PHOTON ACADEMY · IB MATH AI HL CHEATSHEET

Probability Distributions

SL 4.7–4.9 · AHL 4.17

ibtuition.sg · free

FORMULAS YOU MUST KNOW

$0 \le P(X=x) \le 1$ — never negative, never above 1.
Each probability:

$\sum P(X=x) = 1$ — every entry in the table must add to one.
Valid distribution:

Set the sum equal to 1 and solve for the unknown k.
Missing entry:

$E(X) = \sum x\, P(X=x)$ — weight each value by its probability, then add.
Expected value:

$E(\text{gain}) = E(\text{payout}) - \text{cost to play}.$
Net gain (game):

Fair when $E(\text{gain}) = 0$; positive favours the player, negative favours the house.
Fair game:

$X \sim B(n,p)$ — n independent trials, constant p, success/failure each time.
Model:

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

binompdf(n,p,k) returns $P(X=k)$ directly — no need to expand the formula by hand.
On the GDC:

$P(X \le k) =$ binomcdf(n,p,k) — the calculator sums from 0 up to k.
At most k:

$P(X \ge k) = 1 - P(X \le k-1) = 1-$ binomcdf$(n,p,k-1)$.
At least k:

$P(a \le X \le b) =$ binomcdf$(n,p,b)-$ binomcdf$(n,p,a-1)$.
Between:

KEY METHODS

- Add every probability you are given, then subtract from 1 to fill a gap — that missing value is one of the most commonly dropped marks on the whole topic.
- Find any missing k from $\sum P(X=x)=1$ before you touch $E(X)$ — the two steps almost always appear in the same question, and $E(X)$ is wrong if k is.
- $E(X)$ need not be an attainable value — a fair die has mean 3.5. It is a balance point of the distribution, not a prediction of any single outcome.
- "Exactly", "precisely" or "equal to" k successes $\Rightarrow$ binompdf. Read the wording before you type a single number.

EXAM TRAPS TO AVOID

- Your answer for k must itself be a valid probability. If solving gives $k>1$ or $k<0$, re-check the algebra — a probability can never sit outside $[0,1]$.
- $E(\text{payout})$ on its own is not the verdict on a game — always subtract the stake first. An expected payout smaller than the ticket price is still a losing game.
- The binomial needs a fixed n, a constant p and independent trials. Sampling without replacement changes p each draw, so the binomial no longer applies.
- For "at least k" you subtract $P(X \le k-1)$, not $P(X \le k)$ — "at least 2" excludes only 0 and 1. The off-by-one $k-1$ is the classic slip here.

- Translate the words first: "fewer than k" $= P(X \leq k-1)$; "more than k" $= 1 - P(X \leq k)$. A quick sketch of which counts are included stops every mix-up.
- Compute $np(1-p)$ in full first, then take the square root once at the very end for σ. Rounding the variance before rooting throws away accuracy.
- Sketch the bell and shade the region before touching the GDC. The picture tells you which bounds to enter and instantly settles "greater-than vs less-than" confusion.

- The variance carries the extra $(1-p)$ factor; the mean does not. Dropping $(1-p)$ — writing the spread as np — is the single most common error in this sub-topic.
- Enter the standard deviation σ, not the variance σ^2. If the question quotes a variance, take its square root first — a "variance of 16" means you type $\sigma = 4$.
- invNorm always expects the area to the LEFT of x. For a "top 5%" cut-off, enter 0.95; for the "bottom 5%", enter 0.05. Feeding the wrong tail is the number-one inverse-normal error.