More properties of the leakage

- $H(S) = H_\infty(S) = 0$ iff S is a point probability distribution (aka delta of Dirac), i.e., all the probability mass is in one single value
- The maximum value of $H(S)$ and $H_\infty(S)$ is $\log \#S$
- Shannon mutual information is symmetric: $I(S;O) = I(O;S)$
  Namely: $H(S) - H(S|O) = H(O) - H(O|S)$.
  This does not hold for the min-entropy case
- If the channel is deterministic, then $I(S;O) = H(O)$
- If the channel is deterministic, then $C_\infty = C = \log \#O$

Exercises

1. Prove that $I_\infty(S;O) \geq 0$
2. Prove that if all rows of the channel matrix are equal, then $I_\infty(S;O) = 0$
3. Prove that all rows of the channel matrix are equal if and only if $C_\infty = 0$
4. Compute Shannon leakage and Rényi min-leakage for the password checker (the version where the adversary can observe the execution time), assuming a uniform distribution on the passwords

Rényi min-entropy vs. Shannon entropy

$S = \{a, b\}$

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

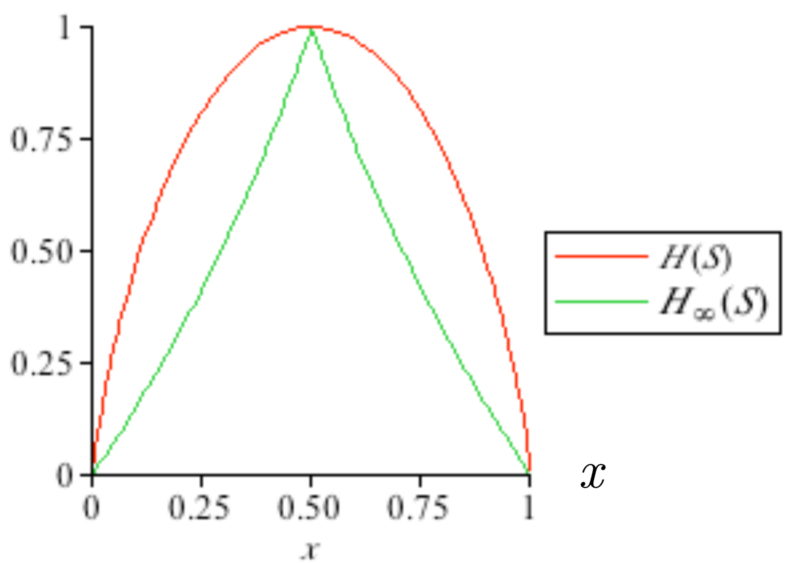

$S = \{a, b, c\}$

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

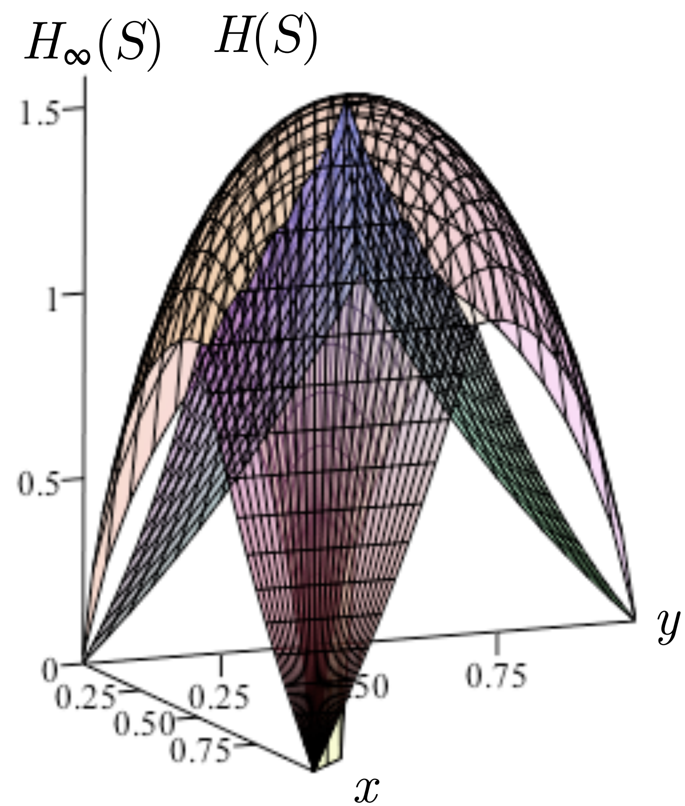

Rényi min entropy and conditional entropy are the log of piecewise linear functions

Shannon capacity vs. Rényi min-capacity

binary channel

a	1−a
b	1−b

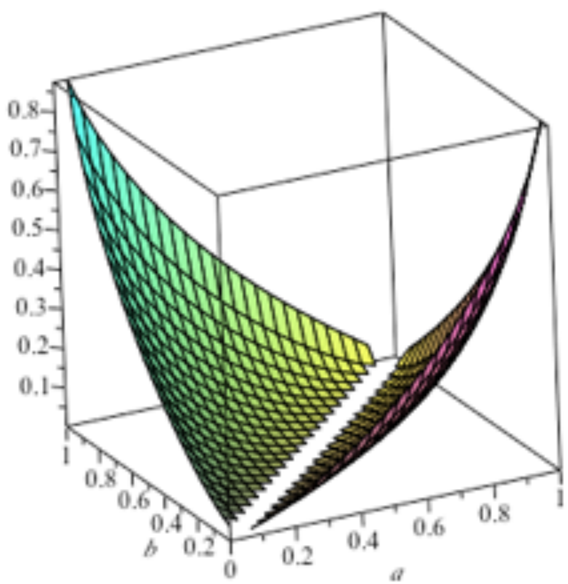

Shannon capacity

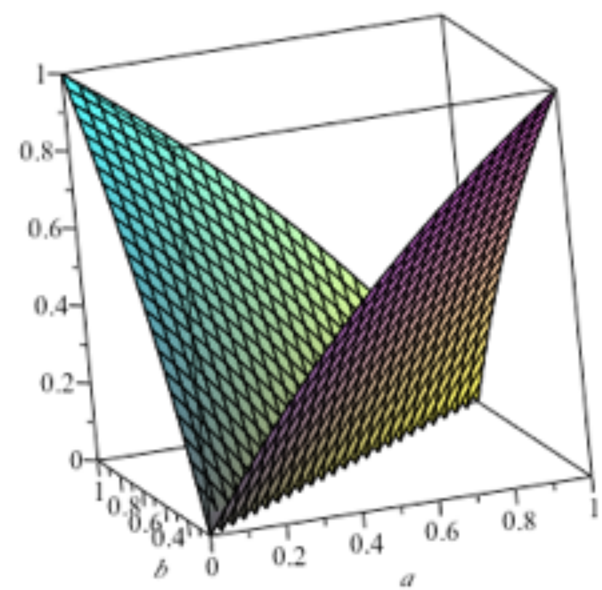

Rényi min-capacity

In general, Rényi min capacity is an upper bound for Shannon capacity

Limitations of min-entropy leakage

- Min-entropy leakage implicitly assumes an operational scenario where adversary $\mathcal{A}$ benefits only by guessing secret S exactly, and in one try.
- But many other scenarios are possible:
  - Maybe $\mathcal{A}$ can benefit by guessing S partially or approximately.
  - Maybe $\mathcal{A}$ is allowed to make multiple guesses.
  - Maybe $\mathcal{A}$ is penalized for making a wrong guess.
- How can any single leakage measure be appropriate in all scenarios?

Notation

- π prior probability
- $x, x_1, x_2 \ldots$ $\mathcal{X}$ secrets
- $x, y_1, y_2 \ldots$ $\mathcal{Y}$ observables
- $w, w_1, w_2 \ldots$ $\mathcal{W}$ guesses
  (they may be different from the secrets)

Gain functions and g-leakage

- We generalize min-entropy leakage by introducing gain functions to model the operational scenario.
- In any scenario, there is a finite set $\mathcal{W}$ of guesses that $\mathcal{A}$ can make about the secret.
- For each guess w and secret value x, there is a gain g(w,x) that $\mathcal{A}$ gets by choosing w when the secret's actual value is x.
- **Definition**: gain function $g : \mathcal{W} \times \mathcal{X} \rightarrow [0, 1]$
- **Example**: Min-entropy leakage implicitly uses

$$g_{id}(w,x) = \begin{cases} 1, \text{ if } w = x \\ 0, \text{ otherwise} \end{cases}$$

g-vulnerability and g-leakage

- Definition: Prior g-vulnerability:

$$V_g[\pi] = \max_w \sum_x \pi[x]g(w,x)$$

"$\mathcal{A}$'s maximum expected gain, over all possible guesses."

- Posterior g-vulnerability:

$$V_g[\pi,C] = \sum_y p(y)\, V_g[p_{X|y}]$$

- g-leakage: $\mathcal{L}_g(\pi,C) = \log V_g[\pi,C] - \log V_g[\pi]$
- g-capacity: $\mathcal{ML}_g(C) = \sup_\pi \mathcal{L}_g(\pi,C)$

The power of gain functions

Guessing a secret approximately.

$g(w,x) = 1 - \text{dist}(w,x)$

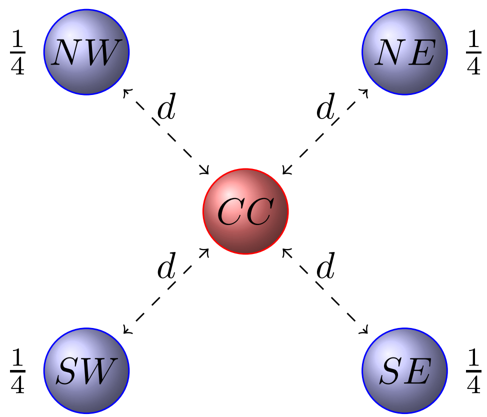

Guessing a property of a secret.

$g(w,x)$ = Is x of gender w?

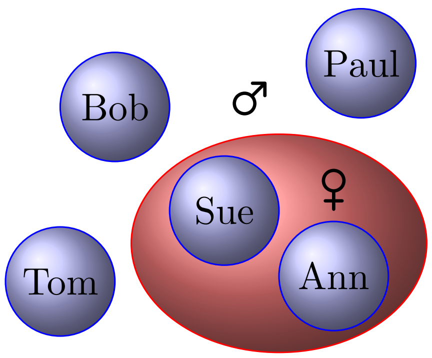

Guessing a part of a secret.

$g(w, x)$ = Does w match the high-order bits of x?

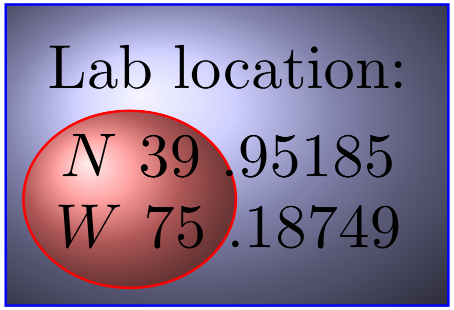

Guessing a secret in 3 tries.

$g_3(w, x)$ = Is x an element of set w of size 3?

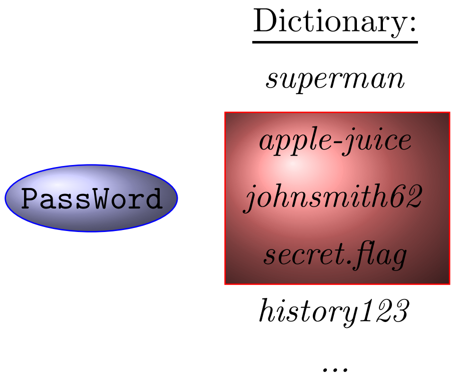

Distinguishing channels with gain functions

- Two channels on a uniformly distributed, 64-bit x:

  A. y = (x or 00000… 0111);

  B. if (x % 8 == 0) then y = x; else y = 1;

  - A always leaks all but the last three bits of x.
  - B leaks all of x one-eighth of the time, and almost nothing seven-eighths of the time.
  - Both have min-entropy leakage of 61 bits out of 64.
- We can distinguish them with gain functions.
- g_8, which allows 8 tries, makes A worse than B.
- g_{tiger}, which gives a penalty for a wrong guess (allowing "⊥" to mean "don't guess") makes B worse.

Robustness worries

- Using g-leakage, we can express precisely a rich variety of operational scenarios.
- But we could worry about the **robustness** of our conclusions about leakage.
- The g-leakage $\mathcal{L}_g(\pi,C)$ depends on both π and g.
  - π models adversary $\mathcal{A}$'s prior knowledge about X
  - g models (among other things) what is valuable to $\mathcal{A}$.
- How confident can we be about these?
- Can we minimize sensitivity to questionable assumptions about π and g?

Capacity results

- **Capacity** (the maximum leakage over all priors) eliminates assumptions about the prior π.
- Capacity relationships between **different** leakage measures are particularly useful.
- **Theorem**: Min-capacity is an upper bound on Shannon capacity: $\mathcal{ML}(C) \geq \mathcal{SC}(C)$.
- **Theorem** ("Miracle"): Min-capacity is an upper bound on g-capacity, for **every** g: $\mathcal{ML}(C) \geq \mathcal{ML}_g(C)$.
  - Hence if C has small min-capacity, then it has small g-leakage under every prior and every gain function.
  - (Note that the choice of g does affect both the prior and the posterior g-vulnerability.)