

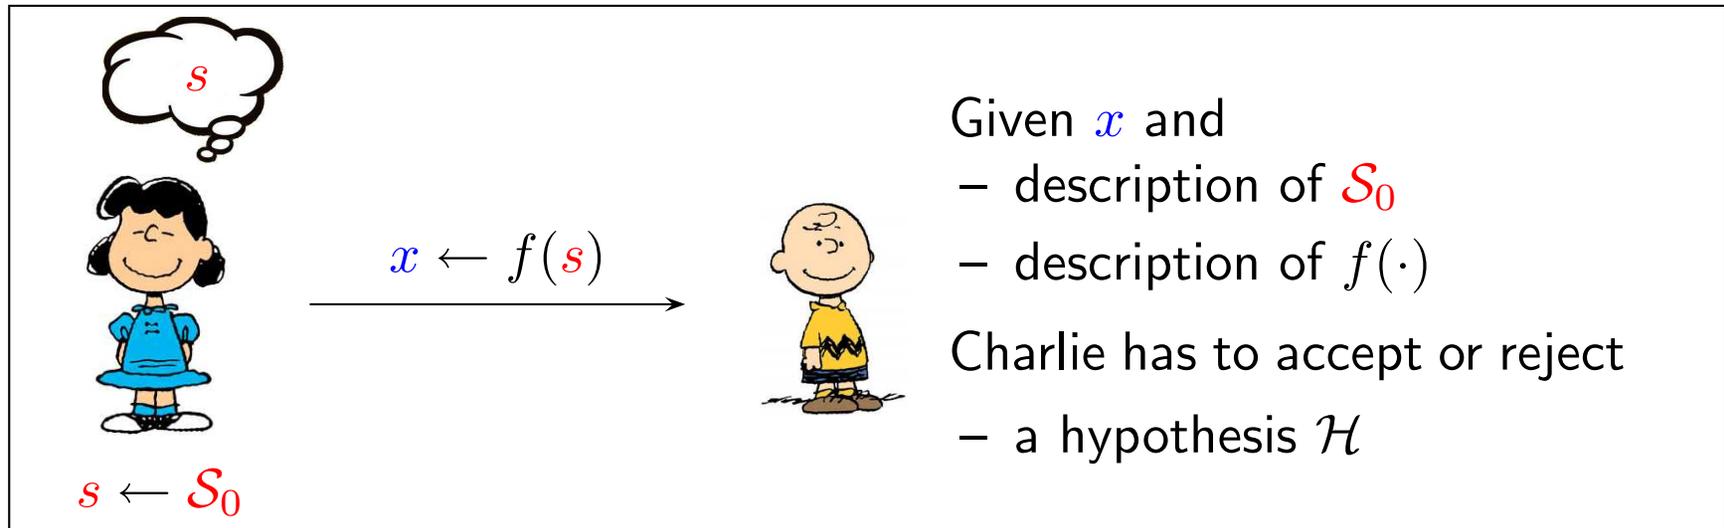
MTAT.07.003 CRYPTOLOGY II

Computational Indistinguishability

Sven Laur
University of Tartu

How to Quantify Secrecy?

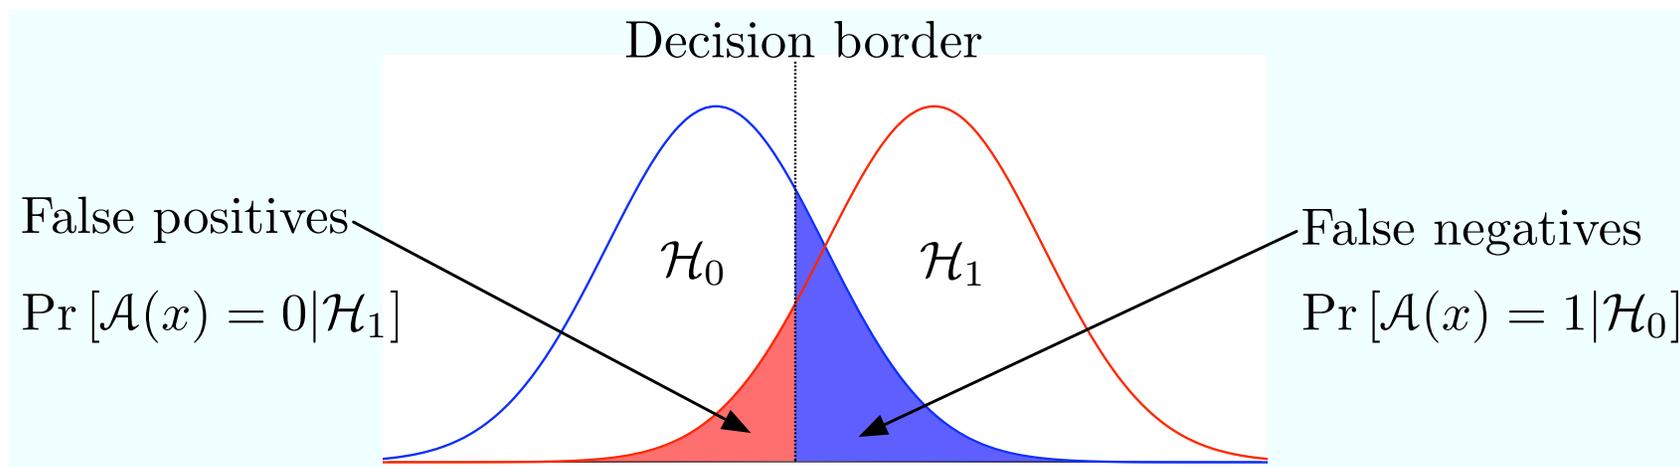
Hypothesis testing scenario



There are several types of hypotheses one can try to resolve:

- ▷ *simple hypotheses* $\mathcal{H} = [s \stackrel{?}{=} s_0]$
- ▷ *complex hypotheses* $\mathcal{H} = [s \stackrel{?}{=} s_0 \vee s \stackrel{?}{=} s_1 \vee \dots \vee s \stackrel{?}{=} s_k]$
- ▷ *trivial hypotheses* that always hold or never hold.

Simple hypothesis testing

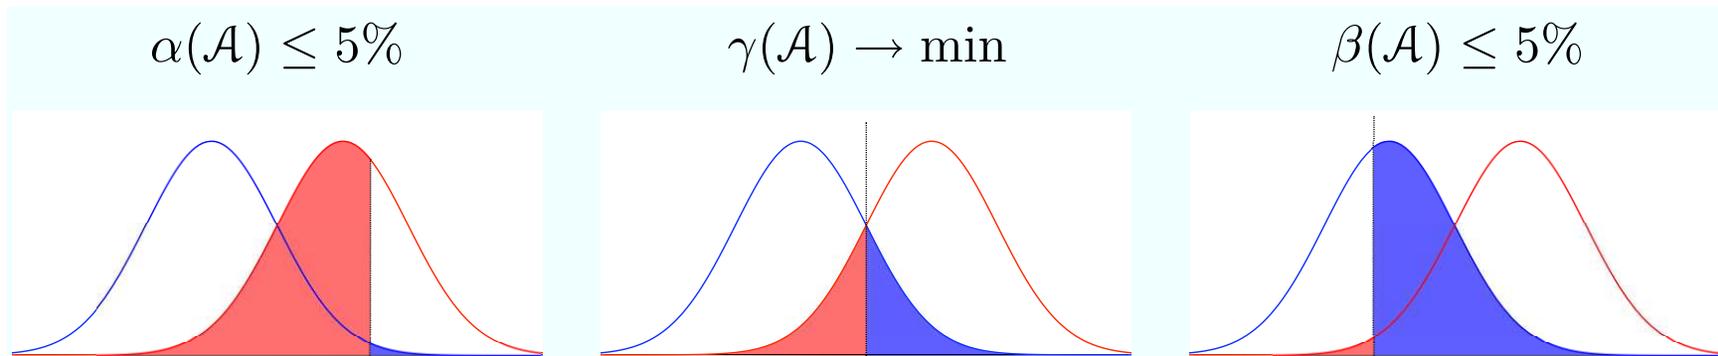


Simple hypothesis \mathcal{H}_0 and \mathcal{H}_1 always determine the distribution of the observable variable $x \leftarrow f(s)$. Consequently, an adversary \mathcal{A} that can choose between two hypothesis \mathcal{H}_0 and \mathcal{H}_1 can do two types of errors:

- ▷ probability of *false negatives* $\alpha(\mathcal{A}) \doteq \Pr[\mathcal{A}(x) = 1 | \mathcal{H}_0]$
- ▷ probability of *false positives* $\beta(\mathcal{A}) \doteq \Pr[\mathcal{A}(x) = 0 | \mathcal{H}_1]$

The corresponding aggregate error is $\gamma(\mathcal{A}) = \alpha(\mathcal{A}) + \beta(\mathcal{A})$.

Various trade-offs



A potential adversary choose between different strategies:

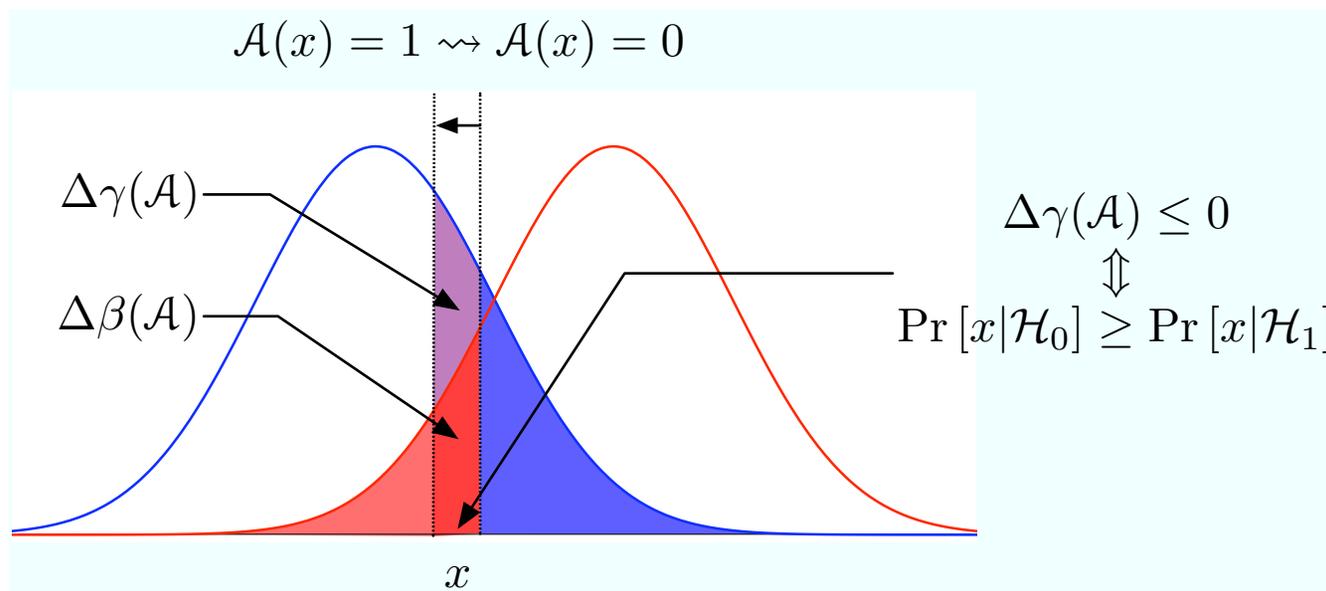
- ▷ Minimise the probability of false positives $\beta(\mathcal{A})$ so that the probability of false negatives $\alpha(\mathcal{A})$ is bounded.
- ▷ Minimise the probability of false negatives $\alpha(\mathcal{A})$ so that the probability of false positives $\beta(\mathcal{A})$ is bounded.
- ▷ Minimise the probability of the aggregate error $\gamma(\mathcal{A}) = \alpha(\mathcal{A}) + \beta(\mathcal{A})$.

Neyman-Pearson theorem

The *likelihood ratio test* described below achieves optimal $\beta(\mathcal{A})$ for any bound $\alpha(\mathcal{A}) \leq \alpha_0$. The aggregate error $\gamma(\mathcal{A})$ is minimised by choosing $\eta = 1$ and using a fair coin to break ties.

$$\mathcal{A}(x) = \begin{cases} 1, & \text{if } \Pr [x|\mathcal{H}_0] < \eta \cdot \Pr [x|\mathcal{H}_1] \\ 0, & \text{if } \Pr [x|\mathcal{H}_0] > \eta \cdot \Pr [x|\mathcal{H}_1] \\ \text{throw a } \rho\text{-biased coin,} & \text{otherwise} \end{cases}$$

The corresponding proof idea



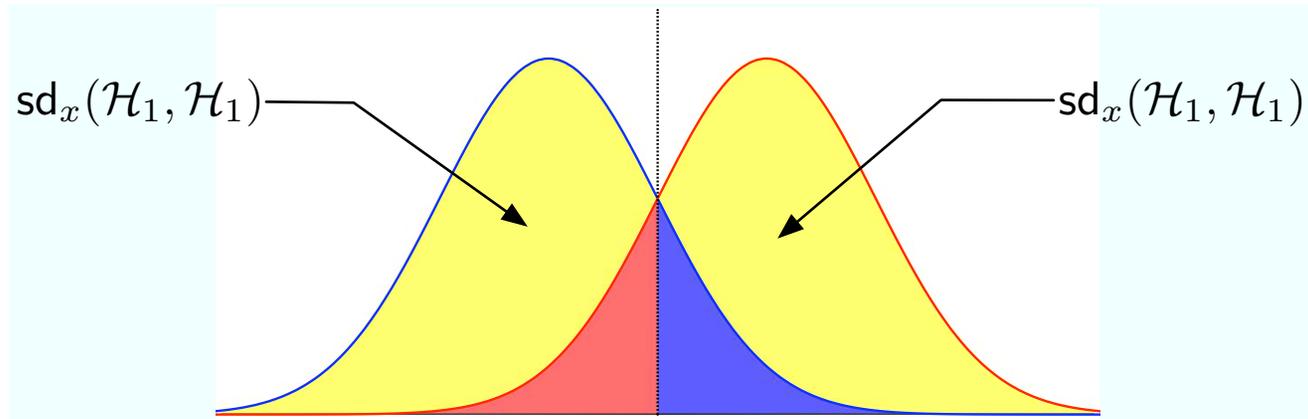
If $\Pr[x|\mathcal{H}_0] \geq \Pr[x|\mathcal{H}_1]$ on for an input x , then by setting $\mathcal{A}(x) = 0$

▷ we decrease the probability of false negatives $\alpha(\mathcal{A})$ by $\Delta\alpha(\mathcal{A})$

▷ we increase the probability of false positives $\beta(\mathcal{A})$ by $\Delta\beta(\mathcal{A})$

By the assumption $\Delta\alpha(\mathcal{A}) \leq \Delta\beta(\mathcal{A})$ and thus the change $\Delta\gamma(\mathcal{A}) \leq 0$.

Statistical distance



Formally, statistical distance is defined as re-scaled ℓ_1 -distance

$$sd_x(\mathcal{H}_1, \mathcal{H}_1) = \frac{1}{2} \cdot \sum_x |\Pr[x|\mathcal{H}_0] - \Pr[x|\mathcal{H}_1]|$$

but there are several other ways how to compute it.

Statistical distance as a limit

For any adversary \mathcal{A} we can define advantage for distinguishing hypotheses:

$$\text{Adv}_{\mathcal{H}_0, \mathcal{H}_1}^{\text{ind}}(\mathcal{A}) = 1 - \gamma(\mathcal{A})$$

$$\Updownarrow$$

$$\text{Adv}_{\mathcal{H}_0, \mathcal{H}_1}^{\text{ind}}(\mathcal{A}) = |\Pr[\mathcal{A}(x) = 1 | \mathcal{H}_0] - \Pr[\mathcal{A}(x) = 1 | \mathcal{H}_1]|$$

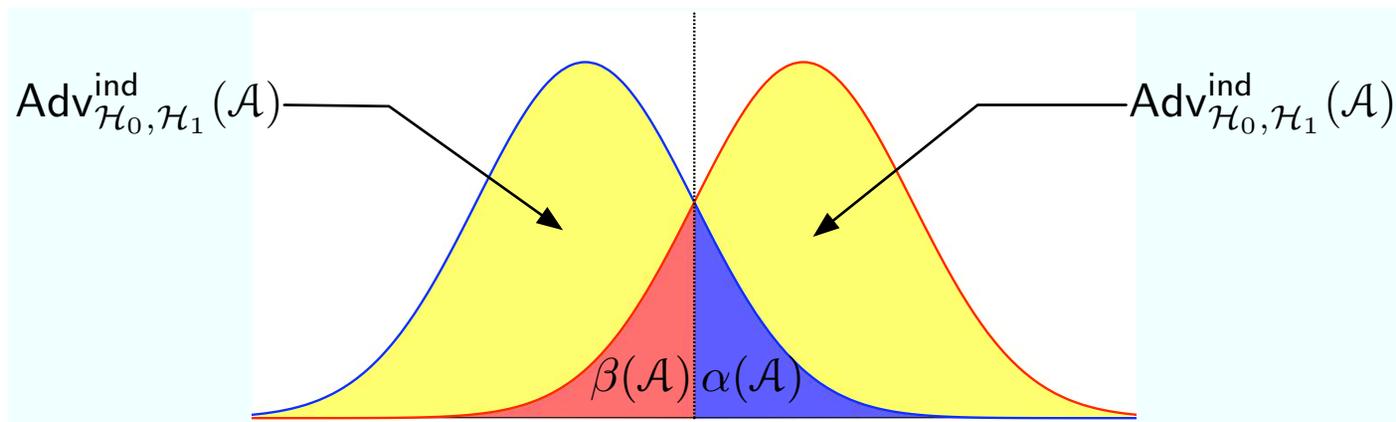
The maximal distinguishing advantage coincides with the statistical distance:

$$\text{sd}_x(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A}} \{\text{Adv}_{\mathcal{H}_0, \mathcal{H}_1}^{\text{ind}}(\mathcal{A})\}$$

$$\Updownarrow$$

$$\text{sd}_x(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A}} \{\Pr[\mathcal{A}(x) = 1 | \mathcal{H}_0] - \Pr[\mathcal{A}(x) = 1 | \mathcal{H}_1]\}$$

The corresponding proof idea



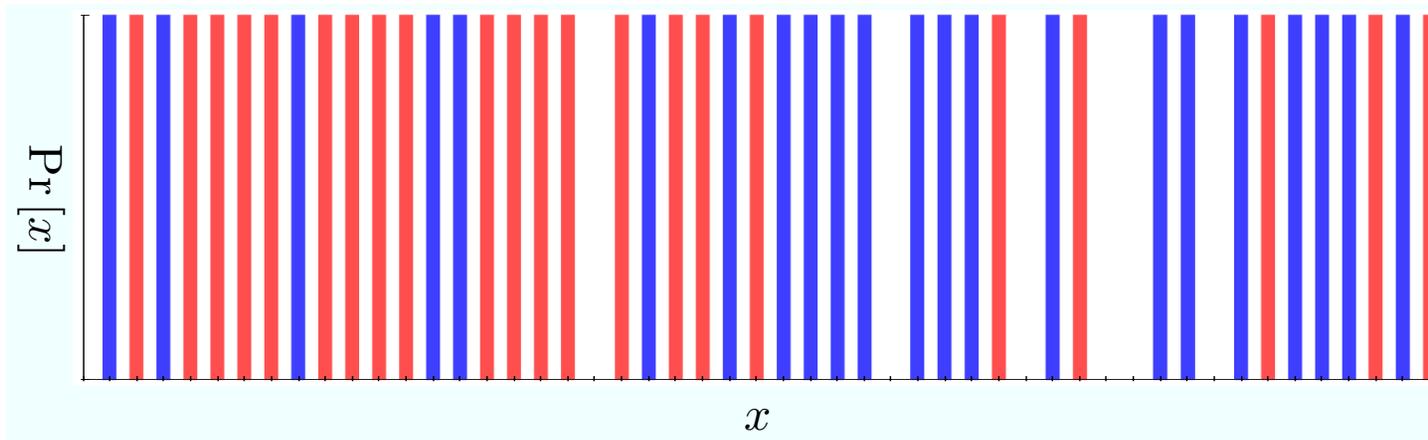
▷ By Neyman-Pearson theorem the optimal adversary behaves as follows

$$\mathcal{A}(x) = \begin{cases} 0, & \text{if } \Pr[x|\mathcal{H}_0] > \Pr[x|\mathcal{H}_1], \\ 1, & \text{if } \Pr[x|\mathcal{H}_0] < \Pr[x|\mathcal{H}_1]. \end{cases}$$

▷ From geometrical considerations $2 \cdot \text{Adv}_{\mathcal{H}_0, \mathcal{H}_1}^{\text{ind}}(\mathcal{A}) = 2 \cdot \text{sd}_x(\mathcal{H}_0, \mathcal{H}_1)$.

What If the Adversary
Is Computationally Bounded?

Infeasibility of statistical distance



The best likelihood ratio test is often infeasible in practice.

- ▷ It is often infeasible to compute $\Pr[x|\mathcal{H}_0]$ and $\Pr[x|\mathcal{H}_1]$.
- ▷ The description of the optimal decision border is too complex to directly hardwire into the description of the distinguishing algorithm.

Instead, we must resort to sub-optimal t -time *distinguishing algorithms*.

Computational distance

Computational distance is defined analogously to the statistical distance:

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A} \text{ is } t\text{-time}} \{ \text{Adv}_{\mathcal{H}_0, \mathcal{H}_1}^{\text{ind}}(\mathcal{A}) \} .$$

As hypotheses uniquely determine observable distributions \mathcal{X}_0 and \mathcal{X}_1 we can talk about indistinguishability of distributions. Distributions \mathcal{X}_0 and \mathcal{X}_1 are (t, ε) -*indistinguishable* if for all t -time algorithms \mathcal{A} :

$$\text{Adv}_{\mathcal{X}_0, \mathcal{X}_1}^{\text{ind}}(\mathcal{A}) = |\Pr [x \leftarrow \mathcal{X}_0 : \mathcal{A}(x) = 0] - \Pr [x \leftarrow \mathcal{X}_1 : \mathcal{A}(x) = 0]| \leq \varepsilon$$

In other terms, the distributions \mathcal{X}_0 and \mathcal{X}_1 are (t, ε) -indistinguishable if

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) \leq \varepsilon .$$

Basic properties of computational distance

Triangle inequality. For any triple of simple hypotheses $\mathcal{H}_0, \mathcal{H}_1, \mathcal{H}_2$:

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_2) \leq \text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) + \text{cd}_x^t(\mathcal{H}_1, \mathcal{H}_2) .$$

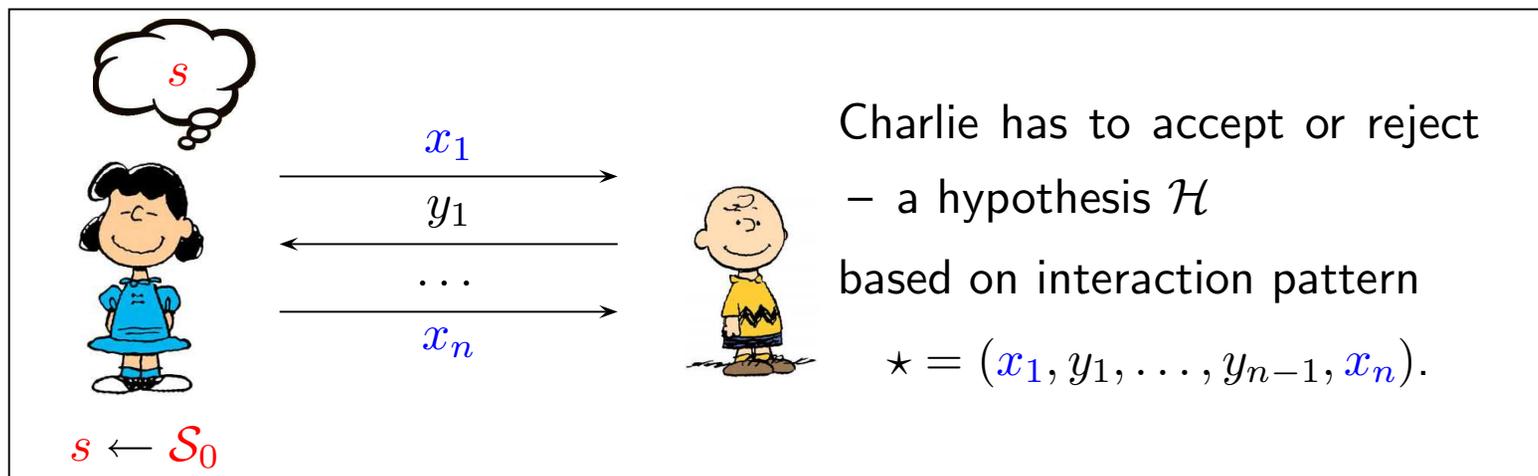
Symmetry. For any two simple hypothesis \mathcal{H}_0 and \mathcal{H}_1 :

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) = \text{cd}_x^t(\mathcal{H}_1, \mathcal{H}_0) .$$

Positively definiteness. For any reasonably large time bound t :

$$\text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) = 0 \quad \Leftrightarrow \quad \text{sd}_x(\mathcal{H}_0, \mathcal{H}_1) = 0 \quad \Leftrightarrow \quad \mathcal{H}_0 \equiv \mathcal{H}_1 .$$

Interactive hypothesis testing



We use analogous notation for computational and statistical distance:

$$\text{cd}_{\star}^t(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A} \text{ is } t\text{-time}} |\Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_0] - \Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_1]| \quad ,$$

$$\text{sd}_{\star}(\mathcal{H}_0, \mathcal{H}_1) = \max_{\mathcal{A}} |\Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_0] - \Pr[\mathcal{A}(\star) = 0 | \mathcal{H}_1]| \quad .$$

These measures also satisfy triangle inequality and other distance axioms.

Security Definitions Based on Indistinguishability

Pseudorandom generators

Model. Let f be a function that stretches m -bit seed s to n -bit string. Then we can consider the following classical hypothesis testing scenario. A t -time adversary \mathcal{A} gets x and must distinguish two hypotheses (*games*):

- ▷ \mathcal{H}_0 : The string x is uniformly chosen over $\{0, 1\}^n$.
- ▷ \mathcal{H}_1 : The string $x \leftarrow f(s)$ for uniformly chosen $s \leftarrow_{\mathcal{U}} \{0, 1\}^m$.

Definition. A function f is (t, ε) -*pseudorandom generator* if $\text{Adv}_f^{\text{prg}}(\mathcal{A}) \leq \varepsilon$ for any t -time adversary \mathcal{A} where the advantage is defined as follows

$$\text{Adv}_f^{\text{prg}}(\mathcal{A}) = |\Pr [x \leftarrow_{\mathcal{U}} \{0, 1\}^n : \mathcal{A}(x) = 0] - \Pr [s \leftarrow_{\mathcal{U}} \{0, 1\}^m : \mathcal{A}(f(s)) = 0]| .$$

Pseudorandom functions

Model. Let \mathcal{F}_{all} denote the set of all functions $f : \mathcal{M} \rightarrow \mathcal{C}$ and let $\mathcal{F} \subseteq \mathcal{F}_{\text{all}}$ be a function family. Then we can consider the following interactive hypothesis testing scenario. A t -time adversary \mathcal{A} that makes at most q calls to the oracle $\mathcal{O}(\cdot)$ in order to distinguish two hypotheses:

- ▷ \mathcal{H}_0 : Oracle chooses $f \xleftarrow{u} \mathcal{F}_{\text{all}}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.
- ▷ \mathcal{H}_1 : Oracle chooses $f \xleftarrow{u} \mathcal{F}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.

Definition. A function family \mathcal{F} is (t, q, ε) -*pseudorandom* if for any t -time adversary \mathcal{A} that makes at most q queries the corresponding advantage

$$\text{Adv}_{\mathcal{F}}^{\text{prf}}(\mathcal{A}) = |\Pr [f \xleftarrow{u} \mathcal{F}_{\text{all}} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0] - \Pr [f \xleftarrow{u} \mathcal{F} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0]| \leq \varepsilon .$$

Pseudorandom permutations

Model. Let \mathcal{F}_{prm} denote the set of all permutations $f : \mathcal{M} \rightarrow \mathcal{M}$ and let $\mathcal{F} \subseteq \mathcal{F}_{\text{prm}}$ be a permutation family. Then we can consider the following interactive hypothesis testing scenario. A t -time adversary \mathcal{A} that makes at most q calls to the oracle $\mathcal{O}(\cdot)$ in order to distinguish two hypotheses:

- ▷ \mathcal{H}_0 : Oracle chooses $f \xleftarrow{u} \mathcal{F}_{\text{prm}}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.
- ▷ \mathcal{H}_1 : Oracle chooses $f \xleftarrow{u} \mathcal{F}$ and for every query x_i replies $y_i \leftarrow f(x_i)$.

Definition. A function family \mathcal{F} is (t, q, ε) -*pseudorandom permutation* if for any t -time adversary \mathcal{A} that makes at most q queries the advantage

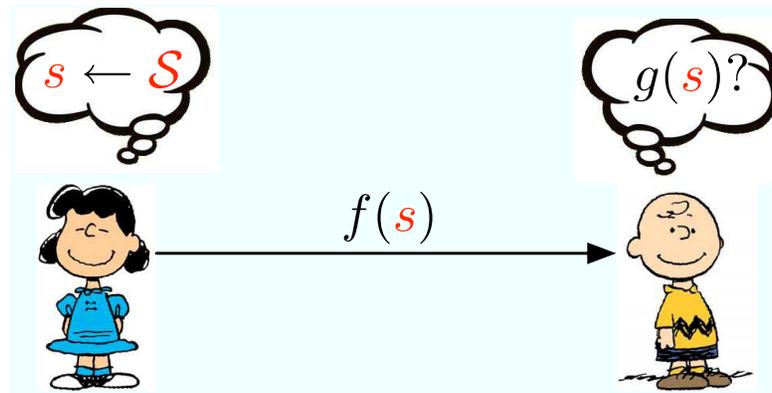
$$\text{Adv}_{\mathcal{F}}^{\text{prf}}(\mathcal{A}) = |\Pr [f \xleftarrow{u} \mathcal{F}_{\text{prm}} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0] - \Pr [f \xleftarrow{u} \mathcal{F} : \mathcal{A}^{\mathcal{O}(\cdot)} = 0]| \leq \varepsilon .$$

Practical implementations

- ▷ **Pseudorandom functions.** Constructing good pseudorandom functions has never been an explicit design goal. Cryptographic hash functions $h : \mathcal{M} \times \mathcal{K} \rightarrow \mathcal{T}$ with implicit or explicit keys are often treated as pseudorandom functions. However, they are also known to contain much more weaknesses than good block ciphers.
- ▷ **Pseudorandom permutations.** Block ciphers are specifically designed to be pseudorandom permutations. This is the most thoroughly studied branch of practical primitive design and we have many good candidates.
- ▷ **Pseudorandom generators.** Stream ciphers are designed to be fast pseudorandom generators. However, we know much more about block ciphers than about stream ciphers. In fact, there is no widely adopted stream cipher standard. There are also more secure constructions based on number theoretical constructions but they are much slower.

Indistinguishability and Guessing Games

Informal definition of semantic security



- ▷ A value $f(s)$ sent to the adversary leaks information.
- ▷ Adversary can try to guess the output of a function $g(s)$.
- ▷ Semantic security is inability to correctly guess the output of $g(\cdot)$.
- ▷ The success of an adversary depends on the functions $f(\cdot)$ and $g(\cdot)$.
- ▷ The success of an adversary depends on the distribution S of secrets.
- ▷ A certain amount of success can be achieved without observing $f(s)$.

The simplest guessing game

Model. Consider the simplest attack scenario:

1. \mathcal{S} is a uniform distribution over two states s_0 and s_1 .
2. \mathcal{H}_0 and \mathcal{H}_1 denote simple hypotheses $[s \stackrel{?}{=} s_0]$ and $[s \stackrel{?}{=} s_1]$.
3. Given $x \leftarrow f(s)$, Charlie must choose between hypotheses \mathcal{H}_0 and \mathcal{H}_1 .

Success bound. The probability of an incorrect guess

$$\begin{aligned}\Pr[\text{Failure}] &= \Pr[\mathcal{H}_0] \cdot \Pr[\mathcal{A}(x) = 1|\mathcal{H}_0] + \Pr[\mathcal{H}_1] \cdot \Pr[\mathcal{A}(x) = 0|\mathcal{H}_1] \\ &= \frac{1}{2} \cdot \left(\underbrace{\Pr[\mathcal{A}(x) = 1|\mathcal{H}_0]}_{\text{False negatives}} + \underbrace{\Pr[\mathcal{A}(x) = 0|\mathcal{H}_1]}_{\text{False positives}} \right) \\ &= \frac{1}{2} + \frac{1}{2} \cdot \text{Adv}_{\mathcal{H}_0, \mathcal{H}_1}^{\text{ind}}(\mathcal{A}) \leq \frac{1}{2} + \frac{1}{2} \cdot \text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1)\end{aligned}$$

Guessing game with a biased coin

Model. Let \mathcal{S} be a distribution over $\{0, 1\}$ such that $\Pr [s \leftarrow \mathcal{S} : s = 0] \leq \frac{1}{2}$ and consider a guessing game \mathcal{G} between a challenger and an adversary \mathcal{A} :

$$\mathcal{G}^{\mathcal{A}} \left[\begin{array}{l} s \leftarrow \mathcal{S} \\ \mathbf{return} [s = \mathcal{A}(f(s))] \end{array} \right]$$

Success bound. For this game, the adversary succeeds with probability

$$\begin{aligned} \Pr [\text{Success}] &= \Pr [\mathcal{H}_0] \cdot \Pr [\mathcal{A} = 0 | \mathcal{H}_0] + \Pr [\mathcal{H}_1] \cdot \Pr [\mathcal{A} = 1 | \mathcal{H}_1] \\ &\leq \Pr [\mathcal{H}_1] \cdot (1 + \Pr [\mathcal{A} = 0 | \mathcal{H}_0] - \Pr [\mathcal{A} = 0 | \mathcal{H}_1]) \\ &\leq \Pr [\mathcal{H}_1] + \Pr [\mathcal{H}_1] \cdot \text{cd}_x^t(\mathcal{H}_0, \mathcal{H}_1) . \end{aligned}$$

Choosing between many values

Model. Let \mathcal{S} be an arbitrary distribution and consider a guessing game

$$\mathcal{G}^{\mathcal{A}} \left[\begin{array}{l} s \leftarrow \mathcal{S} \\ \mathbf{return} [s \stackrel{?}{=} \mathcal{A}(f(s))] \end{array} \right]$$

Success bound. If for all possible states $s_i, s_j \in \text{supp}(\mathcal{S})$ distributions $f(s_i)$ and $f(s_j)$ are (t, ε) -indistinguishable, then for all t -time algorithms

$$\min_{s \in \text{supp}(\mathcal{S})} \Pr [s] - \varepsilon \leq \Pr [\text{Success}] \leq \max_{s \in \text{supp}(\mathcal{S})} \Pr [s] + \varepsilon .$$

The corresponding proof

Let s_* the element with the maximal probability over \mathcal{S} . Then

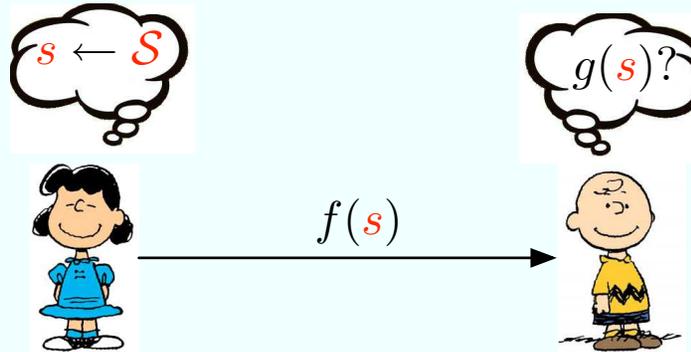
$$\begin{aligned}\Pr[\text{Success}] &= \sum_{s \neq s_*} \Pr[s] \cdot \Pr[\mathcal{A}(f(s)) = s] \\ &\quad + \Pr[s_*] - \sum_{s \neq s_*} \Pr[s_*] \cdot \Pr[\mathcal{A}(f(s_*) = s)] \\ &\leq \Pr[s_*] + \sum_{s \neq s_*} \Pr[s] \cdot \underbrace{|\Pr[\mathcal{A}(f(s)) = s] - \Pr[\mathcal{A}(f(s_*)) = s]|}_{\leq \varepsilon} \\ &\leq \max_{s \in \text{supp}(\mathcal{S})} \Pr[s] + \varepsilon .\end{aligned}$$

The proof of the lower bound is analogous.

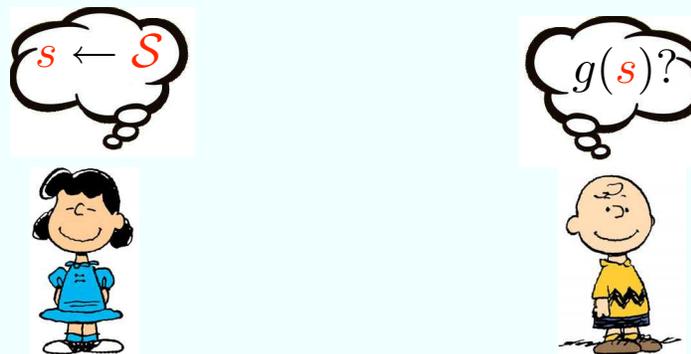
Indistinguishability Implies
Semantic Security

Semantic security

Charlie tries to guess $g(s)$ from the description of \mathcal{S} and $f(s)$.



Charlie tries to guess $g(s)$ solely from the description of \mathcal{S} .



Formal definition

We can define a true guessing advantage

$$\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) = \Pr[\mathcal{G}_0^{\mathcal{A}} = 1] - \Pr[\mathcal{G}_1^{\mathcal{A}} = 1]$$

where the games are defined as follows

$$\begin{array}{l} \mathcal{G}_0^{\mathcal{A}} \\ \left[\begin{array}{l} s \leftarrow \mathcal{S} \\ g_* \leftarrow \mathcal{A}(f(s)) \\ \mathbf{return} [g_* \stackrel{?}{=} g(s)] \end{array} \right. \end{array} \quad \begin{array}{l} \mathcal{G}_1^{\mathcal{A}} \\ \left[\begin{array}{l} s \leftarrow \mathcal{S} \\ g_* \leftarrow \text{argmax}_{g_*} \Pr[g(s) = g_*] \\ \mathbf{return} [g_* \stackrel{?}{=} g(s)] \end{array} \right. \end{array}$$

Obviously, we can express the advantage in more explicit terms

$$\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) = \Pr[s \leftarrow \mathcal{S}_0 : \mathcal{A}(f(s)) = g(s)] - \max_{g_*} \Pr[g(s) = g_*] .$$

Indistinguishability implies semantic security

IND-SEM theorem. If for all $s_i, s_j \in \text{supp}(\mathcal{S})$ distributions $f(s_i)$ and $f(s_j)$ are $(2t, \varepsilon)$ -indistinguishable, then for all t -time adversaries \mathcal{A} :

$$\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) \leq \varepsilon .$$

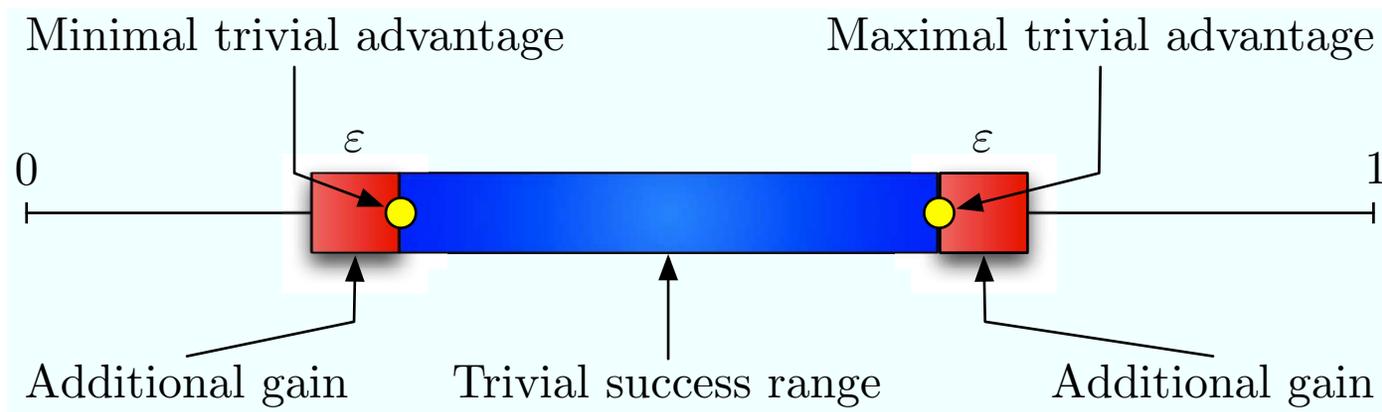
Further comments

- ▷ Note that function g might be randomised.
- ▷ Note that function $g : \mathcal{S} \rightarrow \{0, 1\}^*$ may be extremely difficult to compute.
- ▷ It might be even infeasible to get samples from the distribution \mathcal{S} .
- ▷ The theorem does not hold if \mathcal{S} is specified by the adversary.

Corollary. Under these assumptions, it is *difficult* to predict $f(s)$ from $f(s)$.

Take-home message

As all the results established above are the following graphical form



we can say that

$$\text{Adv}_{\mathcal{G}_0, \mathcal{G}_1}^{\text{ind}}(\mathcal{A}) = |\Pr[\mathcal{G}_0^{\mathcal{A}} = 1] - \Pr[\mathcal{G}_1^{\mathcal{A}} = 1]|$$

bounds the additional gain caused by leakage through published values.