

On definitions of selective opening security

Florian Böhl*, Dennis Hofheinz[†], and Daniel Kraschewski

Karlsruhe Institute of Technology

{Florian.Boehl,Dennis.Hofheinz,Daniel.Kraschewski}@kit.edu

October 29, 2013

Abstract

Assume that an adversary observes many ciphertexts, and may then ask for openings, i.e. the plaintext and the randomness used for encryption, of some of them. Do the unopened ciphertexts remain secure? There are several ways to formalize this question, and the ensuing security notions are not known to be implied by standard notions of encryption security. In this work, we relate the two existing flavors of selective opening security. Our main result is that indistinguishability-based selective opening security and simulation-based selective opening security do not imply each other.

We show our claims by counterexamples. Concretely, we construct two public-key encryption schemes. One scheme is secure under selective openings in a simulation-based sense, but not in an indistinguishability-based sense. The other scheme is secure in an indistinguishability-based sense, but not in a simulation-based sense.

Our results settle an open question of Bellare et al. (Eurocrypt 2009). Also, taken together with known results about selective opening secure encryption, we get an almost complete picture how the two flavors of selective opening security relate to standard security notions.

Keywords: security definitions, selective opening security, public-key encryption

1 Introduction

Security under selective openings. Assume that an adversary observes many ciphertexts, and may then ask for openings of some of them. Do the unopened ciphertexts remain secure? Somewhat surprisingly, security in this setting is not known to be implied by standard security notions for encryption schemes (such as IND-CPA security). In fact, very recently, Bellare et al. [2] showed that a whole class of IND-CPA secure public-key encryption schemes do not achieve a simulation-based notion of selective opening security.

To date, there are two flavors of definitions to capture security under selective openings: simulation-based selective opening security (SIM-SO security, [7, 1]) and indistinguishability-based selective opening security (IND-SO security, [1]). There are indications that SIM-SO and IND-SO security constitute orthogonal requirements. For instance, when looking at selective opening security for commitment schemes, Bellare et al. prove that any statistically hiding commitment scheme is IND-SO secure; however, there are severe limitations on the construction of SIM-SO secure commitment schemes from a number of standard assumptions [1]. Nonetheless, in case of encryption schemes (which are the focus of this paper), no similar result is known.

We will now describe the existing security notions for selective opening security, along with known results.

*Supported by MWK grant “MoSeS”.

[†]Supported by DFG grant GZ HO 4534/2-1.

Simulation-based selective opening security (SIM-SO-CPA). An encryption scheme is called SIM-SO-CPA secure, if anything an adversary can compute from a vector of ciphertexts *and* openings of a subset of these ciphertexts, can also be computed by a simulator that only sees the opened messages (but no ciphertexts at all). SIM-SO-CPA security dates back to Dwork et al. [7], who consider the same security notion for commitments. In the encryption context, SIM-SO-CPA security has been investigated by Bellare et al. [1], who also observe that Goldasser-Micali encryption [9] achieves SIM-SO-CPA security. Later on, several other SIM-SO-CPA secure encryption schemes have been constructed [8, 12, 13].

However, all known SIM-SO-CPA secure encryption schemes are comparatively inefficient: they either encrypt messages bitwise [9, 8], or they are based on assumptions related to Paillier encryption [12, 13]. There is no known efficient SIM-SO-CPA secure encryption scheme based on, say, the DDH problem in a suitable cyclic group. One key difficulty seems to be that SIM-SO-CPA security essentially requires that the encryption is non-committing, such that a ciphertext can be efficiently opened to any message [3, 4, 8] (possibly using a special trapdoor). In fact, Bellare et al. [2] use this property in a clever way to construct an encryption scheme that is IND-CPA secure, but not SIM-SO-CPA secure.

Indistinguishability-based selective opening security (IND-SO-CPA). An encryption scheme is called IND-SO-CPA secure, if no adversary, after given a vector of ciphertexts and openings of a subset of these ciphertexts, can distinguish the unopened messages from fresh messages. There is one subtlety here. Namely, in most applications, the initially received ciphertext vector may contain encryptions of related messages (e.g., encryptions of shares of a secret value). Hence, the “fresh” messages that the adversary must distinguish from the actually encrypted (but unopened) messages must be *conditioned* on the already opened messages. Note that depending on the underlying distribution of message vectors, conditioning on an arbitrary subset of messages can be an inefficient process. In particular, the IND-SO-CPA security experiment may be inefficient.

This subtlety has led to two different IND-SO-CPA variations. *Full IND-SO-CPA* security requires exactly what we have sketched above, with a potentially inefficient security experiment. The problem with full IND-SO-CPA security is that there are no known fully IND-SO-CPA secure encryption schemes.¹

On the other hand, *weak IND-SO-CPA* security requires the above, but only for distributions of message vectors that are efficiently re-samplable. Here, efficiently re-samplable means that the message distribution can be efficiently sampled, even when fixing any subset of messages to a particular value.² The advantage of weak IND-SO-CPA security is that any lossy encryption scheme [16] is already weakly IND-SO-CPA secure [1]. In particular, there are very efficient weakly IND-SO-CPA secure encryption schemes based on standard assumptions. This is also an important advantage over full IND-SO-CPA security for which no realizations are known yet.

The main disadvantage of weak IND-SO-CPA security is that it is obviously only useful in settings in which the joint distribution of all encrypted messages actually is efficiently re-samplable. Many conceivable settings (e.g., when commitments or other non-invertible functions of other messages are encrypted) do not conform to such a re-samplability condition.

The current situation. So far, we can summarize that SIM-SO-CPA as well as (full or weak) IND-SO-CPA security both have advantages and disadvantages. It depends on the concrete setting and requirements which notion is to prefer. However, so far little is known about the *relations* among those notions of selective opening security. A few implications are trivial or at least follow with little effort: full IND-SO-CPA security obviously implies weak IND-SO-CPA security, and

¹We mention that for *commitments*, the situation is less problematic: every statistically commitment scheme is (fully) IND-SO secure [1]. However, a moment of reflection shows that there can be no statistically hiding *encryption* scheme. The closest we can get to statistically hiding encryption is lossy encryption, which is only known to imply *weak* IND-SO-CPA security.

²For instance, the distribution of message tuples (x, x) is efficiently re-samplable, while the distribution (x, g^x) is not (where $x \in \mathbb{Z}_{|\mathbb{G}|}$ is uniform, and $g \in \mathbb{G}$ for some group \mathbb{G} in which discrete logarithms are hard to compute).




Figure 1: Relations of different definitions of selective opening security and IND-CPA. The bold arrows illustrate the results of our work while BDWY11 is the main result of [2]. Crossed arrows symbolize concrete counterexamples and dashed arrows stand for open questions. All other arrows are implications that are pretty much straightforward or follow directly from the already settled relations. Note that the question whether weak IND-SO-CPA security implies SIM-SO-CPA security is settled negatively if a fully IND-SO-CPA secure encryption scheme exists.

it is not hard to see that SIM-SO-CPA security implies weak IND-SO-CPA security. However, otherwise the relation in particular between full IND-SO-CPA security and SIM-SO-CPA security is not known. (We again stress that for *commitments*, the situation is a little different, as sketched above; however, these results do not apply to encryption schemes.)

Our contribution. This paper attempts to fill this gap: we relate full IND-SO-CPA security and SIM-SO-CPA security. Our results show that full IND-SO-CPA security does not imply SIM-SO-CPA security, and vice versa. We give concrete counterexamples, i.e., encryption schemes that are fully IND-SO-CPA secure, but not SIM-SO-CPA secure (and the other way around). In a sense, our results further isolate full IND-SO-CPA security from other notions of encryption scheme security. Thus, there is even less motivation to study full IND-SO-CPA security. Figure 1 depicts the relations of the different flavors of selective opening security to one another and to IND-CPA security.

We now provide some more technical background on our results.

Our first counterexample. We first construct a scheme that is SIM-SO-CPA secure, but not fully IND-SO-CPA secure. The basic idea is to take any SIM-SO-CPA secure scheme, and modify it such that it becomes vulnerable to a full IND-SO-CPA attack (while preserving its SIM-SO-CPA security, of course). Our modification is simple: we add a tuple

$$((g^s u^t)^M, (h^s v^t)) \quad (1)$$

to each ciphertext, where M is the encrypted message, s, t are random exponents, and g, h, u, v are group elements that are part of the public key. In the scheme, $(g, h, u, v) = (g, h, g^\omega, h^\omega)$ is a Diffie-Hellman tuple, such that (1) is a perfectly binding commitment to M . However, during the proof that the modified scheme is still SIM-SO-CPA secure, we will switch (g, h, u, v) to a non-Diffie-Hellman tuple. Then, (1) becomes a perfectly hiding commitment, which can actually be equivocated arbitrarily. (Note that this added commitment really only is an instance of the dual-mode commitment schemes from Damgård and Nielsen [5].) This allows to open ciphertexts in our modified scheme arbitrarily, and shows the modified scheme SIM-SO-CPA secure.

To prove that the modified scheme is not fully IND-SO-CPA secure, we first define a suitable distribution dist of message tuples (x, z) , such that re-sampling dist essentially requires computing a discrete logarithm. Concretely, we define dist such that $(x, z) = (x, g^x)$, resp. $(x, z) = (x, h^x)$ (for uniform x and g, h from the scheme's public key) with probability $1/2$ each. Now suppose an adversary starts off with two ciphertexts, one for x and one for $z = g^x$. He then chooses to open the second ciphertext (for $z = g^x$), which fixes the second component of the ciphertext vector. (However, note that the adversary does not know x at this point.)

Assume, when invoked with the challenge message vector, he then gets a first component y , sampled from dist conditioned on the second component z . By our definition of dist , with a probability of $1/2$, the adversary then does not get $y = x$, but the unique y with $z = h^y$. Note that then, $x = y \cdot \text{dlog}_g h$. Using this relation, the adversary can recognize that the first unopened ciphertext (with commitment $((g^s u^t)^x, (h^s v^t))$) really contained x . This check works *only* if re-sampling occurred, and hence the adversary successfully distinguishes authentic from re-sampled messages. As SIM-SO-CPA security implies IND-CPA security, this counterexample also shows that IND-CPA security does not imply full IND-SO-CPA security.

Our second counterexample. We proceed to construct a scheme that is fully IND-SO-CPA secure, but not SIM-SO-CPA secure. Again, we simply modify an assumed fully IND-SO-CPA secure scheme to make a SIM-SO-CPA attack possible. Concretely, we add a statistically hiding commitment $\text{Com}(M)$ to each ciphertext, where M is the encrypted message. (In fact, we will require non-interactive statistically hiding commitments without any kind of setup, which can be built from collision-resistant hash functions. See Section 4 for details.) This makes the encryption scheme binding (i.e., a public key and a ciphertext form a binding commitment to the message). Hence, applying a general result due to Bellare et al. [2] shows that the scheme is not SIM-SO-CPA secure.

To show that the modified scheme is still fully IND-SO-CPA secure, we show that any IND-SO-CPA adversary A' on the modified scheme can be mapped to an IND-SO-CPA adversary A on the old scheme. The problem for A is that it must present (an internal simulation of) A' with ciphertexts with added commitments $\text{Com}(M_i)$, and later open some of those commitments to the right M_i . In this, A must not know any of the M_i in advance. Our solution to this commitment problem is to embed the $\text{Com}(M_i)$ into A' 's message distribution. (That is, if A' 's message distribution over the M_i is dist' , then A 's message distribution is dist , which is the same as dist' , only with added commitments to the M_i .) Hence, A can go ahead and open all $\text{Com}(M_i)$ -encryptions (and only those) in advance to be able to prepare authentic commitments for A' . The remaining translation between A' 's and A 's IND-SO-CPA experiment is then straightforward.

The technical difficulty in pushing this line of proof through is that by initially opening commitments $\text{Com}(M_i)$ to *all* messages, A may slightly skew a later re-sampling of the M_i . If the used commitment scheme is *perfectly* hiding, this is a non-issue: then, $\text{Com}(M_i)$ reveals no information about M_i , and conditioning on $\text{Com}(M_i)$ does not change the distribution of M_i . However, the most interesting candidates for non-interactive statistically hiding commitment schemes are only statistically, but not perfectly hiding. We thus need to show that conditioning on a statistically hiding commitment does not significantly change a message distribution. This in fact turns out to be surprisingly nontrivial. Specifically, the statement only holds for *bit* messages M_i , but not necessarily for messages, say, from $\{0, 1\}^k$. See Section 4 for details.

Outline. We start by recalling some notation and definitions (including the definitions of selective opening security) in Section 2. We present our counterexamples in Section 3 and Section 4. In Appendix A, we prove a technical lemma that is necessary for our second counterexample.

2 Preliminaries

Notation. For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. Throughout the paper, $k \in \mathbb{N}$ denotes the security parameter. For a finite set \mathcal{S} , we denote by $s \leftarrow \mathcal{S}$ the process of sampling s uniformly from \mathcal{S} . For a distribution X , we denote by $x \leftarrow X$ the process of sampling x from X . For a probabilistic algorithm A , we denote with $y \leftarrow A(x; R)$ the process of running A on input x and with randomness R , and assigning y the result. We let \mathcal{R}_A denote the randomness space of A ; we require \mathcal{R}_A to be of the form $\mathcal{R}_A = \{0, 1\}^r$. We write $y \leftarrow A(x)$ for $y \leftarrow A(x; R)$ with uniformly chosen $R \in \mathcal{R}_A$. If A 's running time is polynomial in k , then A is called probabilistic polynomial-time (PPT). Two sequences of random variables $X = (X_k)_{k \in \mathbb{N}}$ and $Y = (Y_k)_{k \in \mathbb{N}}$ are

computationally indistinguishable (denoted $X \stackrel{c}{\approx} Y$) iff for any PPT algorithm D , the probability $\Pr [D(1^k, X_k) = 1] - \Pr [D(1^k, Y_k) = 1]$ is negligible in k . The statistical distance of X_k and Y_k is defined as $\text{SD}(X_k; Y_k) := \frac{1}{2} \sum_s |\Pr [X_k = s] - \Pr [Y_k = s]|$.

DDH assumption. The *decisional Diffie-Hellman (DDH) assumption* over a group \mathbb{G} (that may depend on the security parameter k) stipulates that

$$(g, g^a, g^b, g^{ab}) \stackrel{c}{\approx} (g, g^a, g^b, g^c),$$

where $g \leftarrow \mathbb{G}$ and $a, b, c \leftarrow [|\mathbb{G}|]$ are uniformly distributed.

PKE schemes. A public-key encryption (PKE) scheme consists of three PPT algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$. Key generation $\text{Gen}(1^k)$ outputs a public key pk and a secret key sk . Encryption $\text{Enc}(pk, M)$ takes a public key pk and a message M , and outputs a ciphertext C . Decryption $\text{Dec}(sk, C)$ takes a secret key sk and a ciphertext C , and outputs a message M . For correctness, we want $\text{Dec}(sk, C) = M$ for all M , all $(pk, sk) \leftarrow \text{Gen}(1^k)$, and all $C \leftarrow \text{Enc}(pk, M)$.

Definition of selective opening security. We present and discuss three definitions for security under selective openings that capture security of an encryption scheme under adaptive attacks. Two definitions are indistinguishability-based, following the IND-SO-COM, resp. IND-SO-ENC definitions by Bellare et al. [1]. These definitions demand that even an adversary that gets to see a vector of ciphertexts cannot distinguish the true contents of the ciphertexts from independently sampled plaintexts. While one of these definitions, called weak IND-SO-CPA here, only considers *efficiently re-samplable* message distributions, the other one, dubbed full IND-SO-CPA, does not restrict the considered message distributions in this way. The third definition, dubbed SIM-SO-CPA by us, resembles the SEM-SO-COM, resp. SEM-SO-ENC definitions from [1] (which in turn follow Dwork et al. [7]). This definition is simulation-based and does not have to cope with different strategies to handle re-sampling.

Definition 2.1 (Efficiently re-samplable). *Let $n = n(k) > 0$, and let dist be a joint distribution over $(\{0, 1\}^k)^n$. We say that dist is efficiently re-samplable if there is a PPT algorithm $\text{ReSamp}_{\text{dist}}$ such that for any $\mathcal{I} \subseteq [n]$ and any partial vector $\mathbf{M}'_{\mathcal{I}} := (M^{(i)})_{i \in \mathcal{I}} \in (\{0, 1\}^k)^{|\mathcal{I}|}$, $\text{ReSamp}_{\text{dist}}(\mathbf{M}'_{\mathcal{I}})$ samples from $\text{dist} \mid \mathbf{M}_{\mathcal{I}}$, i.e., from the distribution dist , conditioned on $M^{(i)} = M'^{(i)}$ for all $i \in \mathcal{I}$.*

Opening oracles. In our definitions of selective opening security we provide the adversary with an *opening oracle* to allow adaptive queries. Such an oracle is a stateful functionality that takes one argument. When queried with a set of indexes, it responds with the corresponding openings. When queried with the string `get queries`, it returns the set of all indexes it has provided openings for since its instantiation.

Definition 2.2 (Weak indistinguishability-based selective opening security). *For a PKE scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$, a polynomially bounded function $n = n(k) > 0$, an opening oracle \mathcal{O} and a stateful PPT adversary A , consider the following experiment:*

Experiment $\text{Exp}_{\text{PKE}, A}^{\text{weak-ind-so}}$

$b \leftarrow \{0, 1\}$

$(pk, sk) \leftarrow \text{Gen}(1^k)$

$(\text{dist}, \text{ReSamp}_{\text{dist}}) \leftarrow A(pk)$

$\mathbf{M}_0 := (M^{(i)})_{i \in [n]} \leftarrow \text{dist}$

$\mathbf{R} := (R^{(i)})_{i \in [n]} \leftarrow (\mathcal{R}_{\text{Enc}})^n$

$\mathbf{C} := (C^{(i)})_{i \in [n]} := (\text{Enc}(pk, M^{(i)}; R^{(i)}))_{i \in [n]}$

$\mathbf{O} := (M^{(i)}, R^{(i)})_{i \in [n]}$

$A^{\mathcal{O}(\cdot)}(\text{select}, \mathbf{C})$

$\mathcal{I} := \mathcal{O}(\text{get queries})$

$\mathbf{M}_1 \leftarrow \text{dist} \mid \mathbf{M}_{\mathcal{I}}$

$out_A \leftarrow A(\text{output}, \mathbf{M}_b)$

return 1 if $out_A = b$, 0 otherwise

We only allow A that always output efficiently re-samplable distributions dist over $(\{0, 1\}^k)^n$ with corresponding efficient re-sampling algorithms $\text{ReSamp}_{\text{dist}}$. We say that PKE is weakly IND-SO-CPA secure, if

$$\text{Adv}_{\text{PKE}, A}^{\text{w-ind-so}}(k) := \Pr \left[\text{Exp}_{\text{PKE}, A}^{\text{weak-ind-so}}(k) = 1 \right] - \frac{1}{2}$$

is negligible.

There are some minor technical differences between Definition 2.2 and the IND-SO-ENC definition from [1]: IND-SO-ENC security universally quantifies over all (efficiently re-samplable) message distributions dist . We let A choose dist instead, e.g., to allow a message distribution that depends on the public key pk . (In fact, otherwise it is not even clear that the resulting definition implies IND-CPA security.) Besides, our definition grants the adversary multiple, possibly adaptive openings, whereas IND-SO-ENC security only allows for a one-shot opening phase. We believe that multiple openings are more realistic in view of a scenario with adaptive party corruptions.

Definition 2.3 (Full indistinguishability-based selective opening security). *For a PKE scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$, a polynomially bounded function $n = n(k) > 0$, a stateful opening oracle \mathcal{O} and a stateful PPT adversary A , we define the experiment $\text{Exp}_{\text{PKE}, A}^{\text{full-ind-so}}$ analogously to $\text{Exp}_{\text{PKE}, A}^{\text{weak-ind-so}}$ but do not require the adversary to provide an algorithm for re-sampling, i.e., $A(pk)$ just outputs a message distribution dist . We say that PKE is fully IND-SO-CPA secure if*

$$\text{Adv}_{\text{PKE}, A}^{\text{s-ind-so}}(k) := \Pr \left[\text{Exp}_{\text{PKE}, A}^{\text{full-ind-so}}(k) = 1 \right] - \frac{1}{2}.$$

is negligible.

Definition 2.3 resembles the IND-SO-COM definition from [1], only for encryption instead of commitments, and with the same syntactic differences as above. (We note that [1] only consider efficiently re-samplable message spaces in their results about encryption schemes. In their results about selective opening secure commitments, the involved message spaces are arbitrary, as in Definition 2.3.)

Definition 2.4 (simulation-based selective opening security). *For a PKE scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$, a polynomially bounded function $n = n(k) > 0$, and a stateful PPT adversary A , consider the following experiments:*

<p>Experiment $\text{Exp}_{\text{PKE}, A}^{\text{sim-so-real}}$ $(pk, sk) \leftarrow \text{Gen}(1^k)$ $\text{dist} \leftarrow A(pk)$ $\mathbf{M} := (M^{(i)})_{i \in [n]} \leftarrow \text{dist}$ $\mathbf{R} := (R^{(i)})_{i \in [n]} \leftarrow (\mathcal{R}_{\text{Enc}})^n$ $\mathbf{C} := (C^{(i)})_{i \in [n]} := (\text{Enc}(pk, M^{(i)}; R^{(i)}))_{i \in [n]}$ $\mathbf{O} := (M^{(i)}, R^{(i)})_{i \in [n]}$ $\text{out}_A \leftarrow A^{\mathcal{O}(\cdot)}(\text{select}, \mathbf{C})$ $\mathcal{I} := \mathcal{O}(\text{get queries})$ return $(\mathbf{M}, \text{dist}, \mathcal{I}, \text{out}_A)$</p>	<p>Experiment $\text{Exp}_S^{\text{sim-so-ideal}}$ $\text{dist} \leftarrow S()$ $\mathbf{M} := (M^{(i)})_{i \in [n]} \leftarrow \text{dist}$ $\text{out}_S \leftarrow S^{\mathcal{O}(\cdot)}(\text{select})$ $\mathcal{I} := \mathcal{O}(\text{get queries})$ return $(\mathbf{M}, \text{dist}, \mathcal{I}, \text{out}_S)$</p>
---	---

We say that the scheme is SIM-SO-CPA secure iff for every adversary A there is a PPT algorithm, the simulator, S such that the distributions induced by the experiments $\text{Exp}_{\text{PKE}, A}^{\text{sim-so-real}}$ and $\text{Exp}_S^{\text{sim-so-ideal}}$ are computationally indistinguishable.

Apart from the differences mentioned above, Definition 2.4 is identical to the SEM-SO-ENC definition from [1].

$\text{Gen}'(1^k)$ $(pk, sk) \leftarrow \text{Gen}(1^k)$ $g \leftarrow \mathbb{G}, h \leftarrow \mathbb{G}$ $\omega \leftarrow [[\mathbb{G}]]$ $u := g^\omega, v := h^\omega$ return $((pk, g, h, u, v), sk)$	$\text{Enc}'(pk', M)$ $((pk, g, h, u, v) := pk')$ $s \leftarrow [[\mathbb{G}]], t \leftarrow [[\mathbb{G}]]$ $C_1 \leftarrow \text{Enc}(pk, M)$ $C_2 := ((g^s u^t)^M, h^s v^t)$ return (C_1, C_2)	$\text{Dec}'(sk, C)$ $(C_1, C_2) := C$ $M := \text{Dec}(sk, C_1)$ return M
--	--	---

Figure 2: PKE' , a scheme which is SIM-SO-CPA but not fully IND-SO-CPA secure

3 SIM-SO-CPA security does not imply full IND-SO-CPA security

We prove by counterexample that there are SIM-SO-CPA secure PKE schemes that are not fully IND-SO-CPA secure. Let $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a PKE scheme with message space $\{0, 1\}^k$ that is SIM-SO-CPA secure³. From PKE we construct a scheme $\text{PKE}' = (\text{Gen}', \text{Enc}', \text{Dec}')$ (see Figure 2) that is still SIM-SO-CPA secure, which is what we prove first, but not fully IND-SO-CPA secure.

For the construction of PKE' (see Figure 2) we use a cyclic DDH group \mathbb{G} of prime order. We assume that the underlying SIM-SO-CPA secure scheme PKE can encrypt elements of \mathbb{G} and \mathbb{G} -exponents.⁴ The idea of our modification is to extend the ciphertext by a “dual-mode” commitment (in the spirit of [5]). If the public key is generated honestly, the commitment is perfectly binding. However, in the course of the proof of Lemma 3.1, we will swap the public key. Thereby we switch to the alternative mode where the commitment is equivocable with the help of a trapdoor. Finally, in the proof of Lemma 3.2, we can use the commitment to show that PKE' is not fully IND-SO-CPA secure.

For a ciphertext $C \leftarrow \text{Enc}'(pk, M)$ under PKE' we write $(M, (r, s, t))$ for the corresponding opening. (r, s, t) resembles the randomness used to generate c : r is the randomness used by Enc and s and t are the coins for the commitment (see Figure 2).

3.1 PKE' is SIM-SO-CPA secure

Lemma 3.1. *PKE' is SIM-SO-CPA secure.*

Proof. Let A' be an adversary for PKE' . Our goal is to construct a simulator S such that $\text{Exp}_{\text{PKE}', A'}^{\text{sim-so-real}}$ and $\text{Exp}_S^{\text{sim-so-ideal}}$ are computationally indistinguishable. Towards this goal we first construct an adversary A that uses A' to attack PKE . Then we show the indistinguishability of $\text{Exp}_{\text{PKE}', A'}^{\text{sim-so-real}}$ and $\text{Exp}_{\text{PKE}, A}^{\text{sim-so-real}}$ and finally use the SIM-SO-CPA security of PKE to obtain S .

The SIM-SO-CPA-real experiment calls A twice, once to obtain the message distribution dist , and once to obtain the output of the adversary after the opening phase. Based on these calls we define A as follows:

Message distribution. A uniformly picks g, h from \mathbb{G} and $\omega_u \neq \omega_v$ from $[[\mathbb{G}]]$. It then computes $u := g^{\omega_u}, v := h^{\omega_v}$ and returns $A'((pk, g, h, u, v))$.

Opening queries. A uniformly picks vectors $\mathbf{S}, \mathbf{T} \leftarrow [[\mathbb{G}]]^n$ of values and computes $C_1^{(i)} := C^{(i)}, C_2^{(i)} := (u^{\mathbf{S}^{(i)}}, v^{\mathbf{T}^{(i)}})$ and $\mathbf{C}' := (C_1^{(i)}, C_2^{(i)})_{i \in [[\mathbf{C}]]}$. Next A constructs an opening oracle \mathcal{O}' that works as follows: If called with an index i , it fetches the corresponding opening $(M, R) := \mathcal{O}(i)$

³Such schemes exist under reasonable assumptions, see [1, 8] for example.

⁴Specifically, in the term $(g^s u^t)^M$ used in Enc' , the message M can be a group element. We thus implicitly assume a suitable encoding of group elements as (nonzero) \mathbb{G} -exponents; depending on \mathbb{G} , this may additionally require application of a collision-resistant hash function H , so that the term becomes $(g^s u^t)^{H(M)}$. We stress that our results do not depend on the used encoding or hash function.

from \mathcal{O} and computes

$$s := \omega_u \omega_v (\mathbf{S}^{(i)} - \mathbf{T}^{(i)} M) / (\omega_v M - \omega_u M)$$

and

$$t := \mathbf{T}^{(i)} - s / \omega_v$$

which yield the opening $(M, (R, s, t))$ for $C'^{(i)}$. Note that we have $(g^s u^t)^M = u^{\mathbf{S}^{(i)}}$ and $h^s v^t = v^{\mathbf{T}^{(i)}}$. A returns $A'^{\mathcal{O}'(\cdot)}(\text{select}, \mathbf{C}')$.

We now provide a sequence of games that shows the computational indistinguishability of $\text{Exp}_{\text{PKE}', A'}^{\text{sim-so-real}}$ and $\text{Exp}_{\text{PKE}, A}^{\text{sim-so-real}}$. **Game 1** is simply the real SIM-SO-CPA experiment with A' and PKE' . In **Game 2** the experiment runs with a modified public key: Let $pk' = (pk, g, h, u, v)$ denote the public key generated by Gen' . The experiment in Game 2 uniformly picks $\omega_u \neq \omega_v$ from $[\#\mathbb{G}]$ and sends the tuple $(pk, g, h, g^{\omega_u}, h^{\omega_v})$ instead of pk' to A' . Every efficient algorithm that could distinguish the distribution generated by Game 1 from that generated by Game 2 with non-negligible probability would win the DDH-experiment with non-negligible probability. In **Game 3** we remove the information about the encrypted message from the commitment part of the ciphertext. For each ciphertext $C = (\text{Enc}(pk, M), ((g^s u^t)^M, h^s v^t))$ in \mathbf{C} the experiment picks s and t uniformly from $[\#\mathbb{G}]$ and replaces C_2 by (u^s, v^t) . If A' wishes to open the ciphertext, the experiment computes an opening as described in the definition of A above using the knowledge of ω_u and ω_v . The distributions of Game 2 and Game 3 are identical: The commitment part of the ciphertext consists of $((g^s u^t)^M, h^s v^t)$ for uniform s and t . Since $\omega_u = \log_g(u) \neq \log_h(v) = \omega_v$, its distribution is identical to⁵ $(g^a M, g^b)$ for uniformly random a and b and hence obviously identical to (u^s, v^t) for random s, t . Similarly we can see that the random values in the openings are still distributed uniformly as well.

The situation in Game 3 is identical to running the SIM-SO-CPA-real experiment with A and PKE . Since A is SIM-SO-CPA secure there is a simulator S such that $\text{Exp}_{\text{PKE}, A}^{\text{sim-so-real}} \stackrel{c}{\approx} \text{Exp}_S^{\text{sim-so-ideal}}$. Altogether we find $\text{Exp}_{\text{PKE}', A'}^{\text{sim-so-real}} \stackrel{c}{\approx} \text{Exp}_{\text{PKE}, A}^{\text{sim-so-real}} \stackrel{c}{\approx} \text{Exp}_S^{\text{sim-so-ideal}}$. Hence S simulates A' which concludes our proof. \square

3.2 PKE' is not fully IND-SO-CPA secure

Lemma 3.2. *PKE' is not fully IND-SO-CPA secure.*

Proof. We construct an adversary A that wins the full IND-SO-CPA experiment with non-negligible probability. Basically, A benefits from the fact that the experiment conditions the distribution of messages dist on the choice of openings \mathcal{I} to sample \mathbf{M}_1 even if this re-sampling could not be done efficiently by A . In the course of this proof we will see that A can therefore utilize the experiment to compute a discrete logarithm which helps A to learn the experiment's choice b .

We now describe the adversary A .

Message distribution. When A receives the public key $pk' = (pk, g, h, u, v)$ it responds with a distribution of tuples $(x, z) \in \mathbb{Z}_{|\mathbb{G}|} \times \mathbb{G}$ determined by the following algorithm:

Distribution dist

$b \leftarrow \{0, 1\}$

$x \leftarrow [\#\mathbb{G}]$

if $b = 0$ then return (x, g^x) otherwise return (x, h^x)

Intuitively, this algorithm draws a random element z from \mathbb{G} and returns either $(\log_g z, z)$ or $(\log_h z, z)$.

Challenge ciphertexts. A receives $\mathbf{C} \leftarrow (\text{Enc}'(pk', x), \text{Enc}'(pk', z))$ for some x and $z = g^x$ or $z = h^x$. Let $(\text{Enc}(pk, x), ((g^s u^t)^x, h^s v^t)) = C^{(1)}$.

⁵Recall that we have assumed an encoding of M that does not map to 0.

Opening queries. A calls $\mathcal{O}(2)$ to open the second component of \mathbf{C} . The return value of this call is of no interest for A here. However, it is important that the value of z is fixed for the re-sampling of \mathbf{M}_1 .

Challenge messages. Finally, A receives a message vector $\mathbf{M}_b = (y, z)$ from the experiment. If

$$(h^s v^t)^y = (g^s u^t)^x \quad (2)$$

then A returns 1 and 0 otherwise.

Analysis. **Game 1** is the full IND-SO-CPA experiment $\text{Exp}_{\text{PKE},A}^{\text{full-ind-so}}$. In **Game 2** the experiment calls $\text{Gen}(1^k)$ to generate the public key $(pk, g, h, u, v) = pk' \leftarrow \text{Gen}(1^k)$ until $g \neq h$ and $gh \neq 1$ before sending pk' to A . The statistical distance of the two distributions of public keys is $\frac{2}{|\mathbb{G}|}$ and hence negligible.

We now analyze the advantage of A in Game 2. By opening the second component of the ciphertext vector A fixes its value, i.e. $z := \mathbf{M}_0^{(2)} = \mathbf{M}_1^{(2)}$. However, since the value of z does not determine whether the first component of \mathbf{M}_b contains the logarithm to base g or to base h , this is decided only when \mathbf{M}_1 is sampled. An adversary A benefits from this re-sampling if $\mathbf{M}_0 = (x = \log_g(z), z)$, $\mathbf{M}_1 = (y = \log_h(z), z)$ and $b = 1$. In this case A learns y and only then⁶ we have that equation 2 holds.

We now show that the advantage of A is non-negligible. We define the three events

- B : The experiment samples $b = 1$.
- $M0$: The experiment samples $\mathbf{M}_0 = (x, g^x)$ (i.e. the first message vector contains a logarithm to base g).
- $M1$: The experiment samples $\mathbf{M}_1 = (y, h^y = z)$ for a fixed z (i.e. the second message vector contains a logarithm to base h).

Let \bar{E} denote the complementary event for an event E . We observe that A outputs 1 if $B \wedge M0 \wedge M1$ and 0 if $\overline{B \wedge M0 \wedge M1}$. Hence

$$\begin{aligned} \Pr \left[\text{Exp}_{\text{PKE},A}^{\text{full-ind-so}} = 1 \right] &= \Pr \left[\overline{B \wedge M0 \wedge M1} \right] + \Pr \left[B \wedge M0 \wedge M1 \right] \\ &\stackrel{(*)}{=} \Pr \left[\overline{B \wedge M0 \wedge M1} \right] + \Pr \left[B \wedge M0 \wedge M1 \right] \\ &= \Pr \left[\overline{B} \right] + \Pr \left[B \right] \Pr \left[M0 \right] \Pr \left[M1 \right] \\ &= \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} = \frac{5}{8}, \end{aligned}$$

where $(*)$ uses that B , $M0$ and $M1$ are independent events. Altogether, the adversary's advantage in Game 2 is

$$\text{Adv}_{\text{PKE},A}^{\text{s-ind-so}} = \Pr \left[\text{Exp}_{\text{PKE},A}^{\text{full-ind-so}} = 1 \right] - \frac{1}{2} = \frac{1}{8}$$

which is non-negligible. □

4 Full IND-SO-CPA does not imply SIM-SO-CPA

4.1 Outline

We will now construct a fully IND-SO-CPA secure PKE scheme that is *not* SIM-SO-CPA secure. To this end, we will start from a fully IND-SO-CPA secure scheme PKE.⁷ We will then add a commitment (to the encrypted message) to each PKE ciphertext, such that the resulting scheme PKE' becomes committing. The result of Bellare et al. [2] then implies that PKE' is not SIM-SO-CPA secure.

⁶Since $g \neq h$ and $gh \neq 1$.

⁷To date, there is no PKE scheme that is known to be fully IND-SO-CPA secure. However, in case no IND-SO-CPA secure scheme exists, of course no separating example can be constructed.

The heart of our argument will thus be to show that PKE' is still fully IND-SO-CPA secure. We will reduce the IND-SO-CPA security of PKE' to that of PKE . Concretely, assume an IND-SO-CPA adversary A' on PKE' . We need to construct an IND-SO-CPA adversary A on PKE . Of course, A will internally run A' and try to map PKE ciphertexts and openings to those of PKE' .

The concrete problem for A is that initially, A' expects a vector of PKE' ciphertexts, which contain commitments to each message. Because these commitments do not appear in PKE ciphertexts, A will have to make up those commitments for A' *without knowing the respective messages*. Later on, however, when A' requests openings, A will have to also open those commitments to messages not known in advance (to A). In other words, A will have to equivocate commitments for A' .

This seems like an insurmountable problem: we need PKE' to be committing, in order to derive (using [2]) that PKE' is not SIM-SO-CPA secure. However, if PKE' is committing, then how could A possibly equivocate commitments? Our solution is to abuse the (possibly inefficient) re-sampling that occurs during the IND-SO-CPA experiment. Namely, observe that statistically hiding commitments can always be equivocated *inefficiently* (at least with high probability). In fact, equivocating a commitment $com = \text{Com}(M; R)$ (with message M and randomness R) can be formulated as re-sampling from the message distribution $(M, R, \text{Com}(M; R))$, conditioned on a fixed value com for the third component. This will essentially allow our adversary A to formulate the necessary equivocations as a re-sampling of suitable message distribution.

4.2 Non-interactive statistically hiding commitments

As a technical tool for our separation, we will require the notion of suitable commitments. To allow for (inefficient) equivocation, we will require that the commitments are statistically hiding. Additionally, for the use in a PKE scheme, the commitments should be non-interactive. Finally, we stress that we do not allow any public parameters (such as a common reference string).

Definition 4.1 (NISHCOMs). *A non-interactive statistically hiding commitment scheme (NISHCOM) is a PPT algorithm Com that takes as input a message $M \in \{0, 1\}$ and outputs a commitment $com \in \{0, 1\}^*$. We require the following properties:*

Statistical hiding. *The statistical distance $\text{SD}(\text{Com}(0); \text{Com}(1))$ is negligible in k .*

Binding. *For every PPT A , the following probability is negligible (in k):*

$$\Pr \left[\text{Com}(0; R_0) = \text{Com}(1; R_1) \mid (R_0, R_1) \leftarrow A(1^k) \right].$$

While one-way functions imply statistically hiding commitments [11], we cannot expect to construct NISHCOMs even from trapdoor one-way permutations [10]. In fact, there can be no NISHCOM that is binding against *non-uniform* adversaries. (The statistical hiding property implies that for each k , there exist many tuples (R_0, R_1) with $\text{Com}(0; R_0) = \text{Com}(1; R_1)$. We can always hardcode one such tuple into a non-uniform A .) However, under specific assumptions, we can construct NISHCOMs:

NISHCOMs from CRHFs. Assume a collision-resistant hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^k$. We stress that H is not keyed but fixed. (In particular, we can only hope for collision-resistance against uniform adversaries.) Instantiated with such an H , Naor and Yung [14], and Damgård et al. [6] yield several constructions of NISHCOMs. For instance, implicit in [14] is the NISHCOM

$$\text{Com}(M; (X, h)) := (H(X), h, h(X) \oplus M)$$

for $M \in \{0, 1\}$, $X \in \{0, 1\}^\ell$ for suitably large ℓ , and a suitable randomness extractor h .

NISHCOMs from fixed groups. Let $(\mathbb{G}_k, g_k, h_k)_{k \in \mathbb{N}}$ be a family of finite groups, one for each value of the security parameter k , along with (fixed) generators g_k, h_k of \mathbb{G}_k . If we assume that the problem of computing $\text{dlog}_{g_k}(h_k)$ is computationally infeasible, then Pedersen's commitment [15] (i.e., $\text{Com}(M; R) := g_k^M h_k^R$) is a NISHCOM that is even perfectly hiding.

$\text{Gen}'(1^k)$ $(pk, sk) \leftarrow \text{Gen}(1^k)$ return (pk, sk)	$\text{Enc}'(pk', M)$ $C \leftarrow \text{Enc}(pk, M)$ $com \leftarrow \text{Com}(M)$ return $C' := (C, com)$	$\text{Dec}'(sk', C')$ $(C, com) := C'$ $M \leftarrow \text{Dec}(sk', C)$ return M
--	--	---

Figure 3: PKE' — a fully IND-SO-CPA, but not SIM-SO-CPA secure PKE scheme

4.3 The separating scheme

We are now ready to describe our scheme. We assume a fully IND-SO-CPA secure scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ with message space $\{0, 1\}$, as well as a NISHCOM Com . In our scheme, depicted in Figure 3, we simply append to each ciphertext a commitment to the encrypted message. This commitment is never checked or opened during execution of the scheme; it only serves as a means to make the scheme committing in the sense of Bellare et al. [2].

4.4 SIM-SO-CPA insecurity of the scheme

First, we note that because of our use of Com , scheme PKE' is a binding CE (“committing encryption”) scheme in the sense of Bellare et al. [2]. Concretely, opening a ciphertext (by releasing the encryption randomness) as an honest encryption in two different ways (i.e., for two different messages) requires breaking the binding property of Com . Hence, we can apply [2, Theorem 4.1]⁸, and we get:

Theorem 4.2. PKE' as depicted in Figure 3 is not SIM-SO-CPA secure.

4.5 Full IND-SO-CPA security of the scheme

The main part of our work is to prove that PKE' is fully IND-SO-CPA secure. As explained above, our intuition will be to use the (potentially inefficient) message re-sampling in the full IND-SO-CPA experiment to equivocate Com commitments.

Theorem 4.3. PKE' as depicted in Figure 3 is fully IND-SO-CPA secure, provided that PKE is fully IND-SO-CPA secure, and Com is a NISHCOM.

Proof. Given an IND-SO-CPA adversary A' on PKE' , we construct an IND-SO-CPA adversary on PKE with roughly the same complexity and success. Concretely, A proceeds as follows:

Message distribution. When invoked with a PKE public key pk , A sets $pk' := pk$ and runs $\text{dist}' \leftarrow A'(pk')$ to obtain an N' -message distribution dist' . Then A creates and outputs its own N -message distribution (for $N := 3N'$) dist as follows:

Distribution dist
 $(M'_i)_{i \in [N']} \leftarrow \text{dist}'$
 $(R_i^{\text{Com}})_{i \in [N']} \leftarrow (\mathcal{R}_{\text{Com}})^{N'}$
 $(com_i)_{i \in [N']} := (\text{Com}(M'_i; R_i^{\text{Com}}))_{i \in [N']}$
return $(M'_1, R_1^{\text{Com}}, com_1, \dots, M'_{N'}, R_{N'}^{\text{Com}}, com_{N'})$

Challenge ciphertexts. When receiving an N -ciphertext vector $(C_i)_{i \in [N]}$, A prepares an N' -ciphertext vector $(C'_i)_{i \in [N']}$ for A' as follows. First, A asks its own IND-SO-CPA experiment for openings of C_3, C_6, \dots, C_N to obtain the commitments com_i (for $i \in [N']$). It then sets $C'_i := (C_i, com_i)$ for all i and hands $(C'_i)_{i \in [N']}$ to D . Note that this results in a challenge

⁸Note that there is an important difference between our SIM-SO-CPA definition and the one from [2]: In [2] the simulator and the adversary are allowed a common *auxiliary input* which is of importance for Theorem 4.1. However, it is easy to verify that all of our proofs concerning SIM-SO-CPA security are still valid in presence of an auxiliary input, which we omitted for the sake of simplicity.

ciphertext for D that is perfectly distributed as in D 's own IND-SO-CPA experiment. Furthermore, because Com is statistically hiding, opening the encrypted commitments does not fix any of the encrypted messages.

Opening queries. When A' wants a ciphertext C'_i opened, A asks for an opening of C_{3i-2} and C_{3i-1} . The opening of C_{3i-2} yields a properly distributed opening of the PKE part C_i of $C'_i = (C_i, \text{com}_i)$. On the other hand, the opening of C_{3i-1} reveals the randomness R_i^{Com} of the corresponding commitment com_i . Together, this forms a perfectly distributed opening of C'_i , which A then hands to A' .

Challenge messages. Finally, when A' is finished asking for openings and requests challenge messages, A does the same and hands the corresponding M'_i (for $i \in [N']$) to A' . When A' outputs a decision bit b' , then A outputs the same bit.

To analyze this A , first note that up to the challenge message, A provides a perfect internal simulation of A' running in its own IND-SO-COM experiment with PKE'. In particular, both challenge ciphertexts and openings are exactly distributed as with PKE'. For the eventual challenge message (and A' 's decision bit), we make the following case distinction:

When A 's experiment tosses $b = 0$ (i.e., no re-sampling). In this case, A eventually obtains the initially sampled plaintext vector with all $M'_i, R_i^{\text{Com}}, \text{com}_i$. In particular, A' gets the messages M'_i just as it would have in its own IND-SO-CPA experiment with PKE'. We get:

$$\Pr \left[\text{Exp}_{\text{PKE}, A}^{\text{full-ind-so}}(k) = 1 \mid b = 0 \right] = \Pr \left[\text{Exp}_{\text{PKE}', A'}^{\text{full-ind-so}}(k) = 1 \mid b = 0 \right]. \quad (3)$$

When A 's experiment tosses $b = 1$ (i.e., re-sampling occurs). In this case, A eventually obtains a plaintext vector that has been re-sampled from dist , conditioned on all opened messages M'_i (along with the corresponding R_i^{Com}), and all commitments com_i . In particular, A' gets a re-sampled message vector that is additionally conditioned on all com_i . This marks a difference to what A' would have gotten in its IND-SO-CPA experiment with PKE': there, A' would have gotten M'_i that are only conditioned on the so far opened messages, but not on all com_i . However, recall that Com is statistically hiding, and thus the distribution of the com_i is statistically close to, say, commitments to all-zero strings. Thus, we will now prove that

$$\Pr \left[\text{Exp}_{\text{PKE}, A}^{\text{full-ind-so}}(k) = 1 \mid b = 1 \right] - \Pr \left[\text{Exp}_{\text{PKE}', A'}^{\text{full-ind-so}}(k) = 1 \mid b = 1 \right]. \quad (4)$$

is negligible in k , using a sequence of Games.

Game 1 is simply the IND-SO-CPA experiment with A and PKE as described above, but with b fixed to 1.

In **Game 2**, we substitute all $\text{com}_i \leftarrow \text{Com}(M'_i)$ by $\text{com}_i \leftarrow \text{Com}(0)$. We stress that during the resampling operation, we still condition on the com_i being output as M_i -commitments. Note that this conditioning operation may fail, e.g., when some M_i has been opened as $M_i = 1$, but com_i lies not in the range of $\text{Com}(1)$. However, this can happen only with negligible probability by the hiding property of Com . Namely, note that for each sampled message vector $(M'_i)_{i \in [N']}$, we can view the whole experiment (including A' 's output) as a probabilistic function of the commitments com_i . If any commitment randomness R_i^{Com} is to be revealed, this randomness can be — inefficiently — generated from com_i and the corresponding M_i . Since Com is statistically hiding, we know that hence, A' 's output does not significantly change compared to Game 1.

In **Game 3**, we no longer condition on the com_i during re-sampling. (Of course, we still condition on the so far opened M'_i .) Lemma A.1 in Appendix A shows that this has no significant effect on the experiment's output. Concretely, note that we can view both Game 2 and Game 3 (including A) as an unbounded algorithm that

- gets a vector $(\text{com}_i)_{i \in [n]}$ of 0-commitments as input,

- then deterministically⁹ selects a message distribution $\widetilde{\text{dist}}$ over $\{0, 1\}^n$ (that internally corresponds to dist' , conditioned on all opened messages),
- and finally gets a sample from either $\widetilde{\text{dist}}$, or $\widetilde{\text{dist}}$ conditioned on all commitments com_i . With a dist -sample, this results in Game 3, whereas with a sample from $\widetilde{\text{dist}} \mid (\text{com}_i)_i$, this yields an execution of Game 2.

Applying Lemma A.1 yields that the output in Game 3 does not significantly differ from that in Game 2. (Somewhat surprisingly, the same statement would not hold if the M_i were not bits but, say, k -bitstrings. See Appendix A for details.) At first glance, it might seem like we only need a non-adaptive version of Lemma A.1, in which the adversary chooses the distribution ahead of time. However, such a non-adaptive Lemma would not be sufficient in our case, because the distribution $\widetilde{\text{dist}}$ depends on the adversary's opening requests and thus may depend on the commitments com_i .

Finally, in **Game 4**, we replace all $\text{com}_i \leftarrow \text{Com}(0)$ again by $\text{com}_i \leftarrow \text{Com}(M'_i)$. Like in Game 2, this has no significant effect on the output of the experiment.

Now note that in Game 4, re-sampled message vectors (M'_i) are no longer conditioned on the com_i , and are hence distributed exactly as in $\text{Exp}_{\text{PKE}', A'}^{\text{full-ind-so}}$ with $b = 1$. Also, commitments and openings are distributed exactly as with PKE' . We obtain (4).

Taking (3,4) together, we get that

$$\text{Adv}_{\text{PKE}, A}^{\text{s-ind-so}}(k) - \text{Adv}_{\text{PKE}', A'}^{\text{s-ind-so}}(k)$$

is negligible, which proves the theorem. \square

References

- [1] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In Antoine Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 1–35. Springer, April 2009.
- [2] Mihir Bellare, Rafael Dowsley, Brent Waters, and Scott Yilek. Standard security does not imply security against selective-opening. Cryptology ePrint Archive, Report 2011/581, 2011. Appears at EUROCRYPT 2012.
- [3] Ran Canetti, Uriel Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *28th Annual ACM Symposium on Theory of Computing*, pages 639–648. ACM Press, May 1996.
- [4] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In Burton S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO'97*, volume 1294 of *Lecture Notes in Computer Science*, pages 90–104. Springer, August 1997.
- [5] Ivan Damgård and Jesper Buus Nielsen. Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 581–596. Springer, August 2002.
- [6] Ivan Damgård, Torben P. Pedersen, and Birgit Pfizmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. *Journal of Cryptology*, 10(3):163–194, 1997.
- [7] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry J. Stockmeyer. Magic functions. In *40th Annual Symposium on Foundations of Computer Science*, pages 523–534. IEEE Computer Society Press, October 1999.
- [8] Serge Fehr, Dennis Hofheinz, Eike Kiltz, and Hoeteck Wee. Encryption schemes secure against chosen-ciphertext selective opening attacks. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 381–402. Springer, May 2010.
- [9] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.

⁹At this point, we can assume without loss of generality that the experiment, including A' , is unbounded, and can thus choose its own random coins deterministically.

- [10] Iftach Haitner, Jonathan J. Hoch, Omer Reingold, and Gil Segev. Finding collisions in interactive protocols - a tight lower bound on the round complexity of statistically-hiding commitments. In *48th Annual Symposium on Foundations of Computer Science*, pages 669–679. IEEE Computer Society Press, October 2007.
- [11] Iftach Haitner, Minh-Huyen Nguyen, Shien Jin Ong, Omer Reingold, and Salil P. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM J. Comput.*, 39(3):1153–1218, 2009.
- [12] Brett Hemenway, Benoit Libert, Rafail Ostrovsky, and Damien Vergnaud. Lossy encryption: Constructions from general assumptions and efficient selective opening chosen ciphertext security. In *ASIACRYPT*, Lecture Notes in Computer Science. Springer, 2011.
- [13] Dennis Hofheinz. All-but-many lossy trapdoor functions. Cryptology ePrint Archive, Report 2011/230, 2011. Appears at EUROCRYPT 2012.
- [14] Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In *21st Annual ACM Symposium on Theory of Computing*, pages 33–43. ACM Press, May 1989.
- [15] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO’91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, August 1992.
- [16] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 554–571. Springer, August 2008.

A A technical lemma

In this section, we show indistinguishability of Game 2 and Game 3. For the sake of completeness, we first briefly restate what we have to show. We are given an unbounded algorithm A that

- gets a vector $(com_i)_{i \in [n]}$ of 0-commitments as input,
- then deterministically selects a message distribution $\widetilde{\text{dist}}$ over $\{0, 1\}^n$,
- and finally gets a sample from
 - either dist ,
 - or dist conditioned on all commitments com_i .

The two possible cases in the final step will result in two different distributions of A 's view, and we have to show that the difference is negligible.

Why conditioning on string commitments leads to problems. However, before we give our formal proof, we want to briefly point out why it is crucial that the commitments have small message space. We illustrate this by a counterexample, i.e., we describe a statistically hiding string-commitment scheme and an algorithm A such that the two sample distributions will significantly differ from each other. Consider some commitment functionality, such that on input $m \in \{0, 1\}^k$ the corresponding commitment c is uniformly random over $\{0, 1\}^k \setminus \{m\}$. Obviously, such a commitment scheme is statistically hiding. However, the algorithm A can choose the distribution $\widetilde{\text{dist}}$ such that the final sample either completely consists of all-zero messages, or it completely equals the initially given commitment vector, each with probability $\frac{1}{2}$. Hence, dist conditioned on the initially given commitment vector will always solely return all-zero messages, whereas without this condition the all-zero vector has only probability $\frac{1}{2}$. Thus, the two sample distributions will have statistical distance $\frac{1}{2}$.

Used notation. Having seen this subtle issue, we give now our formal indistinguishability proof. For a concise presentation, in the following lemma we represent

- the distributions of 0-commitments and 1-commitments by two probability mass functions γ_0 and γ_1 respectively,
- the initially given commitment vector $(com_i)_{i \in [n]}$ by a random variable $\mathbf{C} = (C_i)_{i \in [n]}$, i.e., $\Pr[C_i = c] = \gamma_0(c)$,
- by a family of probability mass functions $\beta_{\mathbf{c}}$ we represent how the message distribution $\widetilde{\text{dist}}$ is generated from the initially given commitment vector,

• and two random variables \mathbf{M} and \mathbf{M}' represent the two possible sample distributions. Moreover, in the lemma we implicitly assume that the distribution $\widetilde{\text{dist}}$ conditioned on the initially given commitments $(\text{com}_i)_{i \in [n]}$ is well defined in the sense that it assigns a non-zero probability to some message vector for which $(\text{com}_i)_{i \in [n]}$ is a possible commitment vector. This corresponds to the assumption that in Theorem 4.3, opening a 0-commitment as a commitment to M'_i does not fail. In particular, this assumption may be violated with at most negligible probability by the statistical hiding property of the commitment.

Lemma A.1. *Fix the following parameters:*

- message space $\{0, 1\}$ and some countable commitment space \mathcal{C}
- a tuple $(\gamma_m)_{m \in \{0, 1\}}$, consisting of two probability mass functions over \mathcal{C}
- a dimension $n \in \mathbb{N}_{>0}$ and a family $(\beta_{\mathbf{c}})_{\mathbf{c} \in \mathcal{C}^n}$ of probability mass functions over $\{0, 1\}^n$

In this setting let some random variables $\mathbf{C} = (C_i)_{i \in [n]} \in \mathcal{C}^n$ and $\mathbf{M} = (M_i)_{i \in [n]} \in \{0, 1\}^n$ and $\mathbf{M}' = (M'_i)_{i \in [n]} \in \{0, 1\}^n$ be given, distributed as follows¹⁰:

$$\begin{aligned} \Pr[\mathbf{C} = \mathbf{c}] &= \prod_{i \in [n]} \gamma_0(c_i) \\ \Pr[\mathbf{M} = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] &= \beta_{\mathbf{c}}(\mathbf{m}) \\ \Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] &= \frac{\beta_{\mathbf{c}}(\mathbf{m}') \cdot \prod_{i \in [n]} \gamma_{m'_i}(c_i)}{\sum_{\mathbf{m} \in \{0, 1\}^n} \beta_{\mathbf{c}}(\mathbf{m}) \cdot \prod_{i \in [n]} \gamma_{m_i}(c_i)} \end{aligned}$$

Further, let $\mu := \text{SD}(\gamma_0; \gamma_1)$ in slight abuse of notation. Now, if $(1 + \sqrt{\mu})^n < 2$, it holds:

$$\text{SD}((\mathbf{C}, \mathbf{M}); (\mathbf{C}, \mathbf{M}')) \leq 2n(\sqrt{\mu} + \mu) + \frac{1}{2 - (1 + \sqrt{\mu})^n} - 1$$

In particular, if μ is negligible in some security parameter k and the dimension n grows only polynomially in k , then the statistical distance $\text{SD}((\mathbf{C}, \mathbf{M}); (\mathbf{C}, \mathbf{M}'))$ is also negligible in k .

Proof. Let $(1 + \sqrt{\mu})^n < 2$. For each $c \in \mathcal{C}$ we define the notations $\bar{\gamma}(c) := \min_{m \in \{0, 1\}} \gamma_m(c)$ and $\varepsilon_c := |\gamma_1(c) - \gamma_0(c)|$. Our proof basically consists of two parts. In the first part we bound the probability for the event that the fraction $\frac{\varepsilon_{C_i}}{\bar{\gamma}(C_i)}$, which can be seen as some “relative distance” between $\gamma_0(C_i)$ and $\gamma_1(C_i)$, is relatively large for any i . This is the undesired case, where we cannot say much about the statistical distance between \mathbf{M} and \mathbf{M}' . In the second part we estimate the statistical distance between (\mathbf{C}, \mathbf{M}) and $(\mathbf{C}, \mathbf{M}')$ conditioned to the event that the “relative distance” $\frac{\varepsilon_{C_i}}{\bar{\gamma}(C_i)}$ is very small for every i , i.e., the distributions $\gamma_0(C_i)$ and $\gamma_1(C_i)$ are almost identical. Put together, this will yield the claimed bound for the statistical distance $\text{SD}((\mathbf{C}, \mathbf{M}); (\mathbf{C}, \mathbf{M}'))$. We start by noting that for all $\nu \in \mathbb{R}_{\geq 0}$ it holds:

$$2\mu = \sum_{c \in \mathcal{C}} \varepsilon_c \geq \sum_{c \in \mathcal{C}: \varepsilon_c > \nu \cdot \gamma_0(c)} \varepsilon_c \geq \nu \cdot \sum_{c \in \mathcal{C}: \varepsilon_c > \nu \cdot \gamma_0(c)} \gamma_0(c) \quad (5)$$

Further note that $\varepsilon_c > \frac{\nu'}{1 + \nu'} \cdot \gamma_0(c)$, if only $\varepsilon_c > \nu' \cdot \bar{\gamma}(c)$, as one can see as follows. Given that $\varepsilon_c > \nu' \cdot \bar{\gamma}(c)$, we have:

$$\frac{\nu'}{1 + \nu'} \cdot \gamma_0(c) \leq \frac{\nu'}{1 + \nu'} \cdot \overbrace{\max_{m \in \{0, 1\}} \gamma_m(c)}^{=\bar{\gamma}(c) + \varepsilon_c} = \frac{\nu' \cdot \varepsilon_c + \nu' \cdot \bar{\gamma}(c)}{1 + \nu'} \stackrel{< \varepsilon_c}{<} \frac{\nu' \cdot \varepsilon_c + \varepsilon_c}{1 + \nu'} = \varepsilon_c$$

¹⁰Formally, we additionally need that $\sum_{\mathbf{m} \in \{0, 1\}^n} \beta_{\mathbf{c}}(\mathbf{m}) \cdot \prod_{i \in [n]} \gamma_{m_i}(c_i) \neq 0$ for every commitment vector \mathbf{c} with $\Pr[\mathbf{C} = \mathbf{c}] > 0$; otherwise the distribution of \mathbf{M}' would not be well-defined. However, for better readability we just omitted stating this as an explicit assumption.

Hence it follows that $\Pr[\varepsilon_{C_i} > \nu' \cdot \bar{\gamma}(C_i)] \leq \Pr\left[\varepsilon_{C_i} > \frac{\nu'}{1+\nu'} \cdot \gamma_0(C_i)\right]$ and we can estimate:

$$\Pr[\varepsilon_{C_i} > \nu' \cdot \bar{\gamma}(C_i)] \leq \Pr\left[\varepsilon_{C_i} > \frac{\nu'}{1+\nu'} \cdot \gamma_0(C_i)\right] = \sum_{c \in \mathcal{C}: \varepsilon_c > \frac{\nu'}{1+\nu'} \cdot \gamma_0(c)} \overbrace{\Pr[C_i = c]}{=\gamma_0(c)} \stackrel{(5)}{\leq} \frac{2\mu(1+\nu')}{\nu'}$$

We set $\nu' := \sqrt{\mu}$ and apply the union bound, which yields:

$$\Pr[\exists i \in [n] : \varepsilon_{C_i} > \sqrt{\mu} \cdot \bar{\gamma}(C_i)] \leq 2n(\sqrt{\mu} + \mu) \quad (6)$$

This concludes the first part of our proof. For the second part, we first rewrite the distribution of the random variable \mathbf{M}' . We just exploit that $\beta_{\mathbf{c}}$ is a probability mass function and hence $\prod_{i \in [n]} \gamma_{m'_i}(c_i) = \sum_{\mathbf{m} \in \{0,1\}^n} \beta_{\mathbf{c}}(\mathbf{m}) \cdot \prod_{i \in [n]} \gamma_{m'_i}(c_i)$ for arbitrary \mathbf{c} and \mathbf{m}' . So we can write:

$$\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] = \frac{\beta_{\mathbf{c}}(\mathbf{m}') \cdot \prod_{i \in [n]} \gamma_{m'_i}(c_i)}{\prod_{i \in [n]} \gamma_{m'_i}(c_i) + \sum_{\mathbf{m} \in \{0,1\}^n \setminus \{\mathbf{m}'\}} \beta_{\mathbf{c}}(\mathbf{m}) \cdot \left(\prod_{i \in [n]} \gamma_{m_i}(c_i) - \prod_{i \in [n]} \gamma_{m'_i}(c_i) \right)}$$

Exploiting once again that $\beta_{\mathbf{c}}$ is a probability mass function, we can estimate the big sum under the fraction line as follows:

$$\left| \sum_{\mathbf{m} \in \{0,1\}^n \setminus \{\mathbf{m}'\}} \beta_{\mathbf{c}}(\mathbf{m}) \cdot \left(\prod_{i \in [n]} \gamma_{m_i}(c_i) - \prod_{i \in [n]} \gamma_{m'_i}(c_i) \right) \right| \leq \max_{\mathbf{m} \in \{0,1\}^n} \left| \prod_{i \in [n]} \gamma_{m_i}(c_i) - \prod_{i \in [n]} \gamma_{m'_i}(c_i) \right|$$

In other words, for every $\mathbf{c} = (c_i)_{i \in [n]}$ and $\mathbf{m}' = (m'_i)_{i \in [n]}$ there exists some $d_{\mathbf{c}, \mathbf{m}'} \in \mathbb{R}$, such that it holds:

$$\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] = \frac{\beta_{\mathbf{c}}(\mathbf{m}') \cdot \prod_{i \in [n]} \gamma_{m'_i}(c_i)}{d_{\mathbf{c}, \mathbf{m}'} + \prod_{i \in [n]} \gamma_{m'_i}(c_i)}$$

and $|d_{\mathbf{c}, \mathbf{m}'}| \leq \max_{\mathbf{m} \in \{0,1\}^n} \left| \prod_{i \in [n]} \gamma_{m_i}(c_i) - \prod_{i \in [n]} \gamma_{m'_i}(c_i) \right|$

Now note that $\left| \prod_{i \in [n]} \gamma_{m_i}(c_i) - \prod_{i \in [n]} \gamma_{m'_i}(c_i) \right| \leq \prod_{i \in [n]} (\bar{\gamma}(c_i) + \varepsilon_{c_i}) - \prod_{i \in [n]} \bar{\gamma}(c_i)$ as a direct consequence of how we constructed $\bar{\gamma}(c_i)$ and ε_{c_i} . Therefore, we can estimate $|d_{\mathbf{c}, \mathbf{m}'}|$ independently of the parameter \mathbf{m}' . In particular, for every $\mathbf{c} = (c_i)_{i \in [n]}$ and $\mathbf{m}' = (m'_i)_{i \in [n]}$ there exists some $d_{\mathbf{c}, \mathbf{m}'} \in \mathbb{R}$, such that it holds:

$$\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] = \frac{\beta_{\mathbf{c}}(\mathbf{m}') \cdot \prod_{i \in [n]} \gamma_{m'_i}(c_i)}{d_{\mathbf{c}, \mathbf{m}'} + \prod_{i \in [n]} \gamma_{m'_i}(c_i)} \quad \text{and} \quad |d_{\mathbf{c}, \mathbf{m}'}| \leq \prod_{i \in [n]} (\bar{\gamma}(c_i) + \varepsilon_{c_i}) - \prod_{i \in [n]} \bar{\gamma}(c_i) \quad (7)$$

In the end we want to show that $\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}]$ is close to $\Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}]$, if only the fractions $\frac{\varepsilon_{c_1}}{\bar{\gamma}(c_1)}, \dots, \frac{\varepsilon_{c_n}}{\bar{\gamma}(c_n)}$ are sufficiently small. Now, if $d_{\mathbf{c}, \mathbf{m}'} = 0$, we would even have that $\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] = \beta_{\mathbf{c}}(\mathbf{m}') = \Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}]$. However, in general it will not be the case that $d_{\mathbf{c}, \mathbf{m}'} = 0$, but we will find a sufficiently small upper bound for $|d_{\mathbf{c}, \mathbf{m}'}|$, given that the corresponding $\varepsilon_{c_1}, \dots, \varepsilon_{c_n}$ are relatively small. Given any $\mathbf{m}' = (m'_i)_{i \in [n]}$ and $\mathbf{c} = (c_i)_{i \in [n]}$ such that $\varepsilon_{c_i} \leq \sqrt{\mu} \cdot \bar{\gamma}(c_i)$ for all i , we can estimate the bound for $|d_{\mathbf{c}, \mathbf{m}'}|$ in (7) as follows:

$$\prod_{i \in [n]} (\bar{\gamma}(c_i) + \varepsilon_{c_i}) - \prod_{i \in [n]} \bar{\gamma}(c_i) \leq ((1 + \sqrt{\mu})^n - 1) \cdot \prod_{i \in [n]} \bar{\gamma}(c_i) \quad (8)$$

For our next steps, we need to formally partition the sample space of the random variable \mathbf{C} :

$$\begin{aligned}\hat{\mathcal{C}} &:= \{\mathbf{c} = (c_i)_{i \in [n]} \in \mathcal{C}^n \mid \exists i \in [n] : \varepsilon_{c_i} > \sqrt{\mu} \cdot \bar{\gamma}(c_i)\} \\ \bar{\mathcal{C}} &:= \{\mathbf{c} = (c_i)_{i \in [n]} \in \mathcal{C}^n \mid \forall i \in [n] : \varepsilon_{c_i} \leq \sqrt{\mu} \cdot \bar{\gamma}(c_i)\}\end{aligned}$$

Taking (7,8) together, we find for every $\mathbf{c} = (c_i)_{i \in [n]} \in \bar{\mathcal{C}}$ and $\mathbf{m}' = (m'_i)_{i \in [n]} \in \{0, 1\}^n$ some $d'_{\mathbf{c}, \mathbf{m}'} \in \mathbb{R}$, such that it holds:

$$\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] = \frac{\beta_{\mathbf{c}}(\mathbf{m}')}{d'_{\mathbf{c}, \mathbf{m}'} + 1} \quad \text{and} \quad |d'_{\mathbf{c}, \mathbf{m}'}| \leq (1 + \sqrt{\mu})^n - 1$$

From this we can infer a lower and an upper bound for $\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}]$ as follows:

$$\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] \geq \frac{\beta_{\mathbf{c}}(\mathbf{m}')}{1 + |d'_{\mathbf{c}, \mathbf{m}'}|} \geq \frac{\beta_{\mathbf{c}}(\mathbf{m}')}{(1 + \sqrt{\mu})^n} \quad (9)$$

$$\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] \leq \frac{\beta_{\mathbf{c}}(\mathbf{m}')}{1 - |d'_{\mathbf{c}, \mathbf{m}'}|} \leq \frac{\beta_{\mathbf{c}}(\mathbf{m}')}{2 - (1 + \sqrt{\mu})^n} \quad (10)$$

Note once again that $\Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] = \beta_{\mathbf{c}}(\mathbf{m}')$ by definition. Thus, the distance between $\Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}]$ and $\Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}]$ is bounded by the following conditions:

$$\begin{aligned}\Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] - \Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] &\stackrel{(9)}{\leq} \beta_{\mathbf{c}}(\mathbf{m}') \cdot \left(1 - \frac{1}{(1 + \sqrt{\mu})^n}\right) \\ -\Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] + \Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] &\stackrel{(10)}{\leq} \beta_{\mathbf{c}}(\mathbf{m}') \cdot \left(\frac{1}{2 - (1 + \sqrt{\mu})^n} - 1\right)\end{aligned}$$

One can show straightforwardly that $\frac{1}{2 - (1 + \sqrt{\mu})^n} - 1 \geq 1 - \frac{1}{(1 + \sqrt{\mu})^n}$, using that $(1 + \sqrt{\mu})^n < 2$ by assumption. Thus, for all $\mathbf{c} = (c_i)_{i \in [n]} \in \bar{\mathcal{C}}$ and $\mathbf{m}' = (m'_i)_{i \in [n]} \in \{0, 1\}^n$ we can estimate:

$$\left| \Pr[\mathbf{M} = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] - \Pr[\mathbf{M}' = \mathbf{m}' \mid \mathbf{C} = \mathbf{c}] \right| \leq \beta_{\mathbf{c}}(\mathbf{m}') \cdot \left(\frac{1}{2 - (1 + \sqrt{\mu})^n} - 1 \right) \quad (11)$$

This concludes the second part of our proof and we can start putting things together. The following equation directly follows by the definitions of statistical distance and conditioned probability:

$$\text{SD}((\mathbf{C}, \mathbf{M}); (\mathbf{C}, \mathbf{M}')) = \frac{1}{2} \sum_{\mathbf{c} \in \mathcal{C}^n} \Pr[\mathbf{C} = \mathbf{c}] \cdot \sum_{\mathbf{m} \in \{0, 1\}^n} \left| \Pr[\mathbf{M} = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] - \Pr[\mathbf{M}' = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] \right| \quad (12)$$

We will split the outer sum into two parts, corresponding to $\mathbf{c} \in \hat{\mathcal{C}}$ and $\mathbf{c} \in \bar{\mathcal{C}}$ respectively. Note that for the inner sum it holds:

$$\sum_{\mathbf{m} \in \{0, 1\}^n} \left| \Pr[\mathbf{M} = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] - \Pr[\mathbf{M}' = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] \right| \leq \underbrace{\sum_{\mathbf{m} \in \{0, 1\}^n} \Pr[\mathbf{M} = \mathbf{m} \mid \mathbf{C} = \mathbf{c}]}_{\text{sums to 1}} + \underbrace{\sum_{\mathbf{m} \in \{0, 1\}^n} \Pr[\mathbf{M}' = \mathbf{m} \mid \mathbf{C} = \mathbf{c}]}_{\text{sums to 1}}$$

Thus, for each $\mathbf{c} \in \hat{\mathcal{C}}$ we can estimate the inner sum of equation (12) by 2. Furthermore, for each $\mathbf{c} \in \bar{\mathcal{C}}$ we can estimate the inner sum's elements of equation (12) just by their maximum and neglect the prefactor $\frac{1}{2}$. Altogether, we get:

$$\text{SD}((\mathbf{C}, \mathbf{M}); (\mathbf{C}, \mathbf{M}')) \leq \Pr[\mathbf{C} \in \hat{\mathcal{C}}] + \max_{\mathbf{c} \in \bar{\mathcal{C}}} \left(\sum_{\mathbf{m} \in \{0, 1\}^n} \left| \Pr[\mathbf{M} = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] - \Pr[\mathbf{M}' = \mathbf{m} \mid \mathbf{C} = \mathbf{c}] \right| \right)$$

Finally, we can estimate $\Pr[\mathbf{C} \in \hat{\mathcal{C}}]$ by (6) and all the rest by (11). It holds:

$$\text{SD}((\mathbf{C}, \mathbf{M}); (\mathbf{C}, \mathbf{M}')) \leq 2n(\sqrt{\mu} + \mu) + \left(\frac{1}{2 - (1 + \sqrt{\mu})^n} - 1 \right) \cdot \max_{\mathbf{c} \in \bar{\mathcal{C}}} \underbrace{\left(\sum_{\mathbf{m} \in \{0, 1\}^n} \beta_{\mathbf{c}}(\mathbf{m}) \right)}_{=1} \quad \square$$