Aalto University
School of Science

Complexity of Statistical Attacks

Céline Blondeau

March 2016

Spring School on Symmetric Cryptography, Bochum, Germany

Outline

Key-Alternating Iterative Cipher

We consider iterative block ciphers

- operating on n-bit messages
- using a master key K
- with round function F using subkeys K_i
- $$y = E_K(x) = F_{K_r} \circ F_{K_{r-1}} \circ \cdots \circ F_{K_1}(x \oplus K_0),$$

  with $F_{K_i}(x) = F(x \oplus K_i)$

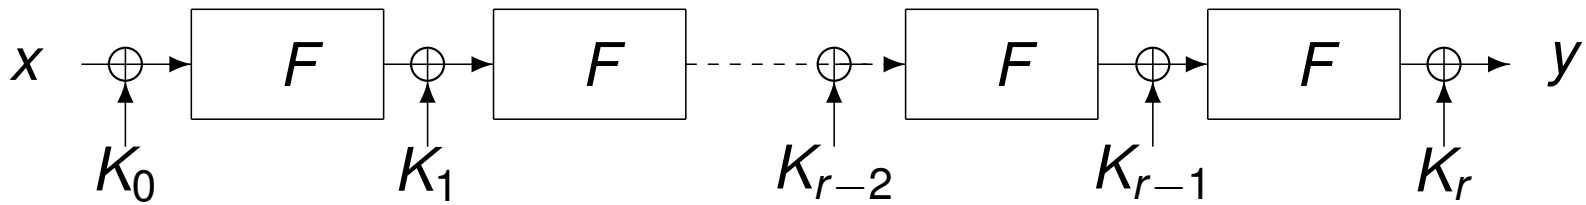

Differential Cryptanalysis [Biham Shamir 90]

Difference between plaintext and ciphertext pairs

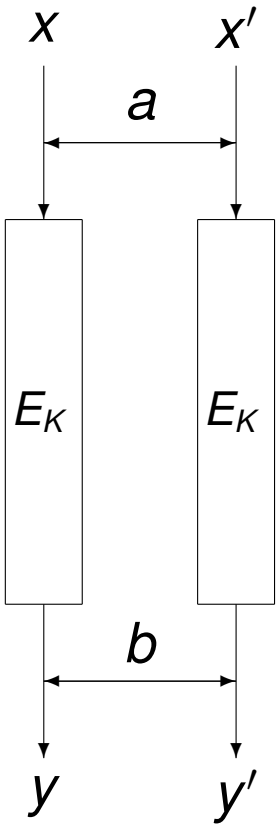

Input difference : a

Output difference : b

Differential probability :

$$\begin{aligned} p_R(K) &= P[a \to b] \\ &= P_{\mathbf{X}}[\, E_K(x) \oplus E_K(x \oplus a) = b \,] \end{aligned}$$

Expected differential probability :

$$\begin{aligned} p_R &= \mathrm{Exp}_{\mathbf{K}}[p_R(K)] \\ &= P_{\mathbf{X},\mathbf{K}}[\, E_K(x) \oplus E_K(x \oplus a) = b \,] \end{aligned}$$

Uniform probability :

$$p_W = 2^{-n}$$

Truncated Differential Attacks [Knudsen 94]

- Set of input differences : $a \in A$
- Set of output differences : $b \in B$
- $$\mathbf{P}[A \to B] = \frac{1}{|A|} \sum_{a \in A} \sum_{b \in B} P[a \to b]$$
- Uniform probability : $p_W = \frac{|B|}{2^n}$

Exercise :

- Why averaging over the input differences

Linear Attacks

Linear relation involving plaintext, key and ciphertext bits

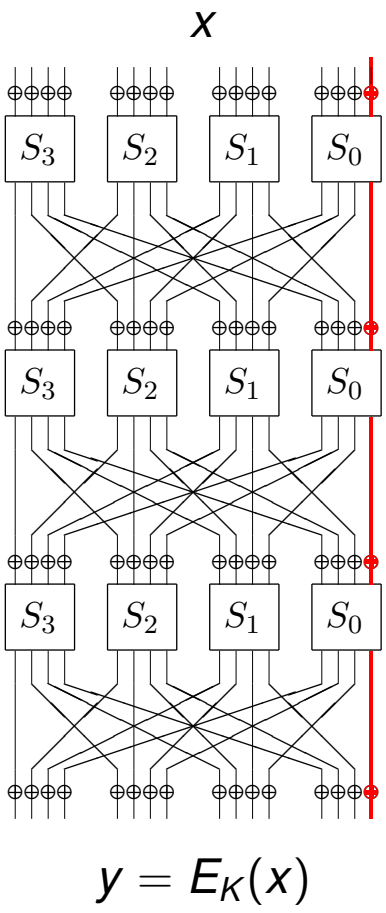

Input mask : u
(Key mask : κ)
Output mask : v

Bias :

$$\varepsilon(K) = 2^{-n}\#\{x \in \mathbb{F}_2^n | u \cdot x(\oplus\kappa \cdot K) \oplus v \cdot y = 0\} - \frac{1}{2}$$

Correlation :

$$\mathrm{cor}_{\mathbf{X}}(u, v)(K) = 2\varepsilon(K)$$

Expected correlation :

$$\mathrm{cor}_{\mathbf{X},\mathbf{K}}(u, v)(K) = \mathrm{Exp}_{\mathbf{X},\mathbf{K}}\left[\mathrm{cor}_{\mathbf{X}}(u, v)(K)\right]$$

Outline

Last Rounds Linear Attack (Matsui's Algorithm 2)

- A linear approximation with masks (u, v) on r rounds
- Partially decrypt the last rounds to find information on the key

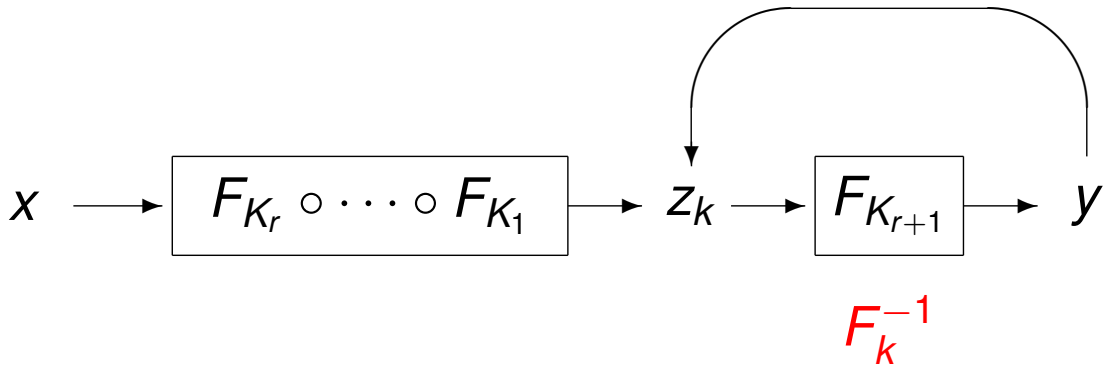

There are 2^κ last round key candidates

Initialize a counter for each key candidate : $T(k)$

Increment $T(k)$ if $u \cdot x \oplus v \cdot z_k = 0$

Last Rounds Differential Attack

- A differential (a, b) on r rounds
- Partially decrypt the last rounds to find information on the key

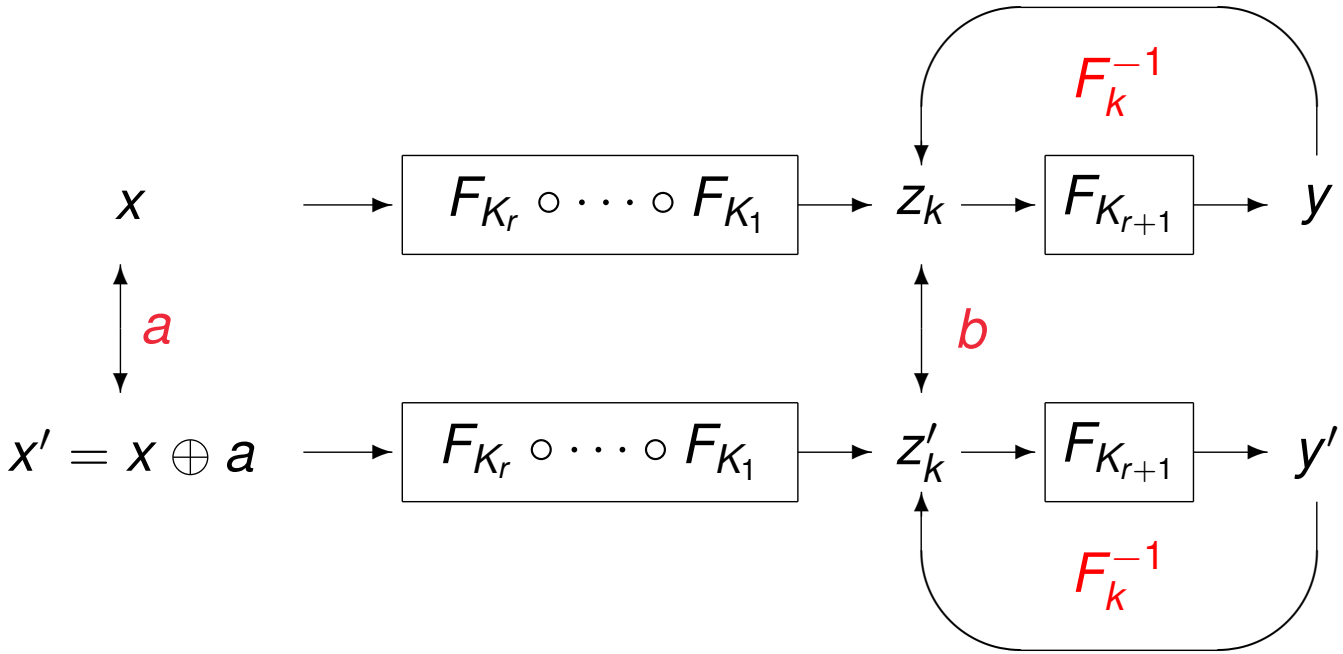

There are 2^κ last round key candidates.
Initialize a counter for each key candidate : $T(k)$
Increment $T(k)$ if $z_k \oplus z_{k'} = b$

Analyzing Phase

- Sort the 2^{κ} key counters $T(k)$ according to their value
- k_R : the good one
- The key candidate corresponding to the good one k_R is among the first ones
- Build the list L of the $\#L$ most likely candidates and try all corresponding master keys (exhaustive search phase)

Data complexity : Number of used plaintexts (denoted by N)

Success probability : Probability that the good key candidate k_R is in the list (denoted by P_S)

The Last-Round Trick

- Wrong Key Assumption in the differential context If on the last round a wrong key candidate is used to decrypt the ciphertext then the differences with a fixed plaintext difference are uniformly distributed

- Wrong Key Assumption in the linear context If on the last round a wrong key candidate is used to decrypt the ciphertext then the linear approximation is equal to zero for "half" of the plaintext/ciphertext pairs

Outline

Hypothesis Testing

- The attacker performs a guess on a subkey K of the cipher and wishes to know whether this guess is correct or not
- There are two possibilities :
  - H_R : k is the correct guess (the right key is denoted by k_R)
  - H_W : k is not the correct guess (a wrong key is denoted by k_W)

The attacker has a certain way of distinguishing the right subkey and a certain amount of plaintext/ciphertext pairs from which he is able to calculate N binary values $X_1, X_2, \ldots, X_N$ which are independent and identically distributed and satisfy

$$p_R = P(X_i = 1|H_R), \qquad p_W = P(X_i = 1|H_W)$$

Hypothesis Testing

- From the samples $X_1, X_2, \ldots, X_N$ the attacker either decides that H_R holds or that H_W is true. Two kind of errors are possible :
  - Non-detection : Occurs if one decides that k is a wrong subkey when H_R holds
  - We denote by α the non-detection error probability
  - False alarm : Occurs if one decides that k is the right subkey when H_W holds
  - We denote by β the false alarm error probability
- In the literature α and β are sometimes denoted by α_0 and α_1

Relation between Terminologies

- To which quantities relate the non-detection and the false-alarm error probability?

Relation between Terminologies

- To which quantities relate the non-detection and the false-alarm error probability?
  - $\alpha = 1 - P_s$
  - $\beta = \frac{\#L - 1}{2^\kappa} \approx \frac{\#L}{2^\kappa}$

Relation between Terminologies

- To which quantities relate the non-detection and the false-alarm error probability?
  - $\alpha = 1 - P_s$
  - $\beta = \frac{\#L - 1}{2^\kappa} \approx \frac{\#L}{2^\kappa}$
- Advantage a : number of key bits “won” in an attack

$$2^{-a} = \beta$$

Sample

- A sample is a set of collected data necessary to measure the involved quantity
- What is a sample in the classical linear context? in the classical differential context?

Sample

- A sample is a set of collected data necessary to measure the involved quantity
- What is a sample in the classical linear context? in the classical differential context?
- For example in a last round key-recovery attack
  - Linear $\mathcal{S} = \{x, z_k\}$
  - Differential $\mathcal{S} = \{(x, x'), (z_k, z'_k)\}$

The Notion of Structure

- Given the 3 input differences $a_1, a_2, a_1 \oplus a_2$, how many pairs can we form with 4 chosen messages?

The Notion of Structure

- Given the 3 input differences $a_1, a_2, a_1 \oplus a_2$, how many pairs can we form with 4 chosen messages?
- 6 pairs
- A structure : a set of messages $(x, x \oplus a_1, x \oplus a_2, x \oplus a_1 \oplus a_2)$
- The number of messages depends on the number of structures, and on the partially encrypted rounds (at the beginning)
- In some cases, I will denote by N_S the number of available samples (number of pairs which can be formed using N plaintexts)
- For the statistical model the notation N corresponds to the number of samples

Scoring Function

- $T(k)$: Scoring function (depends on the data)
- This value is obtained by the analysis of the different samples
- To know the number of plaintexts necessary to the attack, we have to simulate the behavior of this function
- We usually model it using some random variable which follows some well known statistical distribution
- $\mathcal{T}(k)$: random variable associated to $T(k)$

How to Determine the Right Key

Two approaches which lead to similar results :

- We fix the size $\#L$ of the list L of kept key candidates

$$k \in L, k' \notin L \Rightarrow T(k) > T(k')$$

Use order statistic tools : [Selçuk 08]

- We fix the threshold Θ and we keep all keys k with counter $T(k) \geq \Theta$

$$L = \{k, T(k) \geq \Theta\}$$

$$\begin{aligned} \alpha &= Pr(\mathcal{T}_R < \Theta) \\ \beta &= Pr(\mathcal{T}_W \geq \Theta) \end{aligned}$$

Outline

Bernoulli Distribution

- If $\mathcal{T}$ is a random variable with Bernoulli distribution, we have:

$$\Pr(\mathcal{T} = 1) = 1 - \Pr(\mathcal{T} = 0) = p$$

- The probability mass function (pmf or pdf) of this distribution is

$$f(p) = \begin{cases} p & \text{if the condition is fulfilled,} \\ 1 - p & \text{if the condition is not fullfilled} \end{cases}$$

- Equivalently for $j = 0$ or 1 $\quad f(p, j) = p^j (1 - p)^{1-j}$

Binomial Distribution

- A sum of Bernoulli random variables follows a Binomial distribution $\mathcal{T} \sim \mathcal{B}(N, p)$
- The pmf of this distribution is

$$f(N, p, j) = P[\mathcal{T} = j\,] = \binom{N}{j} p^j (1-p)^{N-j}$$

- The cumulative distribution function (cdf) is

$$F(p, i) = P[\mathcal{T} \leq i\,] = \sum_{j=0}^{i} \binom{N}{j} p^j (1-p)^{N-j}$$

- Mean : $Exp[\mathcal{T}] = Np$
- Variance : $Var[\mathcal{T}] = Np(1-p)$

Poisson Distribution

- $\mathcal{T} \sim \mathcal{P}(\lambda)$ if the pmf of $\mathcal{T}$ is

$$f(\lambda, j) = \Pr(\mathcal{T} = j) = \frac{\lambda^j e^{-\lambda}}{j!}$$

- Mean : $Exp(\mathcal{T}) = \lambda$ Variance : $Var(\mathcal{T}) = \lambda$
- Can be applied when we have a large number of events which are rare
- The binomial distribution converges towards the Poisson distribution as the number of trials goes to infinity while the product Np remains fixed, in this case $\lambda = Np$

Normal Distribution

- $\mathcal{T} \sim \mathcal{N}(N, p)$ if the probability density function (pdf) of $\mathcal{T}$ is

$$f(\mu, \sigma^2, x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

  with $\mu = Exp(\mathcal{T})$ and $\sigma^2 = Var(\mathcal{T})$

- The binomial distribution is approximated by a normal distribution for any fixed p (even if p is small) as N is taken to infinity (usually when the product Np is large)

- In this case $\mu = Np$ and $\sigma^2 = Np(1-p)$

Binomial PMF and Approximations

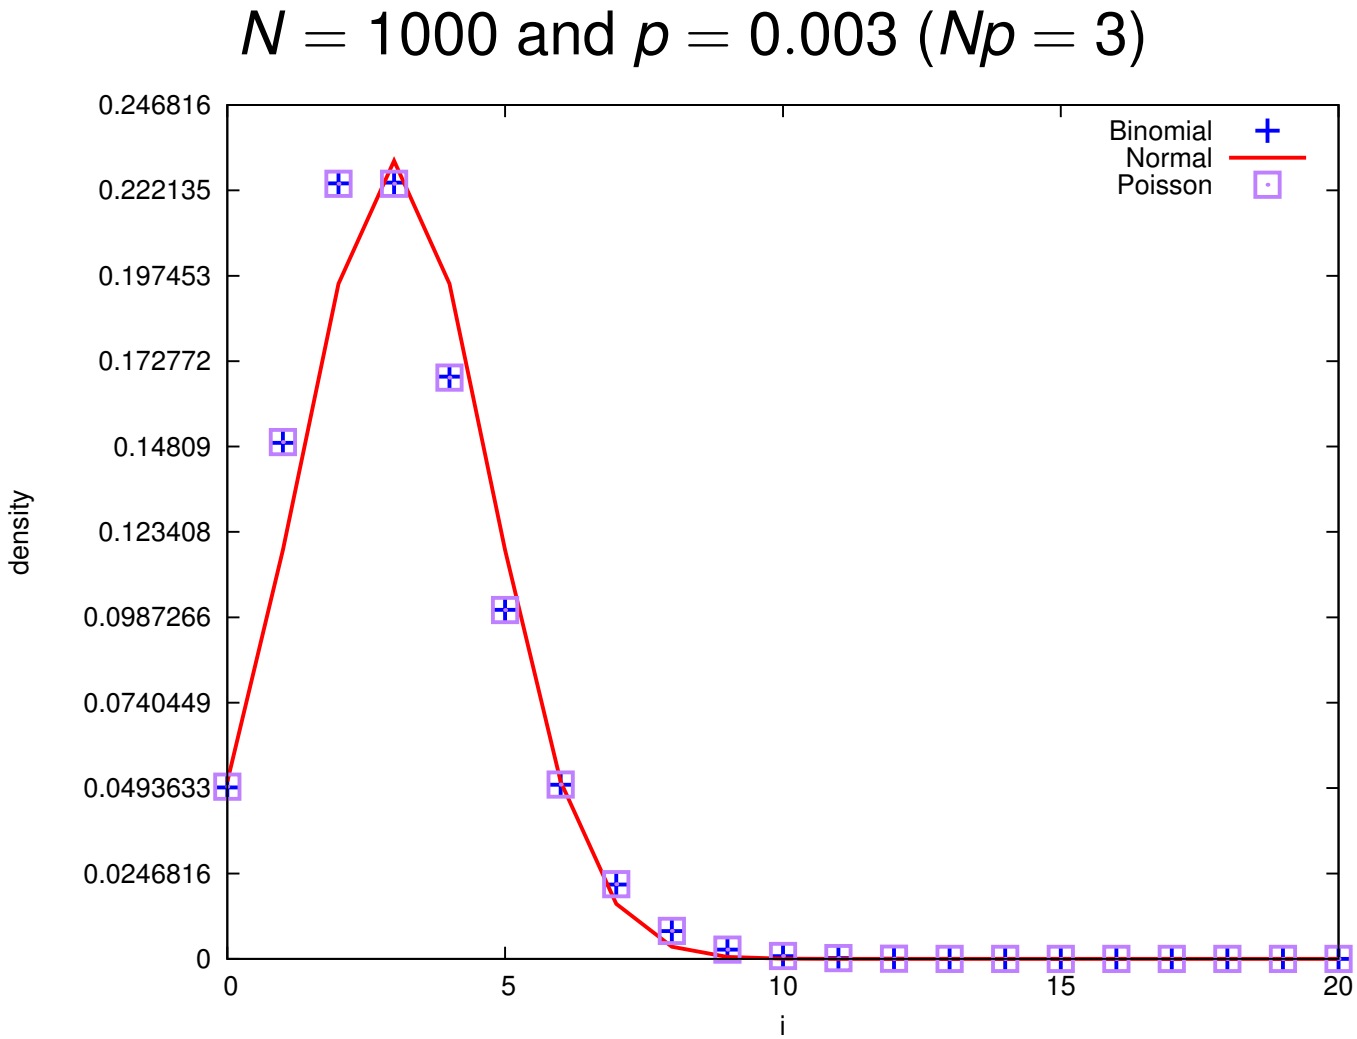

Binomial PMF and Approximations

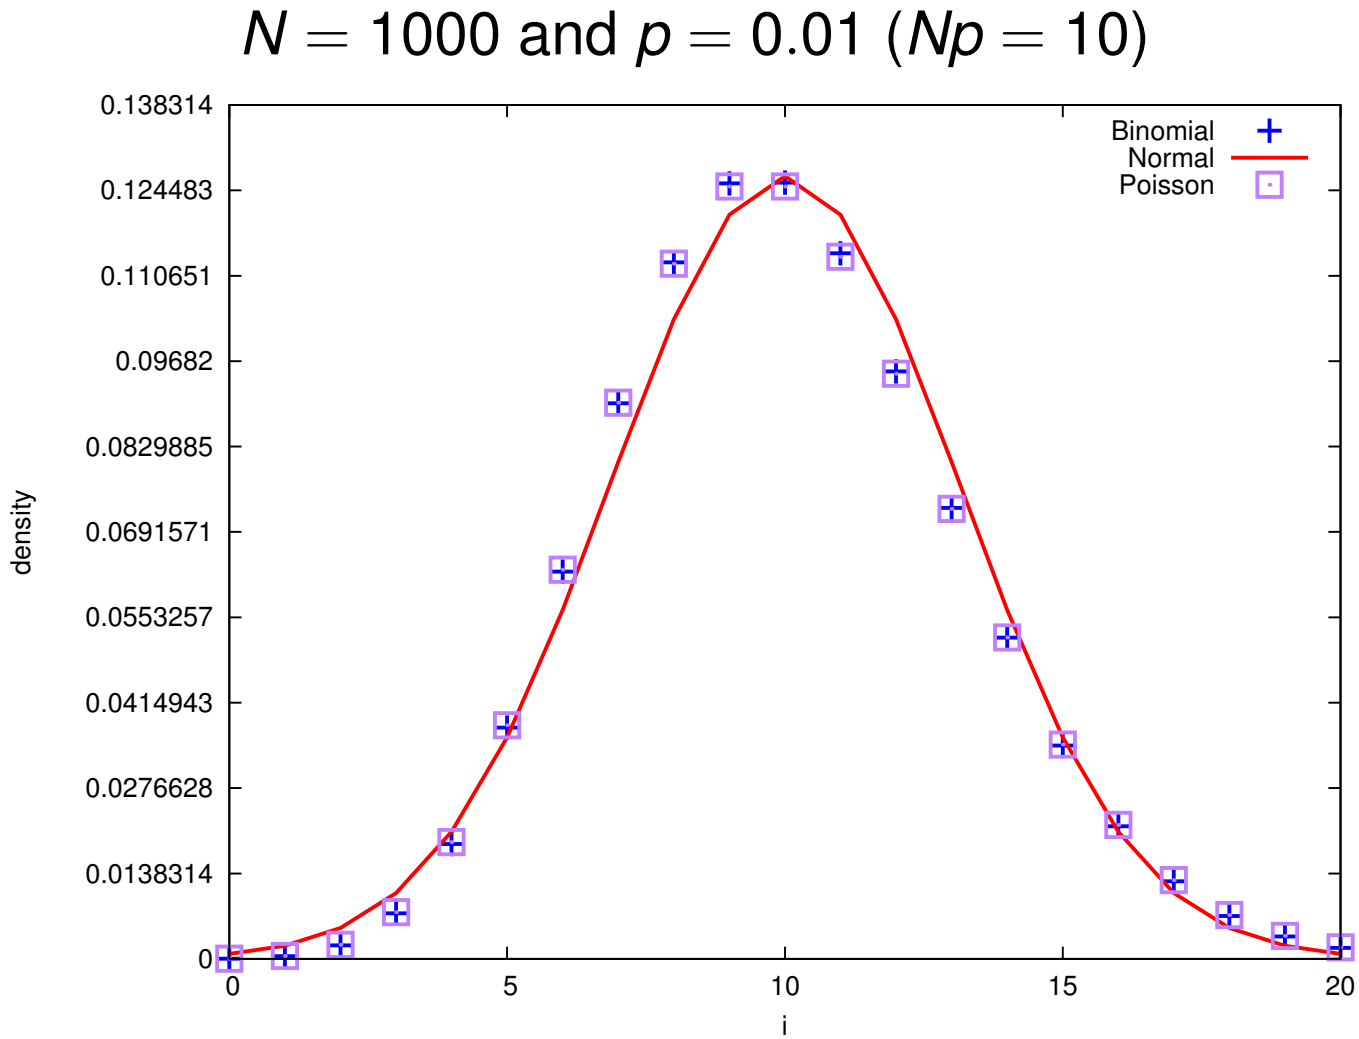

Binomial PMF and Approximations

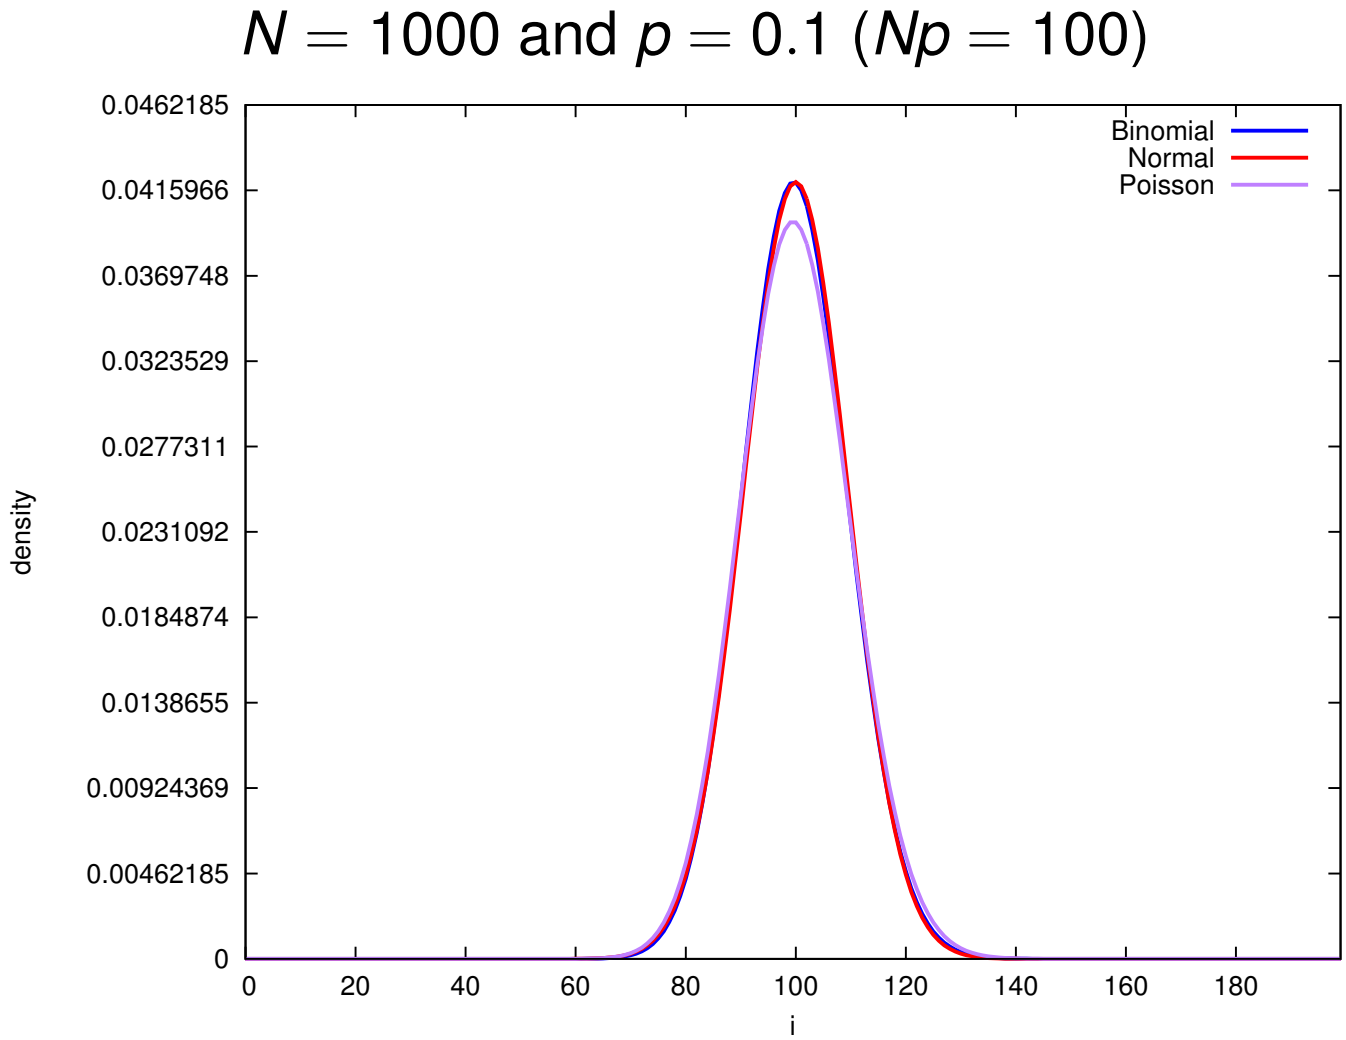

Binomial PMF and Approximations

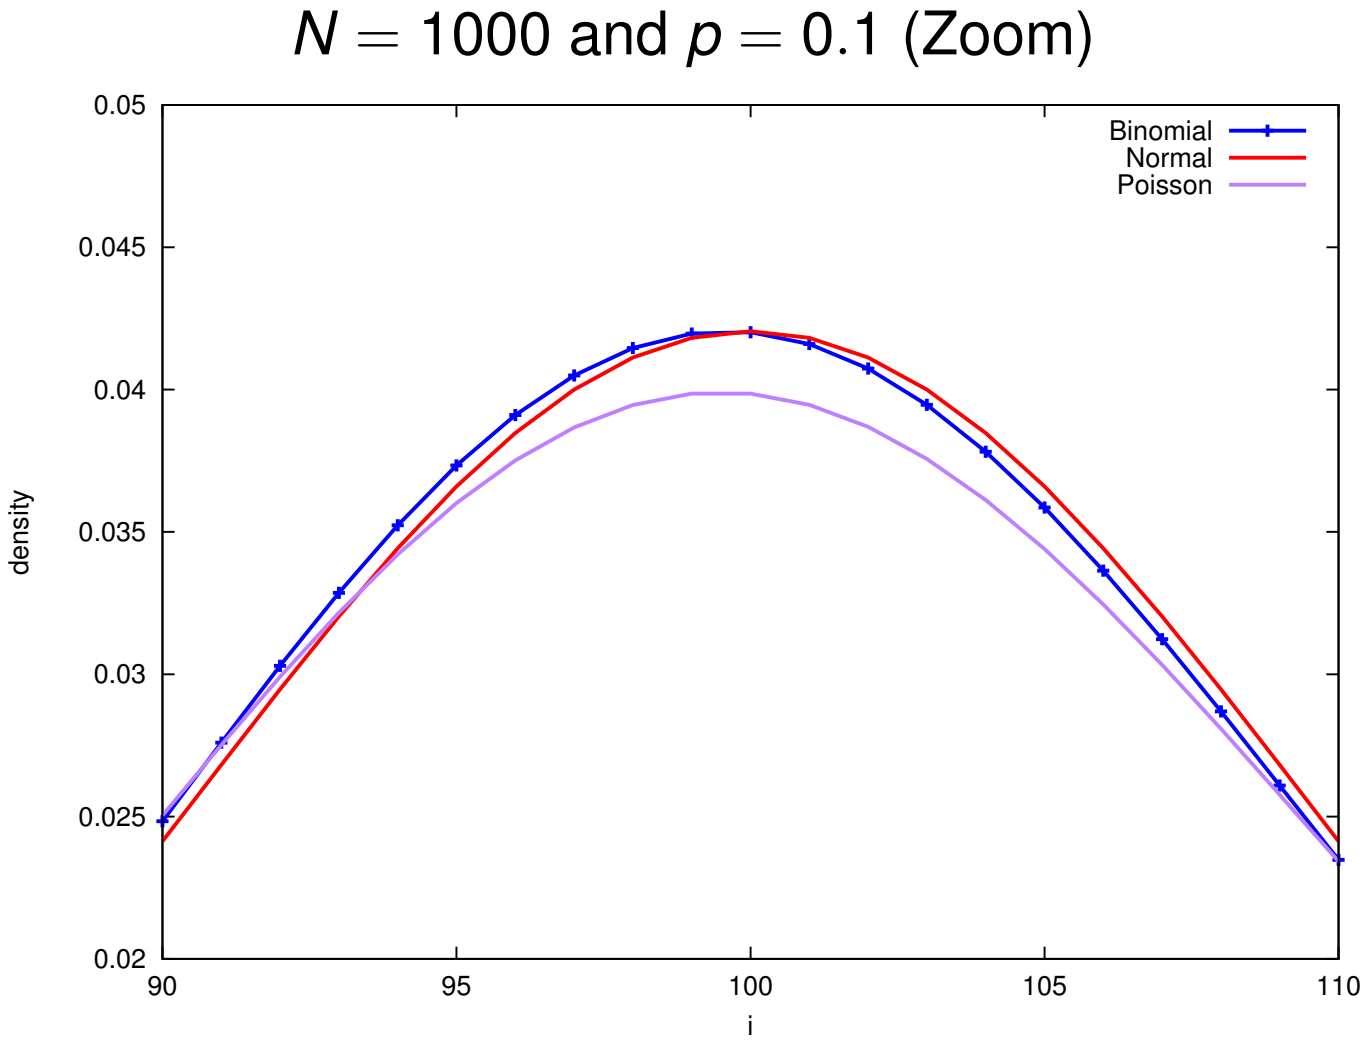

Binomial PMF and Approximations

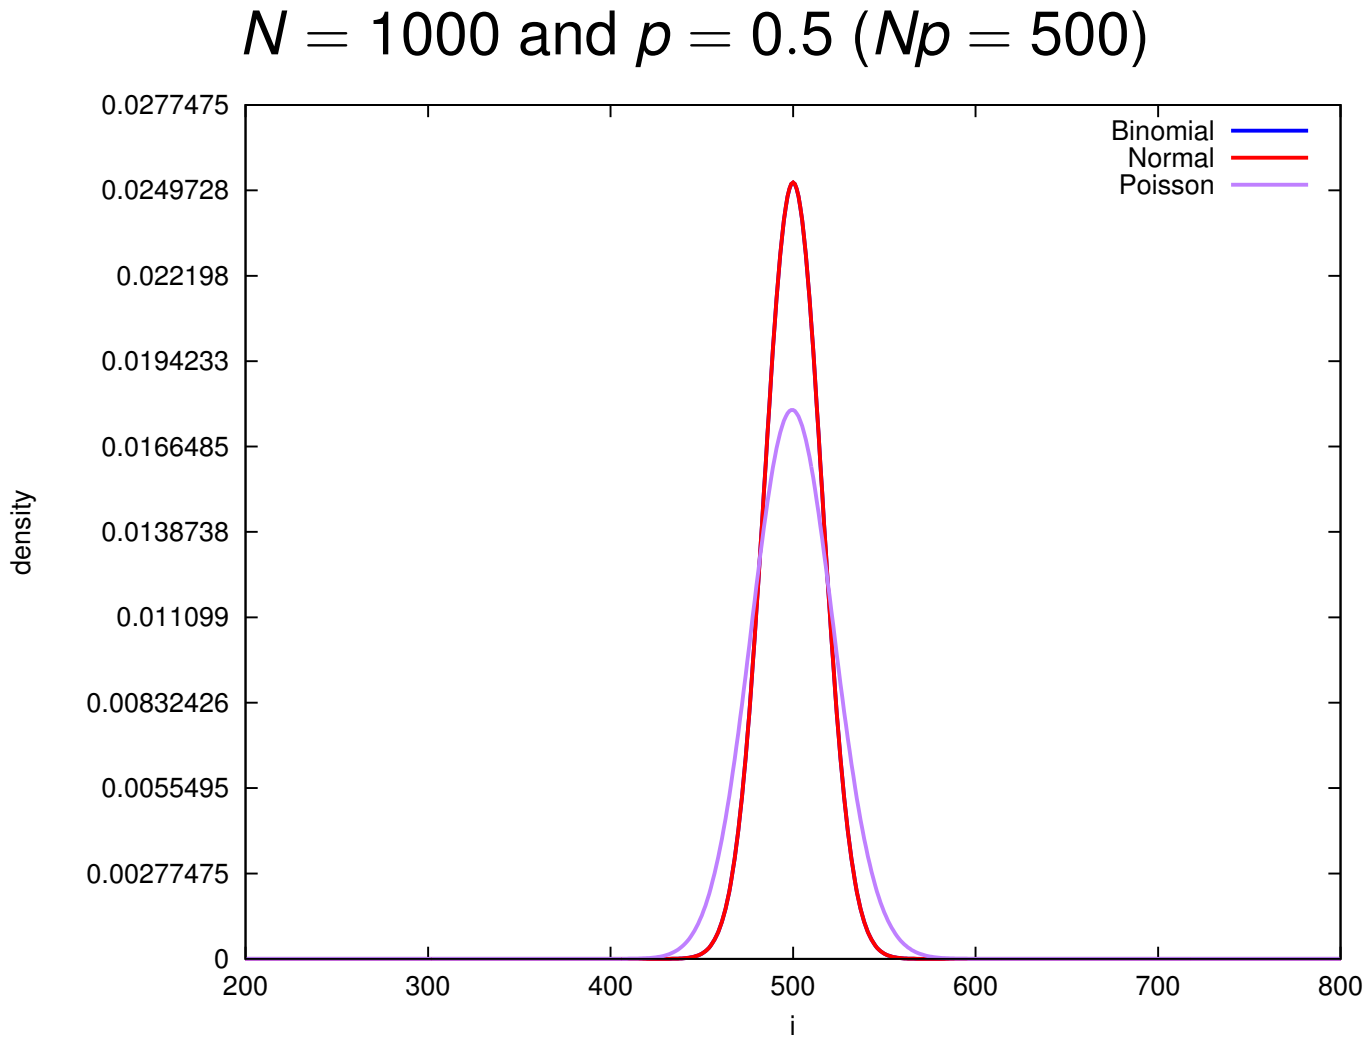

Binomial PMF and Approximations

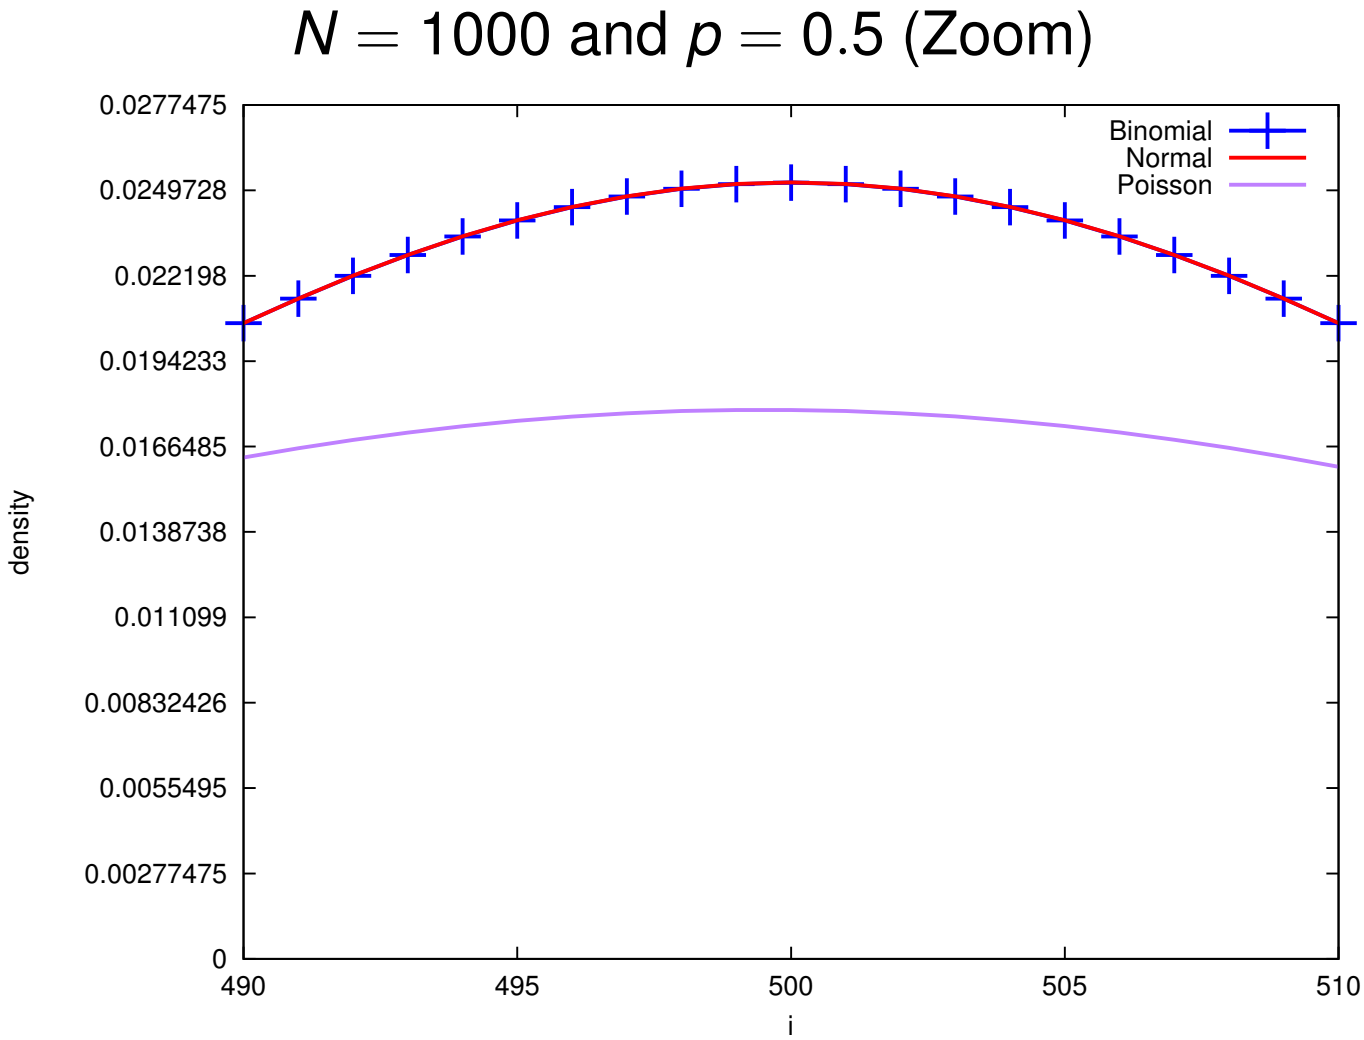

Binomial PMF and Approximations

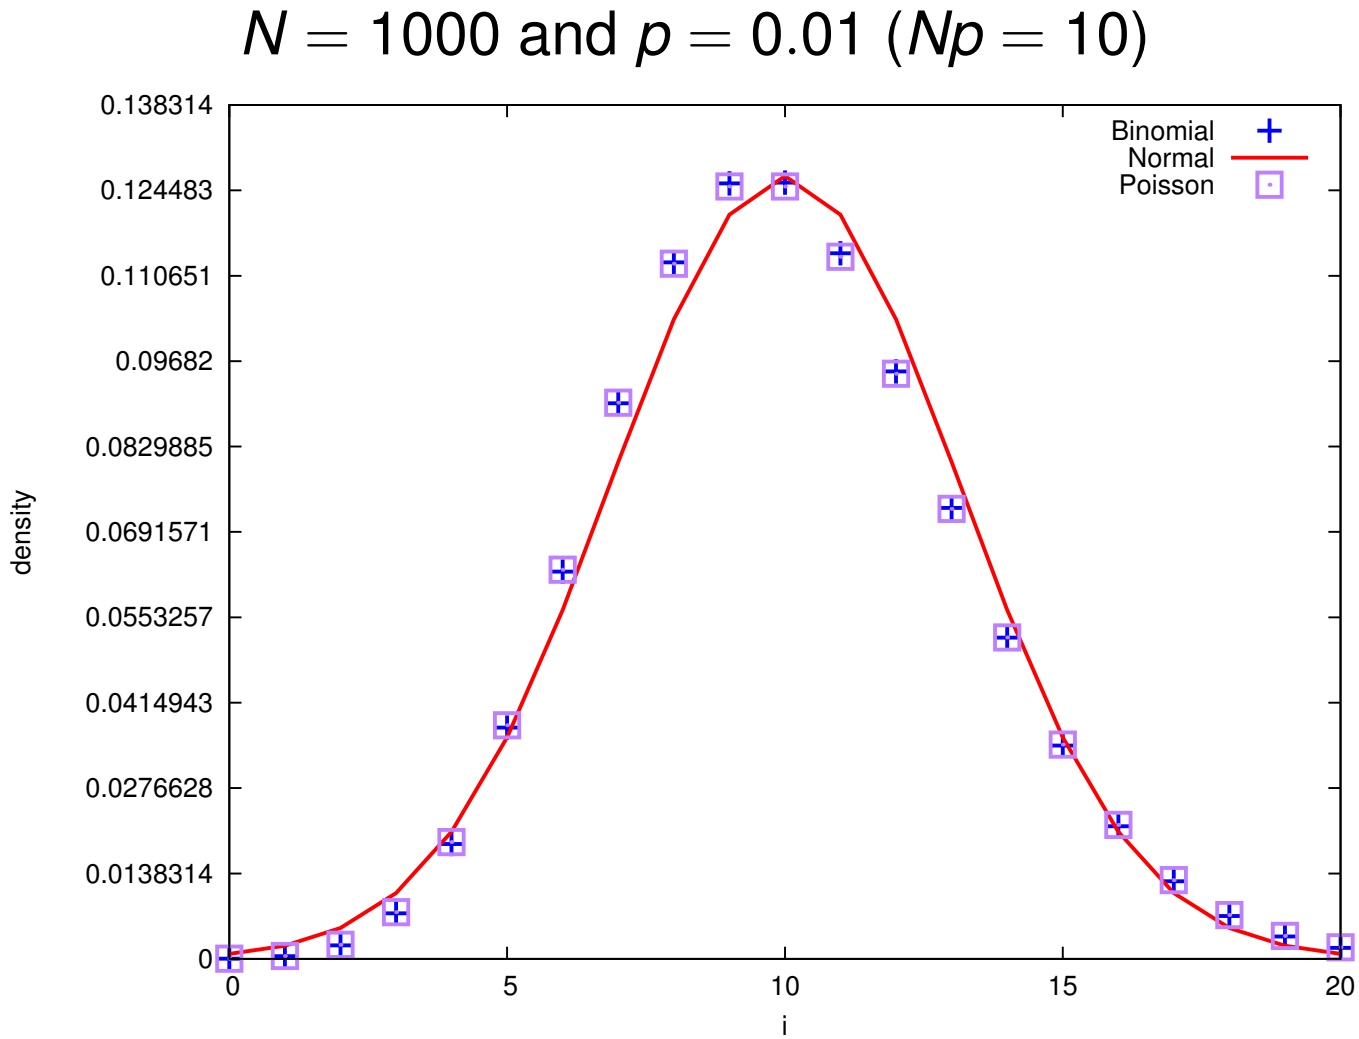

Binomial PMF and Approximations

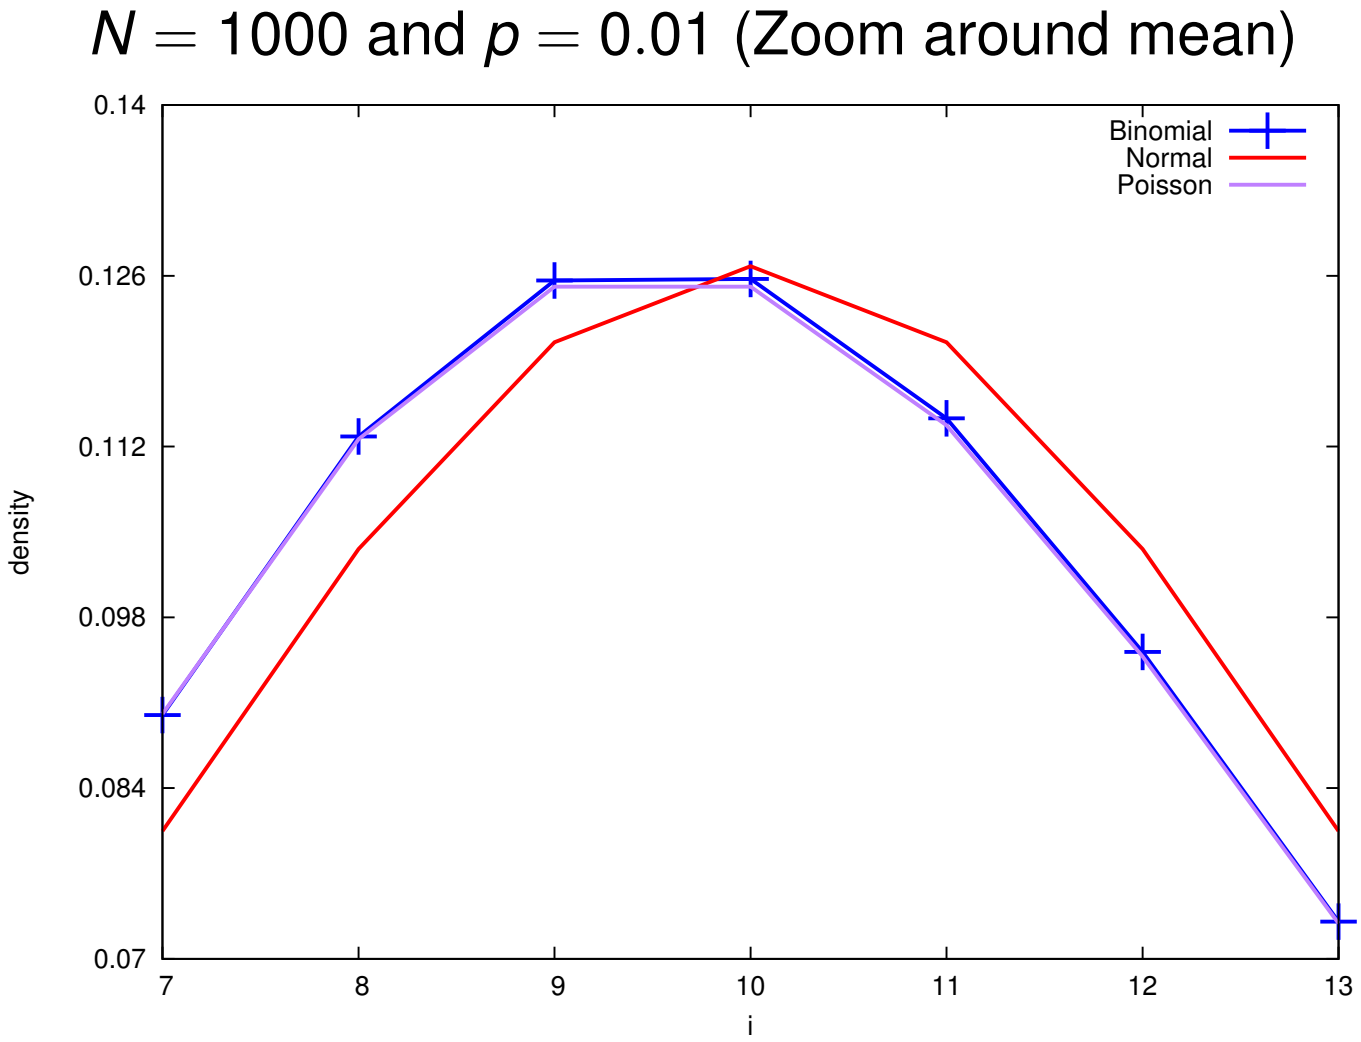

Binomial PMF and Approximations

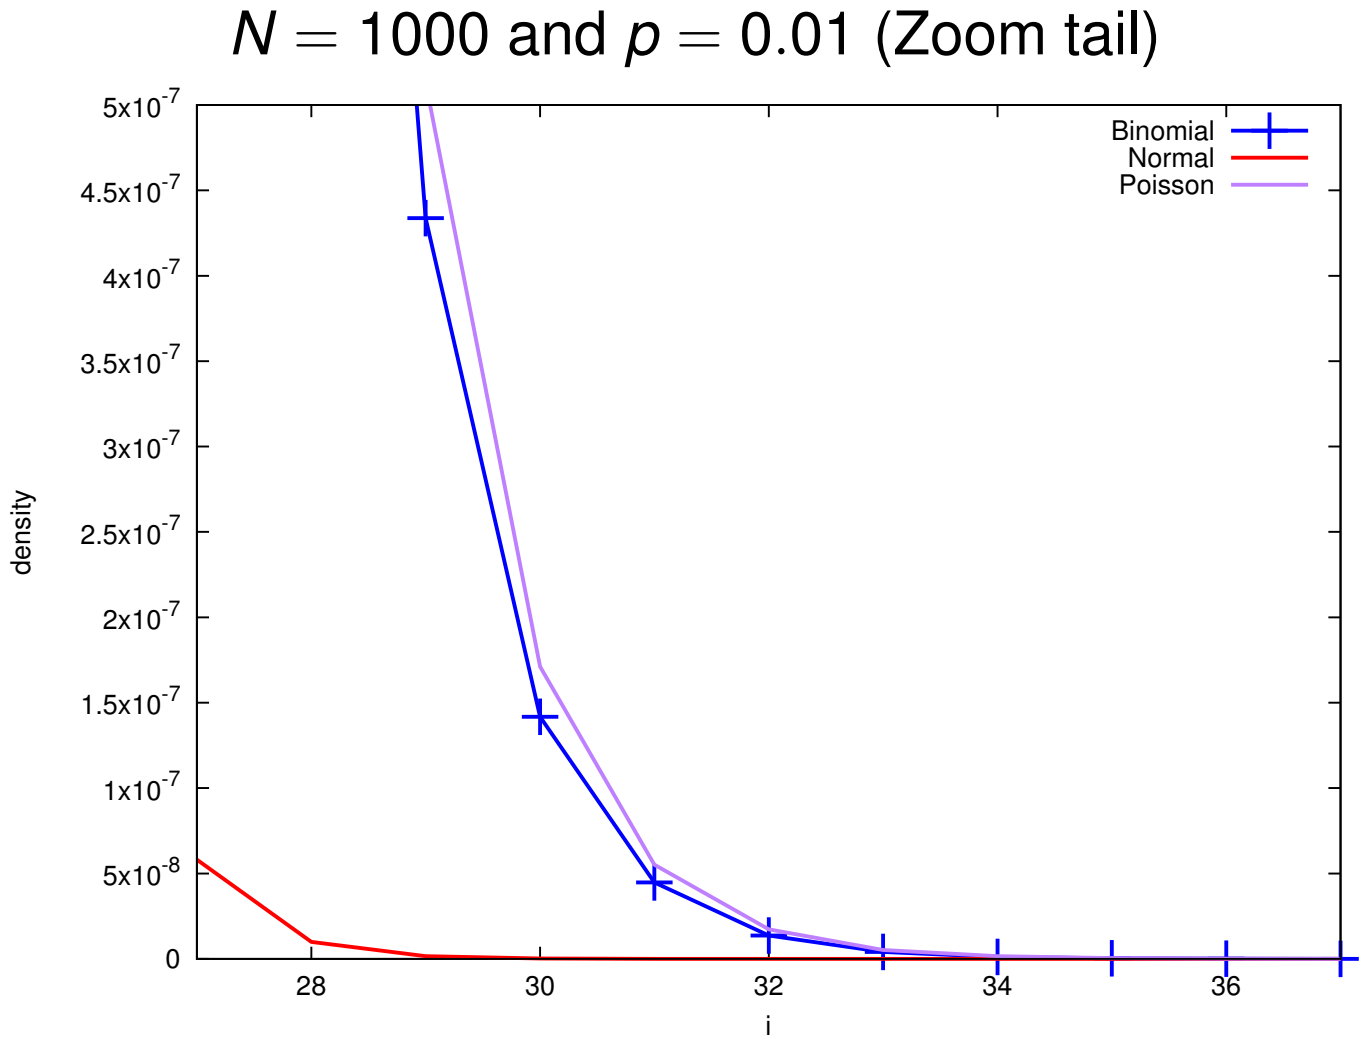

The CDF's

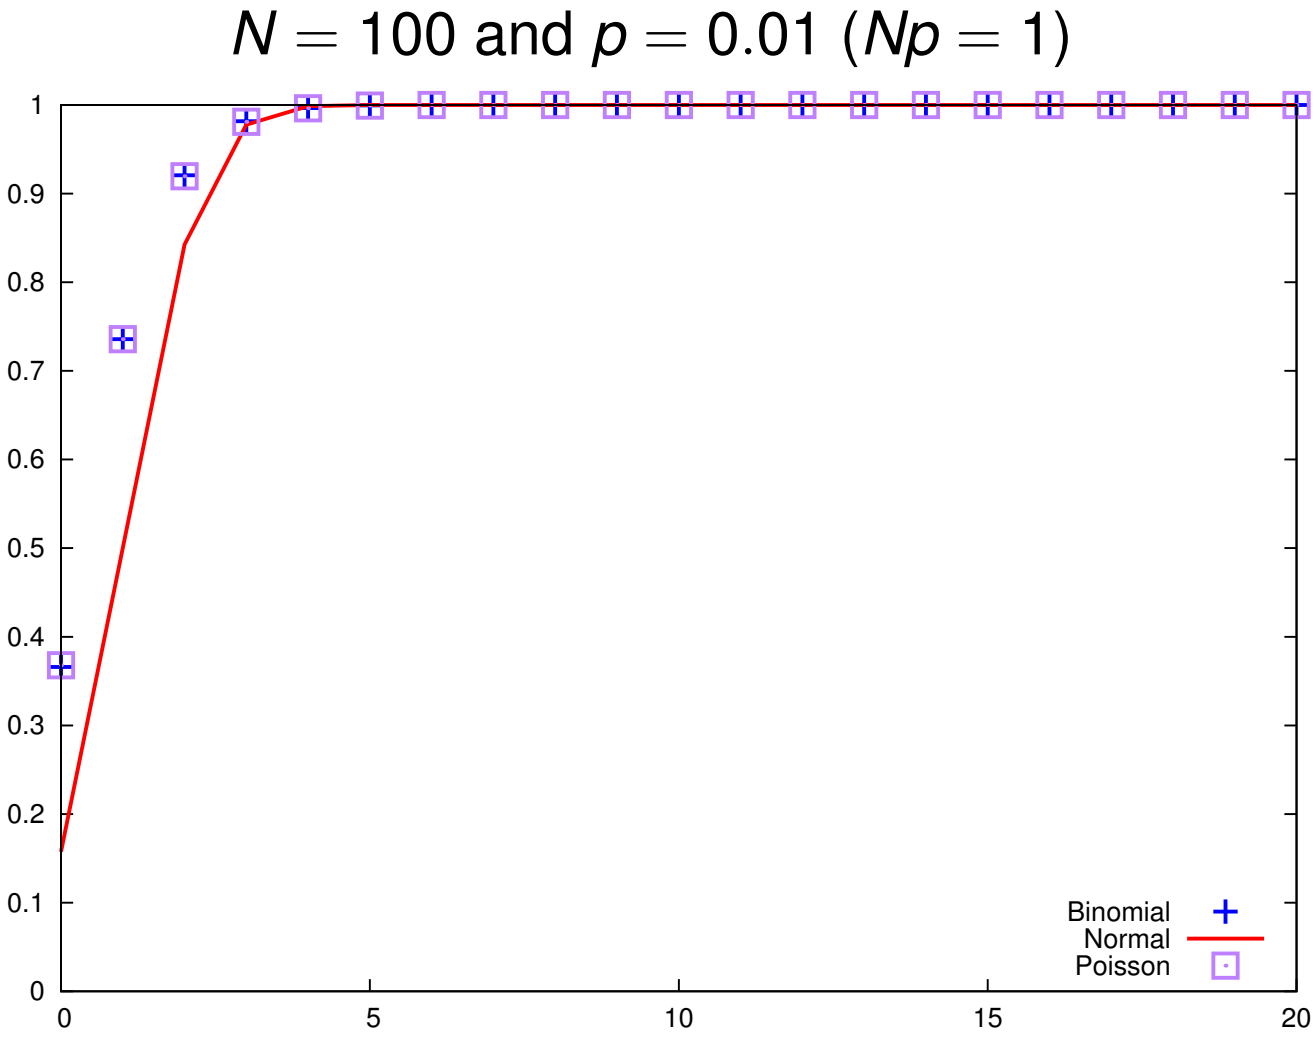
$N = 100$ and $p = 0.01$ ($Np = 1$)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
5
10
15
20
Binomial
Normal
Poisson

Convergence of the Binomial Distribution (Crypto)

- The binomial distribution is “hard” to manipulate
- Classically approximations are used
- In the linear context, the use of the normal distribution is relatively accurate
- In the differential case, we can sometimes use the Poisson distribution
- Exercise : Convergence of a Binomial distribution to a Poisson distribution

Kullback Divergence (1)

- Binomial tail :

$$P\left[\mathcal{T} \leq \Theta\right] = \sum_{i=0}^{\Theta} \binom{N}{i} p^i (1-p)^{N-i}$$

- Relative threshold : $\tau = \Theta / N$
- Kullback-Leibler divergence :

$$\mathrm{Kull}(p||q) = p \ln\left(\frac{p}{q}\right) + (1-p)\ln\left(\frac{1-p}{1-q}\right).$$

- Theorem :

$$P\left[\mathcal{T} \leq \tau N\right] \underset{N\to\infty}{\sim} \frac{p\sqrt{1-\tau}}{(p-\tau)\sqrt{2\pi N\tau}} \cdot 2^{-N\cdot \mathrm{Kull}(\tau||p)}$$

Kullback Divergence (2)

- Exercise :

$$P(\mathcal{T} = \lfloor \tau N \rfloor) \approx \sqrt{\frac{1}{2\pi N(1-\tau)\tau}}\, e^{-N \cdot \mathrm{Kull}(\tau||p)}$$

Kullback Divergence (3)

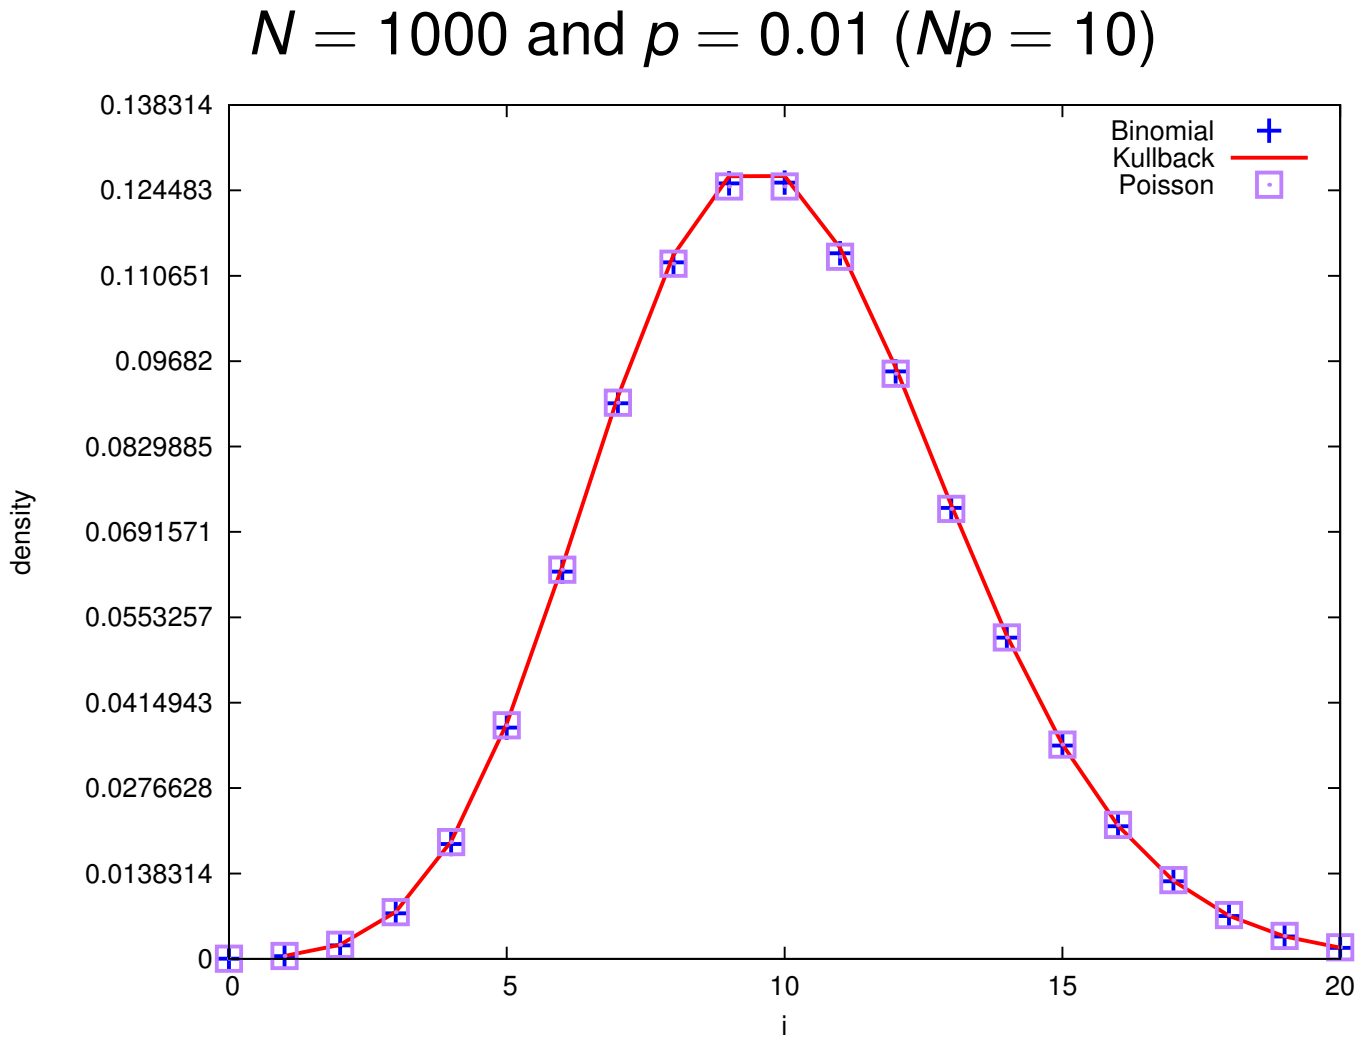

Kullback Divergence (3)

$N = 1000$ and $p = 0.01$ (Zoom around mean)

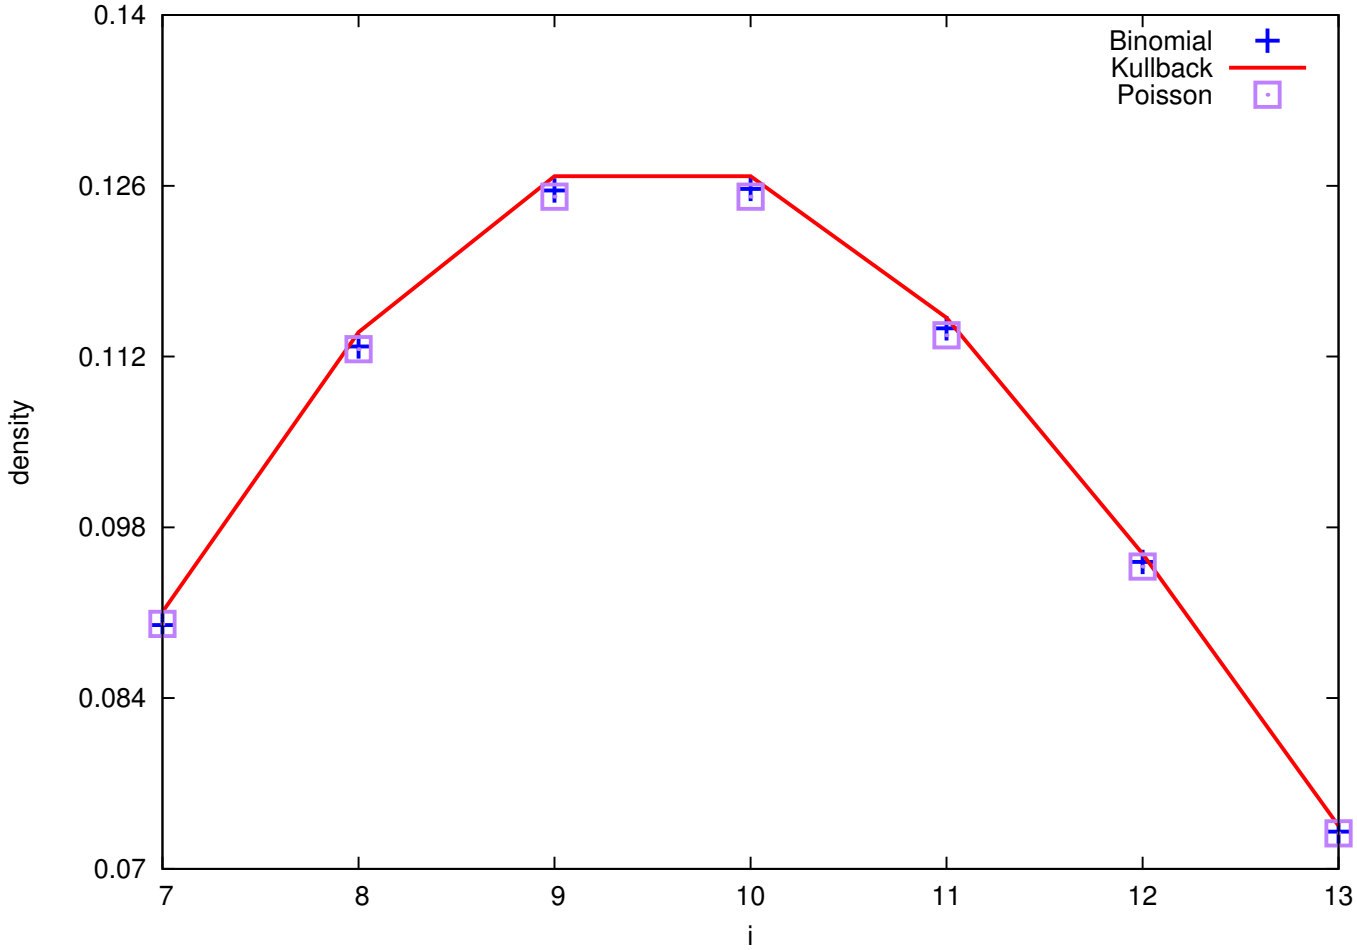

Kullback Divergence (3)

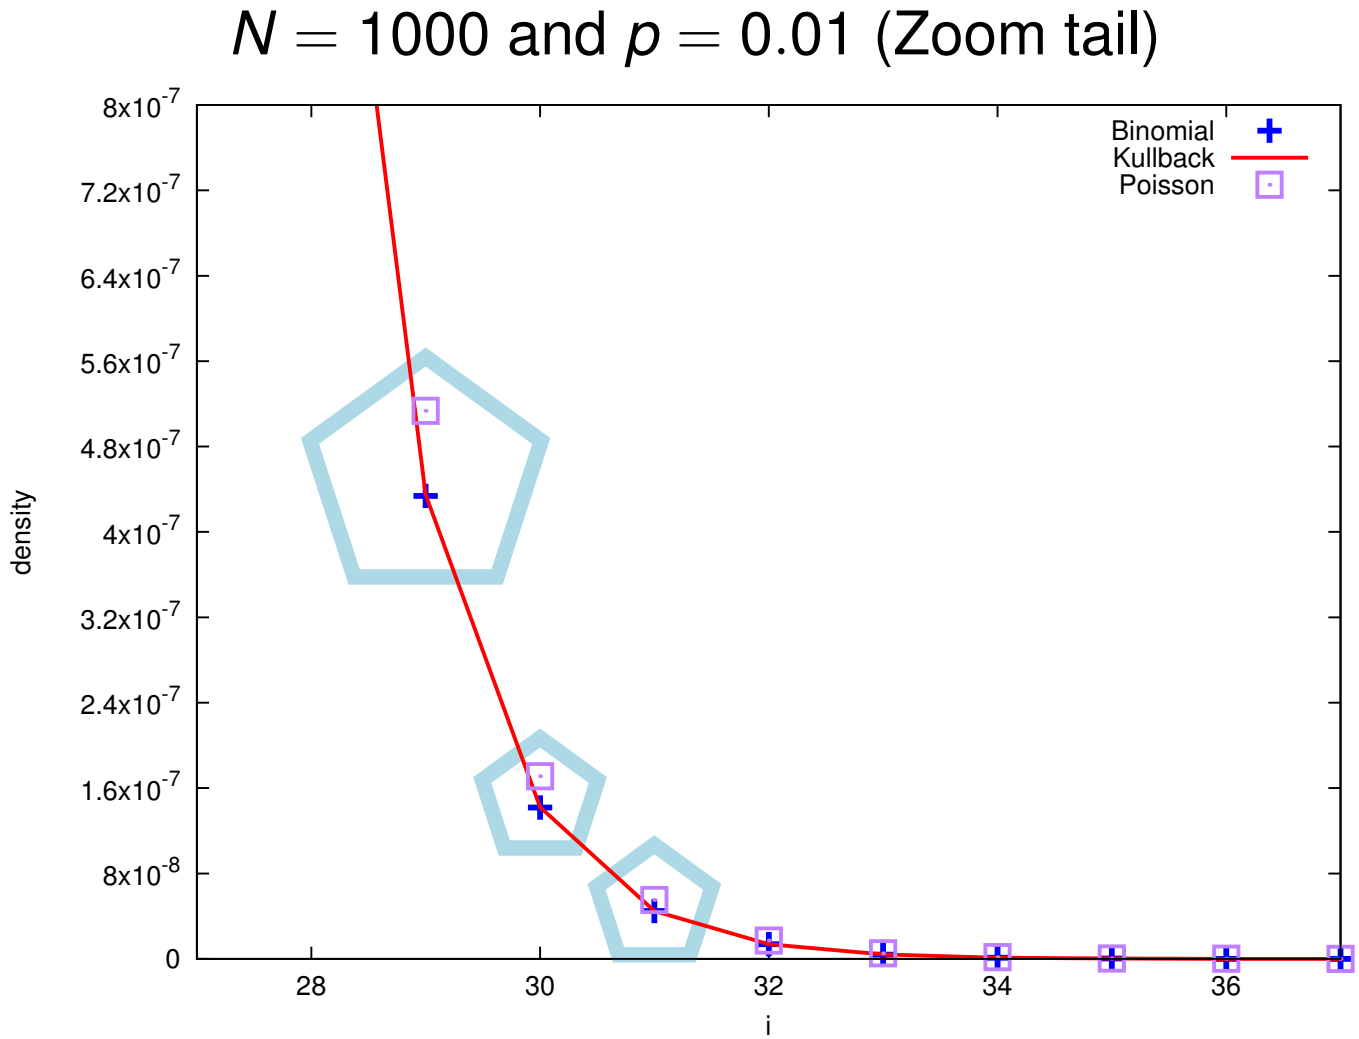

Kullback Divergence (3)

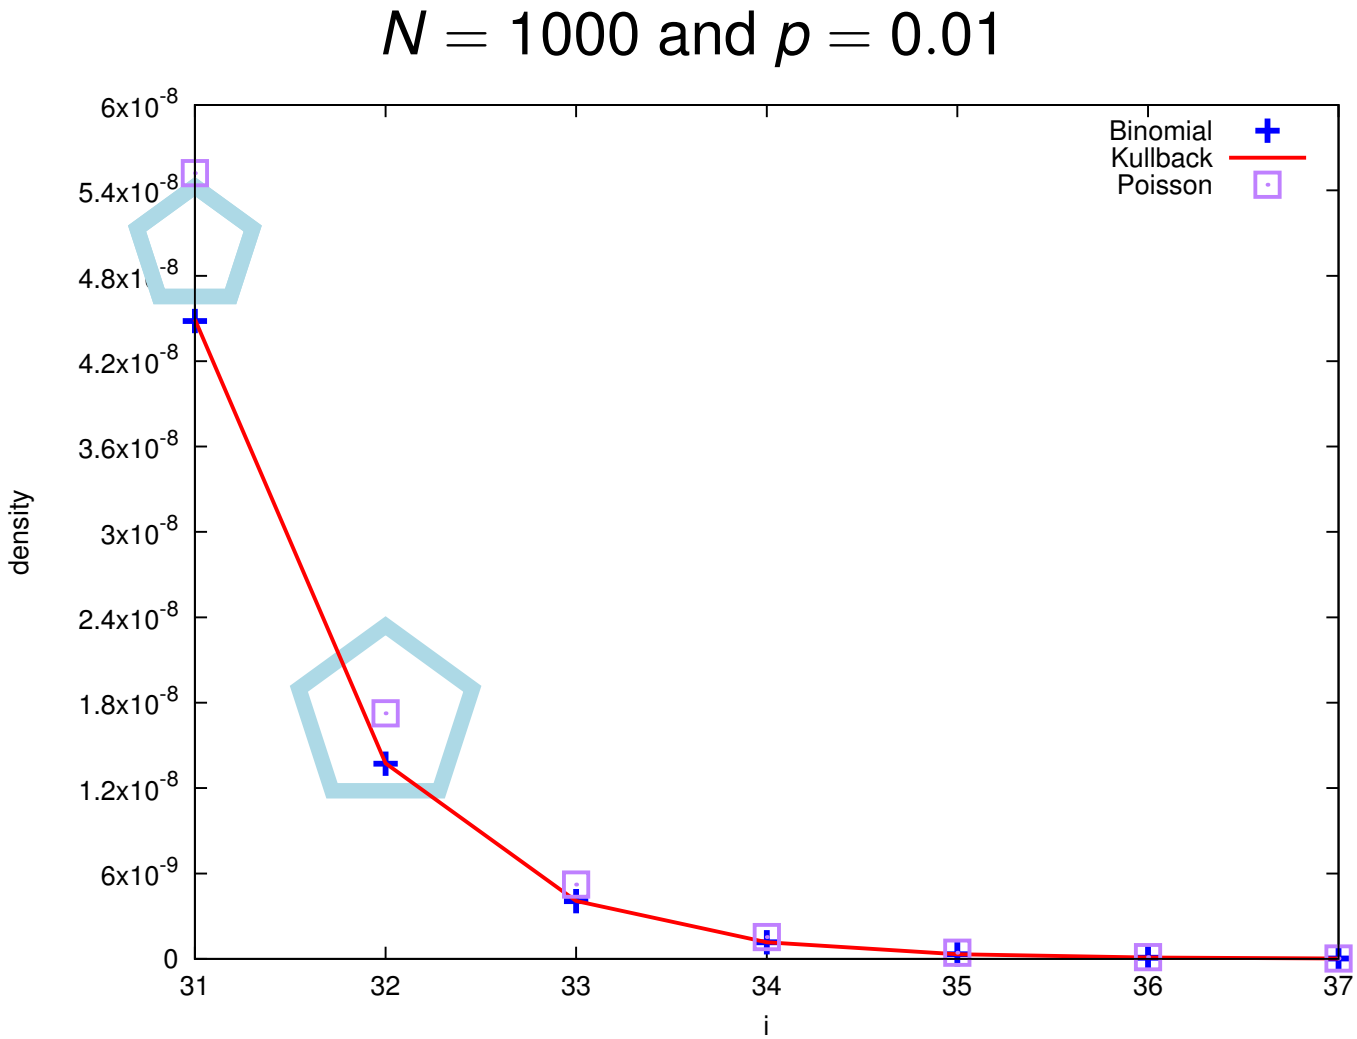

Normal Distribution Facts

- If X follows a normal distribution $\mathcal{N}(\mu, \sigma^2)$, then $\frac{X-\mu}{\sigma^2}$ follows a normal distribution $\mathcal{N}(0,1)$
- Central limit theorem : Suppose $X_1, X_2, \cdots, X_m$ is a sequence of i.i.d. random variables with $E[X_i] = \mu$ and $Var[X_i] = \sigma^2$

  Let $S_m = \frac{X_1 + X_2 + \cdots + X_m}{m}$

  As m approaches infinity, the random variable $\sqrt{m}(S_m - \mu)$ converges in distribution (the cdfs converge) to a normal $\mathcal{N}(0, \sigma^2)$

Attack Models

Which information is collected?

- known ciphertexts

Attack Models

Which information is collected?

- known ciphertexts
- . . .
- known plaintexts (implicitly plaintext-ciphertext pairs)
-
- chosen plaintexts
- chosen plaintexts/ciphertexts
- . . .

Attack Models

Which information is collected?

- known ciphertexts
- . . .
- known plaintexts (implicitly plaintext-ciphertext pairs)
- distinct known plaintexts
- chosen plaintexts
- chosen plaintexts/ciphertexts
- . . .

Attack Models

Which information is collected?

- known ciphertexts
- ...
- known plaintexts (KP)
- distinct known plaintexts (DKP)
- chosen plaintexts (CP)
- chosen plaintexts/ciphertexts
- ...

Outline

The Setting (in the Normal Distribution Case)

- Two random variables $\mathcal{T}_W$ and $\mathcal{T}_R$ such that

$$\mathcal{T}_W \sim \mathcal{N}(\mu_W, \sigma_W^2) \text{ and } \mathcal{T}_R \sim \mathcal{N}(\mu_R, \sigma_R^2),$$

  where $\mathcal{N}$ is the normal distribution with mean μ_W (or μ_R) and variance σ_W^2 (or σ_R^2)

- The scoring value T computed from a sample drawn from either the distribution of $\mathcal{T}_W$ or the one of $\mathcal{T}_R$

- The task is to decide which one of the two

- In some cases, we will identify the scoring value with its associated random variable

Success of the Attack (in the Normal Distribution Case) (1)

- Threshold value Θ :
  - $T < \Theta \Rightarrow T$ is drawn from the distribution of T_W
  - $T \geq \Theta \Rightarrow T$ is drawn from the distribution of T_R

  Assume first $\mu_W < \mu_R$
- We set bounds on the error probabilities

$$P[\, T \geq \Theta \mid T = T_W \,] \leq \beta \text{ and } P[\, T < \Theta \mid T = T_R \,] \leq \alpha,$$

which are satisfied if

$$\mu_W + \sigma_W \Phi^{-1}(1-\beta) \leq \Theta \leq \mu_R - \sigma_R \Phi^{-1}(1-\alpha),$$

where Φ is the cumulative distribution function of the standard normal distribution

Illustration

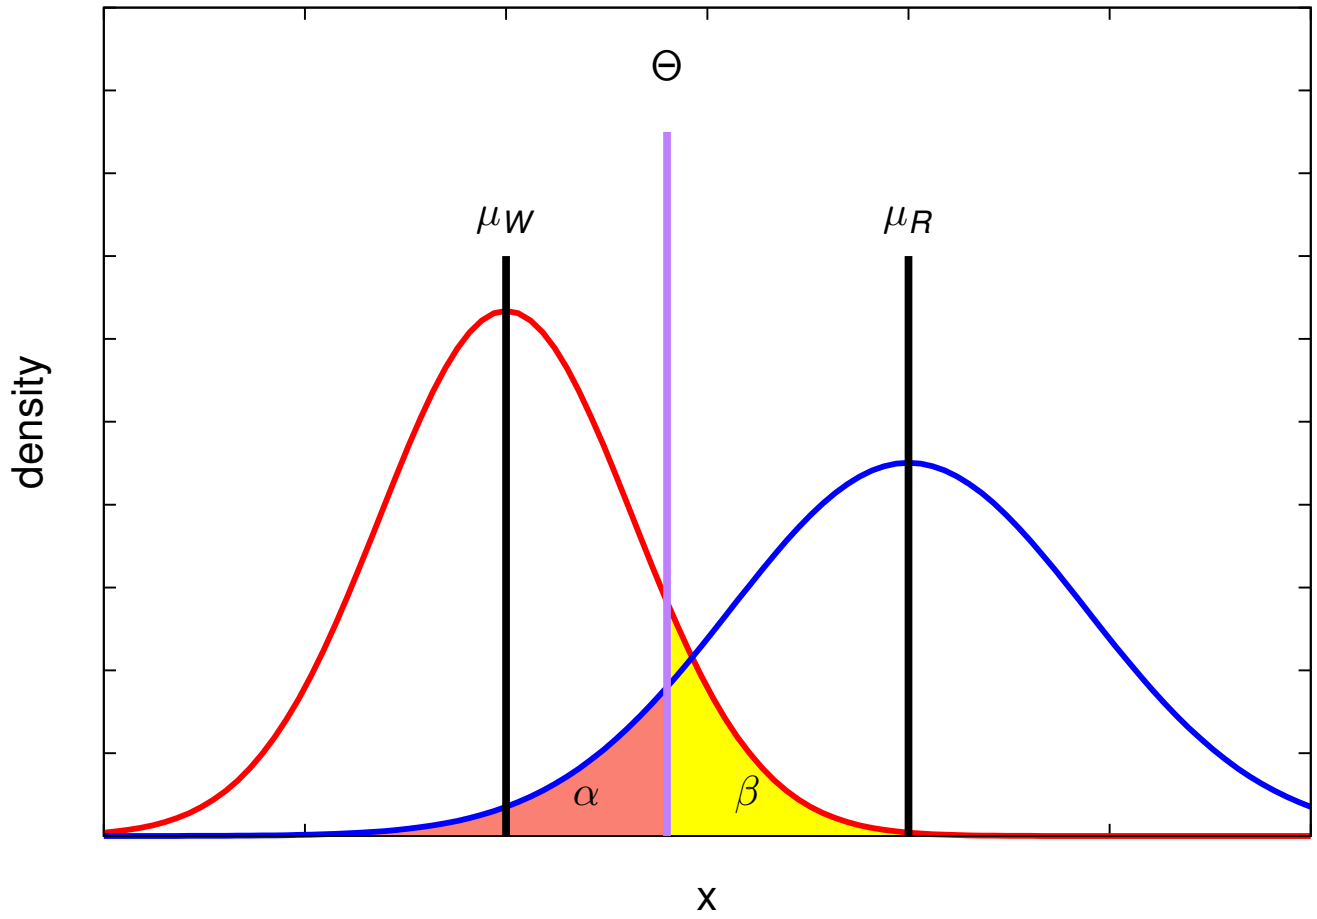

Illustration

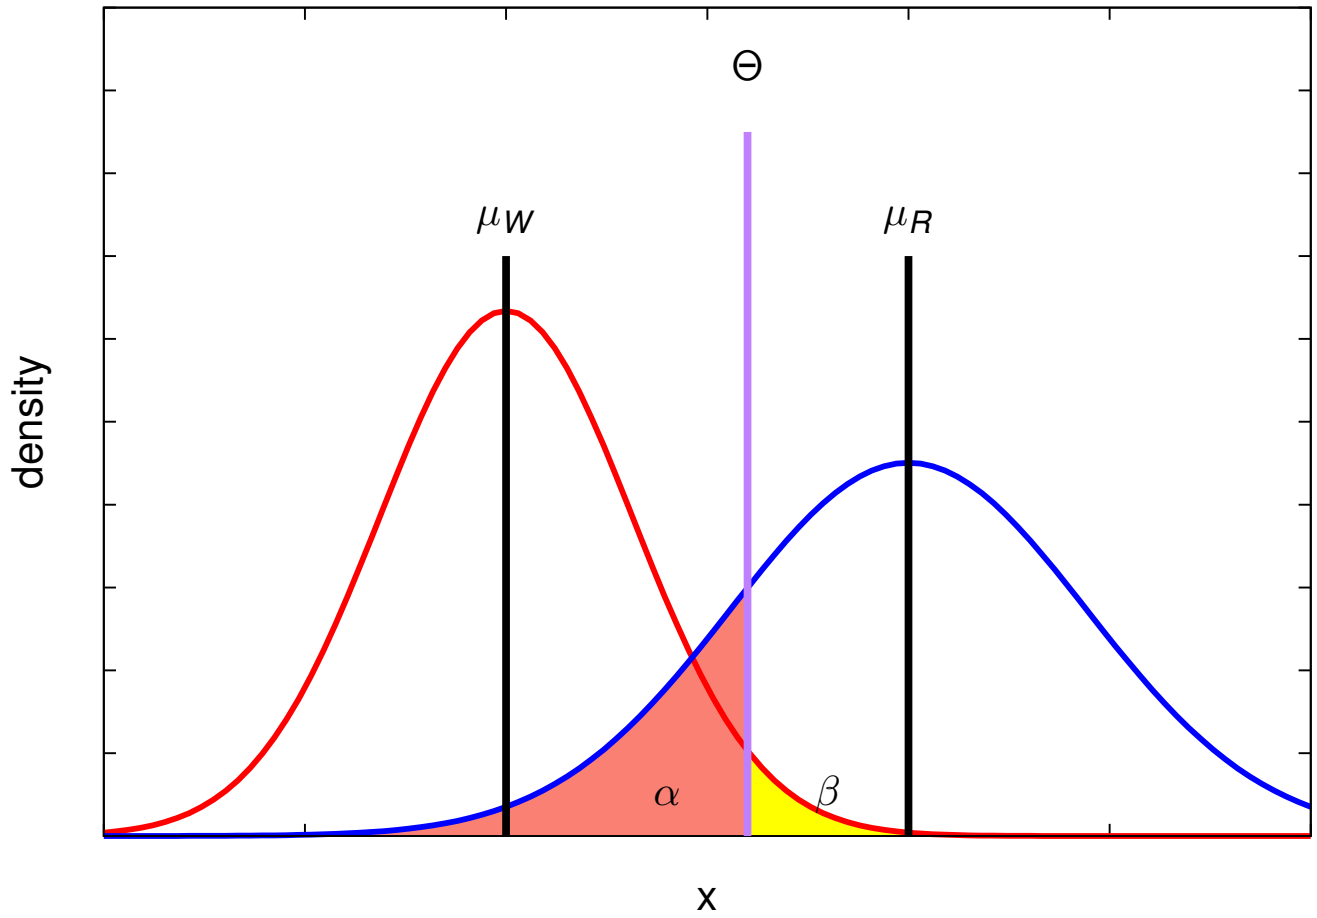

Success of the Attack (2)

- The case $\mu_R < \mu_W$ (red and blue curves) is analogical

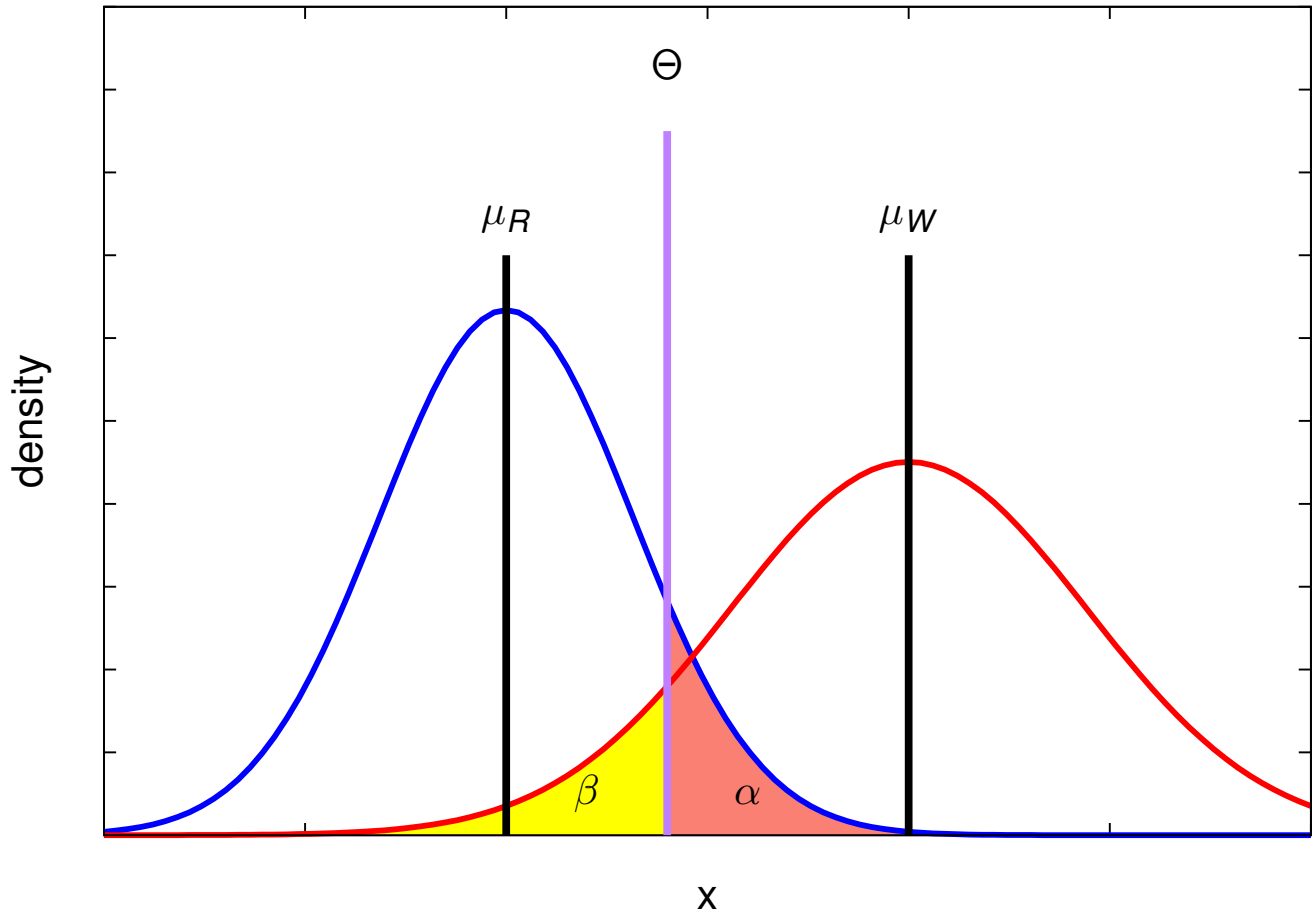

An Algorithm for Finding N and τ (1)

Some properties :

- For a fixed relative threshold $\tau = \Theta/N$, error probabilities decrease when N increases.
- For a fixed N, non-detection error increases with τ.
- For a fixed N, false alarm error decreases when τ increases.

Idea

Dichotomic search for τ.

An Algorithm for finding N and τ (2)

Input : (α, β) and (p_R, p_W)
Output : N and τ the minimum number of samples and the corresponding relative threshold to reach error probabilities less than (α, β).

$\tau_m \leftarrow p_W$ and $\tau_M \leftarrow p_R$.
repeat
$\tau \leftarrow \dfrac{\tau_m + \tau_M}{2}$.
Compute N_{nd} such that $\forall N > N_{\text{nd}},\ P(\mathcal{T}_R < N\tau) \leq \alpha$.
Compute N_{fa} such that $\forall N > N_{\text{fa}},\ P(\mathcal{T}_W \geq N\tau) \leq \beta$.
if $N_{\text{nd}} > N_{\text{fa}}$ **then** $\tau_M = \tau$ **else** $\tau_m = \tau$
until $N_{\text{nd}} = N_{\text{fa}}$.
return N and τ.

Outline

Linear Attack

- One linear approximation: (u, v)
- Empirical correlation: $\hat{cor}(D, k)$
- The distribution of $\hat{cor}(D, k)$ depends on the data D and on the key candidate k
- c: Absolute value of the expected correlation for the right key

$$\mathrm{Exp}_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_R)] = \pm c$$

- KP model

Randomization Hypothesis

- Wrong-key randomization hypothesis: If the key candidate is wrong then $\hat{cor}(D, k_W)$ follows a normal distribution with parameters

$$\begin{aligned} \mathrm{Exp}_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_W)] &= 0 \text{ and} \\ \mathrm{Var}_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_W)] &= \frac{1}{N} \end{aligned}$$

- Hypothesis of right-key equivalence: For all k_R, $\hat{cor}(D, k_R)$ follows a normal distribution with parameters

$$\begin{aligned} \mathrm{Exp}_{\mathbf{D},\mathbf{K}}\left(\hat{cor}(D, k_R)\right) &= \pm c, \\ \mathrm{Var}_{\mathbf{D},\mathbf{K}}\left(\hat{cor}(D, k_R)\right) &= \frac{1}{N}(1 - c^2) \approx \frac{1}{N} \end{aligned}$$

- Since $\mathrm{Exp}_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_R)] = \pm c$ we have $\beta = 2^{-(a+1)}$

Error Probabilities

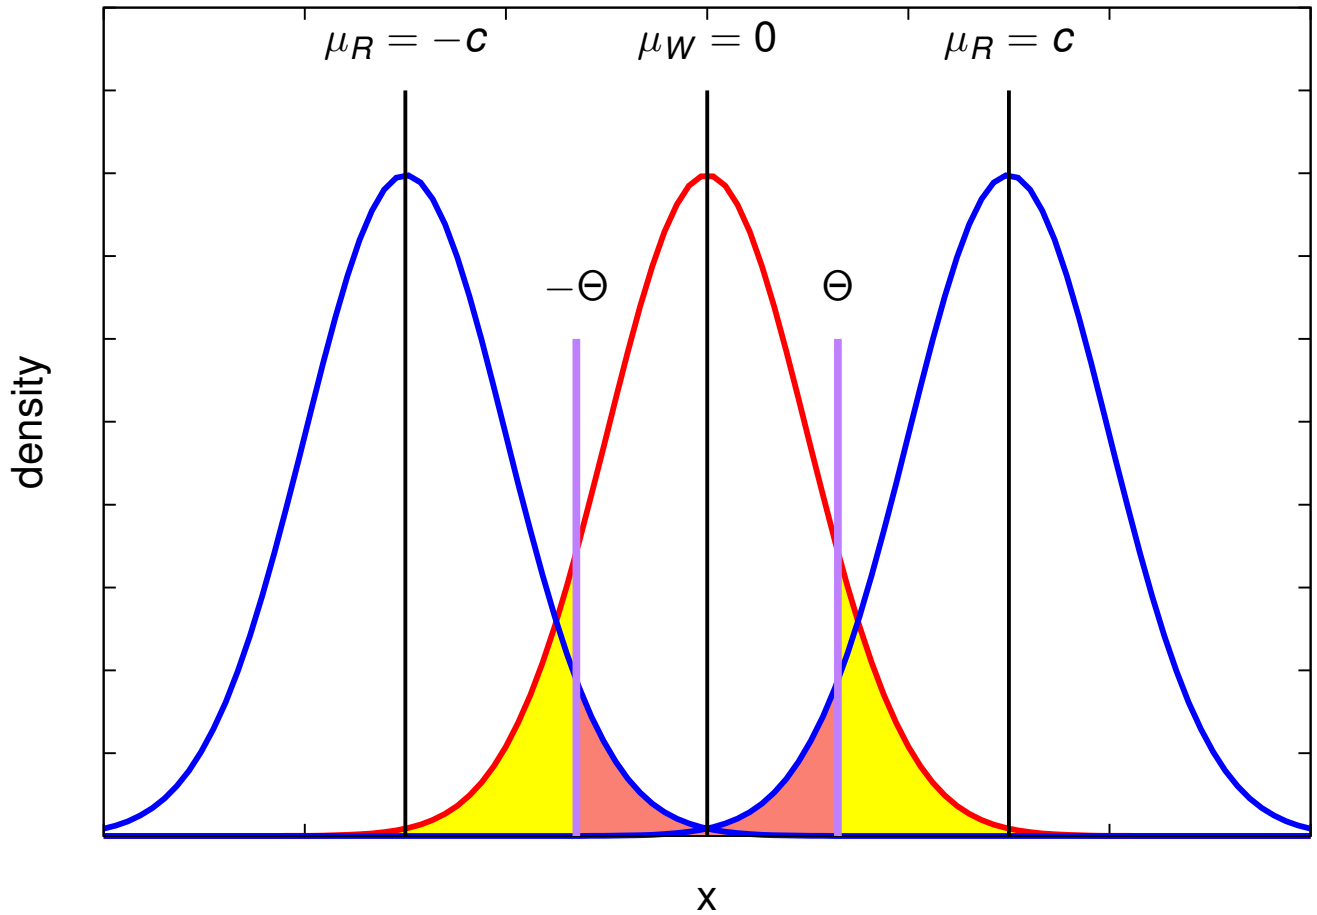

Data Complexity of a Matsui Algorithm-2 Attack (1)

- We use the formula

$$\sigma_R \cdot \Phi^{-1}(1-\alpha) = |\mu_R - \mu_W| - \sigma_W \cdot \Phi^{-1}(1-\beta)$$

- $\beta = 2^{-(a+1)}$
- Data complexity: [Matsui] [Selçuk 08]

$$N \geq \frac{(\varphi_\beta + \sqrt{(1-c^2)} \cdot \varphi_\alpha)^2}{c^2},$$

where $\varphi_\beta = \Phi^{-1}(1-\beta)$ and $\varphi_\alpha = \Phi^{-1}(1-\alpha)$

Adjusting the Wrong-Key Randomization Hypothesis

- [Bogdanov and Tischhauser 13] If the key candidate is wrong then $\hat{cor}(D, k_W)$ follows normal distribution with parameters

$$Exp_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_W)] = 0 \text{ and}$$
$$Var_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_W)] = 1/N + 2^{-n}$$

- Idea
  - For a fixed key k_W, the empirical correlation varies with the data sample meaning that $Var_{\mathbf{D}}[\hat{cor}(D, k_W)] = 1/N$
  - The empirical correlation varies also with the key k_W [Daemen Rijmen 07]

Adjusting the Right-Key Randomization Hypothesis

- Assuming independent round key and linear approximation with a single dominant trail
- ELP: the expected linear potential $ELP = \mathrm{Exp}_{\mathbf{K}}[cor(k)^2]$
- Hypothesis of right-key equivalence: In the context of a linear key-recovery attack of a long-key iterated cipher described in this section, the empirical correlation $\hat{cor}(D, k_R)$ computed from KP sample is approximately normally distributed with parameters

$$\mathrm{Exp}_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_R)] = \pm c \quad \text{and}$$

$$\mathrm{Var}_{\mathbf{D},\mathbf{K}}[\hat{cor}(D, k_R)] = \frac{1}{N} + ELP - c^2$$

Data Complexity of a Matsui Algorithm-2 Attack (2)

- In practice, it is difficult to compute ELP exactly, and therefore it has often been estimated by c^2
- We can show that $ELP - c^2 \geq 2^{-n}$
- By taking $ELP = c^2 + 2^{-n}$ we obtain the following complexity bound

$$N \geq \frac{(\varphi_\beta + \varphi_\alpha)^2}{c^2 - 2^{-n}(\varphi_\beta + \varphi_\alpha)^2}$$

- Other settings: Work in progress [Nyberg Dagstuhl 16]

Outline

χ^2 distribution

- If $X_1, \cdots, X_\ell$ are independent standard normal random variables $X_i \sim \mathcal{N}(0,1)$, then the sum of their squares,

$$Q = \sum_{i=1}^{\ell} X_i^2,$$

is distributed according to the chi-squared distribution with ℓ degrees of freedom

$$Q \sim \chi^2_\ell$$

- The probability density function (pdf) of the chi-squared distribution is

$$f(x;\, \ell) = \begin{cases} \frac{x^{(\ell/2-1)} e^{-\ell/2}}{2^{\ell/2}\Gamma\left(\frac{\ell}{2}\right)}, & x > 0; \\ 0, & \text{otherwise} \end{cases}$$

where $\Gamma(\cdot)$ denotes the Gamma function

Non-Central χ^2

- Let $(X_1, X_2, \ldots, X_i, \ldots, X_\ell)$ be ℓ independent normally distributed random variables $X_i \sim \mathcal{N}(\mu_i, 1)$
- Then the random variable $\sum_{i=1}^{\ell} X_i^2$ is distributed according to the non-central chi-squared distribution with parameters ℓ specifying the number of degrees of freedom and $\lambda = \sum_{i=1}^{\ell} \mu_i^2$ the mean of the means
- λ is sometimes called the non-centrality parameter

Gamma Distribution

- Gamma function:
  - for an integer value n, $\Gamma(n) = (n-1)!$
  - for a real value $t > 0$, $\Gamma(t) = \int_0^\infty x^{t-1} e^{-x}\, dx$
  - $\Gamma(1/2) = \pi$
- pdf of $X \sim \Gamma(\rho, \theta)$:

$$f(x) = \frac{1}{\Gamma(\rho)\theta^\rho} x^{\rho-1} e^{-x/\theta}$$

- $\mathrm{Exp}[X] = \rho\theta$
- $\mathrm{Var}[X] = \rho\theta^2$
- If $X \sim \chi^2(\ell)$, and $c > 0$ then $cX \sim \Gamma(\rho = \ell/2, \theta = 2c)$

Gamma Distribution

$$\Gamma(\rho, \theta)$$

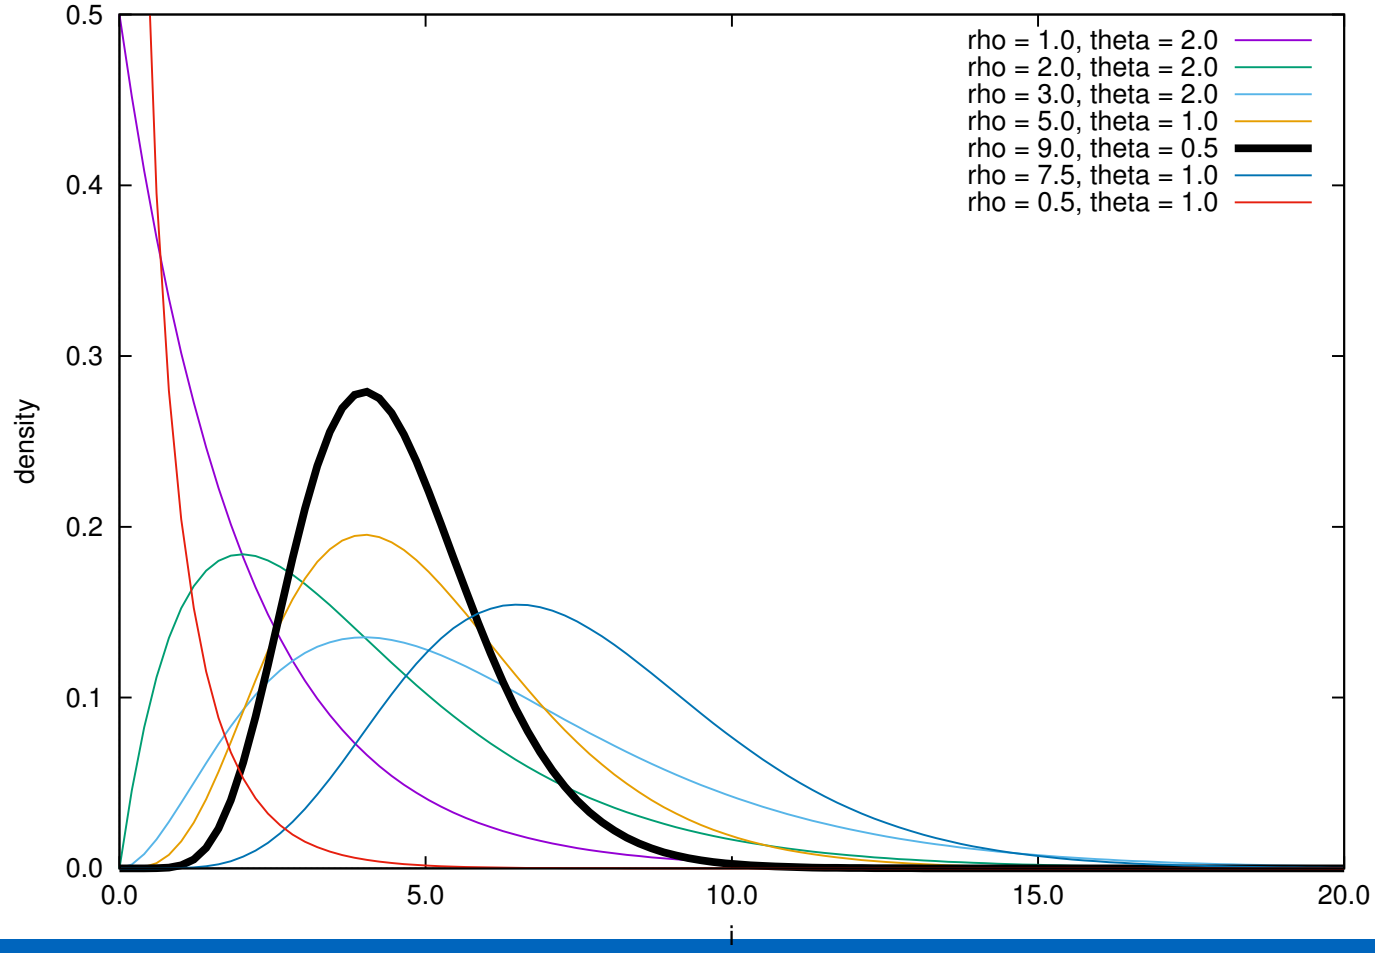

Gamma Distribution

$$\Gamma(\rho, \theta)$$

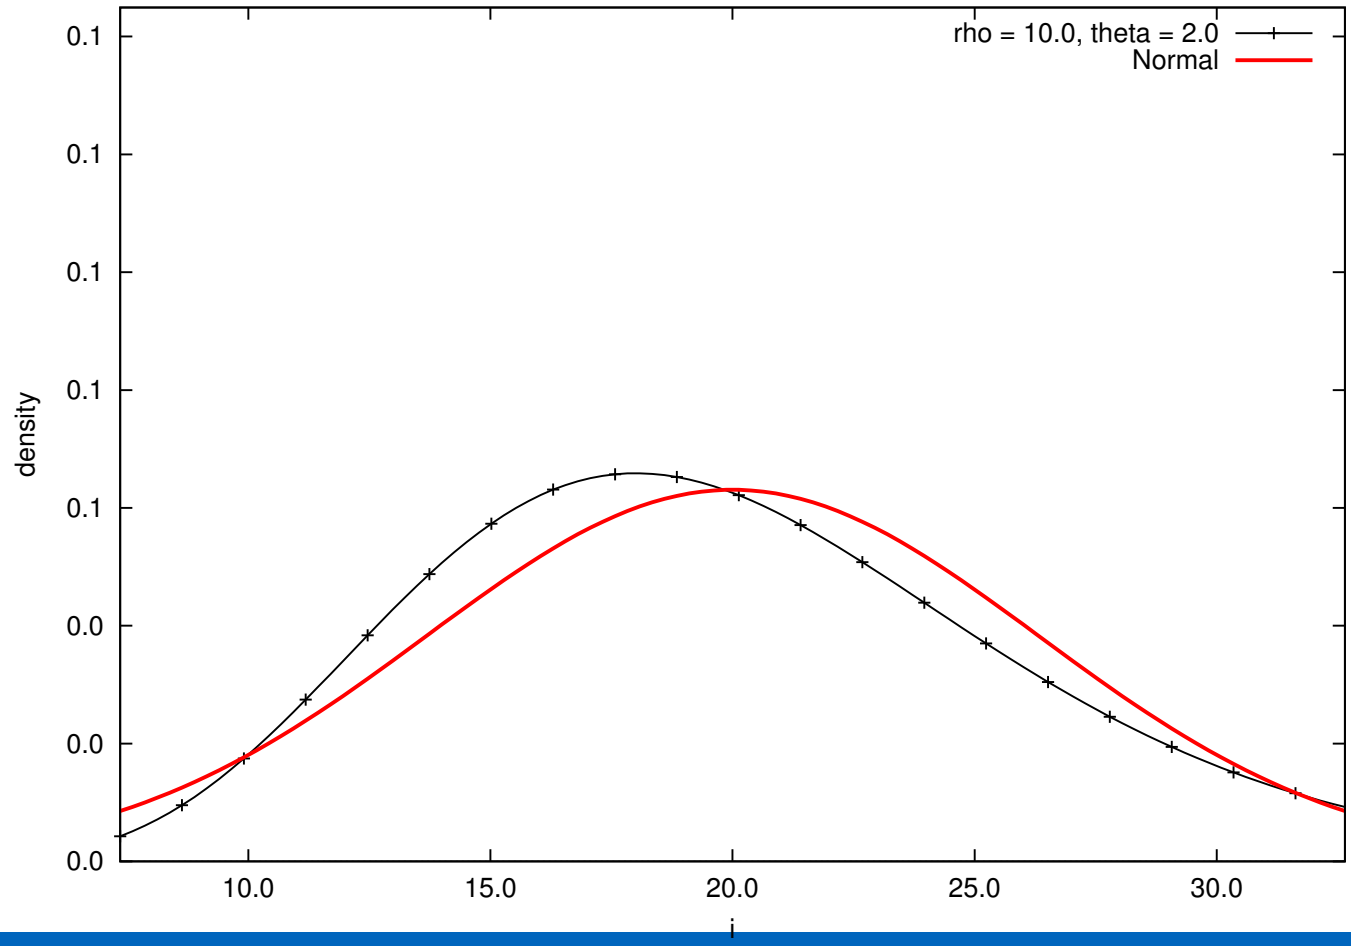

Gamma Distribution

$$\Gamma(\rho, \theta)$$

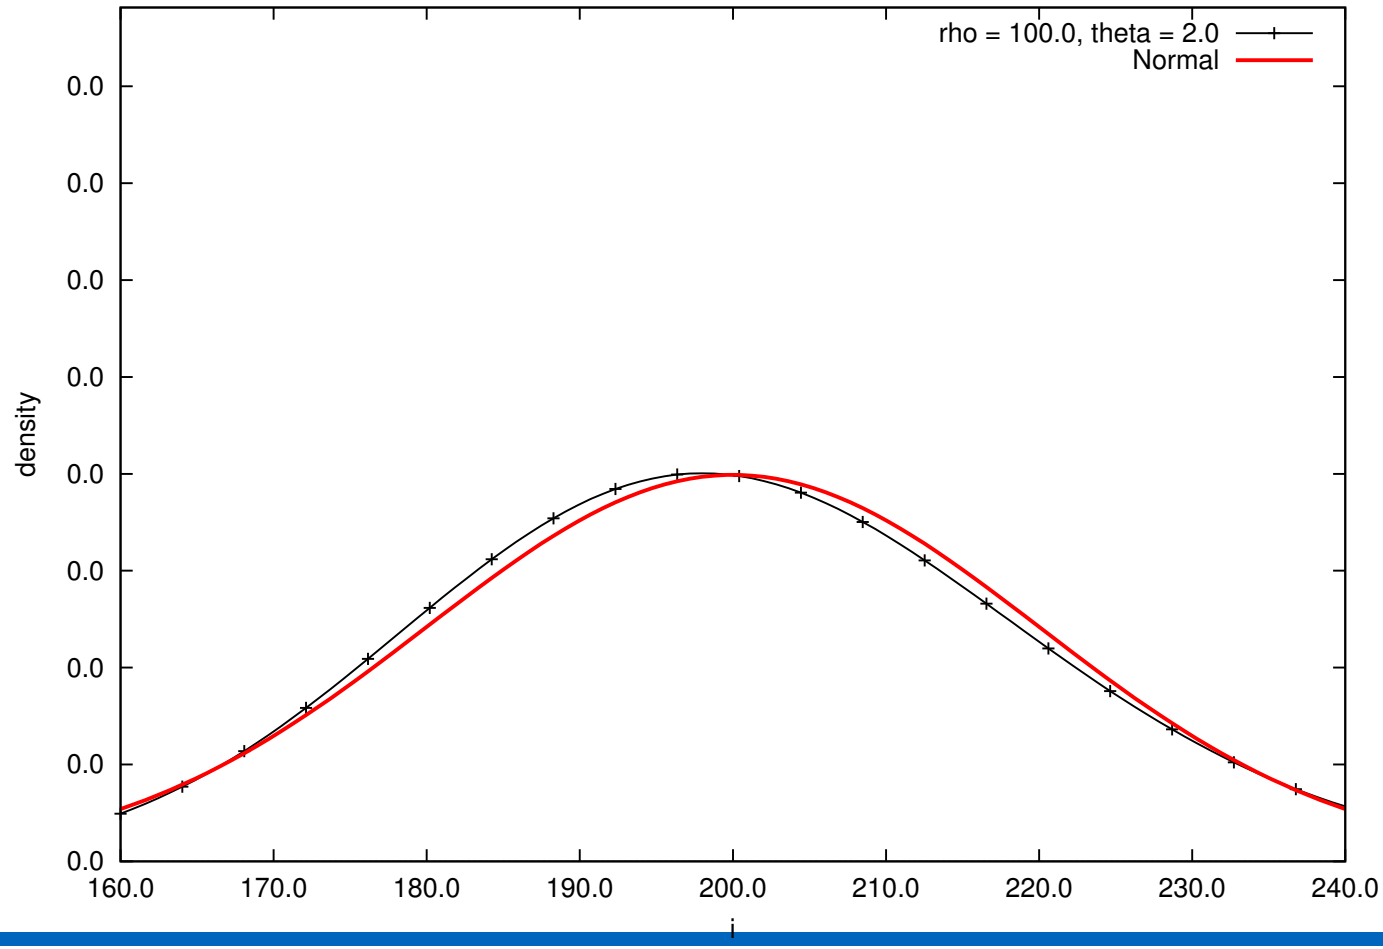

Gamma Distribution

$$\Gamma(\rho, \theta)$$

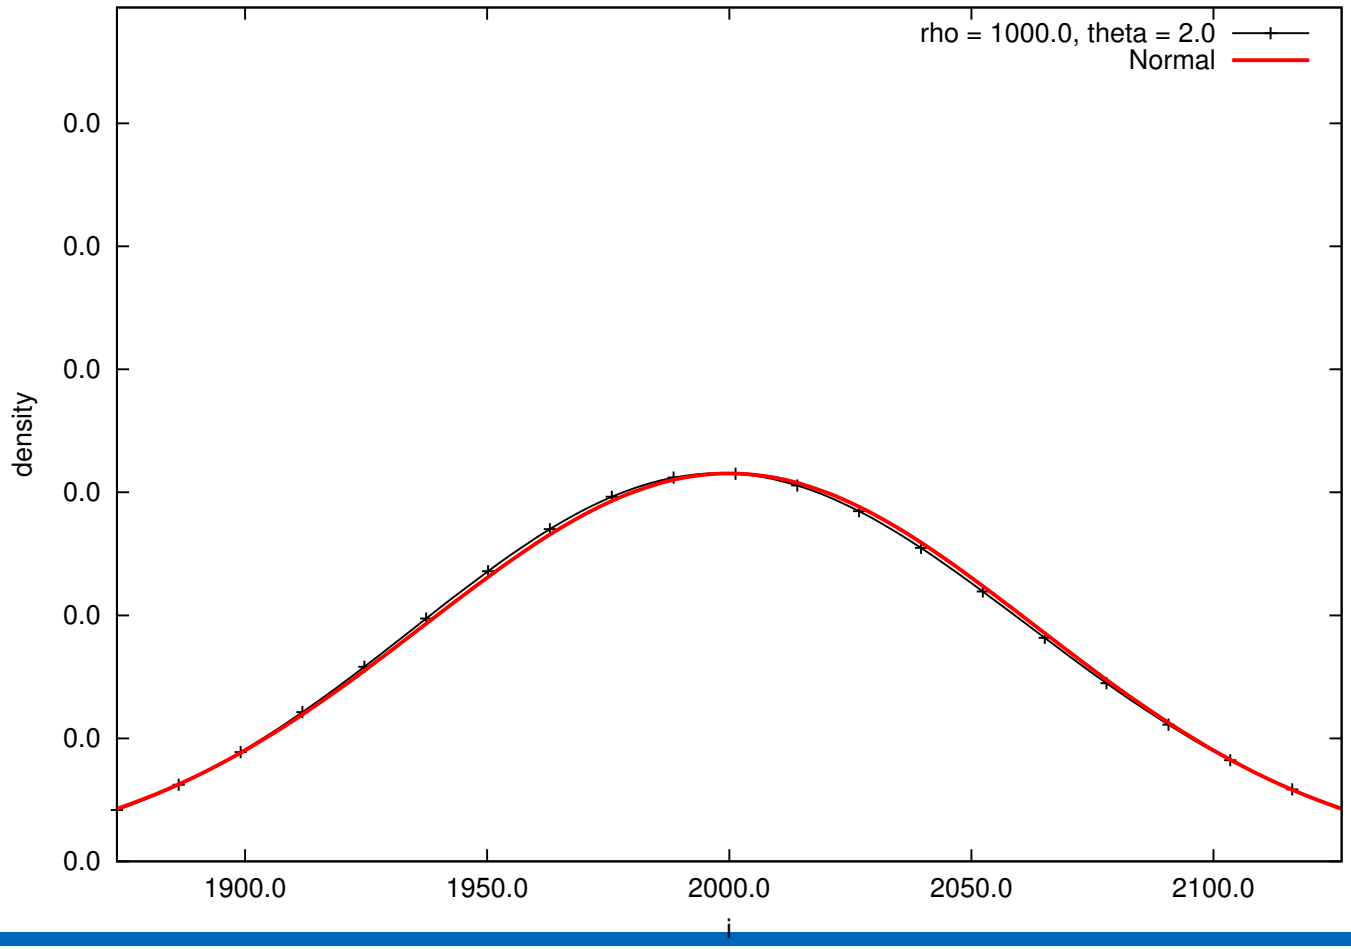

Hypergeometric Distribution

- Probability of i successes in N draws, without replacement, from a finite population of size 2^n that contains exactly $p2^n$ successes, wherein each draw is either a success or a failure
- $\mathcal{X} \sim \mathcal{H}(p2^n, 2^n, N)$
- pmf:

$$f(i) = \frac{\binom{2^n p}{i}\binom{2^n - i}{N - i}}{\binom{2^n}{N}}$$

- $\text{Exp}[X] = Np \qquad Var(X) = Np(1-p)\dfrac{2^n - N}{2^n - 1}$
- $B = \dfrac{2^n - N}{2^n - 1}$

Illustration

- $N = 1000$, $p = 0.5$
- $\mathrm{Exp}[X] = 500$ both in the Binomial and Hypergeometric case
- $2^n = 2000, 1100, 1024$

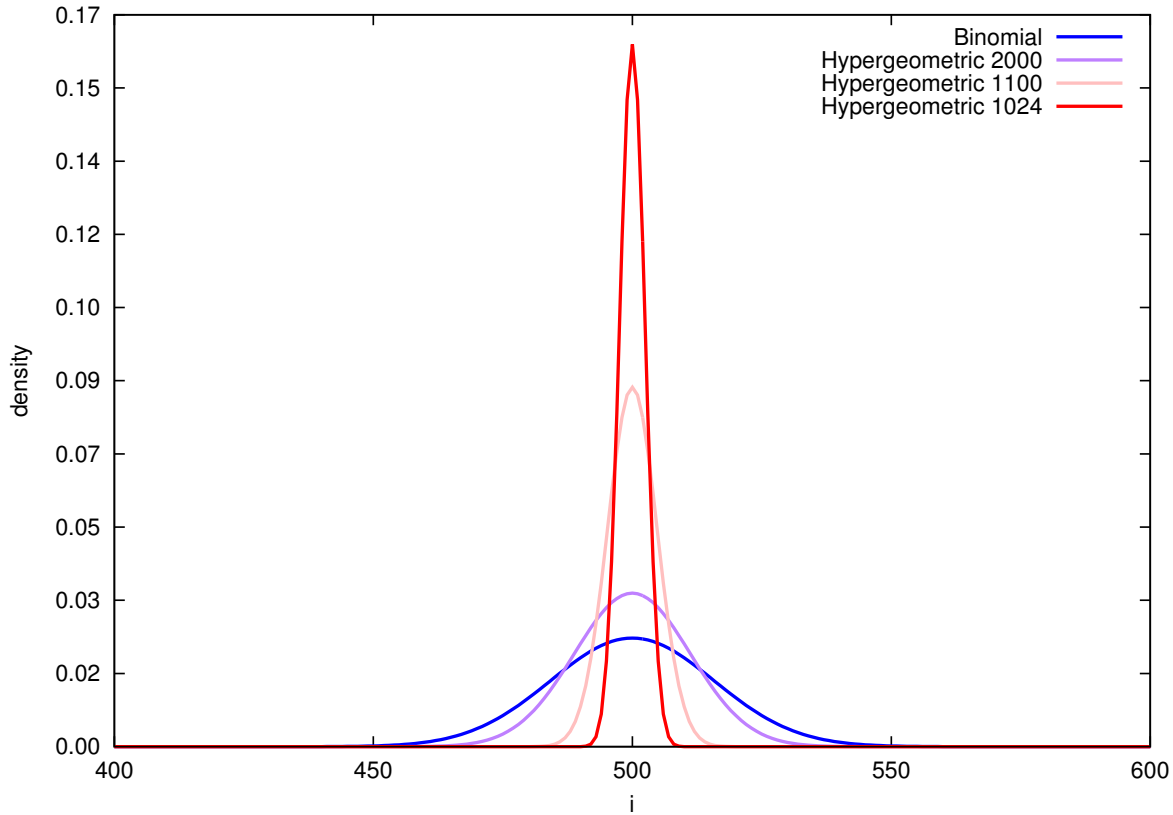

Distinct Plaintexts Attacks

- When performing the attack, we have to make sure that the plaintexts are not repeated
- Appear first for ZC attacks [Bogdanov et al 12]
- KP attacks $\leftrightarrow$ binomial distribution $B = 1$
- DKP attacks $\leftrightarrow$ hypergeometric distribution $B = \frac{2^n - N}{2^n - 1}$
- Variance: $Var(X) = Np(1-p)B$

Outline

Multiple/Multidimensional Linear Attacks

- Set of masks $(u, v) \in U \times V \setminus \{0, 0\}$
- Capacity : $C = \sum_{u \in U} \sum_{v \in V} cor_x^2(u, v)$
- Two approaches
  - Multiple linear attacks [Kaliski and Robshaw 01] [Biryukov et al 04]
  - Multidimensional linear attacks[Hermelin et al 08]
- Two types of tests:
  - χ^2 statistical test
  - LLR statistical test

Zero-Correlation (ZC) Linear Cryptanalysis

[Bogdanov et al 12, 13,14], [Soleimany, Nyberg 13]

The distinguisher takes advantage of linear approximation(s) with no bias

- Single approximation (u, v) with $cor(u, v) = 0$ for all keys
- Multiple approximations:

$$C = \sum_{u \in U, v \in V} cor^2(u, v) = 0,$$

  - multiple ZC: U and V without structure
  - multidimensional ZC: U and V linear (affine) spaces

Multidimensional Linear Attack

- We have 3 output masks $(v||01, v||10, v||11)$ which form a linear space if we include the trivial mask 0
- Instead of measuring the empirical correlations $\hat{cor}(u, v||01), \hat{cor}(u, v||10), \hat{cor}(u, v||11)$
- we can store (assuming that $u \cdot x$ is fixed to 0 for all plaintexts x) the frequency of the two output bits

j	$*\|\|00$	$*\|\|01$	$*\|\|10$	$*\|\|11$
$V[j]$	10	8	7	12

Exercise:
Compute the corresponding correlations

- The second approach is preferred in the multidimensional case

The Scoring Function

- ℓ linear approximations with empirical correlation $\hat{cor}_i$
- Scoring function:

$$T = N \sum_{i=1}^{\ell} \hat{cor}_i^2$$

- In multidimensional linear attacks T is equivalent (Walsh Transform) to:

$$T = \sum_{j=0}^{\ell} \frac{(V[j] - N/(\ell+1))^2}{N/(\ell+1)},$$

where $V[j]$ corresponds to the number of occurrences of the j-th element of the multidimensional distribution

Capacity for a Fixed Key

- Explanation for the multidimensional linear case
- Let $[p_j]_{j=0,\ldots,\ell}$ be a probability distribution with p_j representing the probability of a check in the j-th box

$$p_j = 2^{-n}|\{x \in \mathbb{F}_2^n \mid f(x) = j\}|,$$

  for a given function f

- The capacity is defined by

$$C = \sum_{j=0}^{\ell} \frac{\left(p_j - \frac{1}{\ell+1}\right)^2}{\frac{1}{\ell+1}}$$

Distribution of $T(D)$

- The distribution of $V[j]$ is approximated by a normal distribution with parameters

$$\begin{aligned} \mathrm{Exp}_{\mathbf{D}}(V[j]) &= Np_j \\ \mathrm{Var}_{\mathbf{D}}(V[j]) &= B \cdot Np_j(1-p_j) \approx B \cdot \frac{N}{\ell+1} \end{aligned}$$

$$T = \sum_{j=0}^{\ell} \frac{(V[j] - N\frac{1}{\ell+1})^2}{N\frac{1}{\ell+1}} = B\sum_{j=0}^{\ell} \frac{(V[j] - N\frac{1}{\ell+1})^2}{\mathrm{Var}_D(V[j])}$$

- $B^{-1}T(D)$ follows a non-central χ^2 distribution with ℓ degrees of freedom and non-centrality parameter $B^{-1}NC$
-

$$\begin{aligned} \mathrm{Exp}_D[B^{-1}T(D)] &= \ell + B^{-1}NC \text{ and} \\ \mathrm{Var}_D[B^{-1}T(D)] &= 2\ell + 4B^{-1}NC \end{aligned}$$

Expected Capacity for the Wrong Keys (1)

- The capacity varies with the key
- For a given linear approximation, the expected correlation in the uniform case is $\mathrm{Exp}_{\mathbf{D},\mathbf{K}}[cor_i(D,K)] = 0$
- Up to recently, it was assumed that for ℓ linear approximations the expected capacity

$$\mathrm{Exp}_{\mathbf{D},\mathbf{K}}[C(D,K)] = \sum_i \mathrm{Exp}_{\mathbf{D},\mathbf{K}}[cor_i^2(D,K)]$$

  is equal to 0
- However this is not true

Expected Capacity for the Wrong Keys (2)

- (Wrong-Key Hypothesis - Multiple) The correlations $cor_i(K)$, $i = 1, \ldots, \ell$ over the wrong keys K are i.i.d. and

$$cor_i(K) \sim \mathcal{N}\left(0, 2^{-n}\right)$$

- $2^n C_W(K)$ follows a χ^2 distribution with ℓ degree of freedom
- $C_W(K) \sim \Gamma\left(\frac{\ell}{2}, 2^{1-n}\right)$
- The mean and the variance are :

$$\begin{aligned} \mathrm{Exp}_{\mathbf{K}}[C_W(K)] &= 2^{-n}\ell \\ \mathrm{Var}_{\mathbf{K}}[C_W(K)] &= \frac{2}{\ell}\mathrm{Exp}_{\mathbf{K}}[C_W(K)]^2 = 2^{1-2n}\ell \end{aligned}$$

Combined Mean and Variance

- Combined mean and variance

$$\begin{aligned} \mathrm{Exp}_{\mathbf{D},\mathbf{K}}[T(D,K)] &= Exp_{\mathbf{K}}[Exp_{\mathbf{D}}[T(D,K)]], \\ \mathrm{Var}_{\mathbf{D},\mathbf{K}}[T(D,K)] &= Exp_{\mathbf{K}}[Var_{\mathbf{D}}[T(D,K)]] + Var_{\mathbf{K}}[Exp_{\mathbf{D}}[T(D,K)]] \end{aligned}$$

- For the wrong keys,

$$\begin{aligned} \mathrm{Exp}_{\mathbf{D},\mathbf{K}}\left[T_W(D,K)\right] &\approx B\ell + NC_W \text{ and} \\ \mathrm{Var}_{\mathbf{D},\mathbf{K}}\left[T_W(D,K)\right] &\approx \frac{2}{\ell}\left(B\ell + NC_W\right)^2 \end{aligned}$$

- with $C_W = 2^{-n}\ell$
- and $B = \dfrac{2^n - N}{2^n - 1} \approx 1 - \dfrac{N}{2^n}$ or $B = 1$

The Right Key Case

- Assuming the right-key randomization hypothesis we have

$$\begin{aligned} \mathrm{Exp}_{D,K}\left[T_R(D,K)\right] &= B\ell + NC_R \text{ and} \\ \mathrm{Var}_{D,K}\left[T_R(D,K)\right] &= \frac{2}{\ell}\left(B\ell + NC_R\right)^2, \end{aligned}$$

with $B = \dfrac{2^n - N}{2^n - 1} \approx 1 - \dfrac{N}{2^n}$ or $B = 1$

Experiments Distinct / Non-Distinct Plaintexts

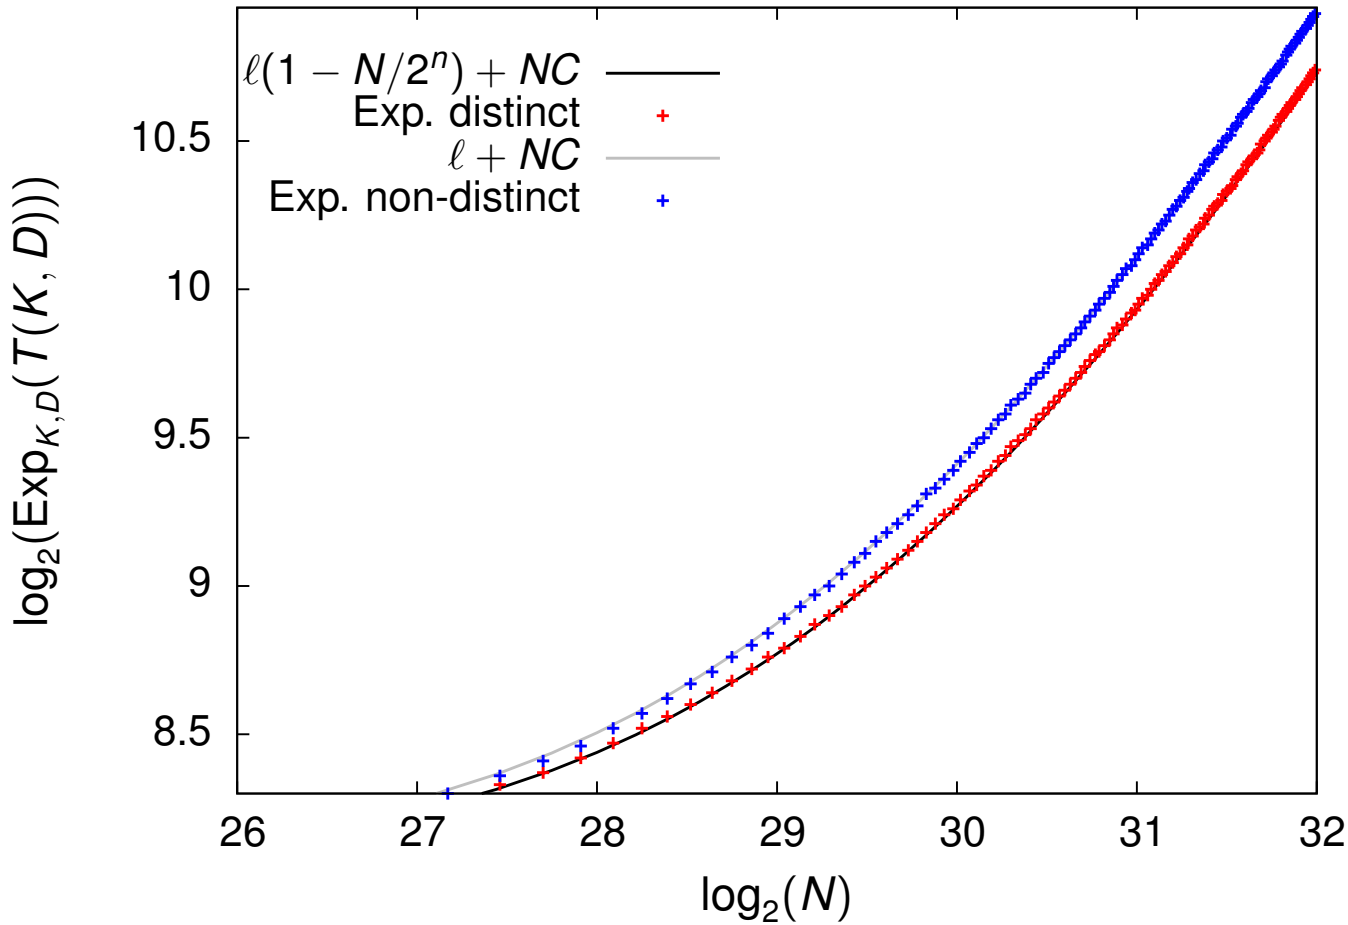

Success Probability of Multiple/Multidimensional Linear Attacks

- Success probability:

$$P_S = 1 - \alpha \approx \Phi\left(\frac{N|C_R - C_W| - \sqrt{2/\ell}(B\ell + NC_W)\varphi_\beta}{\sqrt{2/\ell}(B\ell + NC_R)}\right) \text{ then,}$$

$$N\left(|C_R - C_W| - \sqrt{2/\ell}(C_W\varphi_\beta + C_R\varphi_\alpha)\right) \approx \sqrt{2\ell}B(\varphi_\alpha + \varphi_\beta)$$

- Data complexity for a KP attack ($B = 1$):

$$N^{\text{non-distinct}} \approx \frac{\sqrt{2\ell}(\varphi_\alpha + \varphi_\beta)}{|C_R - C_W| - \sqrt{2/\ell}(C_W\varphi_\beta + C_R\varphi_\alpha)}$$

Outline

LLR Statistical Test

- $p = [p_w]_{w \in W}$: the expected probability distribution vector
  θ: the uniform one
  $q(k)$: the observed one for a key candidate k

- The Neyman-Pearson lemma gives the optimal form of the acceptance region on which is derived the *LLR* method. The optimality requires that both p and θ distributions are known (or at least the values p_w/θ_w).

- For a given number of sample N, the optimal statistical test consists in comparing the following statistic to a fixed threshold

$$LLR(q(k), p, \theta) = N \sum_{w \in W} q_w(k) log(\frac{p_w}{\theta_w})$$

Remarks on the LLR Statistical Test

- For this test “to work” we should have a good estimate of the involved probabilities (correlations)
- In practice this is really complicated to estimate the expected probabilities (correlations):
  - From the linear trails we only obtain a underestimate of the correlations
  - A given key probability differs from the expected one this test can fail

Exercise: Example of LLR test