

Probability, Random Variables, and Probability Distributions

AP Statistics

Random and Non-Random Patterns


Something is **random** if individual outcomes are uncertain but a regular pattern emerges over **many** repetitions. Short-run results look erratic; long-run **relative frequencies** settle down. This long-run stability is what makes probability useful.

Estimating Probabilities Using Simulation

A **simulation** imitates a chance process using random digits or technology. Steps: state the model, assign digits to outcomes, run many trials, and record the proportion of trials meeting the condition. The resulting proportion **estimates** the probability – more trials give a better estimate.

Introduction to Probability

The **probability** of an event is a number from 0 to 1 giving its long-run relative frequency. The **sample space** is the set of all outcomes. For an event A , the **complement** rule: $P(A^c) = 1 - P(A)$. Probabilities of all outcomes sum to 1.



Probability runs from 0 (impossible) to 1 (certain)




A deck of cards is a classic source of probability: 52 equally likely outcomes make the chances easy to count

Mutually Exclusive Events

Two events are **mutually exclusive** (disjoint) if they cannot both happen. Then the **addition rule** simplifies:

$$P(A \text{ or } B) = P(A) + P(B) \quad (\text{if mutually exclusive}).$$

In general, $P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$ – subtract the overlap so it is not counted twice.




A Venn diagram: the overlap is the intersection of two events

Conditional Probability

The **conditional probability** of A given B is

$$P(A | B) = \frac{P(A \text{ and } B)}{P(B)}.$$

It is the chance of A once you know B happened. Two-way tables make these easy: restrict to the row/column for B , then find A 's share.



On a tree diagram, multiply the probabilities along the branches

Independent Events and Unions of Events

Events are **independent** if knowing one does not change the other's probability: $P(A | B) = P(A)$. Then the **multiplication rule** simplifies:

$$P(A \text{ and } B) = P(A) P(B) \quad (\text{if independent}).$$

Independent is not the same as mutually exclusive – mutually exclusive events with nonzero probability are actually **dependent** (if one happens, the other cannot).

		die 2					
die 1	+	1	2	3	4	5	6
	1	2	3	4	5	6	7
	2	3	4	5	6	7	8
	3	4	5	6	7	8	9
	4	5	6	7	8	9	10
	5	6	7	8	9	10	11
	6	7	8	9	10	11	12

the six highlighted cells all give a total of 7

A sample space diagram lists every equally likely outcome

Random Variables and Probability Distributions

A **random variable** assigns a number to each outcome of a chance process.

A **probability distribution** lists each possible value with its probability (they sum to 1). A distribution can be **discrete** (a table of values) or **continuous** (an area-under-a-curve model like the normal).

Mean and Standard Deviation of Random Variables

The **mean (expected value)** of a discrete random variable is the probability-weighted average:

$$\mu_X = E(X) = \sum x_i P(x_i).$$

The **standard deviation** $\sigma_X = \sqrt{\sum (x_i - \mu_X)^2 P(x_i)}$ measures typical spread from the mean. The expected value is the long-run average outcome, not a value you expect on any single trial.

Worked example. A game pays \$5 with probability 0.2 and costs you \$1 (a -1 outcome) with probability 0.8. The expected value is

$$E(X) = 5(0.2) + (-1)(0.8) = 1 - 0.8 = \$0.20,$$

so over many plays you gain about 20 cents per play on average, even though no single play gives exactly that.

Combining Random Variables

When you add or subtract random variables, **means add**: $\mu_{X \pm Y} = \mu_X \pm \mu_Y$. If X and Y are **independent**, **variances add** (even when subtracting):


$$\sigma_{X \pm Y}^2 = \sigma_X^2 + \sigma_Y^2.$$

Take the square root for the standard deviation. Also, scaling: $\mu_{aX+b} = a\mu_X + b$ and $\sigma_{aX+b} = |a|\sigma_X$.

Introduction to the Binomial Distribution

A **binomial** setting (BINS): a fixed number n of **Independent** trials, each with two outcomes (success/failure) and the **same** success probability p . The random variable X = number of successes. Its probability:

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



The binomial distribution, with mean n times p

Parameters for a Binomial Distribution

For a binomial X with n trials and success probability p :

$$\mu_X = np, \quad \sigma_X = \sqrt{np(1-p)}.$$

Use these for “how many successes do we expect, and how much do they vary” questions.

Worked example. A player makes 70% of free throws. In $n = 10$ shots, the probability of exactly 8 makes is

$$P(X = 8) = \binom{10}{8} (0.7)^8 (0.3)^2 = 45 \times 0.0576 \times 0.09 \approx 0.23,$$

and the expected number of makes is $\mu = np = 10(0.7) = 7$, with $\sigma = \sqrt{10(0.7)(0.3)} \approx 1.45$.

The Geometric Distribution

A **geometric** setting is the same as binomial but with **no fixed** n : you keep trying until the **first success**. The random variable Y = the trial of the first success:

$$P(Y = k) = (1-p)^{k-1} p, \quad \mu_Y = \frac{1}{p}.$$

So the expected number of trials until the first success is $1/p$.

Exam tips

- A probability lies in $[0, 1]$; use the **complement** $(1-P)$ and add mutually exclusive events.
- For independent events multiply; for “and/or” use the general addition and conditional rules.
- **Expected value** = $\sum(\text{value} \times \text{probability})$.
- Recognise **binomial** (fixed n , two outcomes, constant p) and **geometric** settings.
- Draw a tree or table for multi-stage problems and multiply along branches.