Quantitative approaches to information protection

Pisa, April 2014

Lecture 2

Information theory: useful concepts

- **Entropy** H(X) of a random variable X
  - A measure of the degree of uncertainty of the events
  - It can be used to measure the vulnerability of the secret, i.e. how “easily” the adversary can discover the secret
- **Mutual information** I(S;O)
  - Degree of correlation between the input S and the output O
  - formally defined as difference between:
    - H(S), the entropy of S ***before*** knowing, and
    - H(S|O), the entropy of S ***after*** knowing O
  - It can be used to measure the leakage:

$$\text{Leakage} = I(S;O) = H(S) - H(S|O)$$

  - H(S) depends only on the prior; H(S|O) can be computed using the prior and the channel matrix

Entropy and Operational Interpretation

In the realm of security, there is no unique notion of entropy.
A suitable notion of entropy should have an **operational interpretation** in terms of the kind of **adversary** we want to **model** , namely:

- the kind of attack, and
- how we measure its success

A general **model of adversary** [Köpf and Basin, CCS'07]:

- Assume an oracle that answers yes/no to questions of a certain form.
- The adversary is defined by the form of the questions, and the measure of success of the attack.
- In general we consider the best strategy for the attacker, with respect to a given measure of success.

Entropy

Case 1:

- The questions are of the form: “is S ∈ P ?”
- The measure of success is: the expected number of questions needed to find the value of S in the attacker’s best strategy

Exercise : guessing a password in case of uniform distribution

Example: $S \in \{ a, b, c, d, e, f, g, h \}$

$$p(a) = p(b) = \frac{1}{4} \qquad p(c) = p(d) = \frac{1}{8} \qquad p(e) = p(f) = p(g) = p(h) = \frac{1}{16}$$

It is possible to prove that the best strategy for the adversary is to split each time the search space in two subspaces with probability masses as close as possible. This gives an almost perfectly balanced tree in terms of masses.

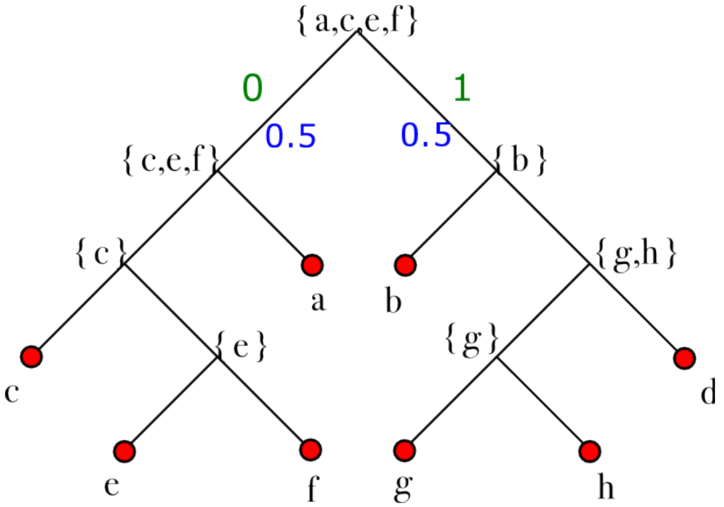

Entropy: Case 1

In the best strategy, the number of questions needed to determine the value of the secret S, when S = *s*, is: **− log *p*(*s*)** (log is in base 2)

This is in case we can construct a *perfectly balanced tree*
In most cases we can only construct an *almost perfectly balanced tree*,
so this formula is an approximation.

hence the **expected number** of question is:

$$H(S) = -\sum_{s} p(s) \log p(s)$$

This is exactly the formula for **Shannon entropy**

Conclusion: For this model of adversary, the degree of protection of the secret, i.e., the degree of difficulty for the adversary to perform his attack, is measured by Shannon entropy

Shannon entropy: information-theoretic int.

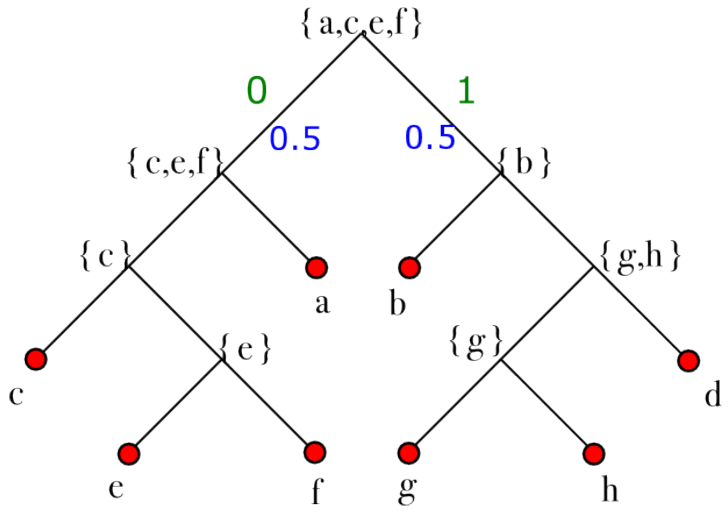

Information-theoretic interpretation:

H(S) is the expected length of the optimal encoding of the values of S

For the strategy in previous example: a: 01 b: 10 c: 000 d: 111 e: 0010 f: 0011 g: 1100 h: 1101

Shannon entropy: properties

In general, the entropy is highest when the distribution is uniform

If $|S| = n$, and the distribution is uniform, then $H(S) = \log n$

$S = \{a, b, c, d, e, f, g, h\} \qquad p(a) = p(b) = \ldots = p(f) = \frac{1}{8}$

$$H(S) \quad = \quad -8\tfrac{1}{8}\log\tfrac{1}{8} \quad = \quad \log 8 \quad = \quad 3$$

$p(a) = p(b) = \frac{1}{4} \qquad p(c) = p(d) = \frac{1}{8} \qquad p(e) = p(f) = p(g) = p(h) = \frac{1}{16}$

$$\begin{aligned} H(S) &= -\textstyle\sum_s p(s)\log p(s) \\ &= -2\tfrac{1}{4}\log\tfrac{1}{4} - 2\tfrac{1}{8}\log\tfrac{1}{8} - 4\tfrac{1}{16}\log\tfrac{1}{16} \\ &= 1 + \tfrac{3}{4} + 1 \\ &= \tfrac{11}{4} \end{aligned}$$

Shannon entropy: properties

The entropy is a concave function of the probability distribution

$S = \{a, b\}$

$p(a) = x \quad p(b) = 1 - x$

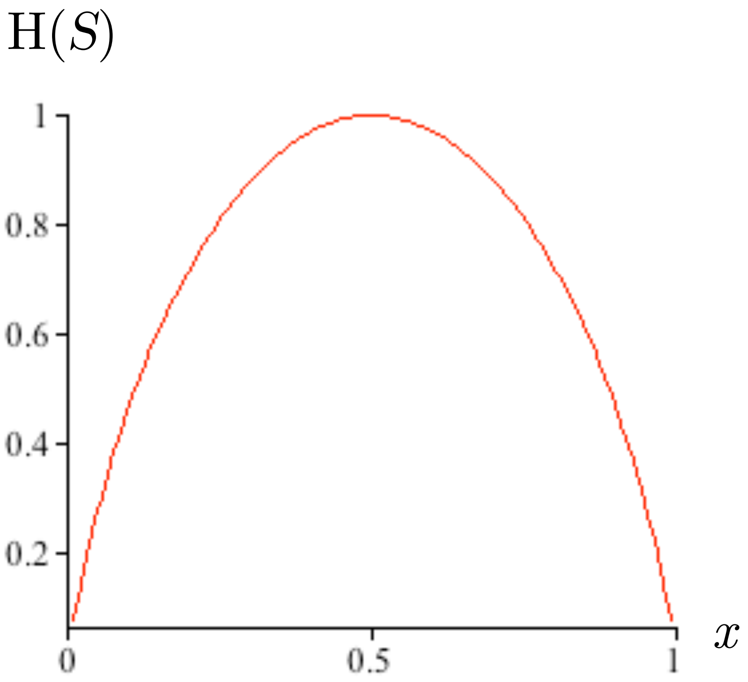

$S = \{a, b, c\}$

$p(a) = x \quad p(b) = y \quad p(c) = 1 - (x + y)$

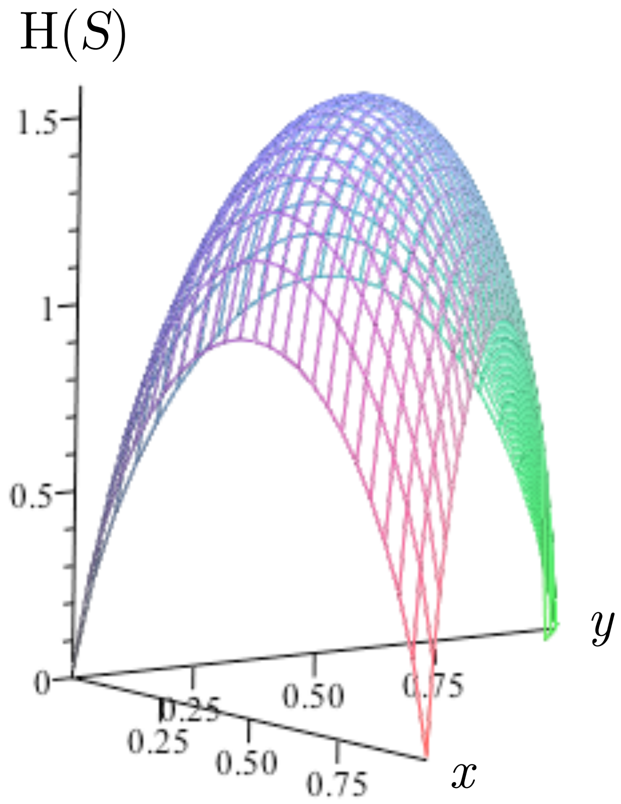

Shannon conditional entropy

An observable o determines a new distribution on S:

$$p(s|o) = p(s)\frac{p(o|s)}{p(o)}$$ Bayes theorem

The entropy of the new distribution on S, given that $O = o$, is:

$$H(S|O = o) \quad = \quad -\sum_{s} p(s|o)\log p(s|o)$$

The conditional entropy is the expected value of the updated entropies:

$$\begin{aligned} H(S|O) &= \sum_{o} p(o)\; H(S|O = o) \\ &= -\sum_{o} p(o) \sum_{s} p(s|o)\log p(s|o) \end{aligned}$$

Shannon mutual information

A priori $$H(S) = -\sum_s p(s) \log p(s)$$

A posteriori $$H(S \mid O) = -\sum_o p(o) \sum_s p(s|o) \log p(s|o)$$

Leakage = Mutual Information $$I(S;O) = H(S) - H(S|O)$$

- In general H(S) ≥ H(S|O)
  - the entropy may increase after one single observation, but in the average it cannot increase

- H(S) = H(S|O) if and only if S and O are independent
  - This is the case if and only if all rows of the channel matrix are the same
  - This case corresponds to strong anonymity in the sense of Chaum

- Shannon capacity C = max I(S;O) over all priors (worst-case leakage)

Example: DC nets. Ring of 2 nodes, b = 1, fair coin

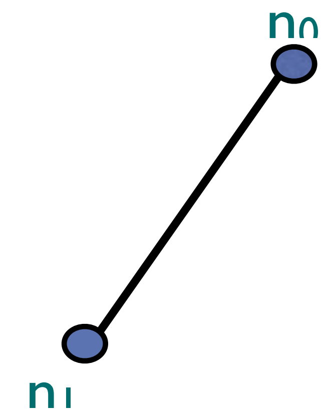

Input S: n_0 , n_1

Output O: the declarations of n_1 and n_0: $d_1 d_0 \in \{01, 10\}$

The entropy of S, as a function of the distribution on

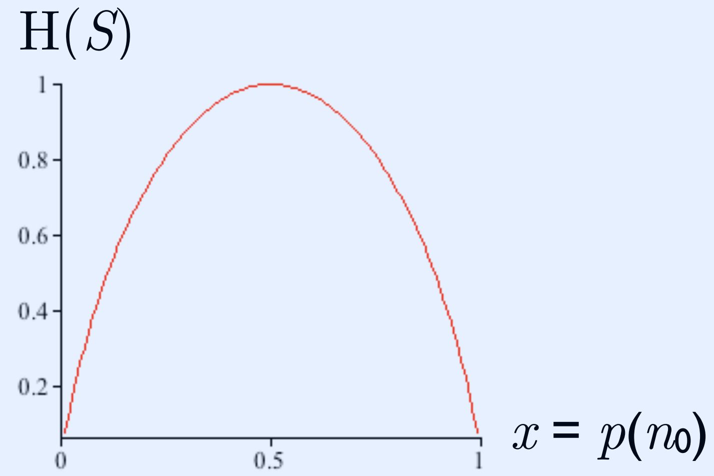

Fair coin: p(0) = p(1) = ½

	01	10
n	½	½
n	½	½

The updated distribution after observation 01

$$
\begin{aligned}
p(n_0|01) &= \frac{p(01|n_0)}{p(01)}\, p(n_0) \\
&= \frac{p(01|n_0)}{p(01|n_0)p(n_0)+p(01|n_1)p(n_1)}\, p(n_0) \\
&= \frac{\frac{1}{2}}{\frac{1}{2}p(n_0)+\frac{1}{2}p(n_1)}\, p(n_0) \\
&= p(n_0)
\end{aligned}
$$

Similarly, after observation 10

$$p(n_0|10) = p(n_0)$$

The entropy of S, as a function of the distribution on

$\mathrm{H}(S)$

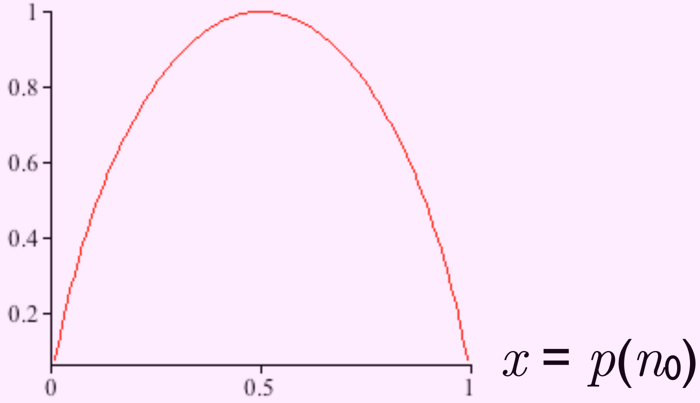

The entropies of S given $O = 01$, and given $O = 10$, as functions of the distribution on S

$\mathrm{H}(S \mid O = 01)$

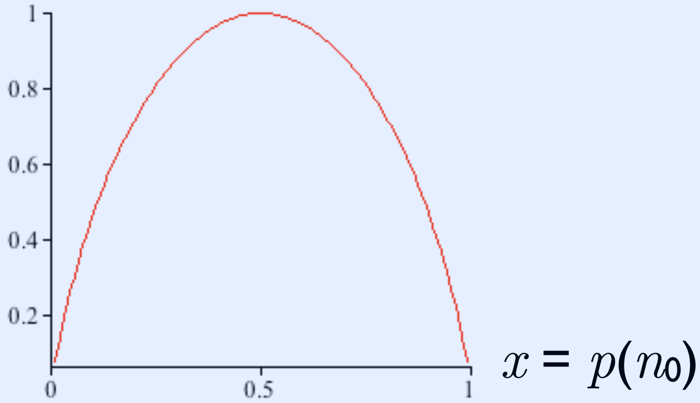

$\mathrm{H}(S \mid O = 10)$

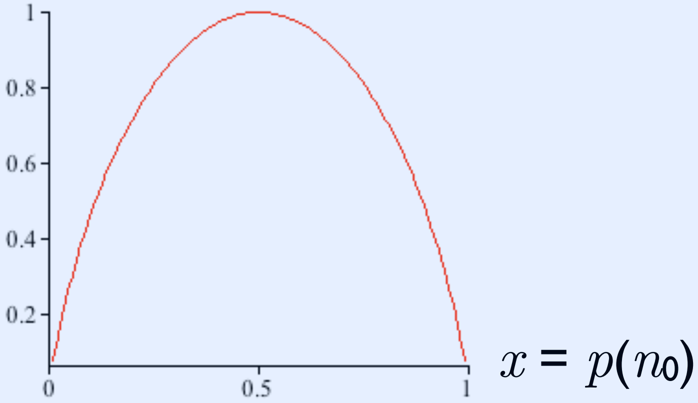

$$H(S|O = 01) \; = \; H(S|O = 10) = H(S)$$

$$\text{Hence } H(S|O) \; = \; p(01)\,H(S|O = 01) + p(10)\,H(S|O = 10) = H(S)$$

Example: DC nets. Ring of 2 nodes, b = 1, biased coin

Input S: n_0 , n_1

Output O: the declarations of n_1 and n_0: $d_1 d_0 \in \{01,10\}$

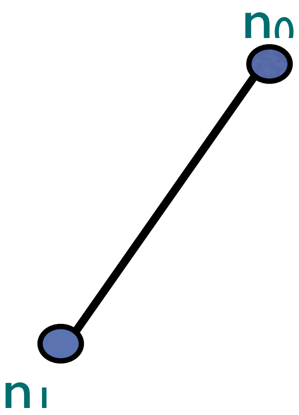

Biased c.: p(0) = ⅔ p(1) = ⅓

	01	10
n	⅔	⅓
n	⅓	⅔

The updated distribution after observation 01

$$\begin{aligned} p(n_0|01) &= \frac{p(01|n_0)}{p(01)}\, p(n_0) \\ &= \frac{p(01|n_0)}{p(01|n_0)p(n_0)+p(01|n_1)p(n_1)}\, p(n_0) \\ &= \frac{\frac{2}{3}}{\frac{2}{3}p(n_0)+\frac{1}{3}p(n_1)}\, p(n_0) \\ &= \frac{\frac{2}{3}}{\frac{2}{3}p(n_0)+\frac{1}{3}(1-p(n_0))}\, p(n_0) \\ &= \frac{2\,p(n_0)}{p(n_0)+1} \end{aligned}$$

The updated distribution after observation 10

$$p(n_0|10) = \frac{p(n_0)}{2-p(n_0)}$$

The entropy of S,

as a function of

the distribution on

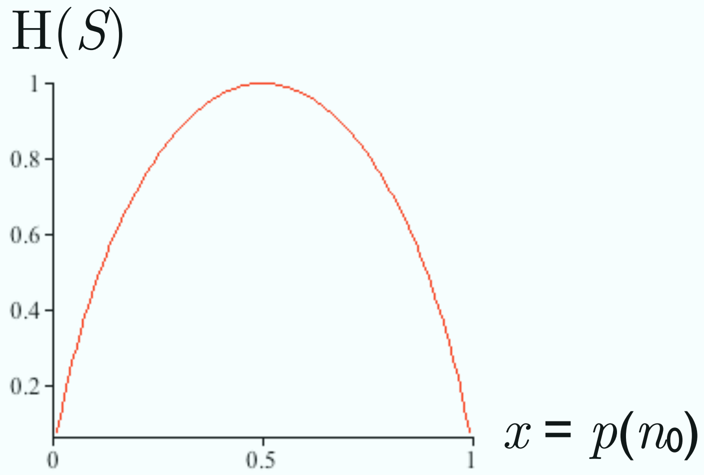

The conditional entropy of S given O as a function of the distribution on S

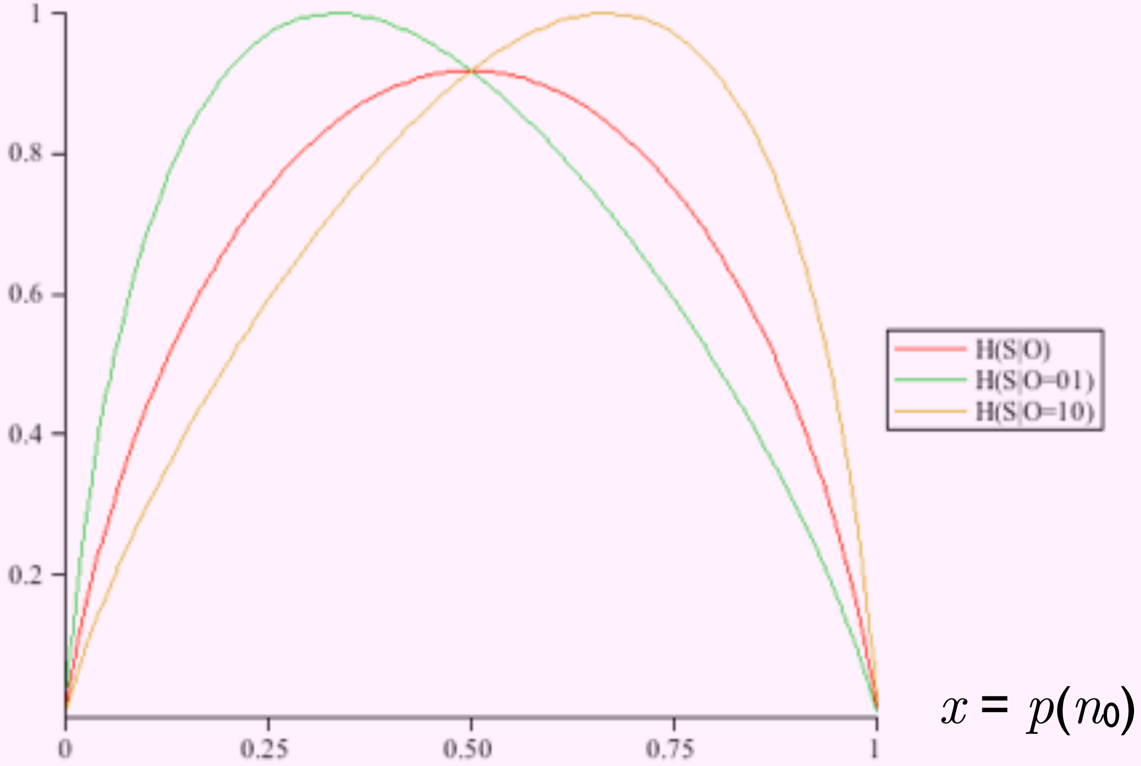

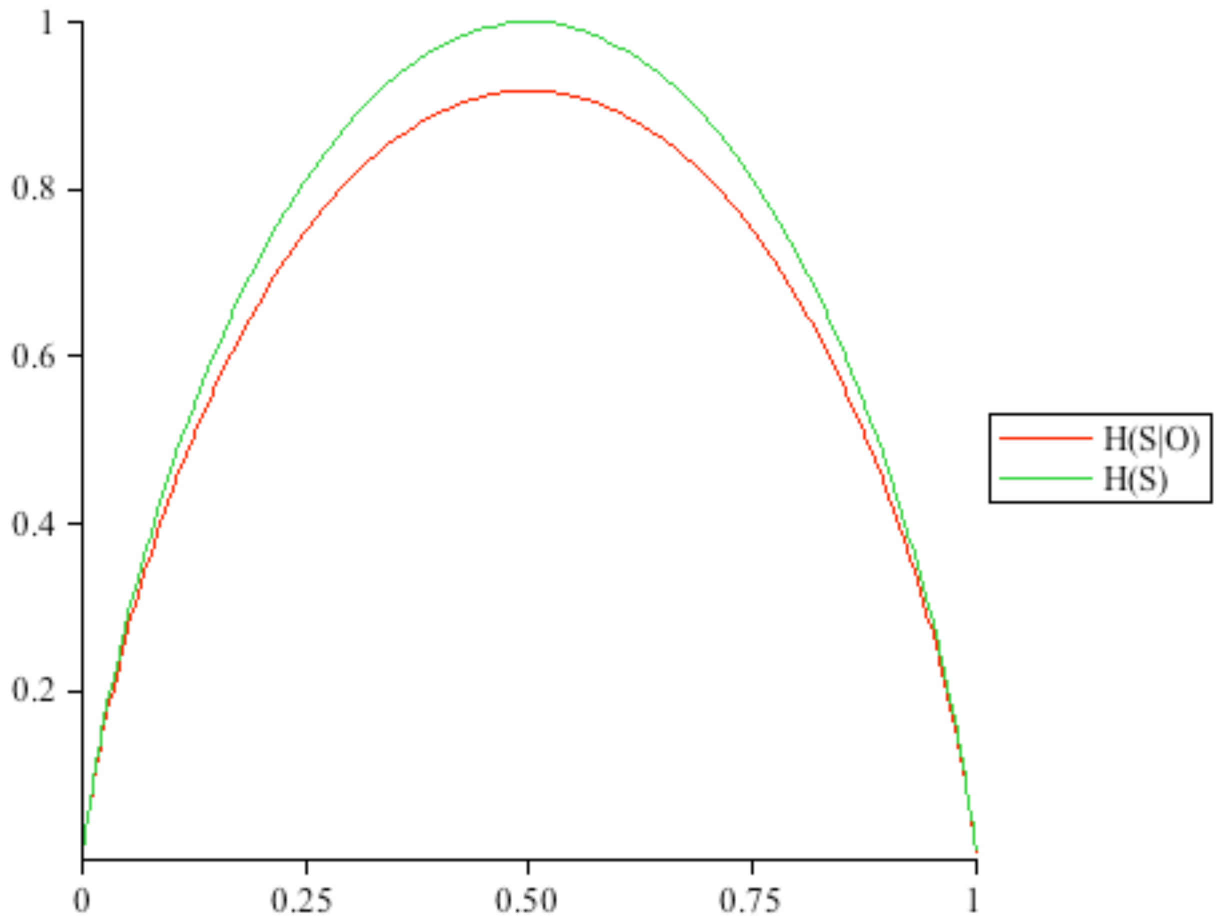

$$\mathrm{I}(S|O) = \mathrm{H}(S) \text{ - } \mathrm{H}(S|O)$$
$$\text{Capacity} = \max \mathrm{I}(S|O)$$

In this example the capacity is about 0.1 bits, and it is obtained when the input distribution is uniform

Exercise 3

- Prove that if the rows of the channel matrix are all equal, then the Shannon Leakage (i.e., Shannon mutual information) is 0.
- Note: The above proof should be relatively easy. The converse is also true, but the proof is much more difficult.

Entropy: Alternative notions

As we argued before, there is no unique notion of vulnerability. It depends on:

- the model of attack, and
- how we measure its success

Consider again the general **model of adversary** proposed by [Köpf and Basin CCS'07] that we saw before:

- Assume an oracle that answers yes/no to questions of a certain form.
- The adversary is defined by the form of the questions and the measure of success.
- In general we consider the best strategy for the adversary, with respect to a given measure of success.

Entropy: Alternative notions

We saw that if

- the questions are of the form: “is S $\in$ P ?”, and
- the measure of success is: the expected number of questions needed to find the value of S in the adversary’s best strategy

then the natural measure of protection is Shannon’s entropy

However, this model of attack does not seem so natural in security, and alternatives have been considered. In particular, the **limited-try attacks**

- The adversary has a limited number of attempts at its disposal
- The measure of success is the probability that he discovers the secret during these attempts (in his best strategy)

Obviously the best strategy for the adversary is to try first the values which have the highest probability

One try attacks: Rényi min-entropy

One-try attacks

- The questions are of the form: “is $S = s$?”
- The measure of success is: $-\log(\max_s p(s))$

The measure of success is Rényi min-entropy: $H_\infty(S) = -\log(\max_s p(s))$

Like in the case of Shannon entropy, $H_\infty(S)$ is highest when the distribution is uniform, and it is 0 when the distribution is a delta of Dirac (no uncertainty).

Towards a notion of leakage based on min-entropy

Leakage = difference between
the a priori vulnerability
and
the a posteriori vulnerability

Leakage = $H_\infty(S) - H_\infty(S \mid O)$

How should we define the conditional min-entropy $H_\infty(S \mid O)$?

Let us recall the conditional entropy in Shannon's case

$$H(S) = -\sum_s p(s) \log p(s)$$ Shannon entropy

An observable o determines a new distribution on S:

$$p(s|o) = p(s)\frac{p(o|s)}{p(o)}$$ Bayes theorem

Define the entropy of the new distribution on S, given that $O = o$, as:

$$H(S|O = o) \quad = \quad -\sum_s p(s|o) \log p(s|o)$$

Define conditional entropy as the expected value of the updated entropies:

$$\begin{aligned} H(S|O) \quad &= \quad \sum_o p(o)\ H(S|O = o) \\ &= \quad -\sum_o p(o) \sum_s p(s|o) \log p(s|o) \end{aligned}$$

Let us try to do the same for the min-entropy case

$$H_\infty(S) = -\log(\max_s p(s))$$ Rényi min-entropy

Define the entropy of the new distribution on S, given that $O = o$, as:

$$H_\infty(S|O=o) \quad = \quad -\log(\max_s p(s|o))$$

Define conditional entropy as the expected value of the updated entropies:

$$\begin{aligned} H_\infty(S|O) &= \sum_o p(o)\ H_\infty(S|O=o) \\ &= -\sum_o p(o)\log(\max_s(s|o)) \end{aligned}$$

However this approach does not work: we would obtain negative leakage!

Conditional min-entropy

Probability of success of an attack on S, given that $O = o$:

$$\Pr_{succ}(S|O=o) \quad = \quad \max_s p(s|o)$$

The expected value of the prob. of success (aka converse of the Bayes risk):

$$\begin{aligned}\Pr_{succ}(S|O) &= \sum_o p(o)\,\Pr_{succ}(S|O=o)\\ &= \sum_o p(o)\max_s\, p(s|o)\\ &= \sum_o \max_s\,(p(o|s)\,p(s))\end{aligned}$$

Now define $H_\infty(S|O) = -\log \Pr_{succ}(S|O)$ [Smith 2009]

Leakage in the min-entropy approach

A priori $H_\infty(S) = -\log \max_s p(s)$

A posteriori $H_\infty(S|O) = -\log \sum_o \max_s (p(o|s) \cdot p(s))$

Leakage = min-Mutual Inf. $I_\infty(S;O) = H_\infty(S) - H_\infty(S|O)$

- In general $I_\infty(S;O) \geq 0$

- $I_\infty(S;O) = 0$ if all rows are the same (but not viceversa)

  Define min-capacity: C_∞ = max $I_\infty(S;O)$ over all priors. We have:
  - $C_\infty = 0$ if and only if all rows are the same
  - C_∞ is obtained on the uniform distribution (but not only)
  - C_∞ = the sum of the max of each column
  - $C_\infty \geq C$