

$$f(x) = \frac{1}{\pi} \cdot \frac{\gamma}{(x - x_0)^2 + \gamma^2}$$







PROBABILITY DISTRIBUTIONS II

*Binomial, Poisson, Normal and
Cauchy–Lorentz Distributions*



From the previous lecture

A quick recap of the key concepts we covered.

Random variable	Distribution	Location	Spread	Shape
 - A quantity whose value is determined by chance. - It can be discrete or continuous. Discrete / Continuous	Discrete (PMF)  Continuous (PDF / CDF)  - Discrete: probabilities for each possible value. - Continuous: described by PDF and CDF. PMF / PDF / CDF	 - Mean (expected value) - Median (50th percentile) - Mode (most probable value) Mean / Median / Mode	 - Variance (σ^2) - Standard deviation (σ) - Full width at half maximum (FWHM) Variance / σ / FWHM	 - Moments describe the shape of a distribution. - Skewness quantifies asymmetry. - Kurtosis quantifies peakedness (tailedness). Moments / Skewness / Kurtosis



These tools help us **describe** and **understand** data and **quantify uncertainty** in measurements.

Roadmap

Today we will study four fundamental probability distributions that are widely used in measurement science and physics.

In this lecture we will:

- Introduce and analyze the **binomial distribution** for successes in repeated trials.
- Explore the **Poisson distribution** for random events occurring in an interval.
- Study the **normal distribution** and understand the central limit theorem.
- Examine the **Cauchy–Lorentz distribution** and its important applications.



1. Binomial Distribution



2. Poisson Distribution



3. Normal Distribution



4. Cauchy–Lorentz Distribution



Key questions we will answer:

- When should we use each distribution?
- What are their main properties?
- How are they connected to real measurements?


Binomial experiment

Independent repeated trials:

- number of trials: n
- success probability: p
- failure probability: $q = 1 - p$

Random variable:

K = number of successes



Key idea: A binomial experiment consists of n independent trials, each with success probability p . The random variable K counts how many successes occur in the n trials.

Binomial distribution

$$P(K = k) = \binom{n}{k} p^k (1 - p)^{n-k}$$

where $k = 0, 1, \dots, n$ and $\binom{n}{k} = \frac{n!}{k!(n-k)!}$



Key idea: The binomial distribution gives the probability of obtaining exactly k successes in n independent Bernoulli trials with success probability p .



Notation: We write $K \sim \text{Bin}(n, p)$ to indicate that K has a binomial distribution with parameters n and p .

Mean and variance

If $K \sim \text{Bin}(n, p)$ with $q = 1 - p$, then:

Mean:

$$\mu = np$$

Variance:

$$\sigma^2 = np(1 - p) = npq$$


Standard deviation:

$$\sigma = \sqrt{np(1 - p)} = \sqrt{npq}$$



Interpretation:

- The mean np is the expected number of successes.
- The variance $np(1 - p)$ measures how much the number of successes typically varies around the mean.



Implication:

As n increases, both the mean and the standard deviation increase, but the distribution becomes relatively less spread out.




Summary: For a binomial random variable $K \sim \text{Bin}(n, p)$:

$$\mu = np, \quad \sigma^2 = np(1 - p) = npq, \quad \sigma = \sqrt{np(1 - p)} = \sqrt{npq}.$$

Binomial distribution in measurements

Suppose a histogram is constructed from N measurements.



For the j -th bin:

$$p_j = P(\text{measurement falls into bin } j)$$



The number of measurements in that bin, N_j , is itself a **random variable**.



$$N_j \sim \text{Binomial}(N, p_j)$$



Takeaway: Even the height of a histogram bin **fluctuates statistically**.

Note: This is why error bars on histogram bins can be estimated using binomial (or Poisson) statistics.

Properties of the Binomial Distribution

If $K \sim \text{Bin}(n, p)$, then the **binomial distribution** has the following important properties.

1. Mean

$$E[K] = np$$

The mean (expected number of successes) is n times the success probability p .

2. Variance

$$\text{Var}(K) = np(1 - p) = npq$$

The variance depends on both n and the product $pq = p(1 - p)$.

3. Standard Deviation

$$\sigma_K = \sqrt{np(1 - p)} = \sqrt{npq}$$

The standard deviation measures the typical spread of K around its mean np .

4. Mode

The mode m (the most probable value) is

$$m = \lfloor (n + 1)p \rfloor$$

If $(n + 1)p$ is an integer, then both m and $m - 1$ are modes.



Here $\lfloor x \rfloor$ denotes the largest integer less than or equal to x .

5. Symmetry

If $p = 0.5$, the distribution is symmetric around $\frac{n}{2}$:

$$P(K = k) = P(K = n - k) \text{ for } k = 0, 1, \dots, n$$




The probabilities are mirrored around the center $\frac{n}{2}$.

6. Shape

The shape depends on p and n .

- $p < 0.5$
Right-skewed
- $p = 0.5$
Symmetric
- $p > 0.5$
Left-skewed



Summary: The binomial distribution is fully determined by n and p . Its mean is np , variance is npq , it is symmetric when $p = 0.5$, and its shape depends on the value of p .

Counting random events



Physical examples:

- photons arriving at a detector;
- radioactive decays;
- cosmic-ray particles;
- detector counts;
- failures occurring in time.



Photon detection



Radioactive decays



Cosmic-ray particles



Detector counts



Failures in time



Common feature:

We count how many **independent events** occur in a **fixed interval**.

Poisson distribution

The Poisson distribution gives the probability of observing exactly k events in a fixed interval.

$$P(K = k) = \frac{\lambda^k e^{-\lambda}}{k!}$$

where

$$k = 0, 1, 2, \dots$$

and λ is the **expected number of events**.



Interpretation:

λ is the average (expected) number of events that occur in the interval.



Properties:

- Events occur independently.
- The probability of more than one event in a very small subinterval is negligible.
- The rate λ is constant over the interval.



Key idea:

The Poisson distribution models the number of random events in a fixed interval when events are **rare** and **independent**.



Note: The Poisson distribution is the standard model for counting random events.

Mean and variance of Poisson distribution

A Poisson distribution with parameter λ has

Mean (expected value):

$$\mu = \lambda$$

and

Variance:

$$\sigma^2 = \lambda$$

Therefore: →

$$\sigma = \sqrt{\lambda}$$



Key message:

Mean = variance



Interpretation:

For a Poisson distribution, the average number of events and the size of its fluctuations are completely determined by λ .




Implication:

If the expected number of events doubles, both the mean and the variance double.


Shape of the Poisson distribution

The shape of the Poisson distribution depends on the value of λ .


$\lambda = 0.2$




$\lambda = 2$



$\lambda = 5$



$\lambda = 12.5$



Key observations:



- For small λ , the distribution is strongly right-skewed.
- As λ increases, the distribution becomes more symmetric.



- For large λ , the distribution approaches a symmetric, bell-shaped form.
- This reflects the approach to the **normal distribution**.



Takeaway:

Small $\lambda \rightarrow$ **asymmetric**; large $\lambda \rightarrow$ **increasingly symmetric**.

Skewness and kurtosis



For a Poisson distribution with parameter λ , the skewness γ_1 and the excess kurtosis γ_2 are

$$\gamma_1 = \frac{1}{\sqrt{\lambda}}$$

and

$$\gamma_2 = \frac{1}{\lambda}$$




Therefore: $\lambda \rightarrow \infty \rightarrow \gamma_1 \rightarrow 0, \gamma_2 \rightarrow 0 \rightarrow \text{Poisson} \rightarrow \text{increasingly Gaussian-like}$



Interpretation:

For small λ , the distribution is positively skewed (long tail to the right) and more peaked than a normal distribution. As λ increases, the distribution becomes more symmetric and flatter.



Key message:

Both skewness and excess kurtosis decrease as λ increases, approaching zero in the limit. The Poisson distribution becomes increasingly similar to the **normal distribution**.



Note: This explains why, for large expected counts (large λ), the Poisson distribution can be well approximated by a normal distribution.

From binomial to Poisson

➔ **Binomial distribution:** the probability of k events in n trials is

$$P(K = k) = \binom{n}{k} p^k (1 - p)^{n-k} \quad \text{for } k = 0, 1, 2, \dots, n$$

Consider the limit: $n \rightarrow \infty$, $p \rightarrow 0$, while $np = \lambda$ remains finite.

➡ **Then** the binomial distribution approaches the **Poisson distribution:**

$$\lim_{\substack{n \rightarrow \infty \\ p \rightarrow 0 \\ np = \lambda}} \binom{n}{k} p^k (1 - p)^{n-k} = \frac{\lambda^k e^{-\lambda}}{k!} \quad \text{for } k = 0, 1, 2, \dots$$



Intuition:

- As n becomes very large, the number of opportunities grows.
- As p becomes very small, each opportunity has a tiny chance p .
- The expected number of events np stays equal to λ .
- In this limit, the binomial distribution converges to **Poisson(λ)**.



Why this matters:




- Many real experiments have a very large number of possible trials (n), but the event probability is tiny (p).
- Poisson distribution is the natural model for counting **rare events** in a fixed interval.



Takeaway: When $n \rightarrow \infty$, $p \rightarrow 0$, and $np = \lambda$ is finite, the binomial distribution converges to **Poisson(λ)**.

Visual comparison: binomial vs. Poisson

Example: $n = 100$, $p = 0.01$ so $np = 1$ (compare with Poisson $\lambda = 1$)



Takeaway: Different models – **almost identical probabilities.**

When n is large and p is small with $np = 1$, the binomial distribution is nearly indistinguishable from Poisson($\lambda = 1$).

When binomial, when Poisson?

Both models describe counts of events, but they are used in different situations.



BINOMIAL

Model for a **fixed number** of trials

Known:

n Number of trials (opportunities)

p Probability of success in each trial

Assumptions:

- There are exactly n independent trials.
- Each trial has the same probability of success p .
- The number of successes K is counted.

Typical situation:



We know in advance how many opportunities we will have.

Example: 10 coin tosses, 100 manufactured items inspected, 20 patients treated.

VS.



POISSON

Model for **events occurring** in time/space

Known:

r Event rate (average number of events per unit time/space)

Δt Observation interval (time or space length)

Assumptions:

- Events occur independently.
- The probability of an event in a very small interval is proportional to the interval length.
- The event rate r is constant.

Expected count:

$$\lambda = r \Delta t$$

Typical situation:



We observe a stream of events for a period of time Δt at an average rate r .

Example: photon counts in a detector, radioactive decays, equipment failures, cosmic-ray hits.



Key takeaway: Use binomial when n and p are known.

Use Poisson when the rate r and the interval Δt are known.

Properties of the Poisson Distribution

If $K \sim \text{Poisson}(\lambda)$, then the following properties hold.

1. Mean and variance

The mean and the variance of a $\text{Poisson}(\lambda)$ random variable are equal to λ .

$$\mathbb{E}[K] = \lambda$$

$$\text{Var}(K) = \lambda$$

Thus the standard deviation is $\sigma = \sqrt{\lambda}$.

2. Additivity

If $K_1 \sim \text{Poisson}(\lambda_1)$ and $K_2 \sim \text{Poisson}(\lambda_2)$ are independent, then

$$K_1 + K_2 \sim \text{Poisson}(\lambda_1 + \lambda_2)$$

More generally:

If $K_1, K_2, \dots, K_n$ are independent with $K_i \sim \text{Poisson}(\lambda_i)$, then

$$\sum_{i=1}^n K_i \sim \text{Poisson}\left(\sum_{i=1}^n \lambda_i\right)$$

3. Splitting (thinning)

If $K \sim \text{Poisson}(\lambda)$ and, independently, each event is classified as "type A" with probability p and "type B" with probability $1-p$, then


- $K_A \sim \text{Poisson}(p\lambda)$
- $K_B \sim \text{Poisson}((1-p)\lambda)$
- K_A and K_B are independent
- $K = K_A + K_B$

Example (splitting):

Customers arrive at a service desk according to a Poisson process with rate $\lambda = 20$ per hour. Independently, 60% of them ask for information (type A) and 40% ask for complaints (type B).

Then:

- $K_A \sim \text{Poisson}(0.6 \cdot 20) = \text{Poisson}(12)$
- $K_B \sim \text{Poisson}(0.4 \cdot 20) = \text{Poisson}(8)$
- K_A and K_B are independent
- $K = K_A + K_B \sim \text{Poisson}(20)$



4. Probability of zero

$$\mathbb{P}(K = 0) = e^{-\lambda}$$


5. Mode

The mode is $\lfloor \lambda \rfloor$.

If λ is an integer, then both $\lambda - 1$ and λ are modes.

6. Shape

For small λ , the distribution is right-skewed. As λ increases, it becomes more symmetric and approaches a normal distribution.



Key takeaway: The Poisson distribution is simple yet powerful: it models random counts in time or space, and has elegant mathematical properties that make it very useful in practice.

$$K \sim \text{Poisson}(\lambda)$$

Example: photon detection

A common physical situation: **photons arrive randomly** at a detector.



Model

Photons are emitted by a light source at random times. The detector counts how many photons arrive during a fixed counting time interval Δt .



Event rate


Let r be the average number of photons detected per unit time (photons per second). We assume r is constant during the measurement.



Expected number of counts

In a time interval Δt , the expected number of detected photons is

$$\lambda = r \Delta t$$



Distribution of counts

Under these conditions (independent arrivals, constant rate), the number of detected photons K in the interval Δt follows

$$K \sim \text{Poisson}(\lambda)$$

where $\lambda = r \Delta t$ is the expected number of counts.



Key message: Photon arrivals are random, independent events with (approximately) constant rate r .

Therefore, the count during a fixed time Δt is **Poisson** distributed with parameter $\lambda = r \Delta t$.



$$K \sim \text{Poisson}(r \Delta t)$$

Slide 19 – Photon counting and shot noise

Since

$$\sigma = \sqrt{\lambda}$$

Relative fluctuation:

$$\frac{\sigma}{\mu} = \frac{\sqrt{\lambda}}{\lambda} = \frac{1}{\sqrt{\lambda}}$$

Large count:

$\lambda \uparrow$




relative noise:

$\frac{1}{\sqrt{\lambda}} \downarrow$



This fundamental fluctuation is called shot noise.




Poisson model

The number of detected photons K in the interval Δt follows

$$K \sim \text{Poisson}(\lambda)$$

with $\lambda = r \Delta t$ (mean count).



Mean: λ Variance: λ SD: $\sqrt{\lambda}$

Slide 20 – Why measure longer?

If

$$\lambda = r \Delta t$$

(r = event rate, Δt = counting time)

Then

$$\sigma = \sqrt{\lambda} = \sqrt{r \Delta t}$$

Relative uncertainty
(relative fluctuation):

$$\frac{\sigma}{\mu} = \frac{1}{\sqrt{\lambda}} = \frac{1}{\sqrt{r \Delta t}}$$

Therefore:

$$\Delta t \uparrow \Rightarrow \text{relative uncertainty } \frac{1}{\sqrt{r \Delta t}} \downarrow$$





Measuring for a longer time **reduces**
the **relative uncertainty** (noise).





The mean number of counts is proportional to time ($\lambda = r \Delta t$), while the standard deviation grows only with $\sqrt{\Delta t}$.

Short counting time → large relative fluctuations



Long counting time → small relative fluctuations



Key message: Longer measurements average out the random fluctuations. The **relative uncertainty** decreases as $1/\sqrt{\Delta t}$.

The normal distribution

The most important continuous distribution in measurement science

Definition

A continuous random variable X has a normal (Gaussian) distribution with parameters μ and $\sigma > 0$, written

$$X \sim \mathcal{N}(\mu, \sigma^2),$$


if its probability density function is

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp \left[-\frac{(x - \mu)^2}{2\sigma^2} \right], \quad -\infty < x < \infty.$$

Parameters

- μ = **mean** (location parameter)
- σ = **standard deviation** (scale parameter), $\sigma > 0$

Shape and interpretation



Key takeaway:

The normal distribution $\mathcal{N}(\mu, \sigma^2)$ is fully determined by two parameters: the mean μ (location) and the standard deviation σ (spread).

Role of μ

Three normal distributions with different means, same standard deviation σ .

Setup


- Same standard deviation: σ
- Different means:

$$\mu_1 < \mu_2 < \mu_3$$

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

Legend

- $\mu = \mu_1$ (leftmost)
- $\mu = \mu_2$ (middle)
- $\mu = \mu_3$ (rightmost)



Key message:

μ determines the **location**.

Role of σ

Three normal distributions with the same mean μ , different standard deviations σ .

Setup

- Same mean: μ (all centered at the same location)
- Different standard deviations:


$$\sigma_1 < \sigma_2 < \sigma_3$$

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

Legend

- $\sigma = \sigma_1$ (small, narrow, tall)
- $\sigma = \sigma_2$ (medium width)
- $\sigma = \sigma_3$ (large, wide, flat)

As σ increases,
the curve gets **wider**
and **lower**.



Small σ (σ_1):
narrow and tall

Medium σ (σ_2):
moderate width

Large σ (σ_3):
wide and flat



Key message:

σ determines the **width** (spread).

Properties of the normal distribution

1 Total area:

$$\int_{-\infty}^{\infty} f(x) dx = 1$$

2 Symmetric:

mean = median = mode = μ

The distribution is symmetric about μ .

3 Skewness:

$$\gamma_1 = 0$$


The distribution is perfectly symmetric.

4 Excess kurtosis:

$$\gamma_2 = 0$$

The distribution has the same "peakedness" as the normal distribution.

Normal distribution



Key takeaway:

The normal distribution is **symmetric**, **unimodal**, has total area **1**, and has **skewness 0** and **excess kurtosis 0**.

Probability as area

What is the probability that X lies between a and b ?


Definition

For a continuous random variable X with probability density function $f(x)$, the probability that X lies between a and b ($a < b$) is

$$P(a < X < b) = \int_a^b f(x)dx$$

Notes

- The probability equals the area under the curve $f(x)$ between a and b .
- Total area under the curve is 1.



Key takeaway:

For a normal distribution, $P(a < X < b) = \int_a^b f(x)dx$, i.e., the probability is equal to the area under the curve between a and b .

Why use standardization?

Different normal distributions have different μ and σ .
Can we describe all of them using one universal distribution?

The problem

Normal distributions can have any mean μ and any standard deviation σ . Their shapes and locations differ.


Goal

Transform any normal variable X into a new variable Z that follows **one universal distribution**.

Why useful?

- One universal distribution $\mathcal{N}(0, 1)$
- Easier to compute probabilities
- Universal tables and software apply to all normal distributions

Idea



By **shifting** (subtracting μ) and **scaling** (dividing by σ), we can make them all look the same.



Yes.

Define the standardized variable

$$z = \frac{x - \mu}{\sigma}$$



Key takeaway: Standardization converts any normal distribution $\mathcal{N}(\mu, \sigma^2)$ into the standard normal distribution $\mathcal{N}(0, 1)$, allowing us to use one universal distribution for all cases.

Standard normal distribution

If $Z = \frac{X - \mu}{\sigma}$, then Z follows the standard normal distribution.

Definition


A random variable Z has a standard normal distribution, denoted by $Z \sim \mathcal{N}(0, 1)$, if it is normally distributed with

$$Z \sim \mathcal{N}(0, 1)$$

Parameters

- Mean: $\mu_Z = 0$
- Standard deviation: $\sigma_Z = 1$

Standard normal distribution $\mathcal{N}(0, 1)$



The standard normal distribution has mean 0 and standard deviation 1.



Key message:

By standardization $z = \frac{x - \mu}{\sigma}$, any normal random variable $X \sim \mathcal{N}(\mu, \sigma^2)$ can be transformed into $Z \sim \mathcal{N}(0, 1)$ with $\mu_Z = 0$ and $\sigma_Z = 1$.

The 68–95–99.7 rule

For a normal distribution, almost all values are close to the mean.

Empirical rule

If $X \sim \mathcal{N}(\mu, \sigma^2)$, then approximately:

■ $P(|X - \mu| < \sigma) \approx 68.3\%$


About 68.3% of the values lie within 1 standard deviation of the mean.

■ $P(|X - \mu| < 2\sigma) \approx 95.4\%$

About 95.4% of the values lie within 2 standard deviations of the mean.

■ $P(|X - \mu| < 3\sigma) \approx 99.7\%$

About 99.7% of the values lie within 3 standard deviations of the mean.



Key takeaway:

For a normal distribution, about 68.3% of values lie within 1σ , about 95.4% within 2σ , and about 99.7% within 3σ of the mean.

Why does the Gaussian appear so often?

Many small independent influences

Many physical measurements are influenced by

$$Y_1, Y_2, \dots, Y_n$$

many small independent contributions.


Measured quantity

The measured quantity is the sum

$$S = Y_1 + Y_2 + \dots + Y_n$$

The central question

What is the distribution of S ?



Central limit theorem

If many independent random contributions with finite variances contribute to a sum, then under broad conditions:

$$S = Y_1 + Y_2 + \cdots + Y_n \rightarrow \text{Normal distribution}$$


as the number of contributions n becomes large.

Interpretation

- No matter what the individual distributions of $Y_1, Y_2, \dots, Y_n$ look like (as long as they have finite variances and are independent), the distribution of their sum S will be approximately Gaussian when n is large.

Important!

The individual variables **do not need to be normally distributed.**



Key message:

The central limit theorem explains why the normal (Gaussian) distribution appears so often in nature and in experimental data.

Visualizing the Central Limit Theorem

As we **average** more independent random variables, the distribution of the average becomes increasingly **Gaussian**.


Setup

Let $X_1, X_2, \dots$ be i.i.d. (random variables) with mean μ and finite variance σ^2 , but not Gaussian.

We look at the average

$$\bar{X}_N = \frac{1}{N} \sum_{i=1}^N X_i$$

and its distribution for increasing N .



What we see

Starting from a strongly non-Gaussian distribution, the distribution of the average $\bar{X}_N$ rapidly approaches a Gaussian shape as N increases.

This is the essence of the Central Limit Theorem.

Consequences for measurement science




Binomial distribution

Number of successes in
 n independent Bernoulli trials


for large n



Gaussian approximation



Approximately normal with
 $\mu = np, \sigma^2 = np(1-p)$




Poisson distribution

Number of events in a fixed
interval with rate λ

for large λ



Gaussian approximation



Approximately normal with
 $\mu = \lambda, \sigma^2 = \lambda$

$X_1, X_2, \dots, X_N$
• • • •


Sample means

Mean of N i.i.d. variables
with mean μ and variance σ^2

for large N



Gaussian distribution



Exactly (asymptotically) normal with
 $\mu = \mu, \sigma^2 = \frac{\sigma^2}{N}$



This is one reason why the normal distribution is central to experimental physics.

Cauchy–Lorentz distribution

Probability density

$$f(x) = \frac{1}{\pi} \frac{\gamma}{(x - x_0)^2 + \gamma^2}$$


Parameters

x_0 = location parameter (center of the peak)

γ = HWHM (half width at half maximum)

Therefore

$$FWHM = 2\gamma$$




Key points

- The Cauchy–Lorentz distribution is centered at x_0 .
- γ is the half width at half maximum (HWHM).
- The full width at half maximum is $FWHM = 2\gamma$.

Gaussian vs. Lorentzian

Gaussian (Normal) distribution

$$f_G(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$




Tails

$\sim e^{-x^2}$
→
rapidly
decreasing tails

Lorentzian (Cauchy–Lorentz) distribution

$$f_L(x) = \frac{1}{\pi} \frac{\gamma}{(x-x_0)^2 + \gamma^2}$$



Tails

$\sim \frac{1}{x^2}$
→
heavy tails



Critical message:

For the Cauchy–Lorentz distribution, the mean and variance are not defined.

- The integrals for the mean diverge.
- The integrals for the variance diverge.
- Only the location parameter x_0 and the scale parameter γ are meaningful.



Key takeaway

The Gaussian has thin (exponential) tails and finite moments of all orders.
The Lorentzian has heavy (power-law) tails and undefined mean and variance.

Physical example: spectral line shape




Many physical systems (e.g., atomic transitions, damped oscillators) produce **Lorentzian line shapes** in the frequency domain.

1. Time domain

Exponentially damped oscillation

$$x(t) = Ae^{-\gamma|t|} \cos(\omega_0 t + \phi)$$




2. Fourier transform

Take the Fourier transform
from time to frequency domain

$$X(\omega) = \mathcal{F}\{x(t)\}$$

3. Frequency domain

Lorentzian spectral profile



Examples

- Natural linewidth of atomic and molecular transitions
- Resonances in scattering and spectroscopy
- Damped oscillators and RLC circuits

In all these cases, energy leaks out **exponentially** in time → the spectrum has a **Lorentzian shape**.



Takeaway

The Lorentz distribution is not a mathematical curiosity — it occurs naturally in physics.


SUMMARY – Key Takeaways

BINOMIAL

- Fixed number n of independent trials
- Known success probability p



Number of successes
in n trials




Mean: np
Variance: $np(1-p)$

POISSON

- Independent random events
- Known average rate λ (events per interval)



Number of events in
a given interval




Mean: λ
Variance: λ

NORMAL

- Continuous variable
- Many independent contributions



Measurement fluctuations
/ averages




Mean: μ
Variance: σ^2

CAUCHY-LORENTZ

- Heavy-tailed continuous distribution
- Resonance / linewidth processes



Spectral profiles
(line shapes)



Mean: not defined
Variance: not defined



Take-home
messages

- Different questions require different distributions—there is **no one-size-fits-all**.
- Binomial and Poisson describe **counts** of discrete events.
- Normal describes **fluctuations and averages** from many small independent effects.
- Cauchy-Lorentz describes **heavy-tailed** phenomena such as linewidths and resonances.
- **Always check the physics: assumptions behind each model matter!**



Statistical models
are tools.
Understanding the
process is essential.