MODULE 1

Random Variables & Probability Distributions

BSM301 · Mathematics III

Track: Internet of Things / Cyber Security / Blockchain Technology · Now with extra worked examples

H T H H T *...one coin toss at a time, building up to a full probability distribution.*

What We'll Cover

1. **Random Variables**
   Discrete vs. continuous — how uncertainty becomes a number

2. **Discrete Distributions**
   Probability mass functions, mean and variance — 2 worked examples

3. **Binomial Distribution**
   Coin tosses → Pascal's triangle → binomial pmf — 2 worked examples

4. **Poisson Distribution**
   The limiting case of Binomial when events are rare — 2 worked examples

5. **Continuous Distributions**
   Exponential & Normal — probability as area under a curve — 3 worked examples

6. **Properties Summary**
   Side-by-side comparison of all four distributions

What Is a Random Variable?

A rule that assigns a number to the outcome of a random experiment

X : Sample Space → Real Numbers. *Example: X = number of heads in 3 coin tosses* $\Rightarrow X \in \{0, 1, 2, 3\}$

DISCRETE

Countable set of values (finite or countably infinite)

- Number of heads in n coin tosses
- Number of corrupted packets per hour
- Number of blocks mined in a day
- Number of failed login attempts

CONTINUOUS

Uncountably infinite values over an interval

- Battery lifetime of an IoT sensor
- Time between successive cyber-attacks
- Block confirmation time on a blockchain
- Sensor calibration error (°C)

Discrete Probability Distribution

Every value of X gets a probability — and the rules are strict

Two Defining Properties

1. $0 \le p(x) \le 1$
   every probability is between 0 and 1
2. $\Sigma\, p(x) = 1$
   probabilities over all outcomes sum to 1

Note: A discrete distribution can also be shown as a table, a formula p(x), or a bar chart — all three describe the same thing.

Example: Rolling a Fair Die

x	1	2	3	4	5	6
p(x)	1/6	1/6	1/6	1/6	1/6	1/6

Check: $p(x) \ge 0$ for all x, and $1/6 \times 6 = 1$ ✓

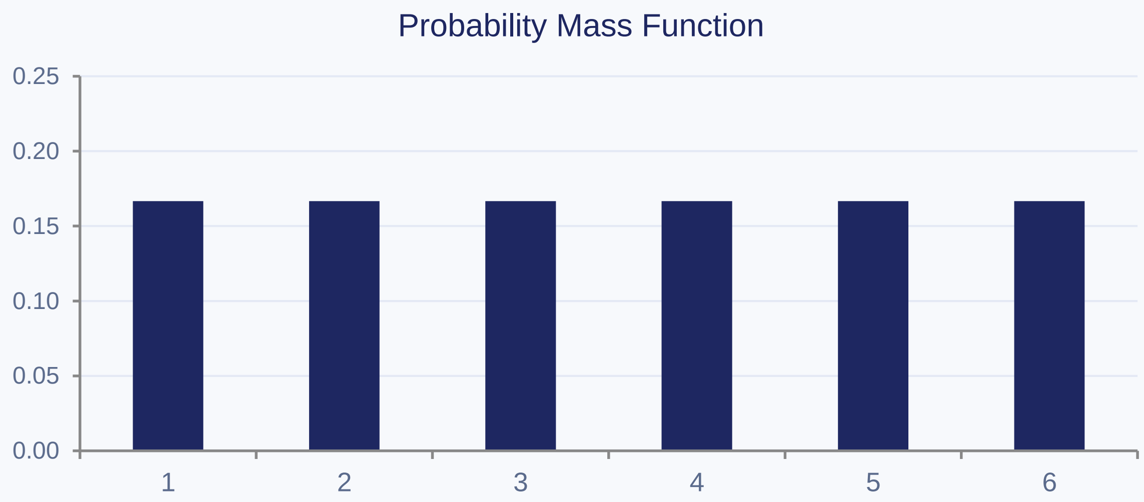

Mean and Standard Deviation

EXAMPLE 1

Mean: $E(X) = \mu = \Sigma\, x \cdot p(x)$

Variance: $Var(X) = \sigma^2 = \Sigma\, x^2 \cdot p(x) - \mu^2$ Standard Deviation: $\sigma = \sqrt{Var(X)}$

Number of Heads in 3 Coin Tosses

x	0	1	2	3
p(x)	0.1250	0.3750	0.3750	0.1250
x·p(x)	0.0000	0.3750	0.7500	0.3750
x^2·p(x)	0.0000	0.3750	1.5000	1.1250

$\mu = E(X) = 1.50$

$\sigma^2 = 3.0000 - (1.50)^2 = 0.7500$

$\sigma = \sqrt{0.7500} \approx 0.866$ heads

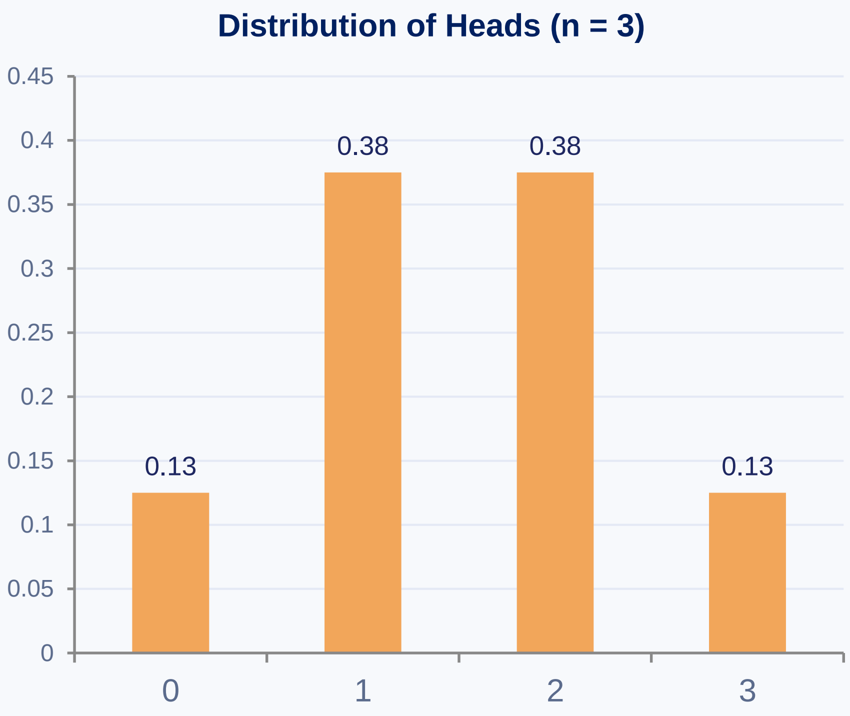

A Second Discrete Example

EXAMPLE 2

The same three steps work for any discrete distribution, not just coin tosses

A help-desk records how many support tickets it resolves per hour. From past data:

x (tickets)	0	1	2	3	4
p(x)	0.10	0.25	0.35	0.20	0.10
x·p(x)	0.00	0.25	0.70	0.60	0.40
x^2·p(x)	0.00	0.25	1.40	1.80	1.60

Check: 0.10+0.25+0.35+0.20+0.10 = 1.00 ✓

Step-by-Step

1. **Mean: $\mu = \Sigma x \cdot p(x) = 1.95$ tickets/hr**
2. **$E(X^2)$: $\Sigma x^2 \cdot p(x) = 5.0500$**
3. **Variance: $\sigma^2 = 5.0500 - (1.95)^2 = 1.2475$**

   $\sigma = \sqrt{1.2475} \approx 1.117$ tickets

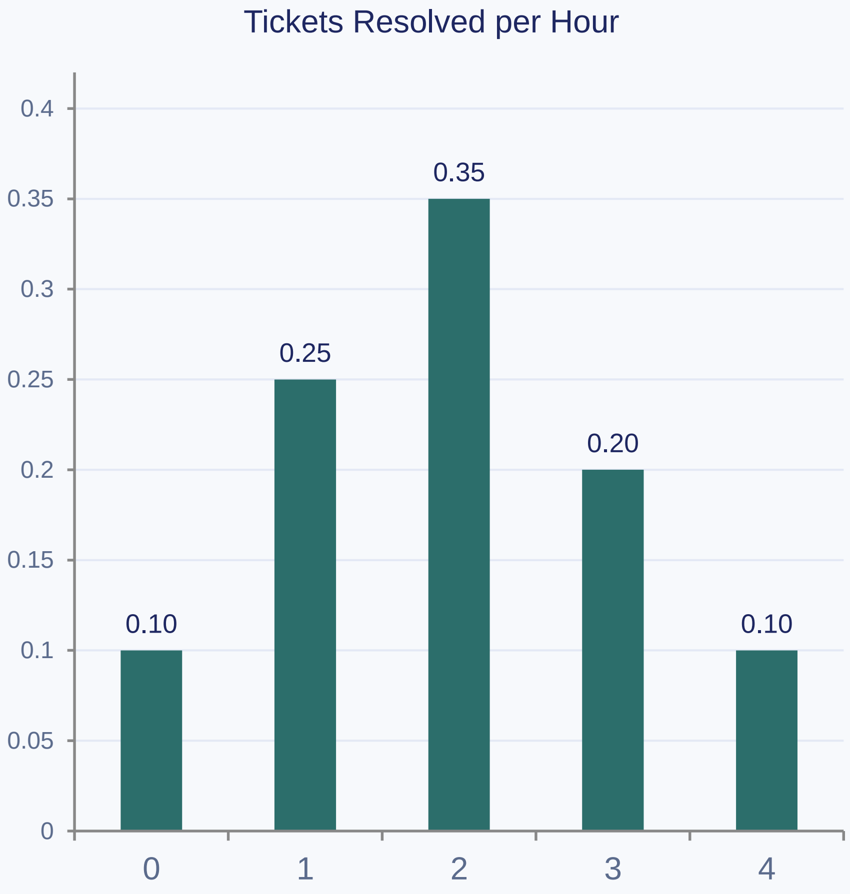

From Coin Tosses to the Binomial Distribution

Toss a coin 3 times and count the heads — every outcome, grouped

k = 0	k = 1	k = 2	k = 3
TTT	**H**TT	**HH**T	**HHH**
	T**H**T	**H**T**H**	
	TT**H**	T**HH**	
1	3	3	1
C(3,0) ways	*C(3,1) ways*	*C(3,2) ways*	*C(3,3) ways*

General pattern: out of 2^n equally likely sequences, exactly C(n, k) of them have k heads.

That count — the binomial coefficient — is exactly what generates the Binomial distribution's shape.

The Binomial Distribution

n independent trials, each with success probability p

$P(X = k) = C(n, k) \cdot p^{k} \cdot (1-p)^{n-k}$ $k = 0, 1, 2, \ldots, n$

n = 10, p = 0.5 — symmetric

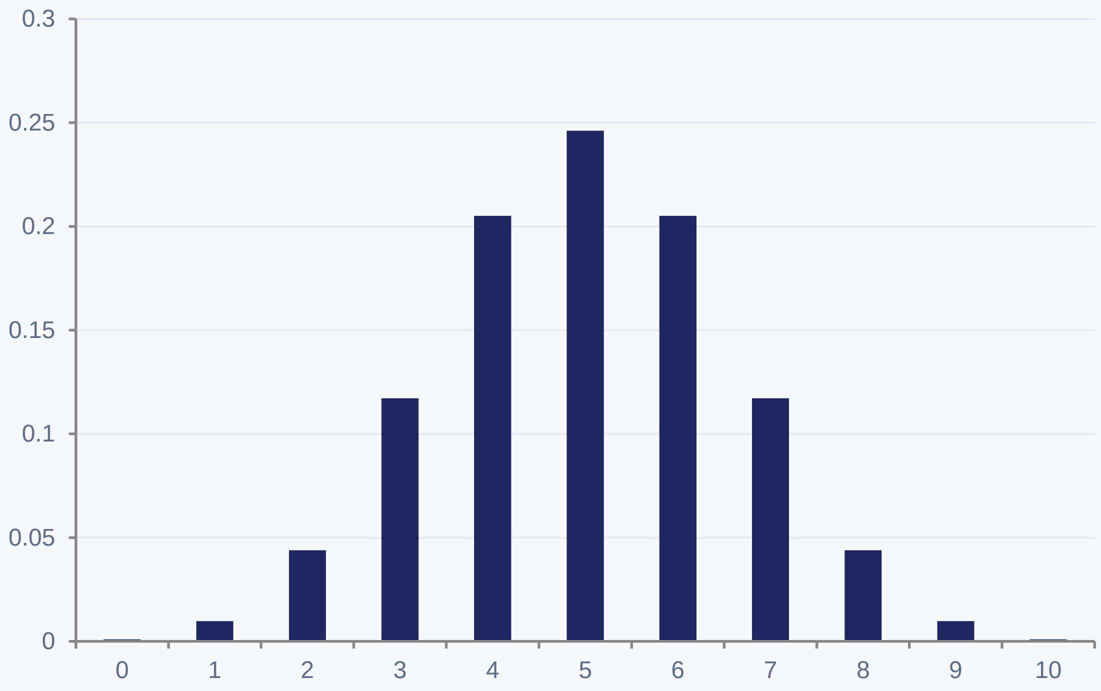

n = 10, p = 0.2 — right-skewed

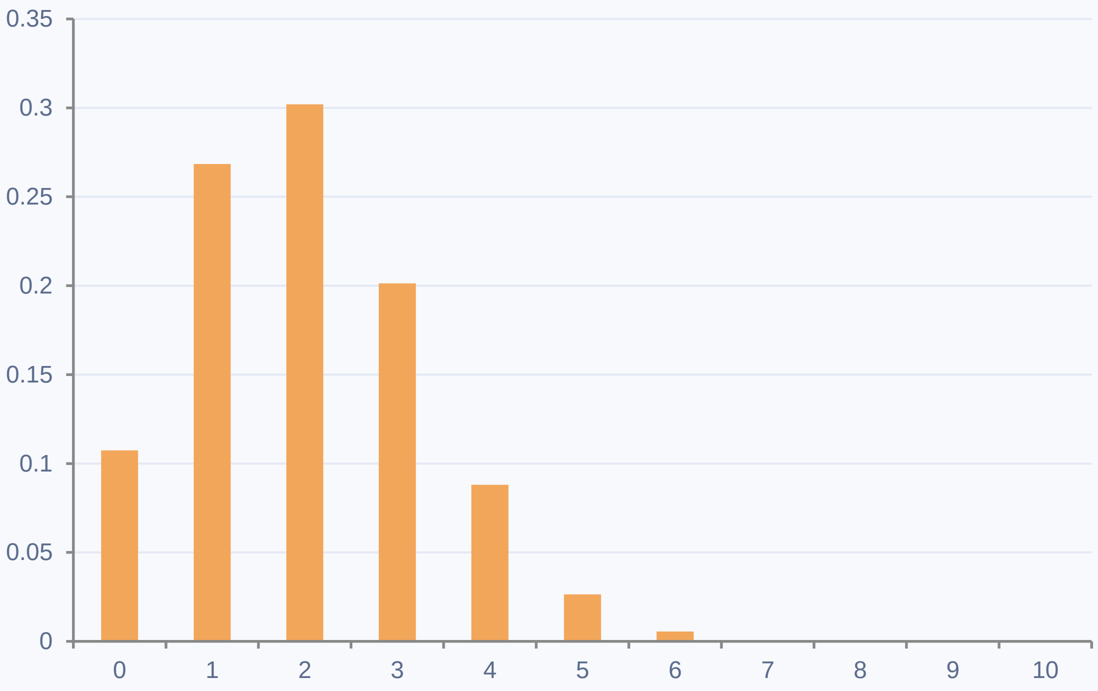

Binomial: Mean, Variance & Properties

EXAMPLE 1

Everything you need to compute with $X \sim \text{Binomial}(n, p)$

Mean

$\mu = E(X) = n \cdot p$

Variance

$\sigma^2 = \text{Var}(X) = n \cdot p \cdot (1-p)$

Std. Deviation

$\sigma = \sqrt{(n \cdot p \cdot (1-p))}$

MGF

$M(t) = (1 - p + p \cdot e^t)^n$

Support

$k \in \{0, 1, 2, \ldots, n\}$

Additivity

Sum of independent Binomials with the same p is Binomial

Worked Example — Packet Transmission

A node sends 50 packets/min; each is corrupted independently with $p = 0.02$.

Expected corrupted packets: $\mu = 50 \times 0.02 = 1.00$ $\sigma^2 = 50 \times 0.02 \times 0.98 = 0.9800$ $\sigma \approx 0.990$ packets

A Second Binomial Example

EXAMPLE 2

Full arithmetic, step by step — quality control: testing IoT devices before shipment

Question: 5 IoT devices are tested. Each independently passes with probability $p = 0.9$. Find P(exactly 4 pass).

Step-by-Step Solution

1. Identify: $n = 5$, $p = 0.9$, $k = 4$
2. Combinations: $C(5,4) = 5$
3. Success term: $p^4 = 0.9^4 = 0.6561$
4. Failure term: $(1-p)^1 = 0.1^1 = 0.1$
5. Multiply all: $5 \times 0.6561 \times 0.1$

$P(X = 4) = 0.3281 \approx 32.8\%$

Full Distribution: X ~ Binomial(5, 0.9)

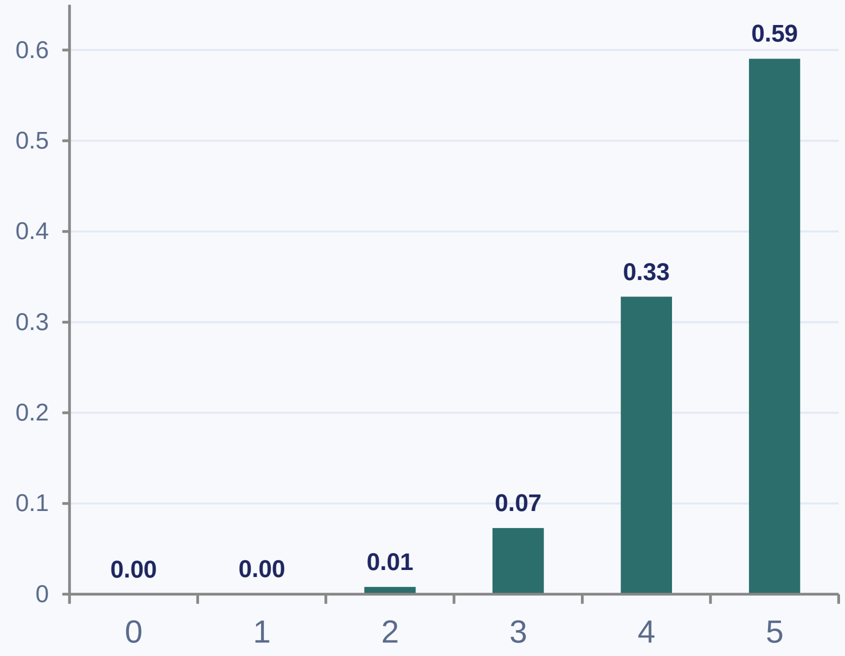

k = 5 (all pass) is most likely since p = 0.9 is high; our k = 4 bar is the second-tallest, matching 0.3281.

From Binomial to Poisson

n → large, p → small, holding n·p = λ fixed

$n = 1000$, $p = 0.01$ ⟶ $n \cdot p = 10 = \lambda$ *The two distributions below are almost indistinguishable.*

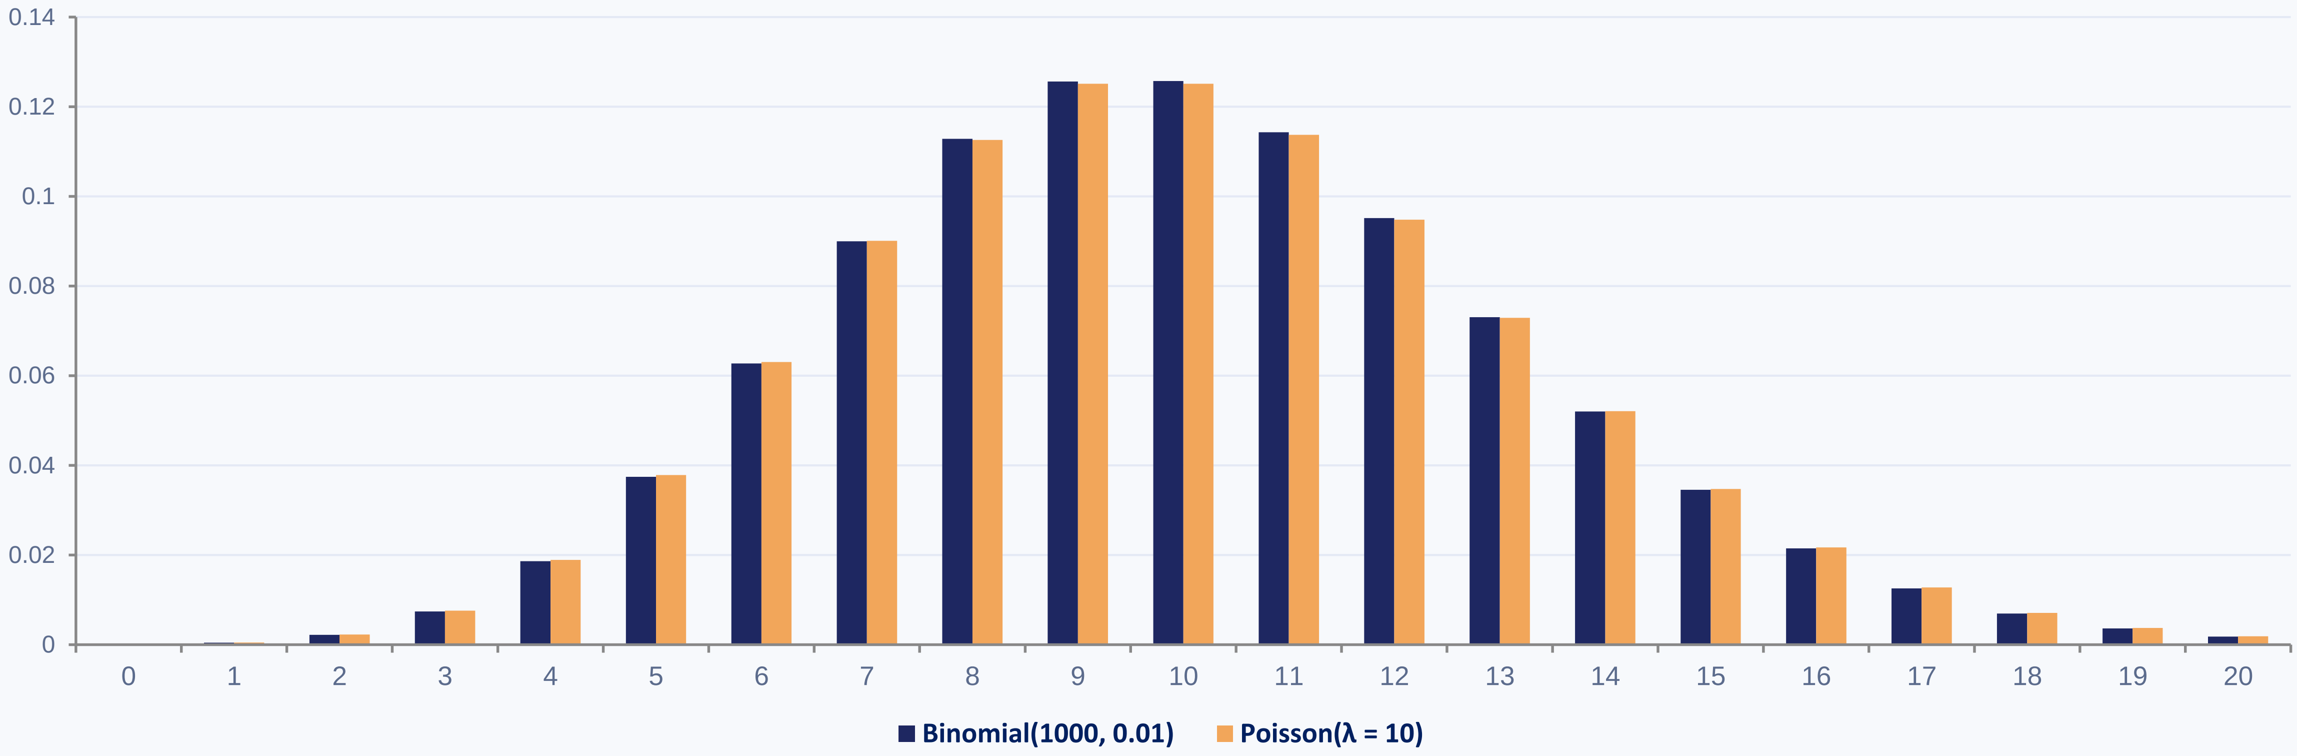

The Poisson Distribution

EXAMPLE 1

Counts of rare events over a fixed interval of time or space

$$P(X = k) = e^{-\lambda} \cdot \lambda^k / k! \qquad k = 0, 1, 2, \ldots$$

Mean: $\mu = E(X) = \lambda$

Variance: $\sigma^2 = \mathrm{Var}(X) = \lambda$ **(mean = variance!)**

MGF: $M(t) = \exp(\lambda \cdot (e^t - 1))$

Support: $k \in \{0, 1, 2, 3, \ldots\}$ **(unbounded)**

Worked Example — Intrusion Alerts

A firewall logs $\lambda = 4$ intrusion attempts/hour on average.

$P(X \geq 2) = 1 - P(0) - P(1) = 0.9084$

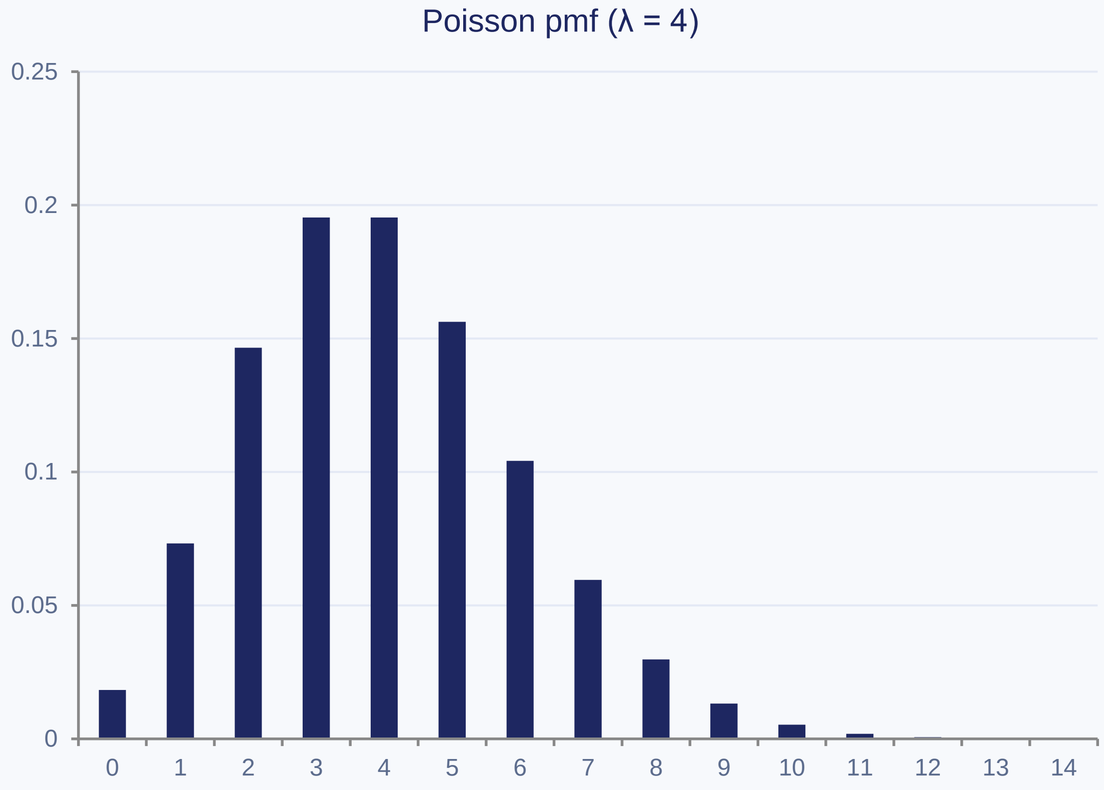

A Second Poisson Example

EXAMPLE 2

Full arithmetic, step by step — blockchain wallet-app logins at a steady average rate

Question: A wallet app receives an average of λ = 3 logins per minute. Find the probability that exactly 5 logins occur in a given minute.

Step-by-Step Solution

1. Identify:
2. $e^{-\nu}$ term:
3. λ^k term:
4. k! term:
5. Combine:

P(X = 5) = 0.1008 ≈ 10.1%

Full Distribution: X ~ Poisson(3)

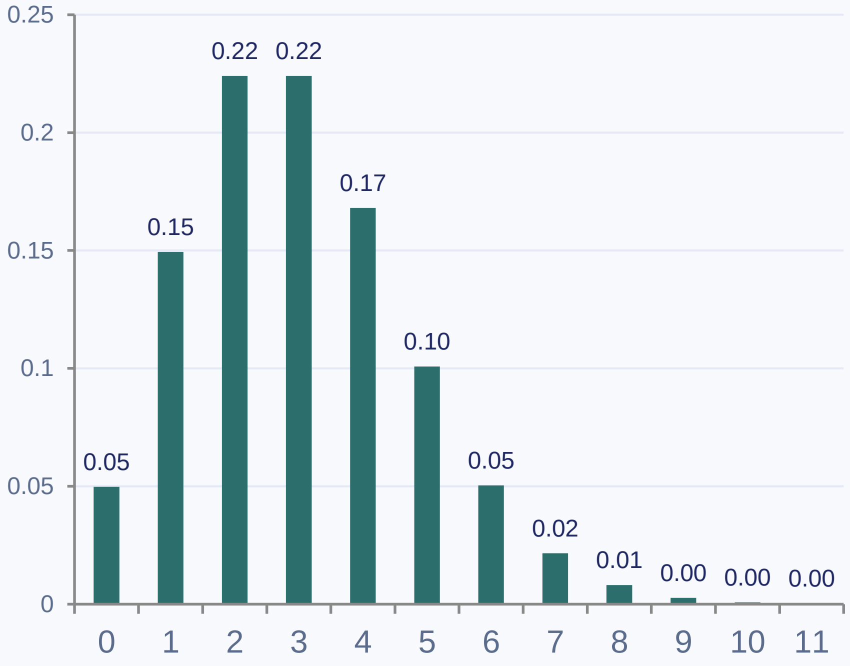

k = 5 sits just past the peak (λ = 3) — matches our computed value.

Continuous Random Variables

Probability becomes area under a curve, not a bar height

Density function f(x)

$f(x) \geq 0$ for all x

Total probability

$\int f(x)\,dx = 1$ (area under the whole curve)

Probability over an interval

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

A single point

$P(X = a) = 0$ — only intervals have positive probability

Example PDF — shaded region = $P(a \leq X \leq b)$

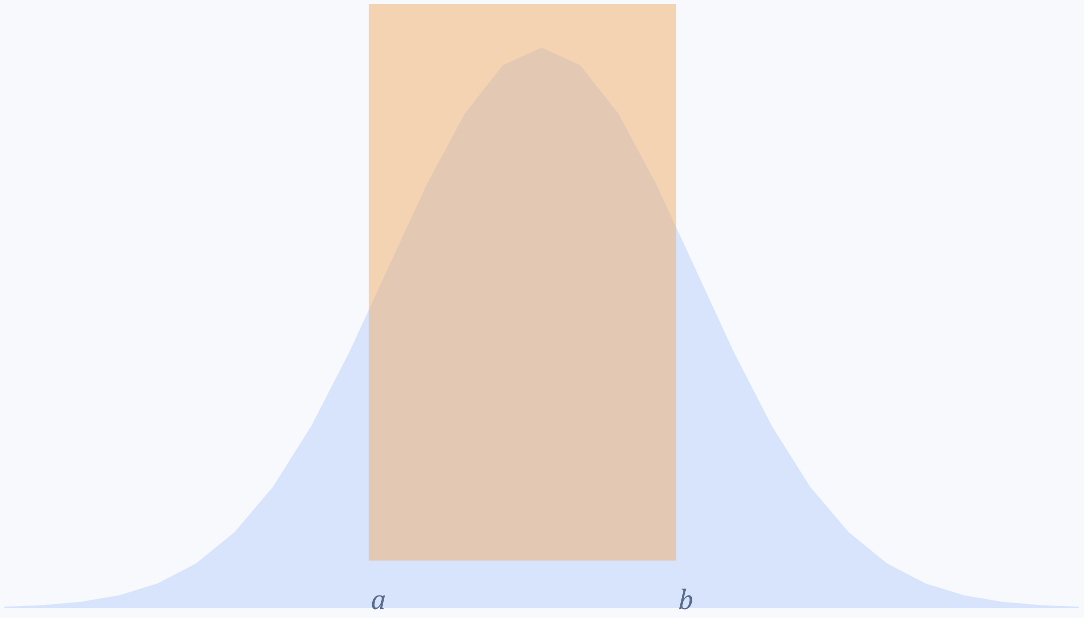

The curve's total area is always exactly 1 — shifting probability to one region always removes it from another.

The Exponential Distribution

Time until the next random event — memoryless waiting

$f(x) = \lambda \cdot e^{-\lambda x}$ for $x \geq 0$ $\mu = 1/\lambda$ $\sigma^2 = 1/\lambda^2$

Sensor battery lifetime, mean = 500 days (λ = 1/500)

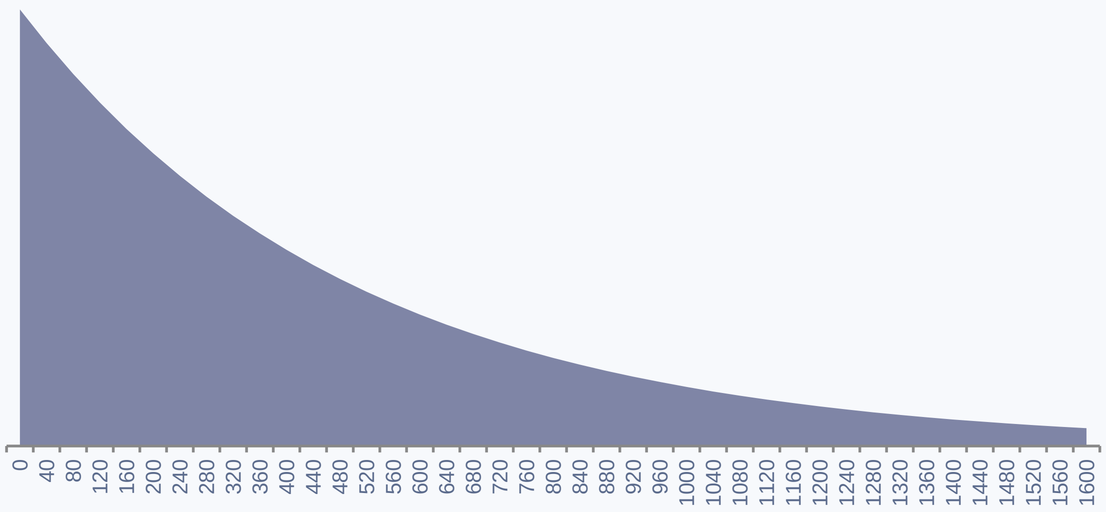

Memoryless Property

$P(X > s + t \mid X > s) = P(X > t)$

A sensor that has already survived 400 days is exactly as likely to survive 200 more days as a brand-new sensor is to survive its first 200 days.

Worked Example

$P(X > 600) = e^{-600/500} = e^{-1.2} = 0.3012$

About a 30.1% chance the sensor survives beyond 600 days.

A Second Exponential Example

EXAMPLE 2

Full arithmetic, step by step — time between phishing emails hitting an inbox

Question: Phishing emails arrive on average every 2 hours. Find the probability the next one arrives within 1 hour.

Step-by-Step Solution

1. Find λ: $\text{mean} = 1/\lambda = 2 \Rightarrow \lambda = 0.5$ / hr
2. We want: $P(X \leq 1) = 1 - e^{-\lambda x}$
3. Plug in: $1 - e^{-(0.5)(1)} = 1 - e^{-0.5}$
4. Compute: $e^{-0.5} = 0.6065$

$P(X \leq 1) = 1 - 0.6065 = 0.3935 \approx 39.3\%$

$f(x) = 0.5 \cdot e^{-0.5x}$ — shaded = $P(X \leq 1)$

Full Distribution: X ~ Exponential(λ = 0.5)

The Normal Distribution

The symmetric bell curve — the most important continuous model

$$f(x) = 1 \,/\, (\sigma\sqrt{2\pi}) \cdot e^{-(x-\mu)^2/2\sigma^2} \qquad -\infty < x < \infty$$

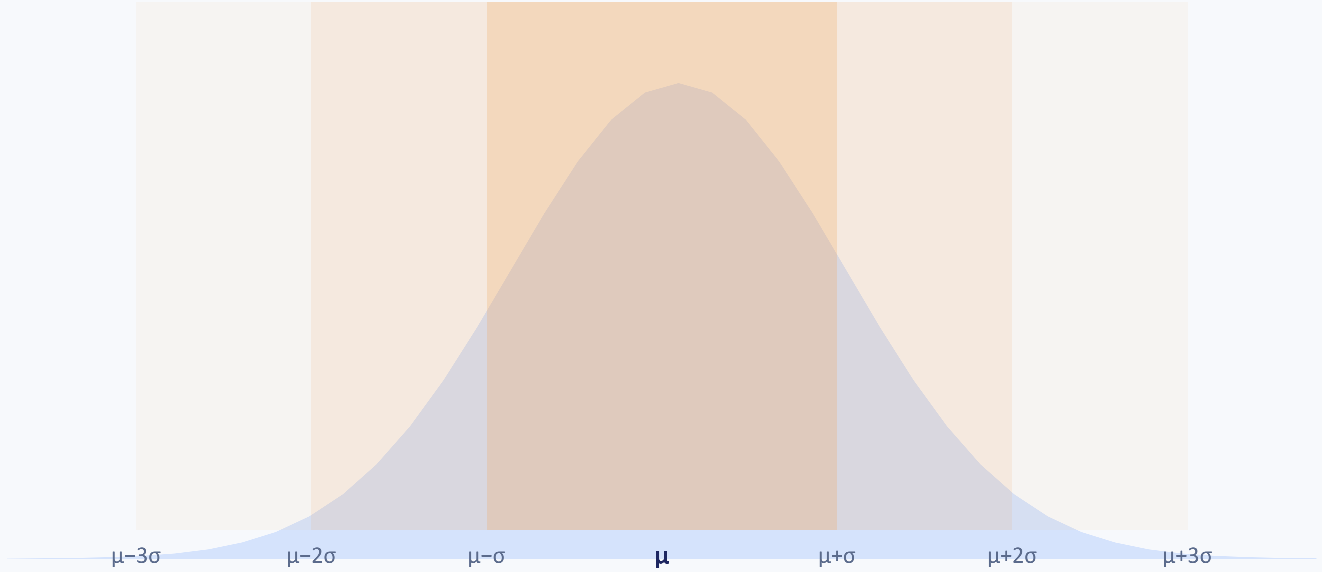

68% / 95% / 99.7% Empirical Rule

Mean = Median = Mode

μ

Symmetry

About the vertical line x = μ

Parameters

μ (location), σ (spread)

MGF

$M(t) = \exp(\mu t + \sigma^2 t^2/2)$

A Worked z-score Example

EXAMPLE

Converting any Normal variable to the standard Normal table

Question: An IoT sensor's readings follow a Normal distribution with $\mu = 50$ and $\sigma = 5$. Find $P(X < 45)$.

Step-by-Step Solution

1. Standardize: $z = (x - \mu) / \sigma$
2. Plug in: $z = (45 - 50) / 5 = -1.00$
3. Look up table: $P(Z < -1.00) = 0.1587$
4. Conclusion: $P(X < 45) = P(Z < -1) = 0.1587$

$P(X < 45) \approx 15.87\%$ (about 1 in 6 readings)

Standard Normal Curve — shaded = $P(Z < -1)$

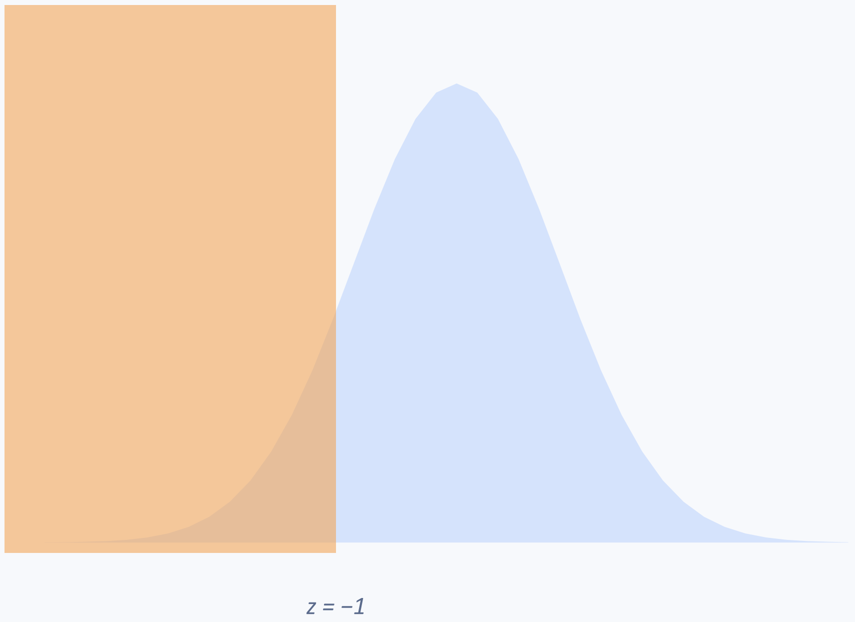

Table gives the area to the left of z — always draw the picture first.

Comparing All Four Distributions

One table to revise the whole module

Distribution	Type	Support	Mean	Variance	Typical Use
Binomial(n,p)	Discrete	0,1,…,n	$n \cdot p$	$n \cdot p(1-p)$	successes in n trials
Poisson(λ)	Discrete	0,1,2,…	λ	λ	rare events per interval
Exponential(λ)	Continuous	$x \geq 0$	$1/\lambda$	$1/\lambda^2$	waiting time to next event
Normal(μ,σ²)	Continuous	$-\infty$ to ∞	μ	σ^2	aggregated / natural measurements

Internet of Things

- Binomial — corrupted packets
- Poisson — sensor failures/hr
- Exponential — battery lifetime
- Normal — calibration error

Cyber Security

- Binomial — flagged logins
- Poisson — intrusion attempts/hr
- Exponential — time between attacks
- Normal — traffic volume noise

Blockchain

- Binomial — valid vs. invalid tx
- Poisson — blocks mined/hour
- Exponential — confirmation wait
- Normal — fee fluctuation

Key Takeaways

1. A random variable turns experiment outcomes into numbers — discrete (countable) or continuous (interval-valued).
2. The Binomial distribution counts successes in n trials; its coefficients literally come from counting coin-toss sequences.
3. The Poisson distribution is the limiting case of Binomial when trials are many and success is rare — mean equals variance.
4. For continuous variables, probability is area under the density curve; the Exponential and Normal distributions are the two workhorses.