

Homework 1

Deadline: January 29; 2026, 11:59pm ET

Instructions

- The solutions must be submitted via Canvas.
- You must typeset your solutions. We suggest using LaTeX or Typst.

Problems

1. (40 Points) Consider the following variant of one-time perfect security, where Eve can obtain two ciphertexts encrypted under the same key, called **two-time perfect security**

We say that an encryption scheme is two-time perfectly secure if $\forall m_{11}, m_{12}, m_{21}, m_{22} \in \mathcal{M}$, the following distributions are identical:

- $\mathcal{D}_1 := \{c_1 := \text{Enc}(k, m_{11}), c_2 := \text{Enc}(k, m_{12}); k \leftarrow \text{KeyGen}()\}$
- $\mathcal{D}_2 := \{c_1 := \text{Enc}(k, m_{21}), c_2 := \text{Enc}(k, m_{22}); k \leftarrow \text{KeyGen}()\}$

Describe an attack demonstrating that one-time pad does **not** satisfy this security definition.

2. (60 Points) Consider an alternative security definition, called **key-privacy security**, that captures the following: a ciphertext should not reveal any information about the *key* that it was encrypted under. In other words, an adversary that tries to guess the key after seeing a ciphertext should do no better than random guessing. We formalize this into a security definition below.

We say that an encryption scheme satisfies *key-privacy security* if for all adversaries \mathcal{A} and for all messages $m \in \mathcal{M}$, we have

$$\Pr_{k \leftarrow \text{KeyGen}()} [\mathcal{A}(\text{Enc}(k, m)) = k] = \frac{1}{|\mathcal{K}|}.$$

- (a) (30 Points) Prove that key-privacy security does **not** imply one-time perfect security. Do so by defining an encryption scheme, proving that it satisfies key-privacy, and then proving that it does not satisfy one-time perfect security.
- (b) (30 Points) Prove that one-time perfect security does **not** imply key-privacy security. Do so by defining an encryption scheme, proving that it satisfies one-time perfect security, and then proving that it does not satisfy key-privacy.