

Relations Among Notions of Non-malleability for Encryption

Rafael Pass¹, Abhi Shelat², and Vinod Vaikuntanathan³

¹ Cornell

² U. Virginia

³ MIT

Abstract. Since its introduction in the early 90’s, the notion of non-malleability for encryption schemes has been formalized using a number of conceptually different definitional approaches—most notably, the “pragmatic” indistinguishability-based approach and the “semantical” simulation-based approach. We provide a full characterization of these approaches and consider their robustness under composition.

Keywords: Public-key Encryption, Non-malleability.

1 Introduction

The basic goal of an encryption scheme is to guarantee the *privacy* of data. A good formalization of privacy is the notion of *semantic security* as defined by Goldwasser and Micali [GM84]. Intuitively, semantic security guarantees that “whatever a polynomial-time machine can learn about a message given its encryption, it can learn even without seeing the encryption.”

When encryption schemes are deployed in more complex environments, the demands for security of encryption grow beyond just the basic privacy requirement. Motivated by practical security requirements, the seminal work of Dolev, Dwork and Naor [DDN00] defined the notion of *non-malleability*—a qualitatively stronger notion of security for encryption schemes. In addition to the normal “privacy” guarantee, non-malleability ensures that it is infeasible for an adversary to *modify* a vector of ciphertexts $\alpha_1, \dots, \alpha_n$ into other ciphertexts of messages which are related to the decryption of $\alpha_1, \dots, \alpha_n$. This stronger notion of security is critical for many practical applications.

Two Formalizations. The notion of non-malleability for encryption schemes has been formalized using two different approaches:

- **The “Semantical” Simulation-based Approach.** The definition presented in the original work of [DDN00] is a so-called “simulation-based” one. The main idea is to capture the requirement that an adversary having access to ciphertexts (and potentially a decryption oracle in case of CCA1/CCA2 attacks), will not be able to “cause more harm” than a simple adversary

that does not see any ciphertexts and does not have access to a decryption oracle. This simulation-based definition of non-malleability is denoted SIM-NME, and like semantic security, the goal of this definition is to capture the “meaning” of non-malleability. As a result, it is often harder to directly prove that a scheme meets the simulation-based definition.

- **The “Pragmatic” Indistinguishability-based Approach.** Bellare et.al. present a “comparison-based” formalization of non-malleability [BDPR98]. This notion does away with the “simulator” used in [DDN00] and instead captures non-malleability through an indistinguishability-style definition. Other indistinguishability-based definitions appear in [BS99, PSV06]. We denote by IND-NME the indistinguishability-based approach to defining non-malleability. The goal of this indistinguishability-based approach is to provide definitions that are easier to “work with.”

Just as Goldwasser and Micali [GM84] show equivalence between simulation-based and indistinguishability-based definitions of secrecy, Bellare and Sahai [BS06] (clarifying [BS99]) show an equivalence between the simulation-based and the indistinguishability-based approach to defining non-malleability. As we discuss later, their proof however makes certain implicit assumptions on the type of encryption schemes used. As far as we know, equivalences for general encryption schemes are not known.

Composition and Invalid Ciphertexts. In practice, encryption schemes must guarantee security also when an adversary receives encryptions of *multiple* messages. It is well known that for the traditional definition of secrecy, “single-message” security implies “multi-message” security – we say that the definition is *closed under composition*. It would be desirable to have a definition of non-malleability that composes (i.e., for which non-malleability for a single message implies non-malleability for multiple messages).

It turns out that this property is highly sensitive to the way non-malleability is formalized. As pointed out by Pass, shelat and Vaikuntanathan [PSV06], there is some ambiguity in the original work of Dolev, Dwork and Naor [DDN00] about how to treat an adversary that sometimes produces *invalid* ciphertexts as part of its output. Whereas the intuitive description of the “spirit” of non-malleability considers an adversary successful if it is able to output ciphertexts that are related to the ciphertexts it receives, the formal definition does not consider an adversary who outputs an invalid ciphertext (even if this event is correlated with the input ciphertexts it receives). It is shown in [PSV06] that for the case of *chosen-plaintext* attacks, this (seemingly minor) issue becomes critical in certain (traditional) applications, and is also essential for proving composability of non-malleability. In both situations a stronger definition, which does not automatically fail an adversary which outputs an invalid ciphertext, is sufficient, whereas the weaker (traditional one) is not. We denote by SIM-NME', IND-NME' these stronger variants of SIM-NME, IND-NME (which are in-line with the definitions of [PSV06, BS06]).

1.1 Our Results

We may thus broadly categorize definitions of non-malleability into two major groups: “simulation-based” and “indistinguishability-based,” and each with two sub-groups: “invalid-allowing” and “invalid-prohibiting.” In this paper we first fully characterize the relationship among the different definitional approaches outlined above. Secondly, we consider the robustness of each of the definitions under a natural (and highly desirable) notion of composition. Our motivation is to clarify the definitional imbroglia surrounding the notions. To so do, we present a unified way of defining non-malleability according the above-mentioned different approaches. We furthermore believe that our definitions provide the simplest and cleanest way to formalize non-malleability according to these approaches.

Relations Between Definitions. Our results are as follows.

1. *The Case of Invalid-Allowing Definitions* For the case of invalid-allowing definitions, we obtain a separation between the simulation-based definition of non-malleability, SIM-NME', and indistinguishability-based definition, IND-NME'. In particular, under CCA1 or CCA2 attacks, SIM-NME' is *strictly* stronger than IND-NME', whereas under CPA attacks they are equivalent.
2. *The Case of Invalid-Prohibiting Definitions* For the case of invalid-prohibiting definitions, the simulation-based definition, SIM-NME is *equivalent* to the indistinguishability-based definition IND-NME, under all attacks (i.e., CPA, CCA1 and CCA2).
3. *The Relation between Invalid-Allowing and -Prohibiting Definitions* The first approach to defining non-malleability is *strictly* stronger than the second one. In fact, this holds under all attacks in the simulation-based notion, and under CCA1 and CPA attacks for the indistinguishability-based notion.

A full characterization of the different definitions is summarized in the table below. The starred results appear in either [DDN00] and/or [BDPR98].

ATTACK	RELATIONSHIPS
CCA2	SIM-NME' > IND-NME' = SIM-NME =* IND-NME =* IND
CCA1	SIM-NME' > IND-NME' > SIM-NME = IND-NME >* IND
CPA	SIM-NME' = IND-NME' > SIM-NME = IND-NME >* IND

Results Concerning Practical Schemes and Restricted Message Spaces. Many practical and efficient encryption schemes only work for *restricted message spaces*. For example, the El Gamal and Cramer-Shoup schemes work for messages that are elements of some finite group. While it seems natural for the above equivalences to also hold for this special class of encryption schemes, we show in §5 that this intuition is not true. In particular, we show that also for the case of CCA2 attacks, SIM-NME is strictly stronger than IND-NME. Thus, somewhat surprisingly,

For restricted message spaces, “simple” IND-CCA2 security does not imply the original semantical (simulation-based) definition of non-malleability.

This stands in sharp contrast to the result of [DDN00, BDPR98] showing that IND-CCA2 indeed is equivalent SIM-NME for the case of full messages spaces.

Why Simulation-Based Non-malleability Is Desirable. Many practical system attacks such as buffer overflows rely on creating a situation in which a process is fed unexpected input. With this in mind, consider an encryption scheme which has been dutifully designed so that an adversary cannot produce a ciphertexts which decrypt to a certain output value (say \perp). A system designer might employ this scheme in a process, and rely on the fact that such inputs cannot be produced by the decrypting algorithm for the correctness of the process.

Now suppose that the adversary might have a way to implement a CCA2 attack. A cryptographer may be content to prove that their encryption scheme is IND-CCA2-secure. However, the systems’ practitioner may require something more. She would like the guarantee that even if the adversary has a decryption oracle, the adversary will be unable to “do any more harm” than if the adversary did not have the decryption oracle. In other words, the adversary will be unable to produce unexpected outputs in this case as well—and so the practitioner’s original assumptions are still valid. In essence, the situation calls for simulation-based security.

Remarks. As shown by Canetti [Can01], a Universally Composable (UC) implementation of an “idealized” encryption functionality \mathcal{F}_{pke} is equivalent to IND-CCA2-secure encryption. Furthermore, the UC definition of security is a semantical notion which provides security under arbitrary concurrent executions; in particular UC security provides security with respect to man-in-the-middle attacks. However, the definition of \mathcal{F}_{pke} allows a corrupted sender to make an honest receiver decrypt a ciphertext to any arbitrary string (and not only those in the domain of the encryption function) *even if this was not possible in a stand-alone setting*; as such UC encryption does not satisfy the above desiderata. We also mention that Goldreich [Gol04] presents a similar semantical (simulation-based) definition of non-malleability, which is equivalent to (simple) indistinguishability under CCA2 attacks; this definition too does not prevent a corrupted sender from making an honest receiver decrypt a ciphertext to any arbitrary string.¹

Additional Equivalences. To further clarify the semantical relation between the various notions, we present additional equivalences for certain restricted encryption schemes: Concisely, a scheme which is IND-NME secure and for which it is possible to efficiently produce a ciphertext which decrypts to every output in

¹ On a high-level, the difference between SIM-NME and the definition of [Gol04] is that in the latter, the simulator is required to output *plaintexts* that are indistinguishable from the messages the adversary encrypts, whereas in the former the simulator must do the same as the adversary and output ciphertexts.

the range of the decryption function is also (multi-message) SIM-NME' secure.² Thus, for encryption schemes satisfying certain technical conditions all the above notions are equivalent. In light of this our separation results might seem “artificial”.³ Note, however, that although these restriction are not implausible, they are far from being satisfied all “practical” encryption schemes. Indeed, whereas RSA-OAEP satisfies them (at the cost of “truncating” the message space), other schemes such as CS1 from [CS98] does not.

Composability of Definitions. The table below summarizes new and known results regarding the composability of of the various definitional approaches. A \checkmark -mark indicates that the definitions composes, **X**-mark indicates it does not, and ? indicates that the status is unknown. Pass, shelat, and Vaikuntanathan [PSV06] show the * result. Gennaro and Lindell [GL03] show the \dagger result. All other results are new in this paper. These new results show that, contrary to folklore belief, *indistinguishability*-based definitions of encryption do *not* necessarily compose in the context of non-malleability.

	SIM-NME'	IND-NME'	SIM-NME	IND-NME
CCA2	?	\checkmark	\checkmark	\checkmark
CCA1	?	\checkmark	X	X
CPA	\checkmark	\checkmark^*	X [†]	X

Related Work. The work of [BS06] (clarifying the original work of [BS99]) provided a comprehensive study of equivalence between indistinguishability-based and simulation-based definitions. Their main results show such an equivalence for the case of *invalid-allowing* definitions. We here note that their result implicitly makes the assumption that the encryption schemes considered have the property that it is “easy” (i.e., there is a prescribed polynomial-time algorithm) to generate invalid ciphertexts. In contrast, we consider general encryption schemes (i.e., without any such restriction). Interestingly, we show that the notions no longer are equivalent when doing so (furthermore, when considering restricted messages spaces, equivalence does not hold even if there exists a prescribed polynomial-time algorithm for generating invalid ciphertexts).

Nevertheless, we emphasize that proof techniques from [BS99, BS06] are useful also when considering general encryption schemes. Indeed, our equivalence proof for the case of *invalid-prohibiting* definitions (i.e., showing that SIM-NME = IND-NME) borrows from their original proof.⁴

² This result generalizes the earlier results by [BS06]. See Section 1.1 for more details.

³ In a sense all separation results can be called either “artificial” or “trivial”—if they are satisfied by known schemes then they are trivial, otherwise they are “artificial”.

⁴ The original published version of their results [BS99] claimed an equivalence between SIM-NME and an indistinguishability-based definition of non-malleability due to [BDPR98]. This claim was later retracted in the new version [BS06] (due to subtleties pointed out by Lindell). We mention, however, that our definition of IND-NME is (seemingly) different from the indistinguishability-based definition of [BDPR98].

We also mention that various other definitions of non-malleability for encryption schemes have been proposed (e.g [BDPR98, BS06, Gol04]). Our goal is not to fully characterize the relative strength of all variants of non-malleability. Rather, we highlight the differences between certain natural definitional approaches (i.e., simulation v.s. indistinguishability, and invalid-allowing v.s. invalid-prohibiting).

2 Definitions

Oracles. In a chosen-plaintext attack (CPA), the oracles O_1, O_2 return the empty string. In a CCA1 attack, the oracle $O_1(\text{PK}, \cdot)$ returns the decryption of its input under public key PK (which is implicit by context). Finally, in a CCA2 attack, both oracles return decryptions with the exception that $O_2(\text{PK}, \mathbf{y}, \cdot)$ returns \perp when queried on a ciphertext contained in \mathbf{y} .

Comparing Definitions. If D1, D2 are two definitions, the notation $D1 > D2$ means that: “Every scheme Π which satisfies D1 also satisfies D2, and if there exists a scheme Π which satisfies D2, then there exists a scheme Π' which also satisfies D2 but does not satisfy D1.” We say that $D1 = D2$ if the set of schemes that satisfy D1 is identical to the set of schemes that satisfy D2.

2.1 Simulation-Based Definitions of Non-malleable Encryption

Definition 1 (SIM-NME' Security). *Define the following two experiments.*

$\text{SIM-NME}'(\Pi, A, k, \ell, r)$ $(\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^k)$ $(M, s) \leftarrow A_1^{O_1}(\text{PK})$ $(m_1, \dots, m_\ell) \stackrel{s}{\leftarrow} M(1^k)$ $\mathbf{y} \leftarrow \text{Enc}(\text{PK}, \mathbf{m})$ $((c_1, \dots, c_r), \sigma) \leftarrow A_2^{O_2}(\mathbf{y}, h(\mathbf{m}), s)$ $d_i = \begin{cases} \text{COPY} & \text{if } c_i \in \mathbf{y} \\ \text{Dec}(\text{SK}, c_i) & \text{o.w.} \end{cases}$ $\text{Output } (M, \mathbf{m}, (d_1, \dots, d_r), \sigma)$	$\overline{\text{SIM-NME}'(\Pi, S, k, \ell, r)}$ $(\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^k)$ <div style="border: 1px solid black; padding: 2px; display: inline-block; margin: 2px;"> $(M, s) \leftarrow S_1(\text{PK})$ </div> $(m_1, \dots, m_\ell) \stackrel{s}{\leftarrow} M(1^k)$ <div style="border: 1px solid black; padding: 2px; display: inline-block; margin: 2px;"> $((c_1, \dots, c_r), \sigma) \leftarrow S_2(h(\mathbf{m}), s)$ </div> $d_i = \begin{cases} \text{COPY} & \text{if } c_i = \text{COPY} \\ \text{Dec}(\text{SK}, c_i) & \text{o.w.} \end{cases}$ $\text{Output } (M, \mathbf{m}, (d_1, \dots, d_r), \sigma)$
--	--

Here M is a Turing machine that samples a vector of $\ell(k)$ messages from a distribution. We say that M is an (p, ℓ) -valid message-sampler if 1) the running-time of $M(1^k)$ is bounded by $p(k)$, and 2) there exists polynomials l_1, l_2, \dots, l_ℓ such that $M(1^k)$ always outputs message sequences $(m_1, \dots, m_{\ell(k)})$ such that $|m_i| = l_i(1^k)$ for all $1 \leq i \leq \ell(k)$.

An encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is SIM-NME'-secure if for polynomials $\ell(k), r(k)$ and $p(k)$, every polynomial-time computable history function $h(\cdot)$, every p.p.t. adversary $A = (A_1, A_2)$ which runs in time $p(k)$ and always outputs a (p, ℓ) -valid message sampler, there exists a p.p.t. algorithm $S = (S_1, S_2)$

that always outputs a (p, ℓ) -valid message sampler, such that the following two distributions are computationally indistinguishable:

$$\left\{ \text{SIM-NME}'(II, A, k, \ell(k), r(k)) \right\}_k \stackrel{c}{\approx} \left\{ \overline{\text{SIM-NME}'(II, S, k, \ell(k), r(k))} \right\}_k \quad (1)$$

We also define a weaker notion of this definition named SIM-NME by requiring that the outputs of the two experiments are indistinguishable only for a certain “restricted” set of adversaries A . Define the following two types of adversaries:

1. *non-copying adversaries*: $A = (A_1, A_2)$ is said to be non-copying if in the above experiment A_2 never outputs a ciphertext c_i , s.t., $c_i \in \mathcal{Y}$.
2. *valid adversaries*⁵: A is said to be valid if in the above experiment A only outputs ciphertexts that are in the range of the encryption function (on input PK), i.e., it holds that for all c_i , there exists an d_i such that $c_i \in \text{Enc}(\text{PK}, d_i)$.

Definition 2 (SIM-NME Security). *An encryption scheme $II = (\text{Gen}, \text{Enc}, \text{Dec})$ is SIM-NME-secure if for polynomials $\ell(k)$, $r(k)$ and $p(k)$, every polynomial-time computable history function $h(\cdot)$, every **non-copying, valid** p.p.t. adversary $A = (A_1, A_2)$ which runs in time $p(k)$ and always outputs a (p, ℓ) -valid message sampler, there exists a p.p.t. algorithm $S = (S_1, S_2)$ that always outputs a (p, ℓ) -valid message sampler, such that the ensembles in equation (1) are indistinguishable to any p.p.t. distinguisher D .*

Single-Message Versus Many-Message Security. We have presented definitions which allow the adversary to see a sequence of encrypted messages. For both the above definitions of non-malleability, a scheme satisfying the definition in the case when $\ell(k) = 1$ (but $r(k)$ is still arbitrary), is said to be *single-message* secure. The question of whether any single-message secure scheme is also (many-message) secure is the question of composability of the definition.

Remarks. Single-message SIM-NME security is a rewriting of the original DDN simulation-based definition of non-malleability. The main difference between our definition and definition of DDN is that we dispense with the relation R and instead use the notion of indistinguishability of the outputs. This difference is inconsequential (since any p.p.t distinguisher can be described as a p.p.t relation and vice versa); however, this draws a parallel to the (upcoming) indistinguishability-based definition of non-malleability, which we term IND-NME. In this way, we emphasize the meaning of this definition: neither a ciphertext of a chosen message or a decryption oracle can substantially alter an adversaries ultimate “behavior.” Given this interpretation, it is also intuitive to see why the valid-adversary is somehow artificial. Moreover this restriction prevents the definition from composing—i.e., it is possible for a scheme to be single-message SIM-NME secure, but not SIM-NME secure. We also remark that our definition of single-message SIM-NME' security is syntactically equivalent to the SNM definition of non-malleability from [BS06].

⁵ This interpretation comes from [DDN00] where they write “ A produces...ciphertexts (f_1, \dots) ...with $f_i \in \text{Enc}(\beta_i)$...”.

2.2 Indistinguishability-Based Definitions

The following definition of non-malleability was introduced in [PSV06] and is syntactically very close to the definition of [BS99, BS06].

Definition 3 (IND-NME' Security [PSV06]). *Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be an encryption scheme and let the random variable $\text{IND-NME}'_b(\Pi, A, k, \ell, r)$ where $b \in \{0, 1\}$, $A = (A_1, A_2)$ and $k, \ell, r \in \mathbb{N}$ denote the result of the following probabilistic experiment:*

$$\begin{aligned} & \text{IND-NME}'_b(\Pi, A, k, \ell, r) : \\ & (\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^k) \\ & ((m_{0,1}, \dots, m_{0,\ell}), (m_{1,1}, \dots, m_{1,\ell}), s) \leftarrow A_1^{O^1}(\text{PK}) \text{ s.t. } |m_{0,i}| = |m_{1,i}| \\ & y_i \leftarrow \text{Enc}(\text{PK}, m_{b,i}) \text{ for } i \in [1, \ell] \\ & (c_1, \dots, c_r) \leftarrow A_2^{O^2}(\mathbf{y}, s) \\ & \text{Output } (d_1, \dots, d_r) \text{ where } d_i = \begin{cases} \text{COPY} & \text{if } c_i \in \mathbf{y} \\ \text{Dec}(\text{SK}, c_i) & \text{otherwise} \end{cases} \end{aligned}$$

$(\text{Gen}, \text{Enc}, \text{Dec})$ is IND-NME'-secure if \forall p.p.t. algorithms $A = (A_1, A_2)$ and for any polynomials $\ell(k)$ and $r(k)$, the following two ensembles are computationally indistinguishable:

$$\left\{ \text{IND-NME}'_0(\Pi, A, k, \ell(k), r(k)) \right\}_k \stackrel{c}{\approx} \left\{ \text{IND-NME}'_1(\Pi, A, k, \ell(k), r(k)) \right\}_k \quad (2)$$

We also introduce a weaker version of this definition, IND-NME, in which, as in the previous section, (2) need only hold for non-copying, valid adversaries A .

Definition 4 (IND-NME Security). *An encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ is IND-NME-secure if \forall non-copying, valid p.p.t. algorithms $A = (A_1, A_2)$ and for any polynomials $\ell(k)$ and $r(k)$, the ensembles in the equation (2) are computationally indistinguishable.*

Single-Message Security. For both the above indistinguishability-based definitions, we obtain the weaker notion of single-message security by restriction attention to the case when $\ell(k) = 1$. We also note that our definition of single-message IND-NME' security is a syntactical rewriting of (and thus equivalent to) the definition of IND-PAX of [BS06].

3 Equivalences Between Definitions

Theorem 1. SIM-NME = IND-NME for all attacks.

The equivalence proof for this theorem uses ideas from Bellare and Sahai [BS99]. Note however that it does not show that SIM-NME' = IND-NME' (as was the goal in Bellare and Sahai's revised paper [BS06]). Let us briefly recall the subtle issue in the original proof in [BS99] (the same issue appears in the revised proof that SIM-NME' = IND-NME' in [BS06]). In one step of the equivalence proof, the SIM-NME simulator must re-encrypt a vector of ciphertexts which the adversary

has produced. If an “aborting” adversary has produced an invalid ciphertext, it is not clear whether the simulator can proceed—in particular, the encryption scheme Π *might not provide* an efficient method available to produce an invalid ciphertext (as was the case in the previous section). The proof does hold, however, for a valid adversary who always produces ciphertexts that are in the range of the Enc function.

In the full version, we present a direct equivalence proof for SIM-NME and IND-NME which is simple and extends to the case of many-message security. Moreover, the proof also leads to the following corollary relating SIM-NME' and IND-NME' used in Theorem 4:

Corollary 1. *If Π is SIM-NME'-secure, then Π is also IND-NME' secure.*

For completeness, we present a proof of the following theorem in the full version which has been partially shown by Dolev, Dwork, and Naor [DDN00].

Theorem 2. $\text{IND-NME}'\text{-CCA2} = \text{IND-NME-CCA2} = \text{IND-CCA2}$.

In the weaker CPA attack, we show that the simulation and indistinguishability definitions for invalid ciphertext-producing adversaries are also equivalent by adapting a simpler version of Thm. 1. This implies that the construction from [PSV06] meets the strongest notion of non-malleability for the CPA attack. The proof appears in the full version.

Theorem 3. *Under a CPA attack, $\text{SIM-NME}' = \text{IND-NME}'$.*

4 Separating the SIM-NME' and IND-NME' Definitions

Theorem 4 (Main Separation). *Under CCA1 or CCA2 attacks, $\text{SIM-NME}' > \text{IND-NME}'$ even for single-message security.*

Corollary 1 shows that SIM-NME' implies IND-NME'. Thus, the main idea for this separation is to design an encryption scheme in which the set of messages for which a ciphertext can be efficiently computed and the range of the decryption function *differ*. As one concrete example below, we design an IND-NME' scheme in which it is nearly impossible for an adversary to produce a ciphertext which decrypts to \perp (i.e., an invalid ciphertext) *unless* it has adaptive access to a decryption oracle.⁶ We show the scheme so constructed meets the IND-NME' definition. However, it does not meet the SIM-NME' definition under a CCA1 or CCA2 attack, because an adversary (with access to a decryption oracle) is able to produce a ciphertext that decrypts to \perp whereas a simulator (without access to a decryption oracle) is unable. Thus, the outputs of the SIM-NME' and IND-NME' experiments will be trivially distinguishable. The general idea behind these type of arguments first appears in [DDN00] and is also used in [BDPR98] to show other separations.

⁶ Another example would be a finite message space, i.e., a message space which includes all strings in $\{0, 1\}^k$ and a scheme in which the range of the decryption function includes one k^2 bit string. We discuss this later in §5.

Proof. Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be an encryption scheme that satisfies IND-NME' under a CCA attack. Consider encryption scheme Π' defined in the figure below. The key property of Π' is that Dec' never outputs \perp unless it is queried with

ENCRYPTION SCHEME Π'
$\text{Gen}'(1^k)$: Run $(\text{PK}, \text{SK}) \leftarrow \text{Gen}(1^k)$. Pick random k -bit string α and set $\text{SK}' \leftarrow (\text{SK}, \alpha)$.
$\text{Enc}'(\text{PK}', m)$: Run $c \leftarrow \text{Enc}(\text{PK}, m)$. Output $(1, 0^k, c)$ as ciphertext.
$\text{Dec}'(\text{SK}', c')$: Parse c' as (b, β, c) where b is a bit, β is a k -bit string. <ol style="list-style-type: none"> 1. If $b = 0$ and $\beta = 1^k$, then output α. 2. If $b = 0$ and $\beta = \alpha$, then output \perp. 3. If $b = 1$ and $\beta = 0^k$, run $m \leftarrow \text{Dec}(\text{SK}, c)$. If the output is \perp, output 0. Otherwise, output m. 4. Otherwise, output 0.

the special “open sesame” string α , and a decryption oracle is necessary to learn the “open sesame” string.

It is easy to see that Π' syntactically is an encryption scheme. The only issue is to argue that Π' is perfectly correct, which follows because perfect correctness only applies to decryption of honestly encrypted messages (which are never invalid ciphertexts).

Claim. $\Pi' = (\text{Gen}', \text{Enc}', \text{Dec}')$ meets the IND-NME'-CCA definition.

Proof. Suppose there exists an adversary A' which breaks the IND-NME'-CCA definition for Π' . Such an adversary can be used to construct an adversary A which breaks the IND-NME'-CCA definition for Π as follows:

The new adversary A simulates $(\text{Gen}', \text{Enc}', \text{Dec}')$ for A' by picking α itself and using the oracles for Dec to answer queries. More precisely, on input a public key PK , A generates a k -bit string α and feeds PK to A' . When A' asks decryption queries, A simulates the Dec' algorithm by using α as the second component of SK' and the decryption oracle in order to compute $\text{Dec}(c, \text{SK})$. When A' produces two challenge messages, A forwards these messages along, and when it receives a challenge ciphertext y , A feeds $(1, 0^k, y)$ to A' . In the case of a CCA2 attack, A again simulates the Dec' function, and when A' finally returns an answer, A echoes it. A perfectly simulates the IND-NME'-CCA game for A' , and thus succeeds with exactly the same probability as A' .

Claim. Π' does not meet the SIM-NME'-CCA definition.

Proof. Consider the relation $R(x, \mathbf{x}, M, s)$ which is 1 if x is \perp and 0 otherwise.

A CCA1 adversary with access to a decryption oracle can satisfy R by making a decryption query on the message $(0, 1^k, 0)$ to get the value α , and then by outputting the ciphertext $(0, \alpha, 0)$.

However, it is not possible for a simulator S without access to the decryption oracle to satisfy R . Such a simulator only has an exponentially small chance of guessing the correct α string necessary to produce \perp . Thus, Π' will not satisfy SIM-NME'-CCA1.

4.1 More Separations for CCA1 and CPA Attacks

We now show that IND-NME' is stronger than IND-NME *when considering weaker CPA and CCA1 attacks*. Recall that IND-NME' and IND-NME are different only in that the former protects against all PPT adversaries, whereas the latter protects against only valid PPT adversaries.⁷ By combining the equivalence from Theorem 1, we also get a separation between IND-NME' and SIM-NME. For CCA2 attacks, they become equivalent (See Thm. 2).

Theorem 5. IND-NME' > IND-NME for CCA1 and CPA attacks even for single-message security.

Corollary 2. IND-NME' > SIM-NME for CCA1 and CPA attacks even for single-message security.

Proof. (Of Corollary 2.) By Theorem 5, IND-NME' > IND-NME for CCA1 and CPA attacks and by Theorem 1, SIM-NME = IND-NME under all attacks.

The main idea for the proof of Theorem 5 is to use the *DDN-lite* transformation [Dwo99, Nao04] to transform an IND-NME-secure encryption scheme into one that remains IND-NME-secure (Claim 4.1), but is vulnerable to an IND-NME' attack (Claim 4.1).

We actually prove a stronger statement which gives us a way to transform an IND-CPA-secure encryption scheme into one that is IND-NME-secure. While this result has been claimed in [Dwo99, Nao04], as far as the authors know, a proof has never been printed. Our proof also shows that the construction also transforms an IND-CCA1 scheme into an IND-NME-CCA1 scheme. The IND-NME'-attack against this scheme is an adaptation of the attack against DDN-lite, given in [PSV06].

Proof. (of Theorem 5) Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be an encryption scheme that is IND-CPA-secure (respectively, IND-CCA1-secure). Let $\Sigma = (\text{Gen}_{sig}, \text{Sign}, \text{Ver})$ be a strongly unforgeable one-time signature scheme. Such a signature scheme can be constructed from one-way functions (The existence of one-way functions, in turn, is implied by the existence of a IND-CPA-secure encryption scheme). We construct a new encryption scheme Π_L from Π and show that Π_L satisfies the IND-NME definition but does not satisfy IND-NME'.

Claim. Π_L meets the IND-NME definition.

⁷ We say that an invalid ciphertext “decrypts” to \perp (Bot) and hence the title of the subsection.

ENCRYPTION SCHEME Π_L

$\text{Gen}'(1^k)$: Run $\text{Gen}(1^k)$ $2k$ times with independent random coins to produce $2k$ pairs of keys $(\text{PK}_b^i, \text{SK}_b^i)$ for $i \in [1, k]$ and $b \in \{0, 1\}$. Let $\text{PK}' = [\text{PK}_b^i]_{i \in [k], b \in \{0,1\}}$ and $\text{SK}' = [\text{SK}_b^i]_{i \in [k], b \in \{0,1\}}$

$\text{Enc}'(m, \text{PK}')$: Run $\text{Gen}_{\text{sig}}(1^k)$ to generate a key-pair $(\text{VKSIG}, \text{SKSIG})$ for the signature scheme. Let VKSIG a k -bit string, and let the i^{th} bit of VKSIG be denoted VKSIG_i .
 Run $c_i \leftarrow \text{Enc}(\text{PK}_i^{\text{VKSIG}_i}, m)$ for $i \in [1, k]$.
 Let $\sigma \leftarrow \text{Sign}(\text{SKSIG}, (c_1, c_2, \dots, c_k))$.
 Output $[(c_1, \dots, c_k), \text{VKSIG}, \sigma]$ as the ciphertext.

$\text{Dec}'(c', \text{SK}')$: Parse c' as $((c_1, \dots, c_k), \text{VKSIG}, \sigma)$.
 If $\text{Ver}(\text{VKSIG}, (c_1, \dots, c_k), \sigma) = \text{REJECT}$, output \perp .
 Otherwise, decrypt the c_i 's with the corresponding secret-keys to get corresponding messages m_i . If all m_i 's are equal, output m_1 , else output \perp .

Proof. First, we show that an encryption scheme $\tilde{\Pi}$, constructed from Π in the following way, meets the IND-CPA (respectively, IND-CCA1) definition (Proposition 1). $\tilde{\Pi} = (\tilde{\text{Gen}}, \tilde{\text{Enc}}, \tilde{\text{Dec}})$ is constructed as follows:

1. $\tilde{\text{Gen}}$ runs k copies of Gen to generate public-keys $\tilde{\text{PK}} = (\text{PK}_1, \text{PK}_2, \dots, \text{PK}_k)$ and corresponding secret-keys $\tilde{\text{SK}} = (\text{SK}_1, \text{SK}_2, \dots, \text{SK}_k)$.
2. $\tilde{\text{Enc}}(m, \tilde{\text{PK}})$ runs $\text{Enc}(m, \text{PK}_i)$ for all $i \in [k]$, with independently chosen randomness, and outputs the vector of k encryptions $[c_1, c_2, \dots, c_k]$.
3. $\tilde{\text{Dec}}(c, \tilde{\text{SK}})$ parses c as $[c_1, c_2, \dots, c_k]$. Let $m_i = \text{Dec}(c_i, \text{SK}_i)$. If all the m_i are the same, output m_1 . Otherwise, output \perp .

Secondly, in Proposition 2, we show that if $\tilde{\Pi}$ is IND-CPA-secure (respectively, IND-CCA1-secure), then Π' is IND-NME-CPA-secure (resp., IND-NME-CCA1-secure). This proof appears in the full version.

Proposition 1. *If Π is IND-CPA-secure (or IND-CCA1-secure), then so is $\tilde{\Pi}$.*

Proof. The proof is a straightforward hybrid argument. The only complication stems from the simulation of the oracle in the CCA1 case. When the adversary asks to decrypt a ciphertext $\hat{c} = (\hat{c}_1, \hat{c}_2, \dots, \hat{c}_k)$, decrypt c_j using the secret-key SK_j (if $j \neq i$) and using the decryption oracle for PK_i (if $j = i$).

Proposition 2. *If $\tilde{\Pi}$ is IND-CPA-secure (respectively, IND-CCA1-secure), then Π_L is IND-NME-secure (respectively, IND-NME-CCA1-secure).*

Claim. Π_L is not IND-NME'-secure under CPA and CCA1 single-message attacks.

Proof. We specify an adversary $A = (A_1, A_2)$ and a distinguisher D such that D distinguishes between $\{\text{IND-NME}'_0(\Pi, A, k, 1)\}$ and $\{\text{IND-NME}'_1(\Pi, A, k, 1)\}$. A works as follows:

1. A_1 outputs two arbitrary messages (m_0, m_1) and no state information.
2. On input ciphertext $c = [(e_1, \dots, e_k), \text{VKSIG}, \sigma]$, let $\text{VKSIG} := b_1 b_2 \dots b_k$. A_2 produces a new ciphertext c' as follows: A_2 generates a new signing key $(\text{SKSIG}', \text{VKSIG}')$. Let $\text{VKSIG}' := b'_1 b'_2 \dots b'_k$. A_2 outputs ciphertext $c' = ((x_1, \dots, x_k), \text{VKSIG}', \sigma')$ where

$$x_i = \begin{cases} e_i & \text{if } b'_i = b_i \\ E_{\text{PK}_i^{b'_i}}(m_0) & \text{otherwise} \end{cases}$$

and σ' is the signature of (x_1, \dots, x_k) under the signing key SKSIG' .

Notice that $\text{NME}_0(\Pi, A, k, 1) = m_0$ and $\text{NME}_1(\Pi, A, k, 1) = \perp$ which can be easily distinguished by a distinguisher D that outputs 0 on m_0 and 1 on \perp .

5 Additional Separations with Finite Message Spaces

Many encryption schemes such as El Gamal, RSA, Cramer-Shoup, and the league of schemes based on elliptic curves and bilinear maps only process messages from a finite message space such as the elements of some group G . In order to capture the security of such systems, Cramer and Shoup [CS98] redefine the encryption primitive to incorporate (a) a key-dependent message space M_{PK} and (b) a p.p.t. message tester algorithm \mathcal{M} that on input $1^k, \text{PK}, \alpha$, determines whether α is an element of the message-space for the security parameter 1^k and the public key PK . The encryption algorithm $\text{Enc} : M_{\text{PK}} \rightarrow \{0, 1\}^*$ now takes an input message from M_{PK} and produces general bit strings, and the decryption algorithm maps $\{0, 1\}^*$ to $M_{\text{PK}} \cup \perp$. The correctness property is only required to hold over the message space.

In this section, however, we note that if the message space is finite, then the previously proven equivalence relationship between the weaker notions of SIM-NME and IND-NME no longer holds. While the particular counter-example that we use for the separation may seem contrived, this separation has practical significance since it runs against our “intuition” about IND-CCA2 security.

The idea behind this separation is as follows. We construct an encryption scheme whose message space includes three elements, $\{0, 1, \chi\}$ where χ is related to the public key PK . Moreover, we make it difficult for an adversary to learn χ unless it has a decryption oracle (notice, the definition for finite message space only requires the message space to be easily decidable, but does not require it to be enumerable.⁸) From this point, the argument is the same. Namely, an adversary with an oracle can produce a (valid) ciphertext decrypting to χ (therefore it is a valid adversary), whereas the simulator can only produce ciphertexts decrypting to 0 or 1. The subtle difference between this argument and

⁸ One could require enumerability of the message space. However, it is unclear such a restriction helps; and it is clear that it needlessly prevents us from using more exotic algebraic structures for encryption.

the one from §4 is that in this one, it is not the simulator’s inability to produce a ciphertext which decrypts to \perp , but rather its inability to *learn a special message in the message space* which provides the separation. In the full message case, there are *no* such special messages since any string can be encrypted. This is the reason that the separation can be extended to valid adversaries.

Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be an IND-NME-secure encryption scheme for general message spaces, and let f be a one-way permutation.⁹

FINITE MESSAGE SPACE ENCRYPTION SCHEME Γ

$\text{Gen}'(1^k)$: Run $\text{Gen}(1^k)$ to generate a key pair (PK, SK) . Pick k -bit random string α and compute $\beta = f(\alpha)$. Set $\text{SK}' = (\text{SK}, \alpha)$ and $\text{PK}' = (\text{PK}, \beta)$. The message tester $\mathcal{M}(m)$ works as follows: if $m \in \{0, 1\}$ or if $f(m) = \beta$, then return 1. Otherwise, return 0. (The messages space consists of $\{0, 1, \alpha\}$).

$\text{Enc}'(m, \text{PK}')$: if $\mathcal{M}(m) = 0$, return an error. Otherwise, run $c \leftarrow \text{Enc}(\text{PK}, m)$ and return $(1, c)$.

$\text{Dec}'(c', \text{SK}')$: Parse c' as (b, c) , and SK' as (SK, α) . If $b = 0$ then output α . Otherwise, output $m \leftarrow \text{Dec}(\text{SK}, c)$.

IND-NME security of the above finite-message space encryption scheme directly follows from the security of $(\text{Gen}, \text{Enc}, \text{Dec})$. In order to violate SIM-NME, the adversary B must be non-aborting. Therefore, the final ciphertext it produces must be in the range of the Enc function (i.e., of the form $(1, c)$). Combined with the one-wayness of f , a simulator not having access to a decryption oracle will not be able to construct a *valid* encryption to the message α .

However, a CCA1 attacker can easily do so by first querying $(0, 0)$ to find α (notice that the attacker can *query* the oracle on invalid ciphertexts, but cannot produce them as final output), and then honestly encrypting α .

6 Special Cases for Equivalence

The separation between SIM-NME' and IND-NME' hinged on the fact that the set of messages for which one can efficiently compute a ciphertext and the range of the decryption procedure differ. When these two sets are made to coincide, a scheme that is IND-NME' secure is also SIM-NME'-secure. Thus, we provide an easy way to prove that a scheme meets the strongest notion of non-malleability. As a corollary, we get that the main construction of [DDN00] achieves the strongest form of security – that is SIM-NME'-security against CCA2 attacks.

Theorem 6. *Any (finite message-space) encryption scheme Π which meets the IND-NME definition and for which there is an efficient algorithm F , which on input (PK, d) where d is a string in the range of Dec , produces a ciphertext c such that $d \leftarrow \text{Dec}(\text{SK}, c)$, also meets the SIM-NME' definition.*

⁹ In fact a one-way function would suffice. We only use a permutation for ease of exposition.

This restriction could easily be added to many schemes by taking the message space to be some set $\{0, 1\}^{\ell(k)}$ for all keys generated by $\text{Gen}(1^k)$ (and by making it easy to generate invalid ciphertext). We note that the RSA-OAEP padding scheme does exactly this.

7 Composition: Many Message Security

In [PSV06], the authors show that IND-NME' security under CPA attacks composes. That is, if an encryption scheme is IND-NME'-secure when the adversary receives one encryption, it will also be IND-NME'-secure in a situation in which the adversary receives many encryptions.

A natural question is whether the same phenomena occurs under stronger CCA1 and CCA2 attacks. In this section, we answer affirmatively as described in the following theorem.

Theorem 7. *A scheme Π meets IND-NME' under attack ATK iff it meets single-message IND-NME' under attack ATK.*

Proof Sketch: The forward implication follows directly. For the reverse direction, we present a routine hybrid argument that uses an adversary $(A_1, A_2), D$ with advantage ϵ to construct a new adversary $(A'_1, A'_2), D$ which breaks the single-message security with advantage η/ℓ^2 .

Define a new experiment $\text{IND-NME}'_{(b_1, \dots, b_\ell)}(\Pi, A, k, \ell)$ indexed by an ℓ -bit string (b_1, \dots, b_ℓ) which is the same as $\text{IND-NME}'_0(\Pi, A, k, \ell)$ except in the fourth line (change is underlined): $y_i \leftarrow \text{Enc}(\text{PK}, \underline{m_{b_i, i}})$ for $i \in [1, \ell]$. Define

$$B(i) = (\overbrace{0, \dots, 0}^{l-i}, \overbrace{1, \dots, 1}^i)$$

and note that $\text{IND-NME}'_0 = \text{IND-NME}'_{B(0)}$ and $\text{IND-NME}'_1 = \text{IND-NME}'_{B(\ell)}$. Because D distinguishes $\text{IND-NME}'_0$ from $\text{IND-NME}'_1$, there exists some $g^* \in [1, \ell]$ such that D distinguishes $\text{IND-NME}'_{B(g^*)}$ from $\text{IND-NME}'_{B(g^*+1)}$ with advantage η/ℓ . This suggests the following adversary: $A_1^O(\text{PK})$ guesses value $g \in [1, \ell]$ and runs $A_1(\text{PK})$ —answering any decryption queries by using its own decryption oracle—and waits to receive the two vector of messages $(m_{0,1}, \dots, m_{0,\ell})$ and $(m_{1,1}, \dots, m_{1,\ell})$. Finally, A' outputs $(m_{0,g}, m_{1,g})$ as its challenge pair and outputs state information containing g and $\mathbf{m}_0, \mathbf{m}_1$.

Adversary $A_2^O(y, \text{STATE}')$, on input an encryption y , first executes the replaced line 4 of experiment $\text{IND-NME}'_{B(g)}$ (described above) with the exception that it uses y for the $(g + 1)$ th encryption: $y_{g+1} \leftarrow y$. This is possible because it receives the messages vectors \mathbf{m}_0 and \mathbf{m}_1 in STATE' .

It then feeds the resulting vector of ciphertexts \mathbf{y} to A_2 to produce another vector of ciphertexts (c_1, \dots, c_ℓ) and uses this vector as its own output. To answer any oracle query c from A_2 , A'_2 uses the following procedure: If $c = y_j$ for any $j \in [1, \ell]$, then return \perp . Otherwise, it uses its own decryption oracle to decrypt c and answers with the returned message.

Notice that $\text{IND-NME}'_0(A'_1, A'_2)$ and $\text{IND-NME}'_{B(g^*)}(A_1, A_2)$ are *syntactically* the same, as are $\text{IND-NME}'_1(A'_1, A'_2)$ and $\text{IND-NME}'_{B(g^*+1)}(A_1, A_2)$. Because A' guesses g^* correctly with probability $1/\ell$, D 's overall advantage in breaking the single-message non-malleability is η/ℓ^2 . \square

One can see here the importance of removing the “valid adversary” restriction for the hybrid argument to work. This follows because the reduction feeds a hybrid distribution to A_2 and, *even if A_2 is itself a valid adversary* for the multi-message experiment, A_2 may produce invalid ciphertexts when it is fed a *hybrid* distribution. Moreover, these \perp values may form the basis for distinguishability in the hybrid experiment. Thus, one cannot guarantee that valid adversaries for the multi-message experiment can be transformed into valid adversaries for the single-message experiment. The separation in Claim 4.1 exploits this issue.¹⁰

SIM-NME and IND-NME Do Not Compose Against CCA1 or CPA Attacks

We now show that (if there exist SIM-NME-secure encryption schemes) there is an encryption scheme Π' that is SIM-NME or IND-NME-secure when the adversary is given *one* ciphertext as the challenge, but there is an adversary A' that completely breaks the IND-NME-security of Π' when given *polynomially many* ciphertexts as challenge.

The encryption scheme Π' is simply the encryption scheme constructed in the proof of Thm. 5 (relying on the DDNLite construction). Thm. 5 showed that Π_L is 1-message IND-NME-secure (and therefore 1-message SIM-NME-secure). The many-message attack against Π' is a simple *covering attack*. (We mention that Gennaro and Lindell [GL03] pointed out that the DDNLite encryption scheme is not secure under many messages. Although they did not include a description of the attack, we believe they had a similar attack in mind.)

Recall that an encryption of a message m under Π' consists of many encryptions of m with respect to a randomly chosen set of k (out of $2k$) public-keys. Given many (roughly $k \log k$) independent encryptions of m , the one can essentially recover an encryption of m under *all* the $2k$ public-keys. This will enable us to construct a completely new encryption of m , and thus break IND-NME' security.

Theorem 8. *Let $\text{atk} \in \{\text{CPA}, \text{CCA1}\}$. If there exists an encryption scheme that is IND- atk secure, then there exists another encryption scheme Π' that is 1-message IND-NME- atk -secure (respectively SIM-NME- atk -secure), but is not even IND-NME-CPA-secure (respectively, SIM-NME-CPA-secure).*

Proof. Omitted

¹⁰ This argument also applies to a different interpretation of “valid adversary” in which one forces the single-message experiment to return 0 when invalid ciphertexts are produced. In this case, when A_2 produces invalid ciphertexts in the hybrid experiments, the value of both hybrid experiments ($b = 0, 1$) will be 0 and the weaker definition will thus be met *even though* there might still be a distinguisher which could have distinguished the output of A'_2 .

SIM-NME and IND-NME Compose Under CCA2 Attacks

Theorem 9. *If an encryption scheme Π is 1-message IND-NME-CCA2-secure, then it is many-message IND-NME-CCA2-secure.*

The proof of this theorem follows from Theorem 2, which shows that under CCA2 attacks, IND-NME and SIM-NME definitions coincide with the IND-NME' definition, and Theorem 7 which shows that IND-NME' composes under a many-message attack.

Acknowledgments. We would like to thank one of the anonymous Crypto referees for thorough and helpful comments.

References

- [BDPR98] Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among notions of security for public key encryption schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, Springer, Heidelberg (1998)
- [BS99] Bellare, M., Sahai, A.: Non-malleable encryption: Equivalence between two notions, and an indistinguishability-based characterization. In: Wiener, M.J. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 519–536. Springer, Heidelberg (1999)
- [BS06] Bellare, M., Sahai, A.: Non-malleable encryption: Equivalence between two notions, and an indistinguishability-based characterization (2006), <http://eprint.iacr.org/2006/228>
- [Can01] Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS, pp. 136–145 (2001)
- [CS98] Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 13–25. Springer, Heidelberg (1998)
- [DDN00] Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM J. Comput. 30(2), 391–437 (2000)
- [Dwo99] Dwork, C.: The non-malleability lectures. Course notes for Stanford CS 359 (1999), <http://theory.stanford.edu/~gdurf/cs359-s99/>
- [GL03] Gennaro, R., Lindell, Y.: A framework for password-based authenticated key exchange. In: Biham, E. (ed.) Advances in Cryptology – EUROCRYPT 2003. LNCS, vol. 2656, pp. 524–543. Springer, Heidelberg (2003)
- [GM84] Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. 28(2), 270–299 (1984)
- [Gol04] Goldreich, O.: Foundations of Cryptography, vol. 2. Cambridge University Press, Cambridge (2004)
- [Nao04] Naor, M.: A taxonomy of encryption scheme security (2004)
- [PSV06] Pass, R., Shelat, A., Vaikuntanathan, V.: Construction of a non-malleable encryption scheme from a any semantically secure one. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, Springer, Heidelberg (2006)