mathbb{G}^2$ does not necessarily correspond to a secret key. Furthermore, $\Pr\left[pk \in \mathcal{K} \backslash \mathcal{V} \mid pk \leftarrow \widetilde{\mathsf{sGen}}(1^k)\right]$ is non-negligible. This is a key property in our proof from Sect. 4.4.

4.5 Realizing Tweaked NCER

Based on key-simulatable PKE. We prove that secure tweaked NCER can be built based on any secure key-simulatable PKE with $\mathcal{K} = \mathcal{V}$ (cf. Definition see Key-Simulatable PKE subsubsection of Sect. 3). Specifically, our construction is based on the separating scheme presented in Sect. 4.4. In addition, we define the fake encryption algorithm so that it outputs two ciphertexts that encrypt two distinct plaintexts rather than the same plaintext twice (implying that ciphertext indistinguishability follows from the **ind-cpa** security of the underlying encryption scheme). More formally, the fake encryption algorithm can be defined as follows. Given $sk = (\alpha, sk_\alpha, r_{1-\alpha})$ and message b, a fake encryption of b is computed by $e^* = (\mathsf{sEnc}_{pk_0}(b), \mathsf{sEnc}_{pk_1}(1-b))$ if $\alpha = 0$ and $e^* = (\mathsf{sEnc}_{pk_0}(1-b), \mathsf{sEnc}_{pk_1}(b))$ otherwise. It is easy to verify that given sk, the decryption of e^* returns b and that e^* is computationally indistinguishable from a valid encryption even given the secret key. Next, we discuss the details of the non-efficient opening algorithm which is required to generate a secret key for a corresponding public key given a fake ciphertext and a message b'. In more details, assuming $sk = (\alpha, sk_\alpha, r_{1-\alpha})$ and $pk = (pk_0, pk_1)$,

$$\mathsf{tOpen}(sk, pk, (e_0^*, e_1^*), b') = \begin{cases} (\alpha, sk_\alpha, r_{1-\alpha}) & \text{if } e_\alpha^* = \mathsf{sEnc}_{pk_\alpha}(b') \\ (1-\alpha, sk_{1-\alpha}, r_\alpha) & \text{else, where } r_\alpha \leftarrow \widetilde{\mathsf{sGen}}^{-1}(pk_\alpha) \\ & \text{and } sk_{1-\alpha} \text{ is a valid secret key} \\ & \text{of } pk_{1-\alpha}. \end{cases}$$

Note that since it holds that $\mathcal{V} = \mathcal{K}$ for the underlying sPKE scheme, there exists a secret key that corresponds to $pk_{1-\alpha}$ and it can be computed (possibly in an inefficient way). Encryption schemes for larger plaintext spaces can be obtained by repeating this basic scheme sufficiently many times.[6] Finally, we note that the scheme is {**ind-tcipher**, **ind-tncer**} secure. Recalling that any simulatable PKE with $\mathcal{K} = \mathcal{V}$ is a key-simulatable PKE [8,9], we conclude that secure tweaked NCER for a binary plaintext space can be built relying on DDH, RSA, factoring and LWE assumptions.

An additional realization based on statistically-hiding $\binom{2}{1}$-OT in presented in the full version [17]. These two implementations support binary plaintext space. Below presented new constructions that support exponential plaintext spaces.

Based on NCER. We show that the DCR based secure NCER of [7] is also a secure tweaked NCER. Let $(p', q') \leftarrow \mathcal{G}(1^n)$ be an algorithm that given a security parameter k returns two random n bit primes p' and q' such that $p = 2p' + 1$ and $q = 2q' + 1$ are also primes. Let $N = pq$ and $N' = p'q'$. Define $(\mathsf{tGen}, \mathsf{tEnc}, \mathsf{tEnc}^*, \mathsf{tDec}, \mathsf{tOpen})$ by,

– tGen, given the security parameter k, run $(p', q') \leftarrow \mathcal{G}(1^n)$ and set $p = 2p' + 1$, $q = 2q' + 1$, $N = pq$ and $N' = p'q'$. Choose random $x_0, x_1 \leftarrow \mathbb{Z}_{N^2/4}$ and a

[6] We note that this construction was discussed in [16] in the context of weak hash proof systems and leakage resilient PKE.

random $g' \in \mathbb{Z}^*_{N^2}$ and compute $g_0 = g'^{2N}$, $h_0 = g_0^{x_0}$ and $h_1 = g_0^{x_1}$. Output public key $pk = (N, g_0, h_0, h_1)$ and secret key $sk = (x_0, x_1)$.

– tEnc, given the public key pk and a plaintext $m \in \mathbb{Z}_N$, choose a uniformly random $t \leftarrow \mathbb{Z}_{N/4}$ and output ciphertext

$$c \leftarrow \mathsf{tEnc}_{pk}(m;t) = \left(g_0^t \bmod N^2, (1+N)^m h_0^t \bmod N^2, h_1^t \bmod N^2\right).$$

– tDec, given the secret key (x_0, x_1) and a ciphertext (c_0, c_1, c_2), check whether $c_0^{2x_1} = (c_2)^2$; if not output $\perp$. Then set $\hat{m} = (c_1/c_0^{x_0})^{N+1}$. If $\hat{m} = 1 + mN$ for some $m \in \mathbb{Z}_N$, then output m; else output $\perp$.

– tEnc^*, given the public key pk, secret key sk and a message m, choose uniformly random $t \leftarrow \mathbb{Z}_{\phi(N)/4}$, compute the fake ciphertext (where all the group elements are computed $\bmod N^2$) $c^* \leftarrow (c_0^*, c_1^*, c_2^*) = ((1+N) \cdot g_0^t, (1+N)^m \cdot (c_0^*)^{x_0}, (c_0^*)^{x_1})$.

– tOpen, given N', (x_0, x_1), a triple (c_0, c_1, c_2) such that $(c_0, c_1, c_2) \leftarrow \mathsf{tEnc}^*_{pk}(sk, m)$ and a plaintext $m^* \in \mathbb{Z}_N$, output $sk^* = (x_0^*, x_1)$, where $x_0^* \leftarrow \mathbb{Z}_{NN'}$ is the unique solution to the equations $x_0^* = x \bmod N'$ and $x_0^* = x_0 + m - m^* \bmod N$. These equations have a unique solution due to the fact that $gcd(N, N') = 1$ and the solution can be obtained employing Chinese Remainder Theorem. It can be verified that the secret key sk^* matches the public key pk and also decrypts the 'simulated' ciphertext to the required message m^*. The first and third components of pk remain the same since x_1 has not been changed. Now $g^{x_0^*} = g^{x_0^* \bmod N'} = g^{x_0 \bmod N'} = g^{x_0} = h_0$. Using the fact that the order of $(1+N)$ in $\mathbb{Z}^*_{N^2}$ is N, we have

$$\left(\frac{c_1}{c_0^{x_0^*}}\right)^{N+1} = \left(\frac{(1+N)^{x_0+m} g_0^{tx_0}}{(1+N)^{x_0^*} g_0^{tx_0^*}}\right)^{N+1}$$
$$= \left((1+N)^{x_0+m-x_0^* \bmod N}\right)^{N+1} = ((1+N)^m)^{N+1} = (1+mN).$$

It is easy to verify that real and fake ciphertexts are computationally indistinguishable under the DCR assumption since the only difference is with respect to the first element (which is an $2N$th power in a real ciphertext and not an $2N$th power in a simulated ciphertext). The other two elements are powers of the first element. Furthermore $sk = (x_0, x_1)$ and $sk^* = (x_0^*, x_1)$ are statistically close since $x_0 \leftarrow \mathbb{Z}_{N^2/4}$ and $x_0^* \leftarrow \mathbb{Z}_{NN'}$ and the uniform distribution over $\mathbb{Z}_{NN'}$ and $\mathbb{Z}_{N^2/4}$ is statistically close.

5 Selective Opening Security for the Sender

In this section we prove **sind-so** is strictly weaker than **ssim-so** security by constructing a scheme that meets the former but not the latter level of security. Our starting point is a lossy encryption scheme $\mathsf{loPKE} = (\mathsf{loGen}, \mathsf{loGen}^*, \mathsf{loEnc}, \mathsf{loDec})$. We then modify loPKE by adding a (statistically hiding) commitment to each ciphertext such that the new scheme, denoted by PKE, becomes committing.

Next, we prove that PKE is **sind-so** secure by showing that the scheme remains lossy and is therefore **sind-so** secure according to [2]. Finally, using the result from [1] we claim that PKE is not **ssim-so** secure. The following theorem is proven in the full version [17].

Theorem 5.1. *Assume there exists a* {**ind-lossy**, **ind-lossycipher**} *secure lossy PKE and a* {**stat-hide**, **comp-bind**} *secure NISHCOM that are compatible. Then, there exists a PKE that is* **sind-so** *secure but is not* **ssim-so** *secure.*

Acknowledgements. Carmit Hazay acknowledges support from the Israel Ministry of Science and Technology (grant No. 3-10883). Arpita Patra acknowledges support from project entitled 'ISEA - Part II' funded by Department of Electronics and Information Technology of Govt. of India. Part of this work was carried out while Bogdan Warinschi was visiting Microsoft Research, Cambridge, UK and IMDEA, Madrid, Spain. He has been supported in part by ERC Advanced Grant ERC-2010-AdG-267188-CRIPTO, by EPSRC via grant EP/H043454/1, and has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement 609611 (PRACTICE).

References

1. 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 (2012)
2. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009)
3. Bellare, M., Waters, B., Yilek, S.: Identity-based encryption secure against selective opening attack. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 235–252. Springer, Heidelberg (2011)
4. Bellare, M., Yilek, S.: Encryption schemes secure under selective opening attack. In: IACR Cryptology ePrint Archive 2009, p. 101 (2009)
5. Pinkas, B., Schneider, T., Smart, N.P., Williams, S.C.: Secure two-party computation is practical. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 250–267. Springer, Heidelberg (2009)
6. Canetti, R., Friege, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: STOC, pp. 639–648 (1996)
7. Canetti, R., Halevi, S., Katz, J.: Adaptively-secure, non-interactive public-key encryption. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 150–168. Springer, Heidelberg (2005)
8. Choi, S.G., Dachman-Soled, D., Malkin, T., Wee, H.: Improved non-committing encryption with applications to adaptively secure protocols. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 287–302. Springer, Heidelberg (2009)
9. Damgård, I.B., 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 (2000)
10. Damgård, I.B., Nielsen, J.B.: Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 581–596. Springer, Heidelberg (2002)

11. Damgård, I.B., Nielsen, J.B.: Universally composable efficient multiparty computation from threshold homomorphic encryption. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 247–264. Springer, Heidelberg (2003)
12. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. J. ACM **50**(6), 852–921 (2003)
13. 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 (2010)
14. El Gamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Trans. Inf. Theor. **31**(4), 469–472 (1985)
15. Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. **28**(2), 270–299 (1984)
16. 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 (2013)
17. Hazay, C., Patra, A., Warinschi, B. Selective opening security for receivers. In: IACR Cryptology ePrint Archive 2015, p. 860 (2015)
18. 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 (2011)
19. 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 (2014)
20. Huang, Z., Liu, S., Qin, B.: Sender-equivocable encryption schemes secure against chosen-ciphertext attacks revisited. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 369–385. Springer, Heidelberg (2013)
21. Jarecki, S., Lysyanskaya, A.: Adaptively secure threshold cryptography: introducing concurrency, removing erasures (extended abstract). In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, p. 221. Springer, Heidelberg (2000)
22. 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 (2002)
23. Ostrovsky, R., Rao, V., Visconti, I.: On selective-opening attacks against encryption schemes. In: Abdalla, M., De Prisco, R. (eds.) SCN 2014. LNCS, vol. 8642, pp. 578–597. Springer, Heidelberg (2014)