CMSC 858K — Advanced Topics in Cryptography February 3, 2004

Lecture 3

Lecturer: Jonathan Katz *Scribe(s): Abheek Anand
Gelareh Taban
Radostina Koleva*

1 Introduction

In the last lecture, we introduced the notion of semantic security and gave a formal definition for semantic security with respect to public-key encryption schemes. It was shown that given the existence of hard-core bits, it is possible to construct a *semantically secure public-key encryption scheme* for messages of length one bit. In this lecture, we introduce the *hybrid technique* and use it to prove that semantic security of the constructed encryption scheme can be extended to polynomially-many messages (of arbitrary polynomial length).

We begin by reviewing the construction of a semantically-secure public-key encryption scheme from a trapdoor permutation. Let $F = (\mathsf{Gen}, \mathsf{Eval}, \mathsf{Invert})$ be a trapdoor permutation family and $H = \{h_k\}$ be a hard-core bit for F. Then we can construct the following public-key encryption scheme $PKE = (\mathsf{KeyGen}, \mathcal{E}, \mathcal{D})$ for the encryption of 1-bit messages:

$$\begin{array}{ll}
\mathsf{KeyGen}(1^k): & (f, f^{-1}) \leftarrow \mathsf{Gen}(1^k) \\
 & PK = (f, h_k) \\
 & SK = f^{-1} \\
 & \\
\mathcal{E}_{PK}(m): & r \leftarrow \{0,1\}^k \\
 & \text{output } \langle f(r), h_k(r) \oplus m \rangle \\
 & \\
\mathcal{D}_{SK}(\langle y, b \rangle): & \text{output } b \oplus h_k(f^{-1}(y))
\end{array}$$

We showed in class last time that the encryption scheme above is semantically secure. In particular, this implies the following theorem:

Theorem 1 *Assuming trapdoor permutations exist, there exists a public-key encryption scheme achieving semantic security (or security in the sense of indistinguishability).*

2 Security for Multiple Messages

The encryption scheme above is semantically secure for messages of length one bit. Would the scheme remain secure if it is applied (in the natural bit-by-bit fashion[1]) to messages of length longer than one bit? Equivalently, is the scheme still secure if it is used to encrypt multiple messages, each of length one bit? If the adversary is able to eavesdrop

[1] Here, encryption of $m = m_1 \cdots m_\ell$ (with $m_i \in \{0,1\}$) is given by $\mathcal{E}_{PK}(m_1) \cdots \mathcal{E}_{PK}(m_\ell)$. The only subtlety here is that *independent* random coins must be used for every invocation of the encryption algorithm.

on these messages, will she obtain extra information (perhaps correlated information about the various messages) and be able to break the semantic security of the encryption scheme?

To model this stronger attack scenario where the adversary can intercept multiple messages via eavesdropping, we introduce the concept of an *encryption oracle* $\mathcal{E}_{PK,b}(\cdot,\cdot)$ which the adversary can query as many times as it wants. This oracle takes as input two messages m_0, m_1 of equal length, and we define $\mathcal{E}_{PK,b}(m_0, m_1) \overset{\text{def}}{=} \mathcal{E}_{PK}(m_b)$ (where new random coins are used to encrypt m_b each time the oracle is invoked). A scheme is secure if the adversary cannot guess the value of the bit b used by the encryption oracle (with much better than probability 1/2). A formal definition follows.

Definition 1 A public-key encryption scheme $PKE = (\mathsf{KeyGen}, \mathcal{E}, \mathcal{D})$ is secure in the sense of left-or-right indistinguishability if the following is negligible (in k) for any PPT adversary:

$$\left|\Pr\left[(PK, SK) \leftarrow \mathsf{KeyGen}(1^k); b \leftarrow \{0,1\} : A^{\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK) = b\right] - 1/2\right|.$$

◇

Theorem 2 *If a public-key encryption scheme $PKE = (\mathsf{KeyGen}, \mathcal{E}, \mathcal{D})$ is semantically secure, then it is also secure in the sense of left-or-right indistinguishability.*

Proof To prove this theorem, the *hybrid argument* is introduced. This technique plays a central role in demonstrating the indistinguishability of complex ensembles based on the indistinguishability of simpler ensembles. However, before we define the technique and prove the more general case, we will show that a semantically secure encryption scheme (for one message) is secure in the sense of left-or-right indistinguishability when *two* messages are encrypted. This is represented by allowing the adversary to have oracle access to the encryption oracle twice.

We will, as usual, perform a proof by contradiction: assume toward a contradiction that PKE is semantically secure but not left-or-right secure. This means that we have a PPT adversary A that can break the PKE in the left-or-right indistinguishability sense with non-negligible probability. Using this adversary we will construct a PPT algorithm that breaks the semantic security of the PKE with non-negligible probability. This is a contradiction as according to the theorem PKE is semantically secure.

In what follows we will let the key generation step be implicit in order to make the notation more readable. Construct adversary $\hat{A}_1$ that can access the encryption oracle just once, and tries to break semantic security as follows:

```
$\underline{\hat{A}_1^{\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK)}$
Run $A(PK)$
At some point $A$ asks for $\mathcal{E}_{PK,b}(m_0, m_1)$
$\hat{A}_1$ queries its own encryption oracle and returns $c \leftarrow \mathcal{E}_{PK,b}(m_0, m_1)$ to $A$
Later, $A$ requests a second encryption $\mathcal{E}_{PK,b}(m'_0, m'_1)$
$\hat{A}_1$ returns $c' \leftarrow \mathcal{E}_{PK}(m'_0)$ to $A$ (i.e., it encrypts $m'_0$ itself)
$\hat{A}_1$ outputs the final output of $A$
```

Since A runs in polynomial time so does $\hat{A}_1$. The probability that $\hat{A}_1$ succeeds is:

$$\begin{aligned}
\mathsf{Succ}_{\hat{A}_1} &\stackrel{\text{def}}{=} \Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k); b\leftarrow\{0,1\} : \hat{A}_1^{\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK)=b\right] \qquad (1)\\
&= \Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k); b\leftarrow\{0,1\} : A^{\mathcal{E}_{PK,b}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK)=b\right]\\
&= \Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,0}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK)=0\right]\times\frac{1}{2}\\
&\quad +\Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK)=1\right]\times\frac{1}{2},
\end{aligned}$$

where we have abused notation and written $A^{\mathcal{E}_{PK,b_1}(\cdot,\cdot),\mathcal{E}_{PK,b_2}(\cdot,\cdot)}$ to indicate that the first time A accesses its oracle, the oracle uses bit b_1, whereas the second time A accesses its oracle, the oracle uses bit b_2.

Similarly, we construct an adversary $\hat{A}_2^{\mathcal{E}_{PK,b}(\cdot,\cdot)}$ that accesses the encryption oracle just once and runs as follows:

```
$\underline{\hat{A}_2^{\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK)}$
Run $A(PK)$
At some point $A$ asks for $\mathcal{E}_{PK,b}(m_0, m_1)$
$\hat{A}_2$ returns $c \leftarrow \mathcal{E}_{PK}(m_1)$ to $A$ (i.e., it encrypts $m_1$ itself)
Later, $A$ requests a second encryption $\mathcal{E}_{PK,b}(m'_0, m'_1)$
$\hat{A}_2$ queries its own encryption oracle and returns $c' \leftarrow \mathcal{E}_{PK,b}(m'_0, m'_1)$ to $A$
$\hat{A}_2$ outputs the final output of $A$
```

Again, $\hat{A}_2$ is clearly a PPT algorithm. The probability that $\hat{A}_2$ succeeds is:

$$\begin{aligned}
\mathsf{Succ}_{\hat{A}_2} &\stackrel{\text{def}}{=} \Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k); b\leftarrow\{0,1\} : \hat{A}_2^{\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK)=b\right] \qquad (2)\\
&= \Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k); b\leftarrow\{0,1\} : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK)=b\right]\\
&= \Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK)=0\right]\times\frac{1}{2}\\
&\quad +\Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,1}(\cdot,\cdot)}(PK)=1\right]\times\frac{1}{2}.
\end{aligned}$$

We now express A's advantage[2] in breaking the left-or-right indistinguishability of the scheme in terms of Equations (1) and (2):

$$\begin{aligned}
\mathsf{Adv}_A &\stackrel{\text{def}}{=} \left|\Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k); b\leftarrow\{0,1\} : A^{\mathcal{E}_{PK,b}(\cdot,\cdot)}(PK)=b\right]-\frac{1}{2}\right|\\
&= \left|\frac{1}{2}\cdot\Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,1}(\cdot,\cdot)}(PK)=1\right]\right.\\
&\quad \left.+\frac{1}{2}\cdot\Pr\left[(PK,SK)\leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,0}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK)=0\right]-\frac{1}{2}\right|
\end{aligned}$$

[2] An adversary's *advantage* in this setting is simply the absolute value of its success probability (i.e., the probability that it correctly guesses b) minus $1/2$.

$$\begin{aligned} = \; & \left| \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,1}(\cdot,\cdot)}(PK) = 1\right] \right. \\ & + \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 0\right] \\ & + \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 1\right] \\ & \left. + \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,0}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 0\right] - 1 \right|, \end{aligned} \tag{3}$$

where we use the fact (from basic probability theory) that

$$\Pr\left[A^{\mathcal{E}_{PK,1}(\cdot,\cdot)\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 0\right] + \Pr\left[A^{\mathcal{E}_{PK,1}(\cdot,\cdot)\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 1\right] = 1.$$

Continuing, we obtain:

$$\begin{aligned} \mathsf{Adv}_A \; \leq \; & \left| \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,1}(\cdot,\cdot)}(PK) = 1\right] \right. \\ & \left. + \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 0\right] - \frac{1}{2} \right| \\ & + \left| \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,1}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 1\right] \right. \\ & \left. + \frac{1}{2} \cdot \Pr\left[(PK,SK) \leftarrow \mathsf{KeyGen}(1^k) : A^{\mathcal{E}_{PK,0}(\cdot,\cdot),\mathcal{E}_{PK,0}(\cdot,\cdot)}(PK) = 0\right] - \frac{1}{2} \right|. \\ = \; & \mathsf{Adv}_{\hat{A}_2} + \mathsf{Adv}_{\hat{A}_1}. \end{aligned}$$

Since (by our initial assumption) Adv_A was non negligible, the above implies that at least one of $\mathsf{Adv}_{\hat{A}_1}$ or $\mathsf{Adv}_{\hat{A}_2}$ must be non negligible. However, this would imply that at least one of $\hat{A}_1$ or $\hat{A}_2$ violate the semantic security of the encryption scheme, contradicting the assumption of the theorem. ■

2.1 The Hybrid Argument

We can generalize the proof technique used above so that it applies to any indistinguishable distributions (the technique is referred to as the "hybrid argument"). We formalize this idea now by first defining *computational indistinguishability.*

Definition 2 Let $\mathcal{X} = \{X_k\}$ and $\mathcal{Y} = \{Y_k\}$ be ensembles of distributions, where for all k, X_k and Y_k are distributions over the same space. $\mathcal{X}$ and $\mathcal{Y}$ are computationally indistinguishable (written $\mathcal{X} \overset{c}{\equiv} \mathcal{Y}$) if the following is negligible (in k) for all PPT A:

$$\left|\Pr\left[x \leftarrow X_k; A(x) = 1\right] - \Pr\left[y \leftarrow Y_k; A(y) = 1\right]\right|. \tag{4}$$

♢

We will sometimes be informal and refer to "distributions" instead of "ensembles of distributions".

As an example of how this notation may be used, we give an equivalent definition of semantic security (for a single bit) in terms of computational indistinguishability. Namely, let

$$X_k \overset{\text{def}}{=} \{(PK, SK) \leftarrow \mathsf{KeyGen}(1^k); C \leftarrow \mathcal{E}_{PK}(0) : (PK, C)\}$$

and

$$Y_k \overset{\text{def}}{=} \{(PK, SK) \leftarrow \mathsf{KeyGen}(1^k); C \leftarrow \mathcal{E}_{PK}(1) : (PK, C)\}.$$

Then encryption scheme $(\mathsf{KeyGen}, \mathcal{E}, \mathcal{D})$ is semantically-secure (for encryption of a single bit) if and only if $\{X_k\} \overset{c}{\equiv} \{Y_k\}$.

Before continuing with our discussion of the "hybrid argument", we note the following useful properties of computational indistinguishability:

Claim 3 *If $\mathcal{X} \overset{c}{\equiv} \mathcal{Y}$ and $\mathcal{Y} \overset{c}{\equiv} \mathcal{Z}$ then $\mathcal{X} \overset{c}{\equiv} \mathcal{Z}$.*

Sketch of Proof (Informal) The proof relies on the *triangle inequality* (namely, the fact that for any real numbers a, b, c we have $|a - b| \leq |a - b| + |b - c|$) and the fact that the sum of two negligible functions is negligible. □

In fact, we can extend this claim as follows:

Claim 4 (Transitivity) *Given polynomially many distributions $\mathcal{X}_1, \ldots, \mathcal{X}_{\ell(k)}$ for which $\mathcal{X}_i \overset{c}{\equiv} \mathcal{X}_{i+1}$ for $i = 1, \ldots, \ell(k) - 1$, then $\mathcal{X}_1 \overset{c}{\equiv} \mathcal{X}_{\ell(k)}$*

Sketch of Proof (Informal) The proof again uses the triangle inequality along with the fact that the sum of a polynomial number of negligible functions remains negligible. □

Note that the claim does *not* hold for a super-polynomial number of distributions.

We now formalize the hybrid argument.

Claim 5 (Hybrid argument) *Let $\mathcal{X}^1, \mathcal{X}^2, \mathcal{Y}^1, \mathcal{Y}^2$ be efficiently sampleable[3] distributions for which $\mathcal{X}^1 \overset{c}{\equiv} \mathcal{Y}^1$ and $\mathcal{X}^2 \overset{c}{\equiv} \mathcal{Y}^2$. Then $(\mathcal{X}^1, \mathcal{X}^2) \overset{c}{\equiv} (\mathcal{Y}^1, \mathcal{Y}^2)$. (Note: if $\mathcal{X} = \{X_k\}$ and $\mathcal{Y} = \{Y_k\}$ are two distribution ensembles, the notation $(\mathcal{X}, \mathcal{Y})$ refers to the distribution ensemble $\{(X_k, Y_k)\}$ where the distribution (X_k, Y_k) is defined by $\{x \leftarrow X_k; y \leftarrow Y_k : (x, y)\}$.)*

Proof Instead of proving this by contradiction, we prove it directly. Let A be an arbitrary PPT algorithm trying to distinguish $(\mathcal{X}^1, \mathcal{X}^2)$ and $(\mathcal{Y}^1, \mathcal{Y}^2)$. We may construct a PPT adversary A_1 trying to distinguish $\mathcal{X}^1$ and $\mathcal{Y}^1$ as follows:

```
$\underline{A_1(1^k, z)}$
Choose random $x \leftarrow X_k^2$
output $A(z, x)$
```

Clearly, A_1 runs in polynomial time (here is where we use the fact that all our distributions are efficiently sampleable). Since $\mathcal{X}^1 \overset{c}{\equiv} \mathcal{Y}^1$ we therefore know that the following must be

[3] A distribution ensemble $\mathcal{X} = \{X_k\}$ is *efficiently-sampleable* if we can generate an element according to distribution X_k in time polynomial in k.

negligible:

$$\begin{aligned}
&\big|\Pr\big[z \leftarrow X_k^1 : A_1(z) = 1\big] - \Pr\big[z \leftarrow Y_k^1 : A_1(z) = 1\big]\big| && (5)\\
&\quad = \big|\Pr\big[z \leftarrow X_k^1; x \leftarrow X_k^2 : A(z,x) = 1\big] - \Pr\big[z \leftarrow Y_k^1; x \leftarrow X_k^2 : A(z,x) = 1\big]\big|\\
&\quad = \big|\Pr\big[x_1 \leftarrow X_k^1; x_2 \leftarrow X_k^2 : A(x_1,x_2) = 1\big] - \Pr\big[y_1 \leftarrow Y_k^1; x_2 \leftarrow X_k^2 : A(y_1,x_2) = 1\big]\big|,
\end{aligned}$$

where the last line is simply a renaming of the variables.

We may similarly construct a PPT algorithm A_2 trying to distinguish $\mathcal{X}^2$ and $\mathcal{Y}^2$ that runs as follows:

```
$\underline{A_2(1^k, z)}$
Choose random $y \leftarrow \mathcal{Y}_k^1$
Output $A(y, z)$
```

Here, since $\mathcal{X}^2 \overset{c}{\equiv} \mathcal{Y}^2$ we know that the following is negligible:

$$\begin{aligned}
&\big|\Pr\big[z \leftarrow X_k^2 : A_2(z) = 1\big] - \Pr\big[z \leftarrow Y_k^2 : A_2(z) = 1\big]\big|\\
&\quad = \big|\Pr\big[y \leftarrow \mathcal{Y}_k^1; z \leftarrow X_k^2 : A(y,z) = 1\big] - \Pr\big[y \leftarrow Y_k^1; z \leftarrow Y_k^2 : A(y,z) = 1\big]\big|\\
&\quad = \big|\Pr\big[y_1 \leftarrow Y_k^1; x_2 \leftarrow X_k^2 : A(y_1,x_2) = 1\big] - \Pr\big[y_1 \leftarrow Y_k^1; y_2 \leftarrow Y_k^2 : A(y_1,y_2) = 1\big]\big|.
\end{aligned}$$

Of course, what we are really interested in is how well A does at distinguishing $(\mathcal{X}^1, \mathcal{X}^2)$ and $(\mathcal{Y}^1, \mathcal{Y}^2)$. We can bound this quantity as follows:

$$\begin{aligned}
&\big|\Pr[x_1 \leftarrow X_k^1; x_2 \leftarrow X_k^2 : A(x_1,x_2) = 1] - \Pr[y_1 \leftarrow Y_k^1; y_2 \leftarrow Y_k^2 : A(y_1,y_2) = 1]\big|\\
&\quad = \big|\Pr[x_1 \leftarrow X_k^1; x_2 \leftarrow X_k^2 : A(x_1,x_2) = 1] - \Pr\big[y_1 \leftarrow Y_k^1; x_2 \leftarrow X_k^2 : A(y_1,x_2) = 1\big]\\
&\qquad + \Pr\big[y_1 \leftarrow Y_k^1; x_2 \leftarrow X_k^2 : A(y_1,x_2) = 1\big] - \Pr[y_1 \leftarrow Y_k^1; y_2 \leftarrow Y_k^2 : A(y_1,y_2) = 1]\big|\\
&\quad \leq \big|\Pr[x_1 \leftarrow X_k^1; x_2 \leftarrow X_k^2 : A(x_1,x_2) = 1] - \Pr\big[y_1 \leftarrow Y_k^1; x_2 \leftarrow X_k^2 : A(y_1,x_2) = 1\big]\big|\\
&\qquad + \big|\Pr\big[y_1 \leftarrow Y_k^1; x_2 \leftarrow X_k^2 : A(y_1,x_2) = 1\big] - \Pr[y_1 \leftarrow Y_k^1; y_2 \leftarrow Y_k^2 : A(y_1,y_2) = 1]\big|
\end{aligned}$$

(where we have again applied the triangle inequality). The last two terms are exactly Equations (5) and (6), and we know they are negligible. Since the sum of two negligible quantities is negligible, the distinguishing advantage of A is negligible, as desired. ■

This is called a "hybrid argument" for the following reason: Looking at the structure of the proof, we introduced the "hybrid" distribution $(\mathcal{Y}^1, \mathcal{X}^2)$ (which is not equal to either of the distributions we are ultimately interested in) and noted that $(\mathcal{X}^1, \mathcal{X}^2) \overset{c}{\equiv} (\mathcal{Y}^1, \mathcal{X}^2)$ and $(\mathcal{Y}^1, \mathcal{X}^2) \overset{c}{\equiv} (\mathcal{Y}^1, \mathcal{Y}^2)$ (this was the purpose of A_1 and A_2, respectively). Applying Claim 3 (which we essentially re-derived above) gives the desired result.

A similar argument can be used for combinations of poly-many ensembles instead of two ensembles, but we omit the details. Furthermore, a corollary of the above is that if $\mathcal{X} \overset{c}{\equiv} \mathcal{Y}$ then polynomially-many copies of $\mathcal{X}$ are indistinguishable from polynomially-many copies of $\mathcal{Y}$. Formally, let $\ell(k)$ be a polynomial and define $\mathcal{X}^\ell = \{X_k^\ell\}$ as follows:

$$X_k^\ell \overset{\text{def}}{=} \overbrace{(X_k, \ldots, X_k)}^{\ell(k)\ \text{times}}$$

(and similarly for $\mathcal{Y}^\ell$). Then $\mathcal{X}^\ell \overset{c}{\equiv} \mathcal{Y}^\ell$.

Strictly speaking, Claim 5 is not quite enough to yield Theorem 2 directly. The problem is the following: recall that if $(\mathsf{KeyGen}, \mathcal{E}, \mathcal{D})$ is a semantically-secure encryption scheme for a single bit then the following ensembles are computationally indistinguishable:

$$X_k \stackrel{\text{def}}{=} \{(PK, SK) \leftarrow \mathsf{KeyGen}(1^k); C \leftarrow \mathcal{E}_{PK}(0) : (PK, C)\}$$
$$Y_k \stackrel{\text{def}}{=} \{(PK, SK) \leftarrow \mathsf{KeyGen}(1^k); C \leftarrow \mathcal{E}_{PK}(1) : (PK, C)\}.$$

But then applying the hybrid argument directly only tells us that the following are indistinguishable:

$$(X_k, X_k) = \left\{ \begin{array}{c} (PK, SK), (PK', SK') \leftarrow \mathsf{KeyGen}(1^k) \\ C \leftarrow \mathcal{E}_{PK}(0); C' \leftarrow \mathcal{E}_{PK'}(0) \end{array} : (PK, C, PK', C') \right\}$$
$$(Y_k, Y_k) = \left\{ \begin{array}{c} (PK, SK), (PK', SK') \leftarrow \mathsf{KeyGen}(1^k) \\ C \leftarrow \mathcal{E}_{PK}(1); C' \leftarrow \mathcal{E}_{PK'}(1) \end{array} : (PK, C, PK', C') \right\};$$

here, encryption is done with respect to two *different* public keys, not a single key as desired. Even so, the hybrid technique is essentially what is used to prove Theorem 2 and we therefore refer to it as such. As a final remark, note that it is crucial in the proof of Theorem 2 that an adversary can generate random encryptions[4] as can be done in any public-key encryption scheme. In particular, an analogue of Theorem 2 does *not* hold for the case of private-key encryption, where an adversary may be unable to generate legal ciphertexts corresponding to an unknown key.

[4] In fact, this directly parallels the requirement in Claim 5 that the distribution ensembles be efficiently sampleable.