Characterization of Security Notions for Probabilistic Private-Key Encryption[*]

Jonathan Katz
Department of Computer Science, University of Maryland
College Park, MD 20742
jkatz@cs.umd.edu

Moti Yung
Department of Computer Science, Columbia University
New York, NY 10027
moti@cs.columbia.edu

Abstract

The development of precise definitions of security for encryption, as well as a detailed understanding of their relationships, has been a major area of research in modern cryptography. Here, we focus on the case of private-key encryption. Extending security notions from the public-key setting, we define security in the sense of both indistinguishability and non-malleability against chosen-plaintext and chosen-ciphertext attacks, considering both non-adaptive (i.e., "lunchtime") and adaptive oracle access (*adaptive* here refers to an adversary's ability to interact with a given oracle even *after* viewing the challenge ciphertext). We then characterize the 18 resulting security notions in two ways. First, we construct a complete *hierarchy* of security notions; that is, for every pair of definitions we show whether one definition is stronger than the other, whether the definitions are equivalent, or whether they are incomparable. Second, we partition these notions of security into two classes (*computational* or *information-theoretic*) depending on whether one-way functions are necessary in order for encryption schemes satisfying the definition to exist.

Perhaps our most surprising result is that security against *adaptive* chosen-plaintext attack is (polynomially) equivalent to security against *non-adaptive* chosen-plaintext attack. On the other hand, the ability of an adversary to mount a (non-adaptive) chosen-plaintext attack is the key feature distinguishing computational and information-theoretic notions of security. These results hold for all security notions considered here.

1 Introduction

The formulation and analysis of definitions of security for encryption has been a fundamental area of modern cryptographic research. Indeed, Shannon's definition of perfect security in the context of private-key encryption [34] and Goldwasser and Micali's formulation of

[*] A preliminary version of this work appeared in [24]. Work by the first author was supported by a DoD/NDSEG Fellowship, and was done while he was in the Dept. of Computer Science, Columbia University.

semantic security in the public-key setting [18] can both arguably be said to mark turning points (each in their own way) in the development of the field. Following the seminal paper of Goldwasser and Micali, much work has focused on refining, analyzing, and extending definitions of computational security for encryption. Some work has focused on giving alternate, but equivalent, characterizations of semantic security [18, 37, 28, 1, 12], most notably in terms of *indistinguishability* (see Section 1.1).[1] Other work includes extending security definitions to the uniform model of computation [11], defining security under stronger adversarial attack models [31, 32, 27, 2, 16], and investigating relations among the various resulting definitions [2, 36, 16]. An alternate definition of security, *non-malleability*, has also been proposed [10] and the relation between this notion and semantic security/indistinguishability investigated [10, 2, 6].

With the exception of [27, 1], the above all focus predominantly on the case of *public-key* encryption. In contrast, rigorous and systematic analysis of definitions of security in the *private-key* setting has been mostly lacking. We are not the first to have noticed this deficiency; see [13, Section 5.5.4]. This relative lack of definitional work is unfortunate since private-key cryptography is the method most often used in practice for the encryption of bulk data. It warrants study separately from public-key cryptography for a number of reasons: from a practical point of view, it introduces concerns not present in the public-key setting (e.g., consideration of *modes of encryption* which are used for efficient encryption of bulk data) while from a theoretical standpoint it has added complications arising from an adversary's potential access to an "encryption oracle" (which is not an issue in the public-key case). As further evidence of the continuing importance of definitional work in this setting, we mention recent research focused on developing even "better" definitions of security for private-key encryption [25, 5, 8, 26, 9, 30, 22] and also the recent resurgence of interest in the design and analysis of modes of encryption (see http://csrc.nist.gov/encryption/modes/), especially those secure under stronger definitions of security.

This paper presents a systematic study of notions of security in the private-key setting. We present definitions of security in the sense of both indistinguishability and non-malleability — under multiple forms of adversarial attack — and characterize the resulting security notions in two ways. First, we explore the relations between these definitions; that is, for every pair of definitions we show whether one definition implies the other (so that any scheme secure under one notion is secure under the other) or whether the definitions are incomparable. In doing so, we construct a complete hierarchy indicating the relative strengths of the various notions. Second, we partition the security notions into two classes which we refer to as *computational* and *information-theoretic*. The first class comprises those definitions which can be satisfied only by assuming the existence of one-way functions, while the second contains notions of security that can be satisfied without any computational assumptions. (We do not claim that our second result is new; however, we believe it represents another useful way to classify notions of security and we are not aware of any previous systematic classification of this type.)

[1] This notion has variously been termed "polynomial security" [18], "find-then-guess security" [1], and "indistinguishability" [2].

1.1 Notions of Security

We begin with an informal overview of the definitions of security considered here; formal definitions appear in Section 2. We consider two types of security goals: *indistinguishability* (IND) [18, 28, 1, 2] and *non-malleability* (NM) [10, 2]. Indistinguishability directly relates to the secrecy afforded by a scheme, and may be viewed as a formalization of the requirement that an adversary not learn any information about a plaintext from the corresponding ciphertext. This definition was introduced [18] as a simpler characterization of semantic security, and the two have been proven equivalent in a variety of settings [18, 28, 11, 1, 36, 16]. Another commonly-used notion of security, left-or-right indistinguishability [1], has also been shown to be equivalent to the notion of indistinguishability considered here.[2]

Non-malleability, in contrast, relates to the resilience of the scheme against ciphertext modifications; secrecy is no longer (explicitly) a concern. Definitions of non-malleability have been proposed for a variety of primitives [10] and, for the case of public-key encryption, simpler characterizations have been given [2, 6]. Informally, a scheme is said to be non-malleable if an adversary — given a challenge ciphertext y representing an encryption of an unknown value x — is unable to generate a second ciphertext y' whose underlying plaintext x' is meaningfully related to x.

Definitions of security for both goals proceed in a common framework [2]. An adversary A is viewed as a pair of algorithms (A_1, A_2) corresponding to two "stages" of an attack. In advance of the adversary's execution, a random key *sk* is chosen; this key is kept hidden from the adversary. In the first stage of the attack, A_1 outputs a distribution over the message space whose format depends on the type of attack under consideration. Next, a message is selected at random according to the distribution and encrypted to give the "challenge" ciphertext. The challenge ciphertext (and any state information output by A_1) are then given as input to A_2, and the "success" of adversary A is determined according to the goal of the attack.

In addition to the goals as outlined above, a notion of security is also characterized by the external resources assumed to be available to an adversary attacking the scheme. Our security notions are labeled according to whether an adversary has possible access to an *encryption oracle* and/or *decryption oracle* during the two stages of its attack. When an adversary never has access to the encryption oracle, we denote this by P0. A *non-adaptive chosen-plaintext attack* indicates that an adversary has access to the encryption oracle during only the *first* stage of its attack ("non-adaptive" here refers to the fact that oracle queries cannot depend on the challenge ciphertext y); we denote this by P1. Finally, an *adaptive chosen-plaintext attack* implies that an adversary has access to an encryption oracle during both the first *and* second stages of its attack. This is denoted by P2. Similar notation is used for decryption oracle access: thus, C0 indicates no access to a decryption oracle, C1 means non-adaptive access during the first stage only [31], and C2 implies adaptive access [32]. We stress that while much previous work has considered C1 and C2 attacks in the public-key setting, the additional complication of an encryption oracle is present in the private-key setting only.

[2] Technically, equivalence between indistinguishability, semantic security, and left-or-right indistinguishability in the private-key setting has only been fully shown for the case of chosen-plaintext attacks [1] (P2 attacks, in the notation developed below). We conjecture that equivalence holds for all attack notions considered here, but leave this for future work.

We consider the eighteen notions of security determined by all possible combinations of goals (IND, NM), encryption oracle access (P0, P1, P2), and decryption oracle access (C0, C1, C2). We note that since the systems performing encryption and decryption are different (and may represent different parties) it makes sense to consider adversaries with different access to the different oracles.

1.2 Summary of Results

As stated earlier, the main contribution of this work is to completely classify the relative strengths of the eighteen notions of security mentioned above and thereby obtain a hierarchy of security notions for private-key encryption. The resulting hierarchy is shown in Figure 1. (For comparison, we have included the known results from the public-key setting [2] in Figure 2.) In the figures, notions A and B boxed together are equivalent; i.e., any scheme meeting notion of security A meets notion of security B and vice versa. A directed path from notion A to notion B means that A is strictly stronger than B; that is, any scheme meeting notion of security A also meets notion of security B but the converse is not necessarily true. If no directed path exists between A and B in either direction then these two notions are incomparable; a scheme meeting either definition of security need not meet the other.

We highlight (informally) those results specific to the private-key setting:

- P1 $\Rightarrow$ P2. This is our most surprising result, and the most challenging to prove. We show that, for all security notions considered here, an encryption scheme secure when an adversary has non-adaptive access to an encryption oracle is also secure when an adversary is given adaptive access to an encryption oracle.

- P0 $\not\Rightarrow$ P1. On the other hand, for all the security notions considered here there exist encryption schemes which are secure when the adversary has no access to an encryption oracle but which are insecure when the adversary has (non-adaptive) access to an encryption oracle. This is to be expected, as security against P0 attacks only requires (informally) the scheme to be secure when used to encrypt a single message: namely, the message that is encrypted to give the challenge ciphertext.

- NM $\not\Rightarrow$ IND. For some notions of security considered here, there exists an encryption scheme secure in the sense of non-malleability which is insecure in the sense of indistinguishability. This is in contrast to the public-key setting where non-malleability always implies indistinguishability [2].

Each of the above results holds unconditionally (however, the first result is vacuous unless one-way functions exist).

Our classification of security notions as either *computational* or *information-theoretic* (described earlier) highlights the adversary's access to an encryption oracle as the specific feature separating these two classes. More precisely, any security notion in which the adversary has no access to an encryption oracle falls into the "information-theoretic" class and is thus achievable without any computational assumptions. In contrast, every security notion in which the adversary has (non-adaptive) access to an encryption oracle lies in the "computational" class, and is achievable if and only if one-way functions exist.

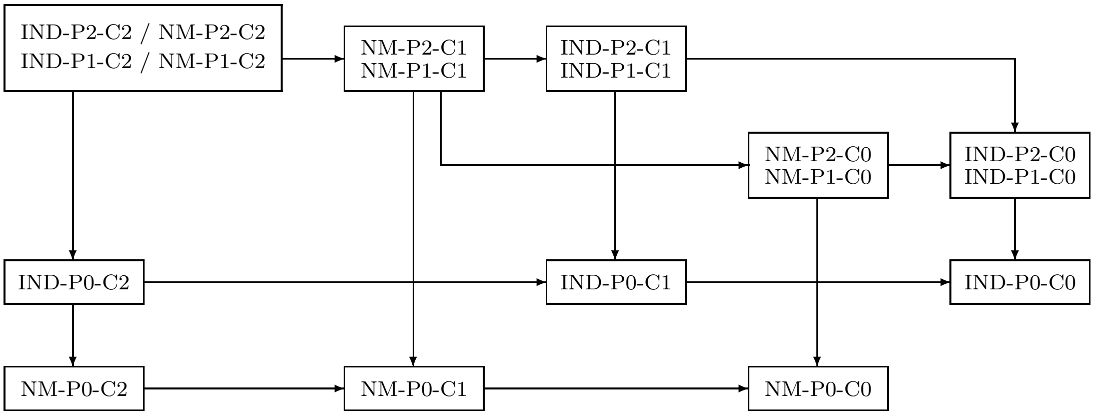

Figure 1: Hierarchy of private-key notions of security.

1.3 Other Related Work

Most previous definitional work related to encryption (referenced earlier) has focused on the public-key setting. There are, however, some notable exceptions. Luby [27, Chapter 14] describes the security notions IND-P1-C0 and IND-P1-C1 and gives constructions of cryptosystems secure with respect to these definitions. Dolev, Dwork, and Naor [10] mention the security notion IND-P1-C2 and present a scheme secure in this sense. Bellare, et al. [1] consider security in the sense of IND-P2-C0 and show concrete security reductions between this notion and various other formulations of indistinguishability. We stress, however, that they do not consider non-adaptive access to an encryption oracle, issues related to non-malleability in the private-key setting, or potential access to a decryption oracle. In fact, all notions of security discussed in [1] are polynomially equivalent to IND-P2-C0.

Our work largely follows the presentation and style of [2] which considers relations between notions of security for public-key encryption and gives a hierarchy of the six resulting security notions in that setting. Indeed, our work was inspired by the open question mentioned in the full version of that work [2, Section 1.6]. We confirm their conjecture that relations analogous to theirs hold in the private-key setting *as long as adaptive chosen-plaintext attacks are assumed.* Our most interesting results, however, arise from consideration of other attacks: namely, non-adaptive chosen-plaintext attacks (which we show — somewhat surprisingly — are equivalent to adaptive chosen-plaintext attacks) and no-plaintext attacks (where different relations hold between the corresponding notions of security).

1.4 Open Questions

As mentioned earlier, a number of new notions of security for private-key encryption have recently been proposed; these include various flavors of unforgeability or "authenticated encryption" [25, 5] (focusing on an adversary's inability to generate new, *valid* ciphertexts) as well as different definitions of "secure channels" [8, 26, 9, 30] (here, the goal is to distill those properties of encryption needed for secure message transmission). New attack models have also been considered [22]. It will be interesting to further develop this research, and

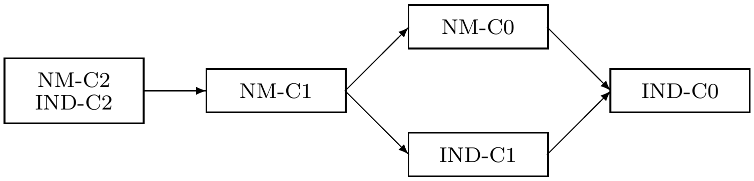

Figure 2: Hierarchy of public-key notions of security [2].

to relate the new definitions and attack models to the ones considered here.

This work considers only the case of stateless encryption schemes; however, new concerns and attacks may arise in the case of *stateful* encryption [4]. It will be useful to characterize security notions in this setting as well, and in fact it seems that our results do not carry over unchanged (for example, it appears that in the case of stateful encryption security against P1 attacks does not imply security against P2 attacks).

2 Preliminaries

We use the standard notation for describing probabilistic algorithms and experiments, following [19]. Denote by $A(x_1, x_2, \ldots; r)$ the result of running probabilistic algorithm A on inputs $x_1, x_2, \ldots$ using randomness r. The notation $y \leftarrow A(x_1, x_2, \ldots)$ denotes the experiment in which r is chosen uniformly at random and y is set equal to the output of $A(x_1, x_2, \ldots; r)$. If $\mathcal{S}$ represents a distribution, then $b \leftarrow \mathcal{S}$ denotes assigning to b an element chosen according to $\mathcal{S}$; when S is a (finite) set, $b \leftarrow S$ simply denotes assigning to b a uniformly-selected element of S. If $p(x_1, x_2, \ldots)$ is a predicate, the notation

$$\Pr\left[x_1 \leftarrow S; x_2 \leftarrow A(x_1, y_2, \ldots); \cdots : p(x_1, x_2, \ldots)\right]$$

denotes the probability that $p(x_1, x_2, \ldots)$ is true after ordered execution of the listed experiments.

We use both "$|$" and "$\circ$" to denote concatenation of strings. A function $\varepsilon : \mathbb{N} \rightarrow [0, 1]$ is *negligible* if for all $c > 0$ there exists an integer N_c such that $\varepsilon(N) \leq N^{-c}$ for all $N > N_c$. For simplicity, we consider only uniform algorithms in this work (both for honest players and for the adversary) and refer to a probabilistic, polynomial time Turing machine (where the running time is measured as a function of the length of its first input) as a "PPT algorithm"; however, all our results extend in the natural way to the case of non-uniform adversaries. Finally, "encryption scheme" refers to a probabilistic, stateless, private-key encryption scheme unless stated otherwise.

We begin with a formal definition of private-key encryption. Our definition is mostly standard, but we stress that we explicitly allow encryption over arbitrary message spaces and, in particular, over message spaces containing different-length messages. This is, in particular, meant to model the practical case of modes of encryption for which variable-size messages may typically be encrypted. We also assume without loss of generality that for a given security parameter k the sender and receiver share a uniformly-distributed key of length k.[3]

Definition 1 A *probabilistic, stateless, private-key encryption scheme* Π is a pair of algorithms $(\mathcal{E}, \mathcal{D})$ defined over a sequence of message spaces $\{\mathcal{M}_k\}_{k \geq 1}$ such that:

- There exists some polynomial p and some collection of index sets $\{I_k\}_{k \geq 1}$, with I_k a non-empty subset of $\{1, \ldots, p(k)\}$, such that $\mathcal{M}_k = \cup_{i \in I_k} \{0,1\}^i$. For a particular k, we call I_k the set of *legal message lengths*. We further require that, given 1^k and ℓ, one can efficiently determine whether $\ell \in I_k$.

- $\mathcal{E}$, the *encryption algorithm*, is a PPT algorithm that takes as input a secret key sk and a message $x \in \mathcal{M}_{|sk|}$ and returns a ciphertext y. (For concreteness, we let $\mathcal{E}_{sk}(x) = \perp$ if $x \notin \mathcal{M}_{|sk|}$.)

- $\mathcal{D}$, the *decryption algorithm*, is a deterministic, poly-time algorithm that takes as input a secret key sk and a ciphertext y and returns either $x \in \mathcal{M}_{|sk|}$ or a special symbol $\perp \notin \mathcal{M}_{|sk|}$.

We require that for all k, for all $sk \in \{0,1\}^k$, for all $x \in \mathcal{M}_k$, and for all y which can be output by $\mathcal{E}_{sk}(x)$, we have $\mathcal{D}_{sk}(y) = x$. ■

We remark that one may extend the definition to consider infinite message spaces and allow, for example, encryption of arbitrary-length messages. (In this case, we must allow $\mathcal{E}$ and $\mathcal{D}$ to take time polynomial in the length of their entire inputs and not merely polynomial in 1^k.) All our results continue to hold as long as $\{I_k\}$ is non-trivial in the following sense: there exists a polynomial p such that, for all k, $I_k \cap \{0,1\}^{\leq p(k)}$ is non-empty.

The above definition only deals with the semantics of private-key encryption. In the next section, we define notions of security for private-key encryption for the cases of both indistinguishability and non-malleability.

A note on concrete security. The definitions given below are all phrased in terms of *asymptotic* security; i.e., we say a scheme is secure if all PPT adversaries have negligible "advantage" in some appropriate sense, where an adversary's running time and advantage are measured as a function of the security parameter k. We note, however, that it is easy to reformulate all our definitions in terms of *concrete* security (as in, e.g., [3, 1]), where one is interested in an explicit bound on the advantage of any adversary expending a specified amount of resources. Although we recognize the importance of concrete security — especially in the case of private-key encryption — we feel that in the current paper a concrete security treatment would obscure the presentation and become overly cumbersome to the reader.[4] We stress that it is straightforward to derive the corresponding "concrete security" versions of all our theorems from the detailed proofs given here.

[3] This is in contrast to the usual definition which allows for an arbitrary *key generation algorithm* $\mathcal{K}$. In the context of private-key encryption, however, it is easy to see that any such scheme may be converted to a scheme in which the sender and receiver simply share the randomness used by $\mathcal{K}$.

[4] Indeed, for a full treatment we would need to consider the number of times an adversary accesses both the encryption and decryption oracles in both the first and the second stages of the experiment!

2.1 Indistinguishability

We refer to Section 1.1 for a discussion of indistinguishability and an intuitive explanation of the type of security it seeks to model. For an adversary $A = (A_1, A_2)$, we may describe the adversary's attack in this context as follows: at the end of the first stage of the attack, A_1 outputs a triple (x_0, x_1, s) consisting of two plaintext messages and state information s. One of x_0 or x_1 is chosen at random and encrypted to give the challenge ciphertext y. In the second stage of the attack, A_2 is given y and s and we say that A succeeds if it correctly determines whether y is an encryption of x_0 or x_1. Different types of attacks are modeled by giving A_1 or A_2 access to an encryption oracle and/or a decryption oracle; this gives rise to multiple security notions. Informally, an encryption scheme is indistinguishable (with respect to a particular type of attack) if every PPT adversary succeeds with probability only negligibly different from 1/2. More formally:

Definition 2 Let $\Pi = (\mathcal{E}, \mathcal{D})$ be an encryption scheme over message space $\{\mathcal{M}_k\}$ and let $A = (A_1, A_2)$ be an adversary. For $X, Y \in \{0, 1, 2\}$ and security parameter k, we define the *advantage* of A as $\mathsf{Adv}_{A,\Pi}^{\text{IND-PX-CY}}(k) \stackrel{\text{def}}{=}$

$$\Big| 2 \cdot \Pr\Big[sk \leftarrow \{0,1\}^k; (x_0, x_1, s) \leftarrow A_1^{\mathcal{O}_1, \mathcal{O}'_1}(1^k); b \leftarrow \{0,1\}; y \leftarrow \mathcal{E}_{sk}(x_b) : \\ A_2^{\mathcal{O}_2, \mathcal{O}'_2}(1^k, s, y) = b \Big] - 1 \Big| ,$$

where:

If $X = 0$ then $\mathcal{O}_1(\cdot) = \varepsilon$ and $\mathcal{O}_2(\cdot) = \varepsilon$
If $X = 1$ then $\mathcal{O}_1(\cdot) = \mathcal{E}_{sk}(\cdot)$ and $\mathcal{O}_2(\cdot) = \varepsilon$
If $X = 2$ then $\mathcal{O}_1(\cdot) = \mathcal{E}_{sk}(\cdot)$ and $\mathcal{O}_2(\cdot) = \mathcal{E}_{sk}(\cdot)$

and

If $Y = 0$ then $\mathcal{O}'_1(\cdot) = \varepsilon$ and $\mathcal{O}'_2(\cdot) = \varepsilon$
If $Y = 1$ then $\mathcal{O}'_1(\cdot) = \mathcal{D}_{sk}(\cdot)$ and $\mathcal{O}'_2(\cdot) = \varepsilon$
If $Y = 2$ then $\mathcal{O}'_1(\cdot) = \mathcal{D}_{sk}(\cdot)$ and $\mathcal{O}'_2(\cdot) = \mathcal{D}_{sk}(\cdot)$.

We insist that $A_1(1^k)$ output $x_0, x_1 \in \mathcal{M}_k$ with $|x_0| = |x_1|$. Furthermore, when $Y = 2$ we insist that A_2 does not ask for the decryption of challenge ciphertext y. We say that Π is secure in the sense of IND-PX-CY if $\mathsf{Adv}_{A,\Pi}^{\text{IND-PX-CY}}(\cdot)$ is negligible for any PPT adversary A. ■

As usual in definitions of this sort [18, 1, 2], we require that $|x_0| = |x_1|$ since encryption does not hide the length of the plaintext.

2.2 Non-Malleability

We refer to Section 1.1 for a discussion of non-malleability and various definitions thereof. In the present work, we extend the definition introduced by Bellare, et al. [2] in the public-key setting. Here, we may describe the adversary's attack as follows: at the end of the first stage of the attack, A_1 outputs a distribution M over messages in the legal message space along with state information s. Two messages $x, \tilde{x}$ are chosen at random according to M, and x is encrypted to give ciphertext y. In the second stage of the attack, A_2 is given y and s and outputs a relation R and a vector of ciphertexts $\vec{y}$. Let $\vec{x}$ correspond to the decryption of ciphertexts in $\vec{y}$ (i.e., if $x[i]$ represents the i^{th} component of vector $\vec{x}$, then $x[i] = \mathcal{D}_{sk}(y[i])$ for $1 \leq i \leq |\vec{x}|$). Informally, we say an encryption scheme is non-malleable if for every PPT A the probability that $R(x, \vec{x})$ is true is at most negligibly different from the probability that $R(\tilde{x}, \vec{x})$ is true.

In the above, both M and R are assumed to be (boolean) circuits, described using some standardized encoding. Since A is a PPT algorithm, note that that M and R may both be evaluated in polynomial time. (More formally, if the running time of A for security parameter k is bounded by $p(k)$ for some polynomial p, then the descriptions of M, R are of length at most $p(k)$ and hence can be evaluated in time at most $p(k)$.)

As above, different attacks are modeled by giving adversaries A_1 or A_2 access to encryption/decryption oracles; this gives rise to multiple notions of security. Formally:

Definition 3 Let $\Pi = (\mathcal{E}, \mathcal{D})$ be an encryption scheme over message space $\{\mathcal{M}_k\}$ and let $A = (A_1, A_2)$ be an adversary. For $X, Y \in \{0, 1, 2\}$ and security parameter k, we define the *advantage* of A as $\mathsf{Adv}_{A,\Pi}^{\text{NM-PX-CY}}(k) \stackrel{\text{def}}{=}$

$$\left| \mathsf{Expt}_{A,\Pi}^{\text{NM-PX-CY}}(k) - \mathsf{Rand}_{A,\Pi}^{\text{NM-PX-CY}}(k) \right|,$$

where $\mathsf{Expt}_{A,\Pi}^{\text{NM-PX-CY}}(k) \stackrel{\text{def}}{=}$

$$\begin{aligned} \Pr \quad & \Big[sk \leftarrow \{0,1\}^k; (M, s) \leftarrow A_1^{\mathcal{O}_1, \mathcal{O}'_1}(1^k); x \leftarrow M; y \leftarrow \mathcal{E}_{sk}(x); \\ & (R, \vec{y}) \leftarrow A_2^{\mathcal{O}_2, \mathcal{O}'_2}(1^k, s, y); \vec{x} = \mathcal{D}_{sk}(\vec{y}) : y \neq \perp \bigwedge y \notin \vec{y} \bigwedge \perp \notin \vec{x} \bigwedge R(x, \vec{x}) \Big], \end{aligned}$$

$\mathsf{Rand}_{A,\Pi}^{\text{NM-PX-CY}}(k) \stackrel{\text{def}}{=}$

$$\begin{aligned} \Pr \quad & \Big[sk \leftarrow \{0,1\}^k; (M, s) \leftarrow A_1^{\mathcal{O}_1, \mathcal{O}'_1}(1^k); x, \tilde{x} \leftarrow M; y \leftarrow \mathcal{E}_{sk}(x); \\ & (R, \vec{y}) \leftarrow A_2^{\mathcal{O}_2, \mathcal{O}'_2}(1^k, s, y); \vec{x} = \mathcal{D}_{sk}(\vec{y}) : y \neq \perp \bigwedge y \notin \vec{y} \bigwedge \perp \notin \vec{x} \bigwedge R(\tilde{x}, \vec{x}) \Big], \end{aligned}$$

and $\mathcal{O}_1(\cdot), \mathcal{O}_2(\cdot), \mathcal{O}'_1(\cdot), \mathcal{O}'_2(\cdot)$ are as in Definition 2 for the corresponding values of X, Y.

In the above, M is a circuit representing a distribution over strings and R is a circuit computing some relation. We require that $|x| = |x'|$ for all x, x' in the support of M. We also require that the vector of ciphertexts $\vec{y}$ output by A_2 be non-empty (but see the remark below). Furthermore, when $Y = 2$ we insist that A_2 does not ask for the decryption of y. We say that Π is secure in the sense of NM-PX-CY if $\mathsf{Adv}_{A,\Pi}^{\text{NM-PX-CY}}(\cdot)$ is negligible for any PPT adversary A. ■

Note that if sampling M results in $x \notin \mathcal{M}_k$ then $\mathcal{E}_{sk}(x) = \perp$ and the adversary's advantage will not increase in this case. For the same reason mentioned in the previous section, we require that all messages in the support of M have the same length. We disallow $y \in \vec{y}$ so the adversary does not get credit for simply "copying".

Alternate definitions of non-malleability. The above definition of non-malleability was adapted directly from the work of [2, 6] in the public-key setting, and seems to most closely capture the intuition of an adversary's "being unable to generate a ciphertext decrypting to a related message" when given an encryption y of some message x. We remark, however, that alternate definitions are also possible and may be more appropriate for certain applications. One possibility is to modify the definition so that A_2 is allowed to output an empty vector $\vec{y}$. (If the adversary must *always* output an empty vector $\vec{y}$ then we obtain a definition similar to "semantic security with respect to relations" as considered by [10].) Notice that such a definition is equivalent to the above definition for $X \in \{1, 2\}$: informally, given some adversary A who (sometimes) outputs an empty $\vec{y}$ and relation R, we may construct an adversary A' who in this case outputs an arbitrary ciphertext $y' \neq y$ (obtained using its encryption oracle, possibly in the first stage) and relation R' such that $R'(x, x') = R(x)$. Clearly, the advantage of A' with respect to Definition 3 is the same as the advantage of A under the modified definition. On the other hand, when $X = 0$ the definitions are *not* equivalent: the encryption scheme constructed in the proof of Theorem 7 gives an example of a scheme secure in the sense of NM-P0-C2 (with respect to Definition 3) which is *not* secure in the sense of NM-P0-C0 (with respect to the modified definition). Furthermore, following [10, Theorem 2.4], it is not hard to show that any scheme which is non-malleable with respect to the modified definition (where the adversary is allowed to output an empty vector of ciphertexts) is also indistinguishable; this is in contrast to what is shown in Theorem 7 with respect to the definition of non-malleability used here.

Another possibility is to remove the requirement in the above definition that $\perp \notin \vec{x}$ (recall, $\vec{x}$ is the vector of plaintexts that result from decrypting $\vec{y}$); in fact, doing so would somewhat simplify the proof of Theorem 5. We have decided against this formulation both in order to parallel previous work in this area [2, 6] and also because the current definition more closely corresponds to our intuitive notion of "producing a ciphertext decrypting to a (meaningful) related message". Again, however, a modified definition in which an adversary may "succeed" even when $\perp \in \vec{x}$ may be more appropriate for certain applications.

3 Relation Between Private-Key Encryption and One-Way Functions

Here, we briefly discuss the relationship between private-key encryption and one-way functions. Although some of these results have previously appeared and others may be considered "folklore", we believe it is instructive to systematically study the relation between one-way functions and each of the definitions presented in the previous section, and to present the results in a unified way.

Much work has focused on exploring the connections between one-way functions and "higher-level" cryptographic primitives such as encryption, authentication, etc. It has been established that one-way functions imply the existence of pseudorandom generators [20], pseudorandom functions [14], message authentication and private-key encryption [15], bit commitment [29], and digital signature schemes [33]. Conversely, all of the above primitives imply the existence of one-way functions [21, 33]. One must be careful, however, in interpreting this last result since, for example, a construction of secure private-key encryption without any computational assumptions is known (namely, the one-time pad [35])!

In fact, the precise statement of [21] is that private-key encryption *in which the encrypted message is longer than the shared key* implies the existence of one-way functions. Our aim in this section is to reformulate this result in terms of the definitions given in Section 2. In this spirit, we present three results which —together with the hierarchy of Figure 1 — completely characterize the relationship between private-key encryption and one-way functions; namely:

- One-way functions are sufficient to construct an encryption scheme secure in the sense of IND-P2-C2. Given the hierarchy of Figure 1, this implies that one-way functions are sufficient for any of the notions considered here.
- The existence of an encryption scheme secure in the sense of IND-P1-C0 implies the existence of one-way functions.
- An encryption scheme secure in the sense of IND-P0-C2 may be constructed without any computational assumptions (we do, however, assume that the adversary makes only polynomially-many queries to the decryption oracle).

Figure 1 shows that all notions considered here are either strictly stronger than (or equivalent to) IND-P1-C0 or strictly weaker than (or equivalent to) IND-P0-C2. Thus, the above results completely partition the present security notions into two classes: (1) *computational* notions of security for which the existence of an encryption scheme satisfying such a notion is equivalent to the existence of a one-way function; and (2) *information-theoretic* notions of security for which an encryption scheme satisfying such a notion may be constructed without any computational assumptions. It thus emerges that the key difference between computational and information-theoretic notions is the adversary's access to an encryption oracle or, equivalently, whether the encryption scheme is required to be secure for the encryption of a single message only (i.e., secure against P0 attacks) or whether it remains secure even when an adversary sees encryptions of multiple messages (e.g., via P1 attacks).

We now formally state and prove the results stated above. We begin with our working definition of a collection of one-way functions [17, 12].

Definition 4 Let $\{D_k\}_{k\geq 1}$ and $\{R_k\}_{k\geq 1}$ be collections of finite sets. We say $F = \{f_k : D_k \rightarrow R_k\}_{k\in\mathbb{N}}$ is a *collection of one-way functions* (informally, a *one-way function*) if:

- There exists a PPT algorithm Sample such that $\mathsf{Sample}(1^k)$ outputs a uniformly-distributed element of D_k. (Thus, "$x \leftarrow D_k$" is equivalent to "$x \leftarrow \mathsf{Sample}(1^k)$".)
- There exists a deterministic, poly-time algorithm Eval such that for all k and all $x \in D_k$, the output of $\mathsf{Eval}(1^k, x)$ is $f_k(x)$.
- For all PPT algorithms A, the following is negligible (in k):

$$\Pr[x \leftarrow D_k; y = f_k(x) : f_k(A(1^k, y)) = y].$$

■

This definition is equivalent to the one which considers a single function f defined over a fixed (infinite) domain D [17].

We first note that encryption schemes secure under the strongest definition of security may be constructed from any one-way function; this was first explicitly noted by Dolev, Dwork, and Naor [10].

Theorem 1 *Assuming the existence of a one-way function, there exists an encryption scheme secure in the sense of* IND-P2-C2.

Proof We repeat the construction given in [10]. Let $\mathcal{F} = \{F^k\}_{k\geq 1}$ be a pseudorandom function family where $F^k = \{F_s : \{0,1\}^k \to \{0,1\}^k\}_{s\in\{0,1\}^k}$ is a collection of functions indexed by a key $s \in \{0,1\}^k$, and let $(\textsc{mac}, \mathsf{Vrfy})$ be a message authentication code. Both of these primitives may be constructed from any one-way function: the existence of one-way functions implies the existence of (length-doubling) pseudorandom generators [20] which in turn imply the existence of pseudorandom functions [14]; the latter may be used to construct secure message authentication codes [15]. We now construct the following encryption scheme $(\mathcal{E}, \mathcal{D})$ over message space $\mathcal{M}_k = \{0,1\}^{\lfloor k/2 \rfloor}$:

```
$\underline{\mathcal{E}_{sk}(m)}$
parse $sk$ as $s_1 \circ s_2$
  with $|s_1| = \lfloor |sk|/2 \rfloor$
$r \leftarrow \{0,1\}^{|s_1|}$
$c = F_{s_1}(r) \oplus m$
$t = \textsc{mac}_{s_2}(r \circ c)$
return $\langle r, c, t \rangle$
```

```
$\underline{\mathcal{D}_{sk}(\langle r, c, t \rangle)}$
parse $sk$ as $s_1 \circ s_2$
  with $|s_1| = \lfloor |sk|/2 \rfloor$
if $\mathsf{Vrfy}_{s_2}(r \circ c, t) = 1$
  return $F_{s_1}(r) \oplus c$
else return $\perp$
```

We briefly sketch the proof that this scheme is secure in the sense of IND-P2-C2. First note that, by security of the message authentication code, the decryption oracle is of no help to the adversary. More formally, with all but negligible probability the decryption oracle will return $\perp$ for all new ciphertexts (i.e., those not returned by the encryption oracle) submitted by the adversary. Thus, it suffices to consider security against an adaptive chosen-plaintext attack (cf. [25]). But it is well-known that this scheme is secure against such an attack [15], since the probability that a nonce r repeats is negligibly small. ■

We now characterize which notions of security require the existence of a one-way function in order to be satisfied. Recall that Impagliazzo and Luby [21] have shown that any encryption scheme where the message is longer than the key implies the existence of a one-way function. We recast their result in terms of the definitions of the previous section. Furthermore, their proof is quite complicated (their proof first constructs a function with false entropy which is then shown to imply the existence of a pseudorandom generator); we give a simpler and more direct proof here which may be of independent interest.

Theorem 2 *The existence of an encryption scheme secure in the sense of* IND-P1-C0 *implies the existence of a one-way function.*

Proof Let $\Pi = (\mathcal{E}, \mathcal{D})$ be an encryption scheme over message space $\{\mathcal{M}_k\}$ which is secure in the sense of IND-P1-C0. For a given k, let $\mathcal{M}'_k \subseteq M_k$ denote some arbitrary pair of distinct, equal-length strings (for concreteness, let $\mathcal{M}'_k = \{0^\ell, 1^\ell\}$ for the smallest $\ell \in I_k$; note that Definition 1 ensures that $\mathcal{M}'_k$ can be found in polynomial time). Slightly abusing notation, we let $\mathcal{M}'_k$ refer to the uniform distribution over $\mathcal{M}'_k$ as well. Consider the modified encryption scheme $\Pi' = (\mathcal{E}', \mathcal{D}')$ over message space $\{(\mathcal{M}'_k)^{2k}\}$ in which encryption of a message $m = m_1 \circ \cdots \circ m_{2k}$ (with $m_i \in \mathcal{M}'_k$) is done by simple concatenation (i.e., $\mathcal{E}'_{sk}(m) = \mathcal{E}_{sk}(m_1) \circ \cdots \circ \mathcal{E}_{sk}(m_{2k})$) and decryption is done in the obvious way. Since Π is secure in the sense of IND-P1-C0, a standard hybrid argument shows that Π' is secure in the sense of IND-P0-C0 (actually, it is also secure in the sense of IND-P1-C0 but we do not use this fact).

Define $F = \{f_k\}_{k \in \mathbb{N}}$ as follows: $f_k(sk, m, \omega) = \mathcal{E}'_{sk}(m; \omega) \circ m$, where $sk \in \{0,1\}^k$, $m \in (\mathcal{M}'_k)^{2k}$, and ω represents the random coins used by $\mathcal{E}'$ when encrypting. We claim that F is a collection of one-way functions. Note that the domain of F is efficiently sampleable since sk and ω are arbitrary strings of the appropriate length and $\mathcal{M}'_k$ is efficiently sampleable. Furthermore, F is efficiently evaluable. Finally, we show that f_k is one-way. Assume toward a contradiction that there exists a PPT algorithm A which inverts f_k with some probability $\delta(k)$. We construct a PPT adversary B attacking Π' as follows:

```
$B_1(1^k)$
$x_0, x_1 \leftarrow (\mathcal{M}'_k)^{2k}$
return $(x_0, x_1, \{x_0, x_1\})$
```

```
$B_2(1^k, \{x_0, x_1\}, y)$
$(sk', x, \omega) \leftarrow A(1^k, y \circ x_0)$
if $\mathcal{E}'_{sk'}(x; \omega) = y$ and $x = x_0$
  return 0
else return 1
```

Clearly, if y is an encryption of x_0, then B_2 outputs 0 with probability exactly $\delta(k)$. On the other hand, if y is an encryption of x_1, then the probability $\delta'(k)$ that B_2 outputs 0 is bounded as follows:

$$\begin{aligned}
&\delta'(k) \\
&\leq \Pr[sk \leftarrow \{0,1\}^k; x_0, x_1 \leftarrow (\mathcal{M}'_k)^{2k}; y \leftarrow \mathcal{E}'_{sk}(x_1) : \exists sk' \text{ s.t. } \mathcal{D}'_{sk'}(y) = x_0] \\
&\leq \sum_{sk'} \Pr[sk \leftarrow \{0,1\}^k; x_0, x_1 \leftarrow (\mathcal{M}'_k)^{2k}; y \leftarrow \mathcal{E}'_{sk}(x_1) : \mathcal{D}'_{sk'}(y) = x_0] \\
&\leq \sum_{sk'} 2^{-2k} = 2^{-k},
\end{aligned}$$

where the last inequality holds because $\left|(\mathcal{M}'_k)^{2k}\right| = 2^{2k}$. The advantage of B is then at least $\delta(k) - 2^{-k}$ and therefore $\delta(k)$ must be negligible, as desired. ■

Finally, we demonstrate that all security notions weaker than IND-P1-C0 can be satisfied without any computational assumptions.

Theorem 3 *(Without any computational assumptions) there exists an encryption scheme secure in the sense of* IND-P0-C2.

Proof We are unaware of any previous explicit statement of this form; however, such a scheme is easy to construct. Consider the following encryption scheme $\Pi = (\mathcal{E}, \mathcal{D})$ defined over message space $\{0,1\}$ (for simplicity, we define the scheme for odd k only):

```
$\mathcal{E}_{sk}(m)$
parse $sk$ as $a \circ b \circ s$,
  where $|a| = |b| = \ell \overset{\text{def}}{=} \lfloor k/2 \rfloor$
$c = m \oplus s$
view $a, b, c$ as elements of $\mathbb{F}_{2^\ell}$
return $\langle c, ac + b \rangle$
```

```
$\mathcal{D}_{sk}(\langle c, t \rangle)$
parse $sk$ as $a \circ b \circ s$,
  where $|a| = |b| = \ell \overset{\text{def}}{=} \lfloor k/2 \rfloor$
view $a, b, c$ as elements of $\mathbb{F}_{2^\ell}$
if $t = ac + b$ return $c \oplus s$
else return $\perp$
```

As in the proof of Theorem 1, decryption oracle queries cannot "help" the adversary since the response to all such queries will be $\perp$ with all but negligible probability. Furthermore, conditioned on all such queries being answered with $\perp$, the adversary has no information about s and therefore no information about m. (More general proofs of a similar result have also been given in the context of authenticated encryption [25, 5].)

We note that Π is secure even against an unbounded adversary, as long as the adversary is limited to only polynomially-many queries to the decryption oracle. ■

4 Relations Among the Notions of Security

In this section, we state and prove the results which give rise to the hierarchy of Figure 1.

4.1 Adaptive Access to an Encryption Oracle: the Case of Indistinguishability

Here, we show that adaptive access to the encryption oracle does not help an adversary in the case of indistinguishability.

Theorem 4 (IND-P1-CY $\Rightarrow$ IND-P2-CY) *If encryption scheme* Π *is secure in the sense of* IND-P1-CY *then* Π *is secure in the sense of* IND-P2-CY, *for* $Y \in \{0, 1, 2\}$.

Proof Let Π be an encryption scheme secure in the sense of IND-P1-CY, and assume we have some adversary A attacking Π in the sense of IND-P2-CY. At a high level, our proof proceeds in the following two stages:

1. We define an oracle $\$_{sk}(\cdot)$ which returns encryptions of *random* plaintext, and show that replacing (in the second stage of A's attack) the "real" encryption oracle $\mathcal{E}_{sk}(\cdot)$ with $\$_{sk}(\cdot)$ does not change A's advantage by more than a negligible amount.

2. We show that Π is secure in the sense of IND-P1-CY even if an adversary additionally has access to $\$_{sk}(\cdot)$ in the second stage of its attack.

When $\mathcal{E}_{sk}(\cdot)$ is replaced with $\$_{sk}(\cdot)$ in the second stage of A's attack, A is simply attacking Π in the sense of IND-P1-CY with the added capability of accessing $\$_{sk}(\cdot)$ in the second stage. Thus, once we have proven the stated claims we will have proved the desired result.

Our oracle $\$_{sk}(\cdot)$ is exactly the "random encryption" oracle introduced by Bellare, et al. in their treatment of private-key encryption [1] which, on input x, returns the encryption of a random plaintext of length $|x|$. More precisely, for $x \in \mathcal{M}_k$, $\$_{sk}(x)$ returns $\mathcal{E}_{sk}(r)$, where $r \leftarrow \{0,1\}^{|x|} \cap \mathcal{M}_k$ (i.e., r is uniformly chosen from elements of $\mathcal{M}_k$ of length $|x|$). If $x \notin \mathcal{M}_k$, then the output of $\$_{sk}(x)$ is $\perp$.

Define $\mathsf{Adv}^{\text{IND-P\$-CY}}_{A,\Pi}(k)$ as the advantage of $A = (A_1, A_2)$ when its access to $\mathcal{E}_{sk}(\cdot)$ in the second stage is replaced with access to $\$_{sk}(\cdot)$. More formally,

$$\begin{aligned}
&\mathsf{Adv}^{\text{IND-P\$-CY}}_{A,\Pi}(k) \overset{\text{def}}{=} \\
&\quad \Big| 2 \cdot \Pr\Big[sk \leftarrow \{0,1\}^k; (x_0, x_1, s) \leftarrow A_1^{\mathcal{E}_{sk}, \mathcal{O}'_1}(1^k); b \leftarrow \{0,1\}; y \leftarrow \mathcal{E}_{sk}(x_b) : \\
&\qquad\qquad A_2^{\$_{sk}, \mathcal{O}'_2}(1^k, s, y) = b \Big] - 1 \Big|,
\end{aligned}$$

where oracles $\mathcal{O}'_1, \mathcal{O}'_2$ depend on the value of Y as in Definition 2. We now prove our first claim from above:

Claim 1 *For $Y \in \{0,1,2\}$, if Π is secure in the sense of* IND-P1-CY *then the following is negligible for any* PPT *adversary A:*

$$\left|\mathsf{Adv}_{A,\Pi}^{\text{IND-P\$-CY}}(k) - \mathsf{Adv}_{A,\Pi}^{\text{IND-P2-CY}}(k)\right|.$$

Proof The proof of this claim is via a hybrid argument where, in the q^{th} hybrid, the adversary's first $q-1$ queries to the encryption oracle are answered correctly, its q^{th} query to the encryption oracle is either answered correctly or "at random" (more precisely, by returning the encryption of a random string), and its remaining queries are answered at random. Let $A = (A_1, A_2)$ be a PPT adversary attacking Π in the sense of IND-P2-CY. We construct an adversary B attacking Π in the sense of IND-P1-CY. Relating the advantage of B to the difference above will give the stated result.

Let $\ell(\cdot)$ be a polynomial bound on the number of queries made by A_2 to its encryption oracle, and let $\{I_k\}$ denote the set of legal message lengths for Π (cf. Definition 1). Without loss of generality, we make the following assumptions: (1) A_2 always makes *exactly* $\ell(k)$ queries to its encryption oracle for a given security parameter k; (2) A_2 never submits to its decryption oracle (if it has access to one) a ciphertext it received from its encryption oracle; and finally, (3) all queries A_2 makes to its encryption oracle are valid messages in $\mathcal{M}_k$. We now define adversary $B = (B_1, B_2)$ as follows:

```
$\underline{B_1^{\mathcal{E}_{sk},\mathcal{O}'_1}(1^k)}$
$c \leftarrow \{0,1\}; q \leftarrow \{1,\ldots,\ell(k)\}$
$(x_0, x_1, s) \leftarrow A_1^{\mathcal{E}_{sk},\mathcal{O}'_1}(1^k)$
$\tilde{y} \leftarrow \mathcal{E}_{sk}(x_c)$
run $A_2^{\mathcal{E}_{sk},\mathcal{O}'_2}(1^k, s, \tilde{y})$ until
  $A_2$ makes $q^{\text{th}}$ query $X_q$ to $\mathcal{E}_{sk}(\cdot)$;
  let $s'$ be the state of $A_2$
$X_\$ \leftarrow \{0,1\}^{|X_q|}$
for $q+1 \le i \le \ell(k)$
  for $j \in I_k$
    $r_{i,j} \leftarrow \{0,1\}^j; y_{i,j} \leftarrow \mathcal{E}_{sk}(r_{i,j})$
return $(X_q, X_\$, (s', c, \{y_{i,j}\}))$
```

```
$\underline{B_2^{\mathcal{O}'_2}(1^k, (s', c, \{y_{i,j}\}), y)}$
run $A_2^{\mathcal{E}_{sk},\mathcal{O}'_2}$ using state $s'$ until it outputs $b$,
  answering current $\mathcal{E}_{sk}(\cdot)$ query with $y$
  and remaining $\mathcal{E}_{sk}(\cdot)$ queries using $\{y_{i,j}\}$
if $b = c$ output 0
else output 1
```

To clarify the description of B_2, when algorithm A_2 makes its i^{th} query to $\$_{sk}(\cdot)$ (for $q+1 \le i \le \ell(k)$), denote this query by X_i; we have B_2 respond to this query with $y_{i,|X_i|}$. We may note the following about the execution of B:

- If y is an encryption of X_q, then B simulates an execution of A in which plaintext x_c was encrypted to give ciphertext $\tilde{y}$, the first q queries of A_2 to its encryption oracle were answered by $\mathcal{E}_{sk}(\cdot)$ and the last $\ell(k) - q$ queries of A_2 to its encryption oracle were answered by $\$_{sk}(\cdot)$.

- If y is an encryption of $X_\$$, then B simulates an execution of A in which plaintext x_c was encrypted to give ciphertext $\tilde{y}$, the first $q-1$ queries of A_2 to its encryption oracle were answered by $\mathcal{E}_{sk}(\cdot)$ and the last $\ell(k)-q+1$ queries of A_2 to its encryption oracle were answered by $\$_{sk}(\cdot)$.

In each case, the simulation of B is perfect (with regard to the stated experiment); in particular, B has no difficulty responding to the decryption oracle queries of A in either stage. It can also be verified easily that B is a PPT algorithm.

For brevity, we let $\Pr_i[A = b'|b]$ denote the probability that A outputs b' when x_b is encrypted and the first i queries of A_2 to its encryption oracle are answered by $\mathcal{E}_{sk}(\cdot)$ and the last $\ell(k) - i$ queries to its encryption oracle are answered by $\$_{sk}(\cdot)$. Furthermore, let $\Pr[B = b'|X_q]$ denote the probability that B outputs b' when X_q is encrypted and let $\Pr[B = b'|X_\$]$ denote the probability that B outputs b' when $X_\$$ is encrypted. (In the above, we drop the explicit dependence of these probabilities on k for convenience.) We then have:

$$\begin{aligned}
&\mathsf{Adv}^{\text{IND-P1-CY}}_{B,\Pi}(k)\\
&\overset{\text{def}}{=} \left|\Pr[B=0|X_q]+\Pr[B=1|X_\$]-1\right|\\
&= \left|\Pr[B=0|X_q]-\Pr[B=0|X_\$]\right|\\
&= \left|\frac{1}{2\ell(k)}\sum_{i=1}^{\ell(k)}\Pr_i[A=0|0]+\frac{1}{2\ell(k)}\sum_{i=1}^{\ell(k)}(1-\Pr_i[A=0|1])\right.\\
&\qquad\left.-\frac{1}{2\ell(k)}\sum_{i=0}^{\ell(k)-1}\Pr_i[A=0|0]-\frac{1}{2\ell(k)}\sum_{i=0}^{\ell(k)-1}(1-\Pr_i[A=0|1])\right|\\
&= \left|\frac{1}{2\ell(k)}\cdot\left(\Pr_{\ell(k)}[A=0|0]-\Pr_{\ell(k)}[A=0|1]\right)\right.\\
&\qquad\left.-\frac{1}{2\ell(k)}\cdot(\Pr_0[A=0|0]-\Pr_0[A=0|1])\right|\\
&\geq \frac{1}{2\ell(k)}\cdot\left|\mathsf{Adv}^{\text{IND-P2-CY}}_{A,\Pi}(k)-\mathsf{Adv}^{\text{IND-P\$-CY}}_{A,\Pi}(k)\right|.
\end{aligned}$$

The proof of the claim concludes by noting that Π is secure in the sense of IND-P1-CY, so that $\mathsf{Adv}^{\text{IND-P1-CY}}_{B,\Pi}(k)$ is negligible. □

We now continue with a proof of our second claim from above:

Claim 2 *For $Y \in \{0,1,2\}$, if Π is secure in the sense of* IND-P1-CY *then, for any* PPT *adversary A,* $\mathsf{Adv}^{\text{IND-P\$-CY}}_{A,\Pi}(k)$ *is negligible.*

Proof Given a PPT adversary A attacking Π in the sense of IND-P\$-CY, we construct a PPT adversary B attacking Π in the sense of IND-P1-CY; furthermore, the advantage of B will be polynomially related to that of A. Since Π is secure in the sense of IND-P1-CY, this will give the result of the claim.

Let $\ell(\cdot)$ be a polynomial bound on the number of queries made by A_2 to $\$_{sk}(\cdot)$, and let I_k be as above. We define adversary $B = (B_1, B_2)$ as follows:

```
$B_1^{\mathcal{E}_{sk}, \mathcal{O}'_1}(1^k)$
$(x_0, x_1, s) \leftarrow A_1^{\mathcal{E}_{sk}, \mathcal{O}'_1}(1^k)$
for $1 \leq i \leq \ell(k)$
  for $j \in I_k$
    $r_{i,j} \leftarrow \{0,1\}^j$
    $y_{i,j} \leftarrow \mathcal{E}_{sk}(r_{i,j})$
return $(x_0, x_1, (s, \{y_{i,j}\}))$
```

```
$B_2^{\mathcal{O}'_2}(x_0, x_1, (s, \{y_{i,j}\}), y)$
run $A_2^{\$_{sk}, \mathcal{O}'_2}(x_0, x_1, s, y)$,
  answering its $\$_{sk}(\cdot)$ queries using $\{y_{i,j}\}$,
  until it outputs $b$
return $b$
```

To clarify the description of B_2: when A_2 makes its i^{th} query to $\$_{sk}(\cdot)$, denote this query by X_i. If $X_i \in \mathcal{M}_k$, B_2 responds with $y_{i,|X_i|}$; otherwise, B_2 responds with $\perp$. Note that this perfectly simulates a response from $\$_{sk}(\cdot)$. Also, B_1 and B_2 have no trouble responding to the other oracle queries of A_1 and A_2. Finally, the running time of B_2 is that of A_2, while B_1 incurs an additional overhead resulting from the $\ell(k)\mu(k)$ queries to its encryption oracle. Since $\ell(k)\mu(k)$ is polynomial, B is a PPT algorithm.

It is easy to see that B provides a perfect simulation for A and hence

$$\mathsf{Adv}_{B,\Pi}^{\text{IND-P1-CY}}(k) = \mathsf{Adv}_{A,\Pi}^{\text{IND-P\$-CY}}(k).$$

The claim follows. □

The preceding claims show that both $\mathsf{Adv}_{A,\Pi}^{\text{IND-P\$-CY}}(k)$ and

$$\left|\mathsf{Adv}_{A,\Pi}^{\text{IND-P\$-CY}}(k) - \mathsf{Adv}_{A,\Pi}^{\text{IND-P2-CY}}(k)\right|$$

are negligible. Thus, $\mathsf{Adv}_{A,\Pi}^{\text{IND-P2-CY}}(k)$ is negligible and Π is secure in the sense of IND-P2-CY. This completes the proof of the theorem. ■

4.2 Adaptive Access to an Encryption Oracle: the Case of Non-Malleability

Intuitively, one might expect that adaptive access to an encryption oracle might help in the context of security in the sense of non-malleability. After all, in this setting an adversary must output a valid ciphertext, and not "just" a bit (as in the case of security in the sense of indistinguishability). However, the following theorem shows that this intuition is wrong.

Before giving the proof of the theorem, we provide a high-level overview. First, we define a notion of security which may be viewed as an "indistinguishability-based" characterization of non-malleability. This definition is based on a similar definition given by Bellare and Sahai [6] in the public-key setting. Our definition (as in [6]) makes use of a "parallel decryption oracle" $\mathcal{D}_{sk}^{\parallel}(\cdot)$ which functions as a standard decryption oracle except that queries to this oracle may be submitted *in parallel*; i.e., on input a vector of ciphertexts $\vec{y} = (y_1, \ldots, y_\ell)$, oracle $\mathcal{D}_{sk}^{\parallel}(\cdot)$ returns $\vec{x} = (x_1, \ldots, x_\ell)$ where $x_i = \mathcal{D}_{sk}(y_i)$.

As in [6], we consider only definitions of security in which the adversary has access to $\mathcal{D}_{sk}^{\parallel}(\cdot)$ in the *second* stage and may query this oracle only *once*. The corresponding notions of security (see below for precise definitions) are denoted IND-PX-C$^{\|}$Y (note, however, that IND-PX-C$^{\|}$2 is equivalent to IND-PX-C2). We stress two differences between our definitions and those of [6]: (1) we require that the adversary access $\mathcal{D}^{\|}_{sk}(\cdot)$ only *after all queries to the encryption oracle have been made* (of course, this only matters when the adversary has access to $\mathcal{E}_{sk}(\cdot)$ in the second stage). Furthermore, (2) let $\vec{y}$ be the query submitted to $\mathcal{D}^{\|}_{sk}(\cdot)$ and let $\vec{x}$ be the corresponding response. Informally, we do not consider the adversary "successful" if $\perp \in \vec{x}$.

Theorem 5 (NM-P1-CY $\Rightarrow$ NM-P2-CY) *If encryption scheme* Π *is secure in the sense of* NM-P1-CY *then* Π *is secure in the sense of* NM-P2-CY, *for* $Y \in \{0, 1, 2\}$.

Proof We first formalize our definition of security in the sense of IND-PX-C$^{\|}$Y. Let $A = (A_1, A_2)$ be an adversary and consider the following experiment:

```
$\underline{\mathsf{Expt}^{\text{IND-PX-C}^{\|}\text{Y}}_{A,\Pi}(k)}$
$sk \leftarrow \{0,1\}^k$
$(x_0, x_1, s) \leftarrow A_1^{\mathcal{O}_1, \mathcal{O}'_1}(1^k)$
$b \leftarrow \{0,1\}; y \leftarrow \mathcal{E}_{sk}(x_b)$
$b' \leftarrow A_2^{\mathcal{O}_2, \mathcal{O}'_2, \mathcal{D}^{\|}_{sk}}(1^k, s, y)$
let $\vec{y}$ be the query submitted to $\mathcal{D}^{\|}_{sk}$
and let $\vec{x}$ be the corresponding response
if $\perp \in \vec{x}$
   $b'' \leftarrow \{0,1\}$; output $b''$
if $\perp \notin \vec{x}$ and $b' = b$ output 1
if $\perp \notin \vec{x}$ and $b' \neq b$ output 0
```

(Oracles $\mathcal{O}_1, \mathcal{O}'_1, \mathcal{O}_2, \mathcal{O}'_2$ are instantiated as in Definitions 2 and 3 depending on the values of X and Y.) As discussed in the paragraph preceding the theorem, we require that A_2 submit only a single query to $\mathcal{D}^{\|}_{sk}(\cdot)$ and furthermore that A_2 not access the encryption oracle $\mathcal{E}_{sk}(\cdot)$ (in case A_2 has access to this oracle) after it has submitted its query to $\mathcal{D}^{\|}_{sk}(\cdot)$. We also require, as usual, that A_1 outputs two messages x_0, x_1 of equal length and that A_2 does not ask for decryption of the challenge ciphertext y. We let

$$\mathsf{Adv}^{\text{IND-PX-C}^{\|}\text{Y}}_{A,\Pi}(k) \stackrel{\text{def}}{=} \left| 2 \cdot \Pr[\mathsf{Expt}^{\text{IND-PX-C}^{\|}\text{Y}}_{A,\Pi}(k) = 1] - 1 \right|$$

and say that Π is secure in the sense of IND-PX-C$^{\|}$Y if $\mathsf{Adv}^{\text{IND-PX-C}^{\|}\text{Y}}_{A,\Pi}(k)$ is negligible for all PPT adversaries A.

We stress that, in contrast to the definition of [6], the definition above does not consider the adversary "successful" in case $\perp \in \vec{x}$. This is so by definition of $\mathsf{Expt}^{\text{IND-PX-C}^{\|}\text{Y}}_{A,\Pi}(k)$: in case $\perp \in \vec{x}$, a random bit b'' is output. Thus (informally), it does not "help" the adversary to submit a query to $\mathcal{D}^{\|}_{sk}$ containing potentially invalid ciphertexts. We do not claim that our definition corresponds to any natural notion of security that one would want to achieve in practice (in fact, it is well-known that decryption oracle queries containing invalid ciphertexts can often help an adversary [7, 23]); rather, the definition is introduced merely to facilitate the proof of the theorem. On the other hand, Claims 3 and 5, below, show that this definition might also be useful as an "indistinguishability-based" characterization of non-malleability.

With the above definition in hand, our proof proceeds in the following stages:

1. We show that if Π is secure in the sense of NM-P1-CY then Π is secure in the sense of IND-P1-C$^{\|}$Y.
2. We show that if Π is secure in the sense of IND-P1-C$^{\|}$Y then Π is secure in the sense of IND-P2-C$^{\|}$Y.
3. We show that if Π is secure in the sense of IND-P2-C$^{\|}$Y then Π is secure in the sense of NM-P2-CY.

The proofs of the first and third claims, above, are similar to the proofs of the analogous claims in [6]; we note, however, that extending their proofs to the private-key setting takes some work (in particular, we must take careful account of access or lack of access to the encryption oracle). The proof of the second claim, above, is similar to the proof of Theorem 4. We begin with a proof of the first claim above.

Claim 3 *If* Π *is secure in the sense of* NM-P1-CY *then* Π *is secure in the sense of* IND-P1-C$^{\|}$Y*, for* $Y \in \{0,1,2\}$.

Proof Note that security in the sense of IND-P1-C$^{\|}$2 is equivalent to security in the sense of IND-P1-C2. Thus, Theorem 8 (below) implies the claim for the case of $Y = 2$. We therefore focus here on the case of $Y \in \{0,1\}$.

Let $A = (A_1, A_2)$ be an adversary attacking Π in the sense of IND-P1-C$^{\|}$Y. We construct the following adversary B attacking Π in the sense of NM-P1-CY:

```
$\underline{B_1^{\mathcal{E}_{sk},\mathcal{O}'_1}(1^k)}$
$(x_0, x_1, s) \leftarrow A_1^{\mathcal{E}_{sk},\mathcal{O}'_1}(1^k)$
$M = \{x_0, x_1\}$
$y'_0 \leftarrow \mathcal{E}_{sk}(x_0)$
$y'_1 \leftarrow \mathcal{E}_{sk}(x_1)$
return $(M, (y'_0, y'_1, s))$
```

```
$\underline{B_2(1^k, (y'_0, y'_1, s), y)}$
run $A_2^{\mathcal{D}_{sk}^{\|}}(1^k, s, y)$ until it submits $\vec{y}$ to $\mathcal{D}_{sk}$
let $\sigma$ be the random coins of $A_2$
let $y' \in \{y'_0, y'_1\}$ be s.t. $y' \neq y$
return $(R_{x_0,x_1,1^k,s,y,\sigma}, y'|\vec{y})$
```

(We assume without loss of generality, above, that $x_0 \neq x_1$.) Relation $R_{x_0,x_1,1^k,s,y,\sigma}$ is defined as:

```
$\underline{R_{x_0,x_1,1^k,s,y,\sigma}(x, \vec{p})}$
parse $\vec{p}$ as $x'|\vec{x}$
let $b \in \{0,1\}$ be s.t. $x = x_b$
  (if no such $b$ exists, return 0)
run $A_2(1^k, s, y)$ using coins $\sigma$
respond to its query to $\mathcal{D}_{sk}^{\|}$ with $\vec{x}$
let $b'$ be the final output of $A_2$
return 1 iff $b' = b$
```

(from now on, we simply write R for convenience). It is clear that R is efficiently computable.

In the description of adversary B, above, ciphertexts y_0', y_1' are necessary for the following technical reason: if A_2 does not submit any parallel decryption query, then $\vec{y}$ is empty; however, B_2 must output some *non-empty* vector of ciphertexts (cf. Definition 3 and the discussion there). Indeed, in the definition of R the first component of the "message vector" is simply ignored and, in particular, R is independent of its second argument if this argument contains only a single message. We remark that the choice to encrypt x_0, x_1 is entirely arbitrary; indeed, it would be simpler (in general) to have B_1 simply choose a single $x \in \mathcal{M}_k$ with $x \notin \{x_0, x_1\}$. However, we use the current formulation because we do not wish to make any assumptions about $\mathcal{M}_k$ and in particular do not wish to assume that $\mathcal{M}_k$ contains more than two messages.

Let $\mathsf{Init}_{A,\Pi}(k)$ denote the following experiment:

$$\left[\begin{array}{l} sk \leftarrow \{0,1\}^k; (x_0, x_1, s) \leftarrow A_1^{\mathcal{E}_{sk}, \mathcal{O}'_1}(1^k); b \leftarrow \{0,1\}; \\ y \leftarrow \mathcal{E}_{sk}(x_b); b' \leftarrow A_2^{\mathcal{D}^{\|}_{sk}}(1^k, s, y); \\ \text{let } \vec{y} \text{ denote the query } A_2 \text{ makes to } \mathcal{D}^{\|}_{sk} \\ \text{and let } \vec{x} \text{ denote the corresponding response} \end{array}\right].$$

By making the appropriate substitutions for B and relation R, we see that:

$$\mathsf{Expt}^{\text{NM-P1-CY}}_{B,\Pi}(k) = \Pr[\mathsf{Init}_{A,\Pi}(k) : \perp \notin \vec{x} \wedge b' = b]$$

and:

$$\begin{aligned} & \mathsf{Rand}^{\text{NM-P1-CY}}_{B,\Pi}(k) \\ & \quad = \quad \frac{1}{2} \cdot \left(\Pr[\mathsf{Init}_{A,\Pi}(k) : \perp \notin \vec{x} \wedge b' = b] + \Pr[\mathsf{Init}_{A,\Pi}(k) : \perp \notin \vec{x} \wedge b' \neq b] \right) \\ & \quad = \quad \frac{1}{2} - \frac{1}{2} \cdot \Pr[\mathsf{Init}_{A,\Pi}(k) : \perp \in \vec{x}]. \end{aligned}$$

Therefore,

$$\begin{aligned} & \mathsf{Adv}^{\text{IND-P1-C}^{\|}\text{Y}}_{A,\Pi}(k) \\ & \quad \stackrel{\text{def}}{=} \quad \left| 2 \cdot \Pr[\mathsf{Expt}^{\text{IND-P1-C}^{\|}\text{Y}}_{A,\Pi}(k) = 1] - 1 \right| \\ & \quad = \quad \left| 2 \cdot \Pr[\mathsf{Init}_{A,\Pi}(k) : \perp \notin \vec{x} \wedge b' = b] + \Pr[\mathsf{Init}_{A,\Pi}(k) : \perp \in \vec{x}] - 1 \right| \\ & \quad = \quad 2 \cdot \left| \mathsf{Expt}^{\text{NM-P1-CY}}_{B,\Pi}(k) - \mathsf{Rand}^{\text{NM-P1-CY}}_{B,\Pi}(k) \right| \\ & \quad \stackrel{\text{def}}{=} \quad 2 \cdot \mathsf{Adv}^{\text{NM-P1-CY}}_{B,\Pi}(k). \end{aligned}$$

Since Π is secure in the sense of NM-P1-CY, $\mathsf{Adv}^{\text{NM-P1-CY}}_{B,\Pi}(k)$ is negligible and hence $\mathsf{Adv}^{\text{IND-P1-C}^{\|}\text{Y}}_{A,\Pi}(k)$ is negligible. This concludes the proof of the claim. □

We now prove the second claim from above.

Claim 4 *If Π is secure in the sense of* IND-P1-C$^{\|}$Y *then Π is secure in the sense of* IND-P2-C$^{\|}$Y*, for $Y \in \{0, 1, 2\}$.*

Proof Note that security in the sense of IND-PX-C$^{\|}$2 is equivalent to security in the sense of IND-PX-C2. Thus, Theorem 4 immediately implies the claim for the case of $Y = 2$. For the case of $Y \in \{0, 1\}$ the proof is almost exactly the same as the proof of Theorem 4 and we therefore omit the details. We do, however, mention one subtlety: it will be crucial in the proof of the analogue to Claim 1 that A_2 make its query to $\mathcal{D}^{\|}_{sk}$ *after* all queries (if any) it makes to the encryption oracle. Of course, this is exactly in accordance with our definition of security in the sense of IND-PX-C$^{\|}$Y. □

We conclude with a proof of the final claim from above.

Claim 5 *If* Π *is secure in the sense of* IND-P2-C$^{\|}$Y *then* Π *is secure in the sense of* NM-P2-CY*, for* $Y \in \{0, 1, 2\}$.

Proof Note that security in the sense of IND-PX-C$^{\|}$2 is equivalent to security in the sense of IND-PX-C2. Thus, Theorem 9 (below) immediately implies the claim for the case of $Y = 2$. We therefore focus on the case of $Y \in \{0, 1\}$.

Let $A = (A_1, A_2)$ be an adversary attacking Π in the sense of NM-P2-CY. We construct the following adversary B attacking Π in the sense of IND-P2-C$^{\|}$Y:

```
$\underline{B_1^{\mathcal{E}_{sk},\mathcal{O}'_1}(1^k)}$
$(M, s) \leftarrow A_1^{\mathcal{E}_{sk},\mathcal{O}'_1}(1^k)$
$x_0, x_1 \leftarrow M$
return $(x_0, x_1, (x_0, x_1, s))$
```

```
$\underline{B_2^{\mathcal{E}_{sk},\mathcal{D}^{\|}_{sk}}(1^k, (x_0, x_1, s), y)}$
$(R, \vec{y}) \leftarrow A_2^{\mathcal{E}_{sk}}(1^k, s, y)$
$\vec{x} = \mathcal{D}^{\|}_{sk}(\vec{y})$
if $(\bot \notin \vec{x} \wedge R(x_0, \vec{x}))$
  return 0
else return 1
```

(We assume here, without loss of generality, that A_2 never outputs $\vec{y}$ with $y \in \vec{y}$.)

Let $\mathsf{Init}_{B,\Pi}(k)$ be the experiment as defined in the proof of Claim 3 (with B substituted for A). Then we have:

$$\Pr[\mathsf{Init}_{B,\Pi}(k) : b' = 0 \bigwedge \bot \notin \vec{x} | b = 0] = \mathsf{Expt}^{\text{NM-P2-CY}}_{A,\Pi}(k)$$

and:

$$\begin{aligned} &\Pr[\mathsf{Init}_{B,\Pi}(k) : b' = 1 \bigwedge \bot \notin \vec{x} | b = 1] \\ &\quad = \Pr[\mathsf{Init}_{B,\Pi}(k) : \overline{R(x_0, \vec{x})} \bigwedge \bot \notin \vec{x} | b = 1] \\ &\quad = 1 - \mathsf{Rand}^{\text{NM-P2-CY}}_{A,\Pi}(k) - \Pr[\mathsf{Init}_{B,\Pi}(k) : \bot \in \vec{x} | b = 1]. \end{aligned}$$

Noting that $\Pr[\mathsf{Init}_{B,\Pi}(k) : \bot \in \vec{x} | b = 1] = \Pr[\mathsf{Init}_{B,\Pi}(k) : \bot \in \vec{x}]$, we then have:

$$\begin{aligned} &\mathsf{Adv}^{\text{NM-P2-CY}}_{A,\Pi}(k) \\ &\quad \stackrel{\text{def}}{=} \left| \mathsf{Expt}^{\text{NM-P2-CY}}_{A,\Pi}(k) - \mathsf{Rand}^{\text{NM-P2-CY}}_{A,\Pi}(k) \right| \\ &\quad = \Big| \Pr[\mathsf{Init}_{B,\Pi}(k) : b' = 0 \bigwedge \bot \notin \vec{x} | b = 0] \\ &\qquad + \Pr[\mathsf{Init}_{B,\Pi}(k) : b' = 1 \bigwedge \bot \notin \vec{x} | b = 1] + \Pr[\mathsf{Init}_{B,\Pi}(k) : \bot \in \vec{x}] - 1 \Big| \end{aligned}$$

$$
\begin{aligned}
&= \left|2 \cdot \Pr[\mathsf{Init}_{B,\Pi}(k) : b' = b \bigwedge \bot \notin \vec{x}] + \Pr[\mathsf{Init}_{B,\Pi}(k) : \bot \in \vec{x}] - 1\right| \\
&= \mathsf{Adv}^{\text{IND-P2-C}^{\|}\text{Y}}_{B,\Pi}(k).
\end{aligned}
$$

Since Π is secure in the sense of IND-P2-C$^{\|}$Y, $\mathsf{Adv}^{\text{IND-P2-C}^{\|}\text{Y}}_{B,\Pi}(k)$ is negligible and hence $\mathsf{Adv}^{\text{NM-P2-CY}}_{A,\Pi}(k)$ is negligible. □

Claims 3–5 imply the result stated in the theorem. ■

4.3 On Non-Adaptive Access to an Encryption Oracle

We show here that non-adaptive access to an encryption oracle can sometime help an adversary break the security of an encryption scheme. Indeed, it is well known that an encryption scheme secure in the sense of IND-P0-C0 is not necessarily secure in the sense of IND-P1-C0; in particular, there exist deterministic schemes secure in the sense of IND-P0-C0 but no deterministic scheme can be secure in the sense of IND-P1-C0. Here, we show a slightly stronger result: namely, that even security in the sense of IND-P0-C2 *against an unbounded adversary* (who is limited only in the number of oracle queries he may ask) does not necessarily imply security in the sense of IND-P1-C0.

Theorem 6 (IND-P0-C2 $\not\Rightarrow$ IND-P1-C0) *There exists an encryption scheme which is secure in the sense of* IND-P0-C2 *(without any computational assumptions) but which is insecure in the sense of* IND-P1-C0*.*

Proof Let Π be the scheme given in the proof of Theorem 3. (Recall, Π encrypts message m deterministically using a one-time pad, and then applies an unconditionally-secure message-authentication code to the resulting ciphertext.) In the proof of Theorem 3 we show that Π is secure in the sense of IND-P0-C2. On the other hand, since Π is deterministic it is clear that it cannot be secure in the sense of IND-P1-C0. ■

4.4 On the Relation between Non-Malleability and Indistinguishability

In contrast to the case for public-key encryption [2], non-malleability does not necessarily imply indistinguishability in the private-key setting. In particular, the next theorem shows that there exists an encryption scheme secure in the sense of NM-P0-C2 (even against an adversary with unbounded running time) which is not secure even in the weakest sense of IND-P0-C0.

Theorem 7 (NM-P0-C2 $\not\Rightarrow$ IND-P0-C0) *There exists an encryption scheme which is secure in the sense of* NM-P0-C2 *(without any computational assumptions) but which is insecure in the sense of* IND-P0-C0*.*

Proof Our construction of the desired encryption scheme Π is simple: we simply send the plaintext in the clear but append an unconditionally-secure MAC. Decryption succeeds only if the given tag on the message is valid. In detail, let Π be the following scheme defined over message space $\{0,1\}^{\lfloor \frac{k-1}{2} \rfloor}$:

```
$\mathcal{E}_{sk}(m)$
let $|sk| = k$, and parse $sk$ as $a|b|s$,
  where $|a| = |b| = \ell \stackrel{\text{def}}{=} \lfloor \frac{k-1}{2} \rfloor$
  and $|s| \in \{0,1\}$
view $a, b, m$ as elements of $\mathbb{F}_{2^\ell}$
return $\langle m, am + b \rangle$
```

```
$\mathcal{D}_{sk}(\langle m, t \rangle)$
let $|sk| = k$, and parse $sk$ as $a|b|s$,
  where $|a| = |b| = \ell \stackrel{\text{def}}{=} \lfloor \frac{k-1}{2} \rfloor$
  and $|s| \in \{0,1\}$
view $a, b, m$ as elements of $\mathbb{F}_{2^\ell}$
if $t \stackrel{?}{=} am + b$ return $m$
else return $\perp$
```

Clearly, Π is not secure in the sense of IND-P0-C0 since the message is sent in the clear. However, we claim that Π *is* secure in the sense of NM-P0-C2 even against an adversary with unbounded running time (but who can make only polynomially-many queries to the decryption oracle). As a sketch of a proof, note that (as in the proof of Theorem 6) all decryption oracle queries are answered with $\perp$ with all but negligible probability. Furthermore, for any vector of ciphertexts $\vec{y}$ output by the adversary we will have $\perp \in \mathcal{D}_{sk}(\vec{y})$ with all but negligible probability. The advantage of any such adversary must therefore be negligible. ■

As noted following Definition 3, the above theorem depends very strongly on the exact definition of non-malleability considered here. See there for further discussion.

4.5 Completing the Picture

Theorems 4–7 are not by themselves enough to completely describe the hierarchy of security notions. To fully determine the hierarchy, we must extend results from the public-key setting [2] to the private-key setting; we do this now.

Theorem 8 (NM-PX-CY $\Rightarrow$ IND-PX-CY) *If encryption scheme* Π *is secure in the sense of* NM-PX-CY *then* Π *is secure in the sense of* IND-PX-CY, *for* $X \in \{1,2\}$ *and* $Y \in \{0,1,2\}$.

Proof This theorem is the analogue of [2, Theorem 3.1]. We stress, however, that in the private-key setting the theorem holds only if $X \neq 0$ (cf. Theorem 7).

Let A be an adversary attacking Π in the sense of IND-PX-CY. We construct an adversary $B = (B_1, B_2)$ attacking Π in the sense of NM-PX-CY as follows (if x is a string, then $\bar{x}$ denotes the ones complement of x):

```
$B_1^{\mathcal{E}_{sk}, \mathcal{O}'_1}(1^k)$
$(x_0, x_1, s) \leftarrow A_1^{\mathcal{E}_{sk}, \mathcal{O}'_1}(1^k)$
$M = \{x_0, x_1\}$
$y_0 \leftarrow \mathcal{E}_{sk}(x_0)$
$y_1 \leftarrow \mathcal{E}_{sk}(x_1)$
return $(M, (s, x_0, x_1, y_0, y_1))$
```

```
$B_2^{\mathcal{O}_2, \mathcal{O}'_2}(1^k, (s, x_0, x_1, y_0, y_1), y)$
$b \leftarrow A_2^{\mathcal{O}_2, \mathcal{O}'_2}(1^k, s, y)$
let $y' \in \{y_0, y_1\}$ be s.t. $y' \neq y$
return $(R, y')$
  where $R(x, x') = 1$ iff $x = x_b$
```

It is not hard to see that $\mathsf{Rand}_{B,\Pi}^{\text{NM-PX-CY}}(k) = 1/2$ so that

$$\mathsf{Adv}_{B,\Pi}^{\text{NM-PX-CY}}(k) = \frac{1}{2} \cdot \mathsf{Adv}_{A,\Pi}^{\text{IND-CX-PY}}(k).$$

Since Π is secure in the sense of NM-PX-CY, $\mathsf{Adv}_{A,\Pi}^{\text{IND-CX-PY}}(k)$ is negligible and the theorem follows. ■

Theorem 9 (IND-PX-C2 $\Rightarrow$ NM-PX-C2) *If encryption scheme Π is secure in the sense of* IND-PX-C2 *then Π is secure in the sense of* NM-PX-C2*, for $X \in \{0,1,2\}$.*

Proof This theorem is the exact counterpart of [2, Theorem 3.3], and we repeat essentially the same proof here for completeness. Let A be an adversary attacking Π in the sense of NM-PX-C2. We define an adversary B attacking Π in the sense of IND-PX-C2 as follows:

```
$B_1^{\mathcal{O}_1,\mathcal{D}_{sk}}(1^k)$
$(M,s) \leftarrow A_1^{\mathcal{O}_1,\mathcal{D}_{sk}}(1^k)$
$x_0, x_1 \leftarrow M$
return $(x_0, x_1, (x_0, s))$
```

```
$B_2^{\mathcal{O}_2,\mathcal{D}_{sk}}(1^k, (x_0, s), y)$
$(R, \vec{y}) \leftarrow A_2^{\mathcal{O}_2,\mathcal{D}_{sk}}(1^k, s, y)$
$\vec{x} = \mathcal{D}_{sk}(\vec{y})$
if $(y \notin \vec{y} \wedge \perp \notin \vec{x} \wedge R(x_0, \vec{x}))$
  return 0
else return 1
```

We may note that the probability that B returns 0 given that y is an encryption of x_0 is exactly $\mathsf{Expt}_{A,\Pi}^{\text{NM-PX-C2}}(k)$ while the probability that B returns 0 given that y is an encryption of x_1 is exactly $\mathsf{Rand}_{A,\Pi}^{\text{NM-PX-C2}}(k)$. Thus,

$$\mathsf{Adv}_{B,\Pi}^{\text{IND-PX-C2}}(k) = \mathsf{Adv}_{A,\Pi}^{\text{NM-PX-C2}}(k)$$

and hence $\mathsf{Adv}_{A,\Pi}^{\text{NM-PX-C2}}(k)$ is negligible. ■

Theorem 10 (IND-P2-C1 $\not\Rightarrow$ NM-P0-C0) *Assuming the existence of a one-way function, there exists an encryption scheme Π which is secure in the sense of* IND-P2-C1 *but which is insecure in the sense of* NM-P0-C0*.*

Proof This theorem is the analogue of [2, Theorem 3.5]. We give here an explicit construction of a scheme which is secure in the sense of IND-P2-C1 (assuming the existence of a one-way function[5]) but insecure in the sense of NM-P0-C0.

Let $\mathcal{F} = \{F^k\}_{k \geq 1}$ be a pseudorandom function family [14] where $F^k = \{F_s : \{0,1\}^k \to \{0,1\}^k\}_{s \in \{0,1\}^k}$ is a finite collection of functions indexed by a key $s \in \{0,1\}^k$; we note that such $\mathcal{F}$ may be constructed assuming one-way functions exist [14, 20]. One may then define [14] the following encryption scheme Π over message space $\mathcal{M}_k = \{0,1\}^k$:

```
$\mathcal{E}_{sk}(m)$
let $|sk| = k$
$r \leftarrow \{0,1\}^k$
return $\langle r, F_{sk}(r) \oplus m \rangle$
```

```
$\mathcal{D}_{sk}(\langle r, c \rangle)$
return $c \oplus F_{sk}(r)$
```

It is easy to verify that this scheme is indeed secure in the sense of IND-P2-C1. Informally, let r be the first half of the challenge ciphertext and let Used denote the event that this same

[5] We note that our assumption is minimal, since the existence of an encryption scheme secure even in the sense of IND-P1-C0 implies the existence of a one-way function [21]. See Section 3.

value is either used (by the encryption oracle) in one of the adversary's queries to $\mathcal{E}_{sk}(\cdot)$ or submitted (by the adversary) in one of the adversary's *non-adaptive* queries to $\mathcal{D}_{sk}(\cdot)$. Clearly, the probability that event Used occurs is negligible. Furthermore, conditioned on the event that Used does not occur, $F_{sk}(r)$ "looks random" to the adversary (by security of the pseudorandom function family) and hence the adversary has negligible advantage in this case.

It is even easier to see that Π is insecure in the sense of NM-P0-C0. To wit, consider the adversary (A_1, A_2) in which $A_1(1^k)$ outputs $M = \{0^k, 1^k\}$ and A_2 — given ciphertext $\langle r, c\rangle$ — outputs $\langle r, c \oplus 1^k\rangle$ along with relation R for which $R(x, y)$ is true iff $x = \bar{y}$. For this adversary, we have $\mathsf{Expt}^{\text{NM-PX-CY}}_{A,\Pi}(k) = 1$ while $\mathsf{Rand}^{\text{NM-PX-CY}}_{A,\Pi}(k) = 1/2$, and thus the advantage of A is not negligible. ■

Theorem 11 (NM-P2-C0 $\not\Rightarrow$ IND-P0-C1, NM-P0-C1) *Assuming the existence of a one-way function, there exists an encryption scheme* Π *which is secure in the sense of* NM-P2-C0 *but which is insecure in the sense of* IND-P0-C1 *and also insecure in the sense of* NM-P0-C1*.*

Proof This is the analogue of [2, Theorem 3.6], although we give a different proof here. We give an explicit construction of a scheme secure in the sense of NM-P2-C0 (assuming the existence of a one-way function; cf. footnote 5 and Theorem 8) which is insecure both in the sense of IND-P0-C1 and in the sense of NM-P0-C1.

Let (MAC, Vrfy) be a secure message authentication code (see [3] for a definition) for which key generation simply consists of choosing a uniformly distributed key of the appropriate length. Furthermore, for concreteness, assume that for keys of length k the scheme authenticates messages of length $2k$. We note that such a scheme may be constructed based on any one-way function [14, 15]. Let $\mathcal{F}$ be a pseudorandom function family as in the proof of Theorem 10. We then define the following encryption scheme Π over message space $\mathcal{M}_k = \{0,1\}^{k/3}$ (for simplicity, we define the scheme only for k a multiple of 3, but it is easy to make the necessary modifications to allow arbitrary k):

```
$\underline{\mathcal{E}_{sk}(m)}$
let $|sk| = k$ and parse $sk$ as $s_1|s_2|v$,
  where $|s_1| = |s_2| = |v| = \frac{k}{3}$
$r \leftarrow \{0,1\}^{k/3}; c = F_{sk_1}(r) \oplus m$
$t \leftarrow \text{MAC}_{s_2}(r|c)$
if $m \overset{?}{=} v$
  return $\langle 1, m, m, t\rangle$
else return $\langle 0, r, c, t\rangle$
```

```
$\underline{\mathcal{D}_{sk}(\langle b, r, c, t\rangle)}$
let $|sk| = k$ and parse $sk$ as $s_1|s_2|v$,
  where $|s_1| = |s_2| = |v| = \frac{k}{3}$
if $b = 1$ return $v$
if $b = 0$ and $\mathsf{Vrfy}_{s_2}(r|c, t) = 1$
  return $F_{sk_1}(r) \oplus c$
otherwise return $\perp$
```

We now examine the security properties of this scheme.

Claim 6 Π *is secure in the sense of* NM-P2-C0*.*

Proof It is not difficult to see that Π is secure in the sense of IND-P2-C0 (an adversary's advantage is negligible unless one of (x_0, x_1) is equal to v, and this occurs with negligible probability). We use this to show that the scheme is secure in the sense of NM-P2-C0. Let A be an adversary attacking Π in the sense of NM-P2-C0; we construct an adversary B attacking Π in the sense of IND-P2-C0.

Adversary B is constructed much the same as in the proof of Theorem 9; we sketch only the differences here. The proof of Theorem 9 only requires that B_2 can decrypt the vector of ciphertexts $\vec{y}$ output by A_2 at the end of the second stage. There, this was easily done since B_2 had access to the decryption oracle. Here, of course, B_2 does not have access to a decryption oracle; however, we show how B_2 can simulate decryption anyway. Let Find denote the event that a ciphertext $\langle 0, r, c, t\rangle$ is ever received in response from the encryption oracle, or as the challenge ciphertext. If Find ever occurs (whether in the first or second stage), B_2 simply outputs a random bit. Otherwise, B_2 simulates decryption of a given ciphertext $\langle b, r, c, t\rangle \in \vec{y}$ as follows:

- If $\langle b, r, c, t\rangle$ was previously output by the encryption oracle on query m, let the message be m.
- If $b = 0$, let the message be a random $v' \in \{0,1\}^{k/3}$ (once such v' is chosen, it is used as the decryption of any other ciphertexts in $\vec{y}$ with $b = 0$).
- Otherwise, let the message be $\perp$.

Assuming Find does not occur, this simulation of the decryption of $\vec{y}$ is perfect unless A was able to forge a tag t on a new "message" $r|c$; this happens with only negligible probability. Finally, noting that the probability that event Find occurs is negligible, we conclude that $\mathsf{Adv}^{\text{IND-P2-C0}}_{B,\Pi}(k)$ is negligible close to $\mathsf{Adv}^{\text{NM-P2-C0}}_{A,\Pi}(k)$ and hence Π is secure in the sense of NM-P2-C0. □

It is easy to show, however, that Π is insecure under a non-adaptive chosen-ciphertext attack.

Claim 7 *Π is insecure in the senses of* IND-P0-C1 *and* NM-P0-C1.

Proof We show an adversary for the case of NM-P0-C1; the case of IND-P0-C1 is even simpler. A_1 submits query $\tilde{y} = \langle 1, r, \bar{r}, t\rangle$ to the decryption oracle and receives in return v. Then, A_1 outputs $M = \{v, \bar{v}\}$. When A_2 gets the challenge ciphertext $y = \langle b, r', c', t'\rangle$ it does the following:

- If $b = 1$, it outputs $\tilde{y}$ and relation R such that $R(x, x')$ is true iff $x' = x'$.
- If $b = 0$, it outputs $\tilde{y}$ and relation R such that $R(x, x')$ is true iff $x' = \bar{x}$.

It should be clear that $\mathsf{Expt}^{\text{NM-P0-C1}}_{A,\Pi}(k) = 1$ while $\mathsf{Rand}^{\text{NM-P0-C1}}_{A,\Pi}(k) = 1/2$, and hence Π is not secure in the sense of NM-P0-C1. □

This completes the proof of the theorem. ■

Theorem 12 (NM-P2-C1 $\not\Rightarrow$ NM-P0-C2) *Assuming the existence of a one-way function, there exists an encryption scheme Π which is secure in the sense of* NM-P2-C1 *but which is insecure in the sense of* NM-P0-C2.

Proof This theorem is the exact counterpart of [2, Theorem 3.7]; we therefore only sketch a proof here and refer the reader to [2] for details. Let $\Pi = (\mathcal{E}, \mathcal{D})$ be a private-key encryption scheme over message space $\{0,1\}^k$ which is secure in the sense of NM-P2-C1; assume furthermore — for convenience — that ciphertexts are strings of length $2k$ (such schemes may be constructed assuming the existence of one-way functions; see [10]). Let $\mathcal{F} = \{F^k\}_{k \geq 1}$ be a pseudorandom function family where $F^k = \{F_s : \{0,1\}^{2k} \to \{0,1\}^k\}_{s \in \{0,1\}^k}$ is a finite collection of functions indexed by a key $s \in \{0,1\}^k$; we note that such $\mathcal{F}$ may be constructed assuming one-way functions exist [14, 20]. We construct scheme $\Pi' = (\mathcal{E}', \mathcal{D}')$ over message space $\{0,1\}^{k/2}$ as follows (for simplicity, we define Π' only for k even, but it is easy to make the necessary changes to allow arbitrary k):

```
$\underline{\mathcal{E}'_{sk}(m)}$
let $|sk| = k$ and parse $sk$ as $s_1|s_2$,
  where $|s_1| = |s_2| = \frac{k}{2}$
$c \leftarrow \mathcal{E}_{s_1}(m)$
return $\langle 0, c, \varepsilon \rangle$
```

```
$\underline{\mathcal{D}'_{sk}(\langle b, c, z \rangle)}$
let $|sk| = k$ and parse $sk$ as $s_1|s_2$,
  where $|s_1| = |s_2| = \frac{k}{2}$
if $(b = 0) \wedge (z = \varepsilon)$ return $\mathcal{D}_{s_1}(c)$
if $(b = 1) \wedge (z = \varepsilon)$ return $F_{s_2}(c)$
if $(b = 1) \wedge (z = F_{s_2}(c))$ return $\mathcal{D}_{s_1}(c)$
otherwise return $\perp$
```

It is easy to see that Π' is not secure in the sense of NM-P0-C2. Namely, consider the adversary (A_1, A_2) for which $A_1(1^k)$ outputs $M = \{x, x'\}$, where x, x' are any distinct messages in $\{0,1\}^{k/2}$. On input challenge ciphertext $\langle 0, c, \varepsilon \rangle$, A_2 submits $\langle 1, c, \varepsilon \rangle$ to its decryption oracle and receives in return a value z such that $z = F_{s_2}(c)$. The final output of A_2 is ciphertext $\langle 1, c, z \rangle$ along with the equality relation. Clearly, the advantage of this adversary is not negligible.

A proof of the following claim exactly follows the proof of [2, Claim 3.15], and we therefore do not repeat it here.

Claim 8 Π' *is secure in the sense of* NM-P2-C1.

This concludes the proof of the theorem. ■

4.6 Obtaining Figure 1

At this point, the reader may — somewhat tediously — convince him- or herself that Theorems 4–12 indeed yield the characterization of Figure 1. To ease this process, we note the following:

- Theorems 4 and 5 show that XXX-P1-CY is equivalent to XXX-P2-CY for either notion of security XXX and any level of chosen-ciphertext attack Y. This generates most of the equivalences (security notions boxed together) in Figure 1. The equivalence between NM-P2-C2 and IND-P2-C2, and the implications of this equivalence, follow from Theorems 8 and 9.

- Theorem 6 indicates that XXX-P1-CY is strictly stronger than XXX-P0-CY for either notion of security XXX and any level of chosen-ciphertext attack Y. Similarly, Theorems 11 and 12 indicate that XXX-PY-C2 is strictly stronger than XXX-PY-C1 which is, in turn, strictly stronger than XXX-PY-C0.

Acknowledgments

We thank an anonymous referee for careful editing and helpful comments.

References

[1] M. Bellare, A. Desai, E. Jokipii, and P. Rogaway. A concrete security treatment of symmetric encryption: analysis of the DES modes of operation. *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, IEEE (1997), pp. 394–403.

[2] M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among notions of security for public-key encryption schemes. *Advances in Cryptology — Crypto ’98*, Lecture Notes in Computer Science, Vol. 1462, H. Krawczyk, ed., Springer-Verlag (1998), pp. 26–45.

[3] M. Bellare, J. Kilian, and P. Rogaway. The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences*, vol. 61, no. 3 (2000), pp. 362–399.

[4] M. Bellare, T. Kohno, and C. Namprempre. Authenticated encryption in SSH: provably fixing the SSH binary packet protocol. *Proceedings of the 9th ACM Conference on Computer and Communications Security*, ACM (2002), pp. 1-11.

[5] M. Bellare and C. Namprempre. Authenticated encryption: relations among notions and analysis of the generic composition paradigm. *Advances in Cryptology — Asiacrypt 2000*, Lecture Notes in Computer Science, vol. 1976, T. Okamoto, ed., Springer-Verlag (2000), pp. 531–545.

[6] M. Bellare and A. Sahai. Non-malleable encryption: equivalence between two notions, and an indistinguishability-based characterization. *Advances in Cryptology — Crypto ’99*, Lecture Notes in Computer Science, vol. 1666, M. Wiener, ed., Springer-Verlag (1999), pp. 519–536.

[7] D. Bleichenbacher. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS#1. *Advances in Cryptology — Crypto ’98*, Lecture Notes in Computer Science, vol. 1462, H. Krawczyk, ed., Springer-Verlag (1998), pp. 1–12.

[8] R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. *Advances in Cryptology — Eurocrypt 2001*, Lecture Notes in Computer Science, vol. 2045, B. Pfitzmann, ed., Springer-Verlag (2001), pp. 453–474.

[9] R. Canetti and H. Krawczyk. Universally composable notions of key exchange and secure channels. *Advances in Cryptology — Eurocrypt 2002*, Lecture Notes in Computer Science, vol. 2332, L.R. Knudsen, ed., Springer-Verlag (2002), pp. 337–351.

[10] D. Dolev, C. Dwork, and M. Naor. Nonmalleable cryptography. *SIAM Journal of Computing*, vol. 30, no. 2 (2000), pp. 391–437.

[11] O. Goldreich. A uniform complexity treatment of encryption and zero-knowledge. *Journal of Cryptology*, vol. 6, no. 1 (1993), pp. 21–53.

[12] O. Goldreich. *Foundations of cryptography: basic tools.* Cambridge University Press, Cambridge, UK, 2001.

[13] O. Goldreich. Draft of a chapter on encryption schemes. Fourth posted version. December 12, 2002. Available at http://www.wisdom.weizmann.ac.il/˜oded/foc-vol2.html.

[14] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, vol. 33, no. 4 (1984), pp. 792–807.

[15] O. Goldreich, S. Goldwasser, and S. Micali. On the cryptographic applications of random functions. *Advances in Cryptology — Crypto '84*, Lecture Notes in Computer Science, vol. 196, G.R. Blakley and D. Chaum, eds., Springer-Verlag (1984), pp. 276–288.

[16] O. Goldreich, Y. Lustig, and M. Naor. On chosen ciphertext security of multiple encryptions. Cryptology ePrint Archive: Report 2002/089 (2002). Available at http://eprint.iacr.org/2002/089.

[17] S. Goldwasser and M. Bellare. Lecture notes on cryptography. August 2001. Available at http://www-cse.ucsd.edu/users/mihir/papers/gb.html.

[18] S. Goldwasser and S. Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, vol. 28, no. 2 (1984), pp. 270–299.

[19] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen message attacks. *SIAM Journal of Computing*, vol. 17, no. 2 (1988), pp. 281–308.

[20] J. Håstad, R. Impagliazzo, L. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal of Computing*, vol. 28, no. 4 (1999), pp. 1364–1396.

[21] R. Impagliazzo and M. Luby. One-way functions are essential for complexity-based cryptography. *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, IEEE (1989), pp. 230-235.

[22] A. Joux, G. Martinet, and F. Valette. Blockwise-adaptive attackers: revisiting the (in)security of some provably secure encryption modes: CMC, GEM, IACBC. *Advances in Cryptology — Crypto 2002*, Lecture Notes in Computer Science, M. Yung, ed., Springer-Verlag (2002), pp. 17–30.

[23] M. Joye, J.-J. Quisquater, and M. Yung. On the power of misbehaving adversaries and security analysis of the original EPOC. *Topics in Cryptology — CT-RSA 2001*, Lecture Notes in Computer Science, vol. 2020, D. Naccache, ed., Springer-Verlag (2001), pp. 208–222.

[24] J. Katz and M. Yung. Complete characterization of security notions for probabilistic private-key encryption. *Proceedings of the 32nd Annual Symposium on Theory of Computing*, ACM (2000), pp. 245–254.

[25] J. Katz and M. Yung. Unforgeable encryption and chosen ciphertext secure modes of operation. *Fast Software Encryption*, Lecture Notes in Computer Science, vol. 1978, B. Schneier, ed., Springer-Verlag (2000), pp. 284–299.

[26] H. Krawczyk. The order of encryption and authentication for protecting communications (or: how secure is SSL?). *Advances in Cryptology — Crypto 2001*, Lecture Notes in Computer Science, vol. 2139, J. Kilian, ed., Springer-Verlag (2001), pp. 310–331.

[27] M. Luby. *Pseudorandomness and cryptographic applications.* Princeton University Press, Princeton, NJ, 1996.

[28] S. Micali, C. Rackoff, and R. Sloan. The notion of security for probabilistic cryptosystems. *SIAM Journal of Computing*, vol. 17, no. 2 (1988), pp. 412–426.

[29] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, vol. 4, no. 2 (1991), pp. 151–158.

[30] C. Namprempre. Secure channels based on authenticated encryption schemes: a simple characterization. *Advances in Cryptology — Asiacrypt 2002*, Lecture Notes in Computer Science, vol. 2501, Y. Zheng, ed., Springer-Verlag (2002), pp. 515–532.

[31] M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen-ciphertext attacks. *Proceedings of the 22nd Annual Symposium on Theory of Computing*, ACM (1990), pp. 427–437.

[32] C. Rackoff and D. Simon. Non-interactive zero-knowledge proof of knowledge and chosen-ciphertext attack. *Advances in Cryptology — Crypto '91*, Lecture Notes in Computer Science, vol. 576, J. Feigenbaum, ed., Springer-Verlag (1991), pp. 433–444.

[33] J. Rompel. One-way functions are necessary and sufficient for secure signatures. *Proceedings of the 22nd Annual Symposium on Theory of Computing*, ACM (1990), pp. 387–394.

[34] C.E. Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, vol. 28 (1949), pp. 656–715.

[35] G.S. Vernam. Cipher printing telegraph systems for secret wire and radio telegraphic communications. *Journal of the American Institute of Electrical Engineers*, vol. 45 (1926), pp. 109–115. See also US patent #1,310,719.

[36] Y. Watanabe, J. Shikata, and H. Imai. Equivalence between semantic security and indistinguishability under chosen ciphertext attacks. *Public Key Cryptography — PKC 2003*, Lecture Notes in Computer Science, vol. 2567, Y. Desmedt, ed., Springer-Verlag (2003), pp. 71–84.

[37] A. Yao. Theory and applications of trapdoor functions. *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, ACM (1982), pp. 80–91.