BP701T · B.Pharm Semester VII · Theory Notes

Biostatistics and Research Methodology

Unit II — Probability and Probability Distributions (9 hours)

Unit II (9 hours)

Pills & Python Notes

From-zero teaching standard: every technical term is defined before use.

Syllabus coverage: Unit II (9 hours).

PILLS & PYTHON

Pharmacy • Data • AI • Healthcare

Instagram @pillsandpython · pillsandpython.com · YouTube @PillsAndPython

CHAPTER 4: Probability — Classical and Statistical, Union, Intersection, Conditional and Marginal

Learning Objectives

After studying this chapter, you will be able to:

- Define probability in simple words and explain what a value between 0 and 1 means.
- Distinguish classical probability from statistical (relative-frequency) probability and say when each one is used.
- Identify the experiment, outcomes, sample space and events in a simple situation.
- Calculate P(A ∪ B) using the addition rule, for mutually exclusive and non-mutually exclusive events.
- Calculate P(A ∩ B) for independent and dependent events.
- Apply the complement rule, especially for "at least one" type of problems.
- Compute conditional probability P(A | B) and explain its meaning in words.
- Read marginal probabilities from a two-way (joint) frequency table.
- Solve pharmacy-style probability problems such as batch acceptance chances and defect rates.

Definitions

- **Experiment (trial):** Any action whose result is not known in advance. Example: checking whether a tablet passes the hardness test.
- **Outcome:** One single possible result of an experiment. Example: "the tablet passes".
- **Sample space (S):** The set of ALL possible outcomes. For tossing a coin once, S = {Head, Tail}.
- **Event (E):** Any collection of outcomes we are interested in. Example: on a die, "getting an even number" = {2, 4, 6}.
- **Probability P(E):** A number from 0 to 1 measuring how likely event E is. 0 = impossible, 1 = certain.
- **Classical probability:** P(E) = (number of favourable outcomes) ÷ (total number of equally likely outcomes).
- **Statistical (empirical) probability:** P(E) = (number of times E happened in many trials) ÷ (total number of trials).
- **Union (A ∪ B):** The event "A or B (or both) happens".
- **Intersection (A ∩ B):** The event "both A and B happen".
- **Complement (A′):** The event "A does not happen". Rule: P(A′) = 1 − P(A).
- **Mutually exclusive events:** Events that cannot happen at the same time. P(A ∩ B) = 0.

- **Independent events:** The happening of one does not change the probability of the other. Then P(A ∩ B) = P(A) × P(B).
- **Dependent events:** The happening of one changes the probability of the other.
- **Conditional probability P(A | B):** The probability of A, given that B has already happened.
- **Marginal probability:** The probability of one event alone, found by adding across a row or down a column of a joint table.

Terminology

Term	Simple meaning
Trial	One repetition of an experiment (one toss, one tablet tested)
Outcome	A single result of a trial
Sample space	The full list of all possible outcomes
Event	The particular result(s) we care about
Favourable outcome	An outcome that makes our event happen
Equally likely	Every outcome has the same chance (like a fair coin)
Relative frequency	Fraction of trials in which the event occurred
Mutually exclusive	Two events that cannot occur together
Independent	One event does not affect the other's chance
Dependent	One event changes the other's chance
Joint probability	Probability that two things happen together, P(A ∩ B)
Exhaustive events	Events that together cover the whole sample space

Simple Explanations

4.1 What probability really means

Technical concept: Probability is a number between 0 and 1 assigned to an event — 0 means impossible, 1 means certain, 0.5 means fifty-fifty.

Simple explanation: Think of it as a "chance meter" from 0% to 100%. A forecast of "80% chance of rain" means the probability of rain is 0.8.

Example: A fair coin has two equally likely outcomes, so P(Head) = 1 ÷ 2 = 0.5.

Diagram:

```
0 ---------- 0.5 ---------- 1
|             |              |
Impossible  Fifty-fifty    Certain
```

Key takeaway: Every probability lies between 0 and 1; closer to 1 means more likely.

4.2 Classical probability

Technical concept: When every outcome is equally likely, P(E) = (favourable outcomes) ÷ (total equally likely outcomes).

Simple explanation: Count how many outcomes make you win, divide by how many outcomes are possible. Works for "fair" situations — a fair die, a well-shuffled deck.

Example: Rolling a fair die, P(getting a 3) = 1 ÷ 6 ≈ 0.1667.

Diagram:

```
Sample space: {1, 2, 3, 4, 5, 6} (6 equally likely)
Event "3":    {3}                 (1 favourable)
P = 1 / 6 ≈ 0.167
```

Key takeaway: Classical probability = count and divide; it needs equally likely outcomes.

4.3 Statistical (empirical) probability

Technical concept: P(E) ≈ (times E occurred) ÷ (total trials). As trials grow, this fraction settles toward the true probability.

Simple explanation: When outcomes are not equally likely, learn the chance from experience: test many tablets, count failures, divide. More tests → more trustworthy number.

Example: A QC inspector checks 1,000 capsules and finds 12 defective (illustrative): P(defective) = 12 ÷ 1,000 = 0.012.

Diagram:

```
10 tablets:    1 defect  → P ≈ 0.10
100 tablets:   9 defects → P ≈ 0.09
1,000 tablets: 12 defects → P = 0.012   (settles down)
```

Key takeaway: Statistical probability = learn from data; it is how real QC estimates defect rates.

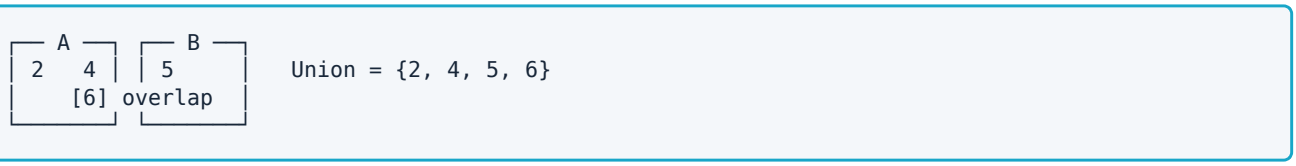
4.4 Union of events — the addition rule

Technical concept: A ∪ B means "A or B (or both)". P(A ∪ B) = P(A) + P(B) − P(A ∩ B); for mutually exclusive events the last term is 0.

Simple explanation: "OR" means add — but outcomes belonging to both events get counted twice, so subtract the overlap once.

Example: Roll a die. A = even = {2,4,6}, P(A) = 3/6. B = greater than 4 = {5,6}, P(B) = 2/6. Overlap = {6}, P = 1/6. P(A ∪ B) = 3/6 + 2/6 − 1/6 = 4/6 ≈ 0.6667. Check: {2,4,5,6} is 4 of 6. ✓

Diagram:



Key takeaway: For "A or B", add the probabilities and subtract the double-counted overlap.

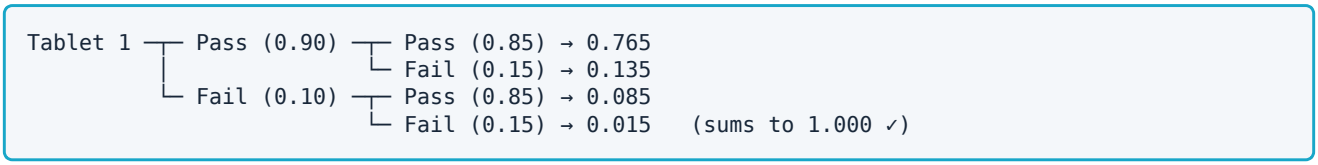
4.5 Intersection of events — the multiplication rule

Technical concept: A ∩ B means "both happen". If independent: P(A ∩ B) = P(A) × P(B). If dependent: P(A ∩ B) = P(A) × P(B | A).

Simple explanation: "AND" means multiply — but only when the events don't influence each other. If the first changes the second's chance, use the changed (conditional) chance.

Example: Two tablets tested independently (illustrative): P(first passes) = 0.90, P(second passes) = 0.85. P(both pass) = 0.90 × 0.85 = 0.765 (90 × 85 = 7,650).

Diagram:



Key takeaway: "A and B" with independent events → multiply; with dependent events → use the conditional probability.

4.6 Complement — the "not" rule

Technical concept: P(A′) = 1 − P(A). Shortcut for "at least one": P(at least one) = 1 − P(none).

Simple explanation: "At least one defective" is hard to count directly (1 or 2 or 3 ...), but "zero defective" is easy — then subtract from 1.

Example: P(a tablet is non-defective) = 0.95, independent (illustrative). P(at least one defective in 10) = $1 - 0.95^{10}$. Steps: $0.95^2 = 0.9025$; $0.95^4 = 0.81450625$; $0.95^8 \approx 0.66342043$; $0.95^{10} \approx 0.59873694$. Answer: $1 - 0.59873694 = 0.40126306 \approx 0.4013$.

Diagram:



Key takeaway: See "at least one" → think complement: 1 − P(none).

4.7 Conditional probability

Technical concept: P(A | B) = P(A ∩ B) ÷ P(B). The condition B shrinks the sample space to outcomes where B happened.

Simple explanation: New information changes the chances. "Given it came from Machine A" — first throw away every non-A tablet, then count inside what remains.

Example: Machine A makes 400 tablets/day with 8 defectives (illustrative). P(defective | Machine A) = 8 ÷ 400 = 0.02.

Diagram:

```
All 1,000 tablets ── "given Machine A" ──→ only its 400 remain
    ┌─ 8 defective → P = 8/400 = 0.02
    └─ 392 good        ("|" = the filter)
```

Key takeaway: P(A | B): narrow the world to B first — B is the denominator.

4.8 Marginal probability

Technical concept: In a two-way table, the marginal probability of one event = its row/column total ÷ grand total ("summing out" the other variable).

Simple explanation: The table shows combinations (day + pass, night + fail ...). For one variable alone — say overall pass rate — add across the other variable and divide by the grand total.

Example: 100 batches by shift and result (illustrative): Day: 50 pass, 5 fail; Night: 35 pass, 10 fail. Marginal P(Pass) = 85 ÷ 100 = 0.85; P(Night) = 45 ÷ 100 = 0.45. Contrast — conditional P(Fail | Night) = 10 ÷ 45 = 2/9 ≈ 0.2222 (row total, not grand total).

Diagram:

```
            Pass  Fail → row total
Day          50     5   → 55
Night        35    10   → 45
             ↓     ↓
            85    15   (column totals); grand total 100
P(Pass) = 85/100 = 0.85
```

Key takeaway: Marginal = margin total ÷ grand total (ignores the other variable); conditional fixes the other variable to one value.

Classifications

Two ways of finding probability

1. **Classical probability** — counted from equally likely outcomes (fair coins, dice, cards).
2. **Statistical (empirical) probability** — learned from repeated trials or past data (QC defect rates, stability data, prescription audits).

Two ways events relate to each other

1. **By overlap:** mutually exclusive (cannot happen together; overlap = 0) vs non-mutually exclusive (can overlap; subtract it in the addition rule).
2. **By influence:** independent (one doesn't change the other's chance; multiply directly) vs dependent (one changes the other's chance; use the conditional probability).

Processes

How to solve any probability word problem (5 steps)

```
Step 1: READ — underline the key event ("at least one defective").
Step 2: IDENTIFY — experiment, sample space, given numbers.
Step 3: CHOOSE by keyword: "or" → addition; "and" → multiplication;
        "not"/"at least one" → complement (1 minus); "given that" → conditional.
Step 4: CHECK — can they overlap? (union) Does one affect the other? (intersection)
Step 5: CALCULATE, then SANITY-CHECK: between 0 and 1? Makes sense?
```

How to read marginal probabilities from a 2 × 2 table

```
Step 1: Write the table with row, column and grand totals.
Step 2: Take the margin total of the event you want ÷ grand total.
Step 3: For a conditional ("Pass GIVEN Day"), divide by the CONDITION's
        row/column total instead.
```

Diagrams and Flowcharts

Which rule do I use?

```
            Does it say "given that"? ──YES──→ CONDITIONAL P(A|B)
                                NO
            Does it say "or"? ──YES──→ UNION: P(A)+P(B) − overlap
                                NO
            Does it say "and"? ──YES──→ INTERSECTION: multiply
                                │         (independent? direct : conditional)
                                NO
            "not" / "at least one"? ──YES──→ COMPLEMENT: 1 − P(opposite)
```

Venn: union, intersection, complement

```
UNION: shade either circle (overlap subtracted once)
INTERSECTION: shade only the overlap lens
COMPLEMENT of A: shade everything OUTSIDE circle A
```

Infographic Concept

INFOGRAPHIC CONCEPT: A single-page "Probability Rule Map". Top: the 0-to-1 chance meter (Impossible → Certain). Left: "OR means ADD" with overlapping Venn circles, overlap subtracted once. Right: "AND means MULTIPLY" with a two-branch tree. Bottom-left: "NOT means 1 MINUS". Bottom-right: "GIVEN means SHRINK FIRST" as a magnifying glass on one table row. Centre: the 5-step checklist. Colour code: blue union, green intersection, orange complement, purple conditional; one mini-example per panel.

Tables

Classical vs statistical probability

Point	Classical	Statistical
Basis	Counting equally likely outcomes	Repeated trials / observed data
Formula	Favourable ÷ total equally likely	Occurrences ÷ total trials
Example	P(Head) = 1/2	P(defective) ≈ 12/1000 (illustrative)
Pharmacy use	Random sampling designs	QC defect rates, stability data

Probability rules at a glance

Situation	Rule	Formula
A or B (any)	Addition (general)	P(A ∪ B) = P(A) + P(B) − P(A ∩ B)
A or B (mutually exclusive)	Addition (simple)	P(A ∪ B) = P(A) + P(B)
A and B (independent)	Multiplication (simple)	P(A ∩ B) = P(A) × P(B)
A and B (dependent)	Multiplication (general)	P(A ∩ B) = P(A) × P(B \
Not A / at least one	Complement	P(A′) = 1 − P(A); 1 − P(none)
A given B	Conditional	P(A \

Worked Examples and Calculation Boxes

CALCULATION BOX 1 — Union

Question: Roll a fair die once. Find P(even OR greater than 4).

```
A = {2,4,6}, P(A) = 3/6;  B = {5,6}, P(B) = 2/6;  A∩B = {6}, P = 1/6
P(A ∪ B) = 3/6 + 2/6 − 1/6 = 4/6 = 2/3 ≈ 0.6667
Check: {2,4,5,6} → 4 of 6 ✓
FINAL ANSWER: 2/3 ≈ 0.6667
```

CALCULATION BOX 2 — Intersection (independent)

Question: Two tablets tested independently (illustrative): P(first passes) = 0.90, P(second passes) = 0.85. Find P(both pass).

```
P(both) = 0.90 × 0.85 = 0.7650   (90 × 85 = 7,650)
FINAL ANSWER: 0.765 (76.5%)
```

CALCULATION BOX 3 — Complement ("at least one")

Question: P(tablet non-defective) = 0.95, independent (illustrative). In 10 tablets, find P(at least one defective).

```
P(at least one) = 1 − 0.95¹⁰
0.95² = 0.9025 → 0.95⁴ = 0.81450625 → 0.95⁸ ≈ 0.66342043
→ 0.95¹⁰ ≈ 0.59873694
P = 1 − 0.59873694 = 0.40126306 ≈ 0.4013
FINAL ANSWER: ≈ 0.4013 (40.1%)
```

CALCULATION BOX 4 — Conditional probability

Question: Machine A: 400 tablets/day, 8 defectives (illustrative). Find P(defective | Machine A).

```
P = 8 ÷ 400 = 2 ÷ 100 = 0.02
FINAL ANSWER: 0.02 (2%)
```

CALCULATION BOX 5 — Marginal vs conditional from a table

Question: Shift × result (illustrative): Day: 50 pass, 5 fail; Night: 35 pass, 10 fail (total 100). Find (a) marginal P(Pass), (b) marginal P(Night), (c) P(Fail | Night).

```
(a) 85 ÷ 100 = 0.85        (b) 45 ÷ 100 = 0.45
(c) 10 ÷ 45 = 2/9 ≈ 0.2222   ← row total, not grand total
FINAL ANSWERS: (a) 0.85  (b) 0.45  (c) ≈ 0.2222
```

Analogies

1 **Cricket prediction.** "70% chance of winning" is a probability of 0.7 — confidence, not certainty. Probability 1 would be "already won".

2 **Thali (union).** A thali with roti OR rice: some have both — don't count those twice. That's subtracting the overlap.

3 **Two locks (intersection).** Both keys must turn: multiply the chances for independent events.

4 **Filtered gallery (conditional).** "Given Diwali" — filter to Diwali photos first, then look inside. "|" is the filter.

5 **Station master (marginal).** Counting all boarding passengers, ignoring which train — summing across the other variable.

Pharmacy Relevance

1 **Batch acceptance.** The chance a batch with a known defect rate passes sampling — the basis of acceptance sampling in QC.

2 **Defect rates.** Defectives ÷ tested is exactly how QC estimates the statistical probability of a defect.

3 **"At least one" in sampling.** The complement rule powers routine in-process checks.

4 **Root-cause analysis.** "Given the defect is capping, what's the probability it came from compression?" — conditional probability is the language of deviations.

5 **Dissolution + uniformity.** The probability a tablet passes both (intersection) decides whether a stability time-point passes.

India Context

- Under Schedule M, Indian QC labs estimate defect probabilities from testing data and compute the chance sampled units meet specifications.
- The Indian Pharmacopoeia's acceptance criteria (e.g. uniformity of weight) make "probability a random tablet complies" a daily working idea.

- India's generic-manufacturing scale means sampling-probability decisions directly affect affordable-medicine supply.
- Hospital prescription audits use relative frequencies (statistical probability applied to drug utilisation).
- Pharmacovigilance-style thinking is conditional: given an adverse event report, how likely is a link to a medicine class? (Concept only; no real data claimed.)

Key Points

- Probability ∈ [0, 1]; 0 = impossible, 1 = certain.
- Classical = favourable ÷ equally likely (fair situations). Statistical = occurrences ÷ trials (real data; more trials = steadier).
- Union "or": add, subtract overlap (0 if mutually exclusive).
- Intersection "and": multiply if independent; use conditional if dependent.
- Complement: P(not A) = 1 − P(A); "at least one" = 1 − P(none).
- Conditional: shrink to B first — B is the denominator. Marginal: margin total ÷ grand total.
- Sanity-check: answer between 0 and 1 and sensible.

Common Confusions

1 **Maths "or" vs spoken "or".** P(A ∪ B) includes both happening — subtract the overlap unless mutually exclusive.

2 **P(A ∩ B) vs P(A | B).** "Both happen" (fraction of everything) vs "A among B-cases" (fraction of B only).

3 **Marginal vs conditional.** 85/100 (grand total) vs 10/45 (row total) — dividing by the wrong total is the classic slip.

4 **Independent vs mutually exclusive.** Different ideas: mutually exclusive = can't co-occur; independent = don't influence each other.

5 **Small samples mislead.** A defect rate from 10 tablets is shaky; from 1,000 it's trustworthy.

Common Mistakes

1 Answers outside [0, 1] (e.g. 0.6 + 0.7 = 1.3 without subtracting overlap).

2 Swapping the rules: addition for "and", multiplication for "or". (OR → add, AND → multiply.)

3 Multiplying directly for dependent events (e.g. two draws without replacement from one small batch).

4 Adding P(1) + P(2) + … for "at least one" instead of 1 − P(none).

5 "Given" with the wrong denominator (factory total instead of Machine A's total).

6 Assuming classical probability where outcomes aren't equally likely.

Mnemonics

- **"OR adds, AND multiplies"** — the two-second rule chooser.
- **"GIVEN shrinks first"** — shrink the world to B, then count.
- **"1 minus is your friend"** — for "not" and "at least one".
- **"Mutually exclusive? Just ADD"** — overlap is zero, so no subtraction.
- **"Margins hold the marginals"** — marginals live in the table margins.

Exam Summary

- Classical = favourable ÷ equally likely; statistical = occurrences ÷ trials.
- P(A ∪ B) = P(A) + P(B) − P(A ∩ B) (last term 0 if mutually exclusive).
- P(A ∩ B) = P(A) × P(B) if independent; P(A) × P(B | A) if dependent.
- P(A′) = 1 − P(A); P(at least one) = 1 − P(none).
- P(A | B) = P(A ∩ B) ÷ P(B). Marginal = margin total ÷ grand total.
- Typical 2-marker: classical vs statistical. Typical 5-marker: union/intersection/complement numerical or a 2 × 2 table (marginal + conditional).

One-Minute Revision

- [0, 1]; classical = count ÷ fair total; statistical = test ÷ trials.
- OR → add (− overlap). AND → multiply (independent) / conditional (dependent).
- NOT / at least one → 1 minus. GIVEN → B becomes the denominator.
- Marginal → margin ÷ grand total. Check overlap and independence before choosing the rule.

Questions

Short-Answer Questions

1 Define classical probability with one example.

2 Define statistical probability and name one pharmacy use.

3 State the addition rule for non-mutually exclusive events.

4 Write the conditional probability formula; explain "|" in words.

5 What is a marginal probability, and where does the name come from?

6 If P(A) = 0.4, find P(not A).

7 Give one pharmacy example each of independent and dependent events.

8 Why is 1 − P(none) the shortcut for "at least one"?

Long-Answer Questions

1 Distinguish classical and statistical probability with examples. Why does statistical probability improve with more trials?

2 State and explain (Venn argument) P(A ∪ B) = P(A) + P(B) − P(A ∩ B); show the mutually exclusive simplification.

3 Explain independent vs dependent events; derive the multiplication rule for each with a pharmacy example.

4 Using a 2 × 2 table (e.g. shift vs pass/fail), explain joint, marginal and conditional probability with full calculations.

Multiple Choice Questions

1 Probability lies between: a) −1 and 1 b) 0 and 1 c) 0 and 100 d) 1 and 10

2 Classical probability = a) occurrences ÷ trials b) favourable ÷ total equally likely c) 1 − P(A) d) P(A) × P(B)

3 If A, B mutually exclusive, P(A ∩ B) = a) P(A)×P(B) b) 0 c) 1 d) P(A)+P(B)

4 If independent, P(A ∩ B) = a) P(A)+P(B) b) P(A)−P(B) c) P(A)×P(B) d) 1−P(A)

5 P(A′) = 1 − P(A) is the: a) addition rule b) multiplication rule c) complement rule d) conditional rule

6 "|" in P(A | B) reads: a) "A or B" b) "A and B" c) "A given B" d) "A not B"

7 Marginal probability uses the: a) condition's row total b) grand total c) smallest cell d) column total only

8 "At least one defective" is easiest via: a) addition b) complement c) conditional d) classical counting

9 P(A)=0.3, P(B)=0.5, mutually exclusive → P(A ∪ B) = a) 0.15 b) 0.8 c) 0.2 d) 1.0

10 Two draws without replacement from one small batch: a) independent b) dependent c) mutually exclusive d) exhaustive

Answer key: 1-b, 2-b, 3-b, 4-c, 5-c, 6-c, 7-b, 8-b, 9-b, 10-b

Viva Questions

1 What does P = 0.5 mean in plain words?

2 When is statistical probability preferred over classical?

3 Why subtract P(A ∩ B) in the addition rule?

4 One-line test for independence?

5 What changes first when you see "given that"?

Application-Based Questions

1 Illustrative defect rate 2%. Find P(at least one defective) in 20 tablets via complement. Show steps.

2 Two lots tested independently, each P(pass) = 0.98. Find P(both pass) and P(at least one fails).

3 From the shift table (Day: 50 pass/5 fail; Night: 35 pass/10 fail), a failed batch is picked. Find P(it came from night shift). (Hint: narrow to failed batches first.)

References

- Bolton (Pharmaceutical Statistics)
- Daniel (Biostatistics)
- Montgomery (Design and Analysis of Experiments)

CHAPTER 5: Probability Distributions — PMF, PDF, Normal, Standard Normal and Z, Binomial, Poisson

Learning Objectives

After studying this chapter, you will be able to:

- Explain what a probability distribution is and what a random variable means.
- Distinguish discrete and continuous distributions with pharmacy examples.
- Define the PMF and use it for a discrete variable.
- Define the PDF and explain why "area under the curve" gives probability for continuous variables.
- Describe the normal distribution: characteristics, parameters μ and σ, and its PDF equation.
- Give pharmaceutical examples of data modelled by the normal distribution.
- Convert a normal value to a Z score, read the standard normal table, and solve "between / above / below" problems.
- State the binomial conditions, write its PMF, and solve binomial problems (including mean and variance).
- State when the Poisson distribution applies, write its PMF, and solve Poisson problems.
- Choose correctly between normal, binomial and Poisson for a pharmacy situation.

Definitions

- **Random variable (X):** A variable whose value is the numerical outcome of a random experiment. Example: X = number of defective tablets in a sample of 10.
- **Probability distribution:** A rule telling us the probability of every possible value of a random variable — the complete "chance map" of X.
- **Discrete random variable:** Takes only separate, countable values (0, 1, 2, ...). Example: number of defective capsules in a strip.
- **Continuous random variable:** Can take any value in an interval, decimals included. Example: exact tablet weight in mg.

- **Probability mass function (PMF):** For a discrete variable, p(x) = P(X = x) — the probability X equals exactly x.
- **Probability density function (PDF):** For a continuous variable, f(x) whose AREA between two points gives the probability of falling between them. f(x) itself is not a probability.
- **Normal distribution:** A continuous, symmetric, bell-shaped distribution; many measurement data follow it approximately.
- **Parameters of the normal distribution:** μ (mean) fixes the centre; σ (standard deviation) fixes the width/spread.
- **Standard normal distribution:** The normal distribution with μ = 0, σ = 1. Its values are Z scores.
- **Z transformation:** Z = (X − μ) ÷ σ — converts any normal value into "standard deviations from the mean".
- **Binomial distribution:** Probability of exactly x "successes" in n fixed independent trials, each with success probability p.
- **Poisson distribution:** Probability of exactly x events in a fixed interval of time/space, at a known average rate λ.

Terminology

Term	Simple meaning
Random variable	A number coming out of a chance experiment
Distribution	The full "value → its chance" map
Discrete	Countable only: 0, 1, 2, ...
Continuous	Any value in a range, decimals allowed
PMF	P(X equals exactly x), for discrete data
PDF	Curve height = density; area = probability, for continuous data
μ (mu)	Population mean; centre of the bell
σ (sigma)	Population SD; spread of the bell
Z score	Standard deviations a value lies from the mean
p (binomial)	Success chance on one trial
n (binomial)	Fixed number of trials
λ (lambda)	Average events per interval (Poisson)
e	≈ 2.71828, used in the Poisson formula

Simple Explanations

5.1 Meaning of a probability distribution

Technical concept: A probability distribution assigns a probability to every possible value (or range) of a random variable; all probabilities total 1.

Simple explanation: It is the complete "chance map": instead of one question ("chance of exactly 2 defectives?"), it answers for 0, 1, 2, 3, … all at once.

Example: X = heads in 2 tosses: P(0) = 0.25, P(1) = 0.50, P(2) = 0.25 (sum = 1.00).

Diagram:

```
X:      0      1      2
P(X):  0.25   0.50   0.25   → adds to 1.00 (always)
```

Key takeaway: A distribution lists every value with its chance; chances total 1.

5.2 Discrete distribution and the PMF

Technical concept: A discrete variable takes countable values; its PMF p(x) = P(X = x) is shown as bars whose HEIGHT is the probability.

Simple explanation: "Discrete" = separate lumps — 0, 1, 2 defective tablets, never 1.5. The PMF table says "chance of exactly x is …", readable straight off the bar.

Example: X = defectives in 3 tablets, each with illustrative p = 0.1 (independent): P(0) = 0.729, P(1) = 0.243, P(2) = 0.027, P(3) = 0.001 (sum = 1.000).

Diagram:



Key takeaway: Discrete = countable; PMF bar height IS P(X = x).

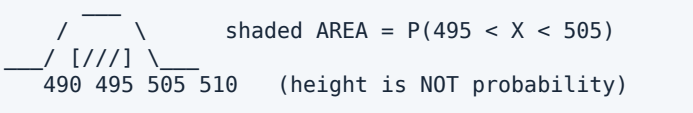
5.3 Continuous distribution and the PDF

Technical concept: A continuous variable takes any value in an interval. Its PDF f(x) is a smooth curve; probability = AREA under it between two points. P(X = one exact value) = 0.

Simple explanation: "Continuous" = smoothly measurable — 500.0, 500.1, 500.13 mg, no smallest step. Infinitely many values → chance of one exact value is zero; ask for ranges instead ("between 495 and 505 mg").

Example: Tablet weights cluster near 500 mg, trailing off both sides — the bell. P(495 < weight < 505) is the shaded area between those marks (≈ 68%, calculated in Box 1).

Diagram:



Key takeaway: Continuous = any value in a range; probability = area; P(exact value) = 0.

5.4 The normal distribution: meaning, characteristics, parameters, equation

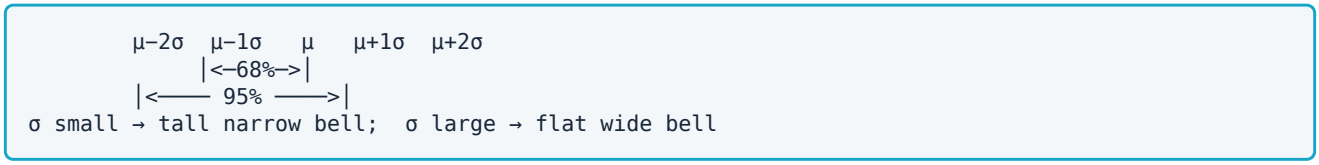
Technical concept: The normal distribution is the symmetric bell-shaped continuous distribution, fully described by two parameters: μ (mean, centre) and σ (SD, spread). Its PDF equation is f(x) = [1 / (σ√(2π))] × e^[−(x−μ)²/(2σ²)].

Simple explanation: Many measurements pile near an average and thin out symmetrically — tablet weights, assay values. Only TWO numbers define the whole bell: μ says where it sits, σ says how wide it is. (You won't compute the equation by hand; tables do that — understand what μ and σ do.)

Characteristics: bell-shaped and symmetric about μ; mean = median = mode (all at centre); total area = 1; tails run to ±infinity; 68–95–99.7 rule (≈68% within μ ± 1σ, ≈95% within μ ± 2σ, ≈99.7% within μ ± 3σ).

Example: Weights with μ = 500 mg, σ = 5 mg (illustrative): ≈68% between 495–505 mg; ≈95% between 490–510 mg. A batch with σ = 2 mg has a tall narrow bell (better control); σ = 8 mg gives a flat wide bell.

Diagram:



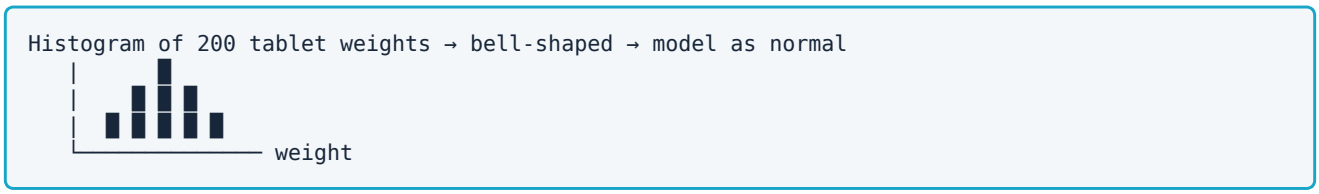
Key takeaway: Normal = symmetric bell; μ = position, σ = width; memorise 68–95–99.7.

5.5 Pharmaceutical examples of normally distributed data

Technical concept: Data are "modelled as normal" when their histogram is approximately bell-shaped.

Simple explanation: Many small random influences (vibration, powder flow, measurement error) add up to a bell. Typical illustrative examples: tablet weights from a press; content uniformity / assay (% of label claim near 100%); dissolution % at a fixed time; capsule fill weights; hardness; disintegration times; repeated HPLC assays. (Generic illustrations — no real study data claimed.)

Diagram:



Key takeaway: Weights, assay, dissolution, hardness — the classic pharmacy bell-curve data.

5.6 Standard normal, Z transformation, and reading the Z table

Technical concept: The standard normal has μ = 0, σ = 1. Z = (X − μ) ÷ σ converts any normal X into "SDs from the mean". The Z table gives the area LEFT of a stated Z: P(Z < z). Then P(Z > z) = 1 − table; P(a < Z < b) = table(b) − table(a); table(−z) = 1 − table(z).

Simple explanation: Infinitely many bells, one printed table — Z is the universal translator. A Z of +2 = "2 SDs above average". Look up Z → get left area → adjust with 1-minus or subtraction.

Example: Weights μ = 500, σ = 5 (illustrative). X = 510 → Z = (510−500)/5 = 2.0. As per the standard table: left of 1.64 → 0.9495, so P(Z > 1.64) = 0.0505; left of 1.96 → 0.9750, so P(−1.96 < Z < 1.96) = 0.9750 − 0.0250 = 0.95.

Diagram:

```
X = 510 mg ──Z=(X−μ)/σ──→ Z = +2.0 ──table──→ left area 0.9772
(μ=500, σ=5)              (mean 0, SD 1)      P(above) = 1 − 0.9772
Table: row "1.6" + column ".04" → 0.9495 for Z = 1.64
```

Key takeaway: Z = (X − μ)/σ; table gives left area; everything else is 1-minus or subtraction.

5.7 Binomial distribution

Technical concept: Needs four conditions: fixed n trials, two outcomes per trial, constant p, independent trials. PMF: P(X = x) = C(n,x) × p^x × (1−p)^(n−x), C(n,x) = n!/[x!(n−x)!]. Mean = np, variance = np(1−p).

Simple explanation: "Binomial" = two names (success/failure). Count successes in fixed independent tries: defectives in 10 tablets, passes in 20 dissolution tests.

Example: 10 tablets, illustrative p = 0.1 each, independent. P(exactly 2 defective) = C(10,2) × 0.1^2 × 0.9^8 ≈ 0.1937 (full steps in Box 4).

Diagram:

```
[✓] fixed n  [✓] two outcomes  [✓] constant p  [✓] independent → BINOMIAL
```

Key takeaway: Fixed n + two outcomes + constant p + independence = binomial; PMF gives P(exactly x).

5.8 Poisson distribution

Technical concept: Counts events in a fixed interval of time/space at average rate λ (no fixed n). PMF: P(X = x) = e^(−λ) × λ^x ÷ x!. Mean = variance = λ.

Simple explanation: For rare random happenings: colonies per plate, particles per vial, defects per m^2 of coating. No "n" — just an average rate λ.

Example: Colonies at λ = 3 per plate (illustrative): P(exactly 2) = e^(−3) × 3^2/2! ≈ 0.224; P(0) = e^(−3) ≈ 0.0498 (full steps in Box 6).

Diagram:

```
Binomial: fixed TRIALS (n tablets, pass/fail each)
Poisson:  fixed INTERVAL at rate λ (3 colonies/plate — varies plate to plate)
```

Key takeaway: Events in time/space at rate λ = Poisson; mean = variance = λ.

Classifications

Discrete vs continuous distributions

1 **Discrete** — counted values; PMF gives exact probabilities as bar heights. Examples: binomial, Poisson (defectives in a sample; colonies per plate).

2 **Continuous** — measured values; PDF gives probabilities as areas. Examples: normal, standard normal (tablet weight; assay value).

The three distributions by question type

1 **Normal** — "probability a MEASUREMENT falls in a range?" (weights, assay, dissolution).

2 **Binomial** — "exactly x successes in n fixed trials?" (defectives in a sample).

3 **Poisson** — "exactly x events at rate λ per interval?" (colonies per plate).

Processes

How to solve a Z problem (6 steps)

```
Step 1: Write μ, σ, and the target X or range.
Step 2: Convert: Z = (X − μ) ÷ σ (show subtraction and division).
Step 3: Sketch the bell, mark Z, SHADE the wanted area.
Step 4: Look up left-tail area(s) in the Z table.
Step 5: Combine: below → table; above → 1 − table; between → upper − lower.
Step 6: State answer; sanity-check (0–1? matches the shading?).
```

How to read the Z table

```
Split Z into row + column (1.64 → row "1.6", column ".04").
Read the entry = area LEFT of Z (1.64 → 0.9495 as per standard table).
Negative z: table(−z) = 1 − table(z) (−1.00 → 0.1587).
```

Binomial or Poisson? (decision steps)

```
Fixed n trials, two outcomes each? → BINOMIAL (check constant p, independence).
Average rate λ per interval, no fixed n? → POISSON.
A measurement with mean and SD? → NORMAL.
```

Diagrams and Flowcharts

PMF vs PDF

```
DISCRETE: bar HEIGHT = probability (read directly)
CONTINUOUS: shaded AREA = probability (height is not)
```

Which distribution?

```
        COUNTED (0,1,2..) + fixed n     → BINOMIAL  C(n,x)·p^x·q^(n−x)
        MEASURED (mg, %) + μ, σ         → NORMAL    Z = (X−μ)/σ
        EVENTS at rate λ per interval   → POISSON   e^(−λ)·λ^x/x!
```

Infographic Concept

INFOGRAPHIC CONCEPT: A one-page "Distribution Chooser". Top: three signposts — COUNTED (binomial: tablets), MEASURED (normal: bell of weights), RATE (Poisson: plate colonies). Middle: each formula in large type with tick-box conditions and a pharmacy icon. Bottom: the Z "translator" strip (510 mg → Z = +2.0) with 68–95–99.7 bands under the bell. Side: "PMF = bar height you read; PDF = area you shade". Colours: blue discrete, orange continuous, green Z.

Tables

PMF vs PDF

Point	PMF (discrete)	PDF (continuous)
Gives	P(X = exactly x) directly	Density; probability = area
Height	IS the probability	Is NOT the probability
P(exact value)	Can be > 0	Always 0
Total	Sums to 1	Area = 1

Binomial vs Poisson

Point	Binomial	Poisson
Key numbers	n (trials), p (per-trial chance)	λ (rate per interval)
PMF	C(n,x)·p^x·(1−p)^(n−x)	e^(−λ)·λ^x/x!
Mean / variance	np / np(1−p)	λ / λ
Example	Defectives in 10 tablets	Colonies per plate

Normal fact sheet

Item	Detail
Shape	Symmetric bell; mean = median = mode = μ
Parameters	μ (position), σ (width)
PDF	$f(x) = [1/(\sigma\sqrt{(2\pi)})]\cdot e^{[-(x-\mu)^2/(2\sigma^2)]}$
68–95–99.7	±1σ → 68%; ±2σ → 95%; ±3σ → 99.7%
Standard form	μ = 0, σ = 1; Z = (X−μ)/σ
Key Z (std table)	±1.00 → 68.26%; ±1.96 → 95%; ±2.58 → 99%

Worked Examples and Calculation Boxes

CALCULATION BOX 1 — Z: probability between two values

Question: Weights normal, μ = 500 mg, σ = 5 mg (illustrative). Find P(495 < X < 505).

```
Z₁ = (495−500)/5 = −1.00 → left area 0.1587 (std table)
Z₂ = (505−500)/5 = +1.00 → left area 0.8413
P = 0.8413 − 0.1587 = 0.6826
FINAL ANSWER: 0.6826 (68.26%) — matches the 68% rule ✓
```

CALCULATION BOX 2 — Z: probability above a value

Question: Same distribution. Find P(X > 512 mg).

```
Z = (512−500)/5 = 2.40 → left area 0.9918 (std table)
P(above) = 1 − 0.9918 = 0.0082
FINAL ANSWER: 0.0082 (0.82%) — tiny, as expected far in the tail ✓
```

CALCULATION BOX 3 — Z in reverse: find X

Question: Same distribution. Below what weight do 99% of tablets fall?

```
Z with 0.99 left area ≈ 2.33 (std table)
X = μ + Zσ = 500 + 2.33 × 5 = 500 + 11.65 = 511.65
FINAL ANSWER: 511.65 mg
```

CALCULATION BOX 4 — Binomial: exactly x

Question: 10 tablets, independent, illustrative p = 0.1 each. Find P(exactly 2 defective).

```
Conditions ✓ → P(X=2) = C(10,2) × 0.1² × 0.9⁸
C(10,2) = 90/2 = 45;  0.1² = 0.01
0.9⁸ = 0.43046721  (0.9²=0.81 → 0.9⁴=0.6561 → square it)
P = 45 × 0.01 × 0.43046721 = 45 × 0.0043046721 = 0.1937102445 ≈ 0.1937
FINAL ANSWER: ≈ 0.1937 (19.37%)
```

CALCULATION BOX 5 — Binomial: second problem

Question: 20 ampoules, independent, illustrative p = 0.05 each. Find P(exactly 3 defective).

```
P(X=3) = C(20,3) × 0.05³ × 0.95¹⁷
C(20,3) = 6840/6 = 1140;  0.05³ = 0.000125
0.95¹⁷ ≈ 0.41812035  (0.95⁸≈0.66342043 → square → ×0.95)
P = 1140 × 0.000125 × 0.41812035 = 1140 × 0.00005226504 ≈ 0.05958215
FINAL ANSWER: ≈ 0.0596 (5.96%)
```

CALCULATION BOX 6 — Poisson: exactly x

Question: Colonies at λ = 3 per plate (illustrative). Find (a) P(exactly 2), (b) P(0).

```
e^(−3) ≈ 0.04978707
(a) P(2) = 0.04978707 × 9/2 = 0.04978707 × 4.5 = 0.22404182 ≈ 0.2240
(b) P(0) = 0.04978707 × 1 = 0.0498
FINAL ANSWERS: (a) ≈ 0.2240  (b) ≈ 0.0498
```

CALCULATION BOX 7 — Poisson: mean = variance

Question: Particles at λ = 1.5 per mL (illustrative). Find P(exactly 1 in 1 mL); state mean, variance.

```
P(1) = e^(−1.5) × 1.5 = 0.22313016 × 1.5 = 0.33469524 ≈ 0.3347
Mean = λ = 1.5; Variance = λ = 1.5
FINAL ANSWER: ≈ 0.3347; mean = variance = 1.5
```

Analogies

1. **Mark-sheet.** A distribution lists every value with its chance, like a mark-sheet lists every student with marks.
2. **Idli vs dosa batter.** PMF = separate idli you count; PDF = batter spread on a tawa — you measure the area poured between two marks.
3. **Currency conversion.** Z converts every measurement "currency" into one universal currency so a single table works.
4. **Railway timetable.** You don't compute the train's speed — you look it up. Z in, probability out.
5. **Raindrops on tiles (Poisson).** No fixed n; just an average per tile — like colonies per plate.

Pharmacy Relevance

1. **Tablet weights.** The textbook normal example; Z converts specification limits into out-of-limit probabilities.
2. **Content uniformity / assay.** % of label claim clusters near 100%; normal probabilities estimate the fraction outside limits.
3. **Defect counting (binomial).** "How many defectives in n?" — acceptance sampling's core question.
4. **Microbial/particle counts (Poisson).** Colonies per plate, particles per vial — no fixed n, just a rate.
5. **Dissolution data.** % dissolved at a fixed time, commonly modelled as normal.
6. **Process capability.** From μ and σ, predict the out-of-specification fraction before mass testing.

India Context

- Under Schedule M, Indian QC labs treat weight, hardness and assay data as approximately normal; control charts monitor processes daily.
- The Indian Pharmacopoeia's acceptance criteria (uniformity of weight/content) make Z-score thinking a routine QC skill.
- India's generic sector relies on acceptance sampling — the binomial calculation behind lot accept/reject decisions.
- Sterile facilities' settle/contact plate counts are classic Poisson data (λ = average per plate).
- This BP701T training equips pharmacists to read stability data, dissolution profiles and QC charts.

Key Points

- Distribution = chance map; chances total 1.

- Discrete → PMF (bar height = probability). Continuous → PDF (area = probability); P(exact value) = 0.
- Normal: symmetric bell; μ = centre, σ = spread; 68–95–99.7.
- Z = (X−μ)/σ; table gives left area; above = 1 − table; between = upper − lower.
- Binomial: fixed n, two outcomes, constant p, independent → C(n,x)p^x(1−p)^(n−x); mean np, variance np(1−p).
- Poisson: rate λ → e^(−λ)λ^x/x!; mean = variance = λ.
- Choose: measurements → normal; successes in n trials → binomial; events at a rate → Poisson.

Common Confusions

1 **PMF vs PDF.** Bar height IS probability (PMF); only AREA is probability (PDF) — the chapter's top error.

2 **P(exact value) = 0** for continuous variables — always use a range with the Z table.

3 **μ, σ vs $\bar{x}$, s.** Population (true bell) vs one sample's statistics.

4 **Z vs table value.** 1.96 is a position; 0.9750 is the area left of it.

5 **Binomial vs Poisson.** Fixed n ("in 10 tablets") → binomial; rate per interval ("3 per plate") → Poisson.

6 **e^(−λ)** — the Poisson PMF starts with it; e^(+λ) or dropping it breaks the answer.

Common Mistakes

1 Binomial with non-independent trials (e.g. no-replacement draws from a tiny batch).

2 Poisson for fixed-n two-outcome data (that's binomial).

3 table(a) − table(b) reversed → negative "probability". Always upper minus lower.

4 Forgetting 1-minus for P(Z > z) / P(X > value).

5 Reverse-Z without multiplying: X = μ + Zσ, not μ + Z.

6 Computing the normal PDF equation by hand in exams — it's for understanding μ and σ, not arithmetic.

Mnemonics

- **"PMF = Points, PDF = Field"** — discrete dots vs continuous area.
- **"X minus mean over SD"** — the Z chant.
- **"68–95–99.7"** — the bell's postcode.
- **"BIN needs four to win"** — fixed n, Binary outcomes, Independent, constant p.
- **"λ is the average"** — Poisson's mean and variance in one number.
- **"1.96 for 95"** — ±1.96 encloses 95% (standard table).

Exam Summary

- Discrete: PMF, bar height = P(X=x). Continuous: PDF, area = probability; P(exact) = 0.
- Normal: bell, μ centre, σ spread, 68–95–99.7; PDF equation (roles of μ, σ).
- Standard normal μ=0, σ=1; Z=(X−μ)/σ; left-tail table; 1-minus and subtraction.
- Binomial: 4 conditions; C(n,x)p^x(1−p)^(n−x); mean np, variance np(1−p).
- Poisson: e^(−λ)λ^x/x!; mean = variance = λ.
- Typical 2-marker: define PMF/PDF; binomial conditions; Z formula. Typical 5-marker: full Z numerical; binomial "exactly x"; Poisson numerical.

One-Minute Revision

- Distribution: every value + chance; total 1.
- PMF: height = probability. PDF: area = probability; exact → 0.
- Normal: μ, σ, 68–95–99.7. Z = (X−μ)/σ; above = 1 − table; between = upper − lower.
- Binomial: C(n,x)p^x q^(n−x); mean np, var npq. Poisson: e^(−λ)λ^x/x!; mean = var = λ.
- Measurements → normal; n trials → binomial; rate → Poisson.

Questions

Short-Answer Questions

1 Define probability distribution and random variable, one example each.

2 Discrete vs continuous variable — one pharmacy example of each.

3 Define PMF and PDF in one line each.

4 Why is P(X = exact value) = 0 for continuous variables?

5 List four characteristics of the normal distribution.

6 Name the normal parameters and what each controls.

7 Write the Z formula; what does a Z score mean?

8 State the four binomial conditions.

9 Write the Poisson PMF; state its mean and variance.

Long-Answer Questions

1 Explain probability distributions; distinguish discrete/continuous with PMF/PDF and one pharmacy example each.

2 Describe the normal distribution: shape, characteristics, parameters, 68–95–99.7, PDF equation (roles of μ, σ), three pharmacy examples.

3 Explain Z transformation and the standard normal; how to read the Z table; find P(Z<z), P(Z>z), P(a<Z<b) with one numerical illustration each.

4 Binomial: conditions, PMF, mean, variance. Solve: n = 12, p = 0.2 (illustrative), find P(exactly 3). Show steps.

Multiple Choice Questions

1 A distribution assigns: a) one probability b) a probability to every value c) a mean only d) a sample size

2 PMF is for: a) continuous b) discrete c) normal curves d) Z tables

3 P(X = exactly 500 mg), continuous: a) 0.5 b) PDF height c) 0 d) 1

4 Normal: mean, median, mode are: a) all different b) equal at centre c) equal to σ d) undefined

5 Normal parameters: a) n, p b) λ, e c) μ, σ d) $\bar{x}$, s

6 Z = +2.0 means: a) 2 units above mean b) 2 SD above mean c) twice the mean d) 2nd percentile

7 P(−1.96 < Z < 1.96) ≈ (std table): a) 0.68 b) 0.90 c) 0.95 d) 0.99

8 NOT a binomial condition: a) fixed n b) two outcomes c) constant p d) normal data

9 Poisson λ = 4 → variance: a) 2 b) 4 c) 16 d) 0.25

10 Colonies per plate → : a) normal b) binomial c) Poisson d) Z

Answer key: 1-b, 2-b, 3-c, 4-b, 5-c, 6-b, 7-c, 8-d, 9-b, 10-c

Viva Questions

1 Random variable in one sentence?

2 PMF or PDF — which gives probability as bar height?

3 What do μ and σ do to the bell?

4 Why convert X to Z before the table?

5 Table gives left or right area?

6 Get P(Z > 1.64) from 0.9495?

7 Poisson over binomial — when?

Application-Based Questions

1 Capsule weights normal, μ = 250 mg, σ = 4 mg (illustrative). Find P(246 < X < 254) and P(X > 256). Show Z steps and table use.

2 15 tablets, independent, illustrative p = 0.08 visual defect. Find P(exactly 1 defective) and the mean number.

3 Settle plates average λ = 2 colonies (illustrative). Find P(exactly 3) and P(0).

4 Explain to a trainee why the normal curve's height at 500 mg is NOT P(weight = 500 mg), and how to find P(495 < weight < 505) correctly.

References

- Bolton (Pharmaceutical Statistics)
- Daniel (Biostatistics)
- Montgomery (Design and Analysis of Experiments)

CHAPTER 6: Sampling Distributions — t, F and Chi-Square

Learning Objectives

After studying this chapter, you will be able to:

- Explain what a sampling distribution is and how it differs from a single sample's distribution.
- Define the t statistic, write its equation, and calculate it from sample data.
- Describe the t distribution, explain degrees of freedom, and state why df shapes the t curve.
- Read and interpret the t table (one-tailed and two-tailed) for given df and α.
- State the applications of the t distribution in pharmacy.
- Define the F statistic, write its equation, and calculate it from two sample variances.
- Describe the F distribution and read/interpret the F table using numerator and denominator df.
- Define the chi-square statistic, write its equation, and calculate it from observed and expected counts.
- Describe the chi-square distribution and read/interpret its table.
- State the applications of the chi-square distribution (goodness of fit, independence).
- Choose correctly between t, F and chi-square for a pharmacy problem.

Definitions

- **Statistic:** A number calculated from sample data ($\bar{x}$, s^2, t, F, χ^2).
- **Sampling distribution:** The probability distribution of a statistic over ALL possible samples of the same size. Example: the distribution of $\bar{x}$ values from thousands of 10-tablet samples.
- **t statistic:** $t = (\bar{x} - \mu_0) \div (s / \sqrt{n})$ — how many standard errors the sample mean lies from a claimed mean μ_0.
- **Standard error of the mean:** $s \div \sqrt{n}$ — typical bounce of sample means; shrinks as n grows.
- **t distribution:** Sampling distribution of the t statistic; symmetric and bell-shaped like normal but heavier-tailed when samples are small.
- **Degrees of freedom (df):** Independent pieces of information; for one-sample t, df = n − 1. Larger df → t curve closer to normal.

- **F statistic:** $F = s_1^2 \div s_2^2$ (larger variance on top) — compares two sample variances as a ratio.
- **F distribution:** Sampling distribution of F; non-negative, right-skewed; needs two df values (df_1 numerator, df_2 denominator).
- **Chi-square statistic:** $\chi^2 = \Sigma (O - E)^2 \div E$ over all categories (O = observed, E = expected counts).
- **Chi-square distribution:** Sampling distribution of χ^2; non-negative, right-skewed; shape set by df.
- **Significance level (α):** Cut-off probability (commonly 0.05); a statistic more extreme than the table's critical value is "statistically significant".

Terminology

Term	Simple meaning
Population	The whole group (all tablets in the batch)
Sample	The small group actually tested
Statistic	Number from the sample ($\bar{x}$, s^2, t, F, χ^2)
Parameter	Number describing the population (μ, σ^2) — usually unknown
Sampling distribution	Distribution of the statistic across repeated samples
Standard error	SD of sample means: $s \div \sqrt{n}$
Critical value	The table value your statistic is compared against
Goodness of fit	Does the observed pattern match the expected one?
Test of independence	Are two categorical variables related?

Simple Explanations

6.1 Meaning of sampling distributions

Technical concept: The sampling distribution is the probability distribution of a statistic (like $\bar{x}$) over every possible sample of size n. It describes how the statistic varies from sample to sample.

Simple explanation: One sample gives one average; a second sample of 10 tablets gives a slightly different average. Plot the averages of thousands of samples — that plot is the sampling distribution of the mean. It tells us what "ordinary chance bounce" looks like, so we can spot a suspiciously far value.

Example: Batch true mean μ = 500 mg (illustrative). Sample means of n = 16 form their own bell centred at 500 mg with spread $\sigma \div \sqrt{n}$. A sample mean of 505 mg can then be judged against this bounce.

Diagram:

```
Population (μ=500, σ=8) → samples of 16 → x¯: 502, 498, 505, ...
→ sampling distribution of x¯: bell at 500, spread σ/√n
```

Key takeaway: A sampling distribution shows a statistic's chance bounce across repeated samples.

6.2 The t statistic, its equation, and the t distribution

Technical concept: When σ is unknown (real life) and the sample is small, we estimate σ with s. The t statistic standardises the gap: $t = (\bar{x} - \mu_0) \div (s \div \sqrt{n})$. Its sampling distribution — the t distribution — is symmetric and bell-shaped like the normal curve but with heavier tails at small df, merging into the normal curve as df grows.

Simple explanation: t asks: "how many standard errors is my sample mean from the claim?" Top = the gap (seen minus claimed); bottom = typical bounce (standard error). t = 2.5 means "2.5 bounces away" — far. The extra wobble from estimating σ with a small sample fattens the tails — the t curve is the normal curve's cautious cousin; with big samples the caution fades.

Example: Claim μ_0 = 500 mg; sample n = 16, $\bar{x}$ = 505 mg, s = 8 mg (illustrative). SE = 8 ÷ √16 = 8 ÷ 4 = 2. t = (505 − 500) ÷ 2 = 2.5. With df = 15, the two-tailed 5% critical value is 2.131 (standard t table) — |2.5| > 2.131, so significant. Note the shrinkage with df: two-tailed 5% critical values are 2.571 (df = 5), 2.042 (df = 30), 1.96 (normal) — t approaches normal as df grows.

Diagram:

```
t = (x̄ − μ₀) / (s/√n)  ← gap ÷ typical bounce
t curve vs normal (same centre): t spreads wider at small df
```

Key takeaway: t = gap ÷ standard error; big |t| = sample disagrees with claim; t → normal as df grows.

6.3 Degrees of freedom and their relevance to t

Technical concept: df = independent values free to vary. For one-sample t, df = n − 1 (once the mean is fixed, the last observation is forced). df selects the t curve: small df → wide, cautious curve; large df → near-normal.

Simple explanation: If the mean of 4 numbers is 10 and you know three (8, 12, 9), the fourth MUST be 11 — only 3 are free. df is the dial setting the t curve's width.

Example: n = 16 → df = 15. n = 6 → df = 5, a much wider curve.

Diagram:

```
8, 12, 9 free; 4th FORCED to 11 → df = 3 = n − 1
df = 5  → wide curve, critical 2.571
df = 30 → near-normal, critical 2.042
```

Key takeaway: df = n − 1 for one-sample t; df dials the t curve's width.

6.4 Reading and interpreting the t table

Technical concept: Rows = df, columns = tail areas (one- or two-tailed α). The entry is the critical value cutting off that tail area. If |t calculated| > critical value, the difference is significant at α.

Simple explanation: Pick the df row, pick the α column (check one- vs two-tailed: "different from" = two-tailed, "greater than" = one-tailed), read the critical value, compare.

Example: df = 15, two-tailed α = 0.05 → 2.131 (standard table). Our t = 2.5 > 2.131 → significant. Also: df = 10, one-tailed 0.05 → 1.812.

Diagram:

```
row df=15  →  column "two-tail 0.05"  →  2.131
|t| = 2.5 > 2.131 → significant at 5%
```

Key takeaway: Row = df, column = α (mind one- vs two-tailed); beyond the table value = significant.

6.5 Applications of the t distribution

Technical concept: t is for testing MEANS with small samples (roughly n < 30) and unknown σ — the standard laboratory situation.

Simple explanation: QC rarely tests 100 tablets for a comparison — usually 6, 10 or 20. Uses: (1) one-sample t-test — batch mean assay vs label claim; mean weight vs 500 mg; (2) paired t-test — same tablets before/after coating; (3) two-sample t-test — two machines, analysts or methods compared on means. (Full test procedure belongs to hypothesis testing, elsewhere in the course.)

Diagram:

```
Small n + σ unknown + question about a MEAN → t
n = 10, x̄ = 99.2%, s = 1.1% → t = (99.2 − 100)/(1.1/√10)
```

Key takeaway: Small samples + unknown σ + means = t distribution.

6.6 The F statistic, its equation, and the F distribution

Technical concept: $F = s_1^2 \div s_2^2$ with the LARGER variance on top (so $F \geq 1$); $df_1 = n_1 - 1$ (numerator), $df_2 = n_2 - 1$ (denominator). Its sampling distribution — the F distribution — takes only values ≥ 0, skews right, and its shape is fixed by the pair (df_1, df_2).

Simple explanation: While t compares averages, F compares spreads: "is Machine A more variable than B?" Near 1 = equally consistent; large = one much more variable. The curve starts at zero (a ratio can't be negative) with a long right tail.

Example: Two dissolution testers (illustrative): $s_1^2 = 48$ ($n_1 = 5$), $s_2^2 = 12$ ($n_2 = 21$). $F = 48 \div 12 = 4.0$, $df_1 = 4$, $df_2 = 20$. The 5% critical value is 2.87 (standard F table) — 4.0 > 2.87, so the variances differ significantly.

Diagram:

```
F = larger s² / smaller s²     F near 1 → similar; F large → one more variable
F curve: starts at 0, right skew; needs (df₁, df₂)
```

Key takeaway: F = bigger ÷ smaller variance; needs two df; right-skewed from 0.

6.7 Reading and interpreting the F table

Technical concept: The F table gives upper-tail (α = 0.05) critical values. Columns = df_1 (numerator), rows = df_2 (denominator). F calculated > table value → variances differ significantly.

Simple explanation: A map grid reference: column = top variance's df, row = bottom variance's df; their crossing is the critical value. Always put the larger variance on top, since the table is upper-tail.

Example: $df_1 = 4$ (column), $df_2 = 20$ (row), $\alpha = 0.05 \rightarrow 2.87$ (standard table). Our $F = 4.0 > 2.87 \rightarrow$ significant.

Diagram:

```
column df₁=4 × row df₂=20 → 2.87
F = 4.0 > 2.87 → variances differ significantly at 5%
```

Key takeaway: Column = numerator df, row = denominator df; beyond the table value = significant.

6.8 The chi-square statistic and the chi-square distribution

Technical concept: $\chi^2 = \Sigma (O - E)^2 \div E$ over all categories (O = observed, E = expected counts). Its sampling distribution — the chi-square distribution — is non-negative and right-skewed; shape set by df. Goodness of fit: $df = k - 1$. Test of independence (r × c table): $df = (r - 1)(c - 1)$.

Simple explanation: Chi-square is the "does the pattern match?" statistic. Expected 25 each of 4 capsule colours; counted 20, 30, 25, 25. Square each gap, divide by expectation, add: small χ^2 = matches; large χ^2 = something's off. Squared gaps can't be negative → curve starts at 0 with a right tail.

Example: $\chi^2 = (20-25)^2/25 + (30-25)^2/25 + 0 + 0 = 1 + 1 = 2$ (illustrative). $df = 4 - 1 = 3$; 5% critical value 7.815 (standard table) — $2 < 7.815$, pattern fits.

Diagram:

```
per category: (O − E)² / E, then ADD → χ² = 2 here
χ² curve: starts at 0, right skew; df = k − 1 (fit) or (r−1)(c−1) (independence)
```

Key takeaway: $\chi^2 = \Sigma(O-E)^2/E$ on counts; small = as expected, large = mismatch.

6.9 Reading and interpreting the chi-square table

Technical concept: Rows = df, columns = upper-tail α. Entry = critical value cutting off area α in the right tail. χ^2 calculated > table value → significant mismatch/association. Always right-tail (no two-tailed fuss).

Simple explanation: df row → α column (usually 0.05) → critical value. Beyond it = "the pattern differs significantly".

Example: df = 3, α = 0.05 → 7.815; our $\chi^2 = 2 < 7.815$ → fits. For a 2 × 2 table (df = 1): 3.841.

Diagram:

```
row df=3 × column α=0.05 → 7.815
χ² = 2 < 7.815 → not significant (fits well)
```

Key takeaway: Row = df, column = upper-tail α; beyond = significant.

6.10 Applications of the chi-square distribution

Technical concept: Chi-square is for COUNT data in categories, not measurements. Two classic uses:

1 **Goodness of fit** — "does the observed pattern match the expected one?" (Capsule colours: $\chi^2 = 2$, df = 3 → fits.)

2 **Test of independence** — "are two categorical variables related?" (Shift vs result table: $\chi^2 \approx 3.35$, df = 1, critical 3.841 → no significant association at 5%. Full steps in Box 4.)

3 **Test of homogeneity** — "do groups share the same distribution?" (e.g. defect-type patterns across three batches.)

Diagram:

```
COUNT data → chi-square
one variable vs expected → GOODNESS OF FIT (df = k − 1)
two variables vs each other → INDEPENDENCE (df = (r−1)(c−1))
```

Key takeaway: Counts in categories → chi-square: fit (one variable) or independence (two variables).

Classifications

The three sampling distributions by question type

1 **t** — MEANS with small samples, σ unknown. "Batch mean vs claim?" df = n − 1.

2 **F** — two VARIANCES compared. "Machine A more variable?" F = bigger s^2 ÷ smaller s^2; $df_1 = n_1 - 1$, $df_2 = n_2 - 1$.

3 **Chi-square** — COUNT patterns. "Pattern matches expectation?" (df = k − 1) / "two variables related?" (df = (r−1)(c−1)).

Processes

One-sample t comparison (6 steps)

```
1. Write claim μ₀; sample n, x̄, s.
2. SE = s ÷ √n.   3. t = (x̄ − μ₀) ÷ SE.
4. df = n − 1; pick one-/two-tailed α (0.05).
5. Read critical value (row = df, column = α).
6. |t| > critical → significant; else not.
```

Comparing two variances with F (5 steps)

```
1. F = larger s² ÷ smaller s² (F ≥ 1).
2. df₁ = n₁ − 1 (numerator), df₂ = n₂ − 1 (denominator).
3. Read critical value: column = df₁, row = df₂, α = 0.05.
4. F > critical → variances differ significantly.
5. State in words.
```

Chi-square test (6 steps)

```
1. Observed counts O per cell.   2. Expected counts E
   (claimed pattern, or row × column ÷ grand total).
3. Per cell: (O − E)² ÷ E.   4. Sum: χ² = Σ(O−E)²/E.
5. df = k − 1 (fit) or (r−1)(c−1) (independence); read critical value.
6. χ² > critical → significant. (Check: Σ(O − E) = 0.)
```

Diagrams and Flowcharts

The three curves

```
t: symmetric, fat tails at small df (−t … 0 … +t)
F, χ²: start at 0, right skew (only the right tail matters)
```

t, F or chi-square?

```
Question about a MEAN, small n, σ unknown → t = (x̄−μ₀)/(s/√n), df = n−1
Question comparing two SPREADS            → F = s₁²/s₂² (bigger top), df₁, df₂
Question about COUNTS in categories       → χ² = Σ(O−E)²/E
```

Infographic Concept

INFOGRAPHIC CONCEPT: A one-page "Which Sampling Distribution?" poster. Three columns: balance scale (t — mean vs claim), two bells of different widths (F — comparing spreads), bar chart of capsule colours vs expected (chi-square — comparing patterns). Each column: the question, the formula in large type, the df rule, and a mini table-reading strip (t: row=df, col=α; F: col=df_1, row=df_2; chi-square: row=df, col=α). Banner: "Small n + unknown σ → t; two spreads → F; counts → chi-square". Bottom: curve sketches with critical markers. Colours: blue t, green F, orange chi-square.

Tables

t vs F vs chi-square

Point	t	F	Chi-square
For	Means (small n, σ unknown)	Two variances	Count patterns
Statistic	$(\bar{x} - \mu_0) \div (s/\sqrt{n})$	$s_1^2 \div s_2^2$ (larger top)	$\Sigma(O-E)^2/E$
Range/shape	$-\infty$ to $+\infty$, symmetric	0 to $+\infty$, right skew	0 to $+\infty$, right skew
df	$n - 1$	$df_1 = n_1 - 1$, $df_2 = n_2 - 1$	$k-1$ or $(r-1)(c-1)$
Table	Row = df, col = α	Col = df_1, row = df_2	Row = df, col = α
Typical 5% critical	2.131 (df=15, two-tail)	2.87 (df_1=4, df_2=20)	7.815 (df=3); 3.841 (df=1)
Pharmacy use	Batch mean vs claim	Precision comparison	Defect patterns; associations

Decision table

Data	Question	Use
Measurements, small n, σ unknown	Mean vs claim?	t

Data	Question	Use
Measurements	Two groups' means differ?	t (two-sample)
Measurements	One method more variable?	F
Counts in categories	Pattern matches expectation?	Chi-square (fit)
Counts in 2-way table	Variables related?	Chi-square (independence)
Measurements, large n or σ known	Mean vs claim?	Z, not t

Worked Examples and Calculation Boxes

CALCULATION BOX 1 — t statistic

Question: Claim μ_0 = 500 mg. Sample n = 16: $\bar{x}$ = 505 mg, s = 8 mg (illustrative). Calculate t; test two-tailed at 5%.

```
SE = s ÷ √n = 8 ÷ 4 = 2
t = (505 − 500) ÷ 2 = 5 ÷ 2 = 2.5
df = 16 − 1 = 15
Critical (df=15, two-tail 0.05) = 2.131 (standard t table)
|2.5| > 2.131 → significant at 5%
FINAL ANSWER: t = 2.5, df = 15; mean differs significantly from 500 mg.
```

CALCULATION BOX 2 — F statistic

Question: Tester 1: s_1^2 = 48, n_1 = 5; Tester 2: s_2^2 = 12, n_2 = 21 (illustrative). Calculate F; test at 5%.

```
F = 48 ÷ 12 = 4.0   (larger on top)
df₁ = 5 − 1 = 4;  df₂ = 21 − 1 = 20
Critical (df₁=4, df₂=20, 0.05) = 2.87 (standard F table)
4.0 > 2.87 → significant at 5%
FINAL ANSWER: F = 4.0 (df 4, 20); testers differ in variability.
```

CALCULATION BOX 3 — Chi-square goodness of fit

Question: 100 capsules, expected 25 per colour; observed 20, 30, 25, 25 (illustrative). Calculate χ^2; test at 5%.

```
(20−25)²/25 = 25/25 = 1;  (30−25)²/25 = 25/25 = 1;  others 0
χ² = 1 + 1 + 0 + 0 = 2
df = 4 − 1 = 3;  critical (df=3, 0.05) = 7.815 (standard table)
2 < 7.815 → not significant
FINAL ANSWER: χ² = 2, df = 3; pattern fits expectations.
```

CALCULATION BOX 4 — Chi-square test of independence

Question: Shift × result (illustrative): Day 50 pass/5 fail; Night 35 pass/10 fail (total 100). Test association at 5%.

```
E = row × column ÷ 100:
E(day,pass) = 55×85/100 = 46.75;  E(day,fail) = 55×15/100 = 8.25
E(night,pass) = 45×85/100 = 38.25; E(night,fail) = 45×15/100 = 6.75
(46.75+8.25+38.25+6.75 = 100 ✓)

(O−E)²/E:  (3.25)²/46.75 = 10.5625/46.75 ≈ 0.226
           10.5625/8.25 ≈ 1.280;  10.5625/38.25 ≈ 0.276;  10.5625/6.75 ≈ 1.565
χ² = 0.226 + 1.280 + 0.276 + 1.565 = 3.347 ≈ 3.35
df = (2−1)(2−1) = 1;  critical (df=1, 0.05) = 3.841 (standard table)
3.35 < 3.841 → not significant at 5%
FINAL ANSWER: χ² ≈ 3.35, df = 1; no significant shift–result association.
```

Analogies

1. **Poll of polls.** One poll = one sample; the spread of many polls = the sampling distribution — chance bounce made visible.
2. **Umpire's margin.** "How far outside the line, in ball-widths?" = t: "how far from the claim, in standard-error widths?"
3. **Free bus seats.** 4 bags must total a fixed weight; 3 placed freely, the 4th forced → df = n − 1.
4. **Tabla players (F).** Same average tempo, different wobble — F = wobblier ÷ steadier.

Pharmacy Relevance

1. **Batch mean vs claim (t).** Small-sample check of mean assay/weight against specification — the everyday QC t-test.
2. **Method/analyst comparison (t).** Do two analysts or old vs new HPLC method give different mean assays?
3. **Precision comparison (F).** New method's variance vs existing method's — significantly larger F = less precise.
4. **Equipment qualification (F).** Two dissolution testers/balances compared for consistency.
5. **Defect patterns (chi-square).** Goodness of fit: do observed defect types match the historical pattern?

India Context

- Indian QC labs routinely work with n = 6–20, making t (not Z) the everyday tool for batch-mean decisions.
- Method transfer between Indian sites uses variance comparisons to match precision between laboratories.
- Schedule M defect logs (by type, shift, machine) are ready-made chi-square tables for deviation investigations.
- Indian Pharmacopoeia uniformity/assay checks on samples lead naturally to t-based "close enough to claim?" decisions.
- BP701T teaches t, F, chi-square as the core small-sample toolkit because Indian QC so often decides from limited data.

Key Points

- Sampling distribution = a statistic's distribution over repeated samples = "normal bounce".
- $t = (\bar{x} - \mu_0)/(s/\sqrt{n})$; small n, σ unknown; $df = n - 1$; $|t| >$ table → significant.
- t curve: symmetric, fat tails at small df → normal as df grows.
- F = larger $s^2 \div$ smaller s^2; $df_1 = n_1 - 1$, $df_2 = n_2 - 1$; 0 to ∞, right skew; table: col = df_1, row = df_2.
- $\chi^2 = \Sigma(O-E)^2/E$ on COUNTS; 0 to ∞, right skew; df = k−1 (fit) or (r−1)(c−1) (independence); table: row = df, col = α.
- Choose: means + small n → t; two spreads → F; category counts → chi-square.

Common Confusions

1 **t vs Z.** σ known / large n → Z; σ unknown + small n → t. Lab data usually needs t.

2 **t statistic vs t distribution.** 2.5 is the statistic (one number); the curve is the distribution.

3 **df ≠ n.** One-sample t uses $n - 1$ — wrong row otherwise.

4 **F needs two df, in order.** df_1 = numerator (larger variance); swapping reads the wrong entry.

5 **Chi-square needs COUNTS.** Percentages or means in $(O-E)^2/E$ are invalid.

6 **2 × 2 independence df = 1**, not 3. ((2−1)(2−1).)

Common Mistakes

1 df = n instead of $n - 1$ for t.

2 Using Z's 1.96 for a small-sample mean test instead of the t critical value.

3 Smaller variance on top in F, then comparing $F < 1$ to an upper-tail table.

4 Swapping df_1/df_2 in the F table.

5 Chi-square on percentages instead of raw counts.

6 Wrong E in independence tests: E = row × column ÷ grand total.

Mnemonics

- **"t is for Tiny"** — small n, unknown σ → t.
- **"F is a Fraction of variances"** — bigger on top.
- **"Observed minus Expected, squared over Expected"** — the chi-square chant.
- **"Column-df-one, row-df-two"** — F table orientation.
- **"Counts, not amounts"** — chi-square eats frequencies only.

Exam Summary

- $t = (\bar{x} - \mu_0)/(s/\sqrt{n})$, df = n−1; symmetric, fat tails → normal; row = df, col = α; $|t|$ > critical → significant.
- F = larger s^2 ÷ smaller s^2; $df_1 = n_1 - 1$, $df_2 = n_2 - 1$; right skew; col = df_1, row = df_2; F > critical → variances differ.
- $\chi^2 = \Sigma(O-E)^2/E$; right skew; row = df, col = upper α; χ^2 > critical → significant.
- df: fit = k−1; independence = (r−1)(c−1); E = row × col ÷ total.
- Uses: t → small-sample means; F → precisions; chi-square → count patterns/associations.
- Typical 2-marker: define sampling distribution; write an equation; state df rules. Typical 5-marker: full t calculation + table verdict; full chi-square numerical.

One-Minute Revision

- Sampling distribution: statistic's bounce across samples.
- $t = (\bar{x} - \mu_0)/(s/\sqrt{n})$; df = n−1; $|t|$ > table → significant.
- F = bigger ÷ smaller variance; df_1, df_2; col = df_1, row = df_2.
- $\chi^2 = \Sigma(O-E)^2/E$ on counts; df = k−1 or (r−1)(c−1).
- F, χ^2: from 0, right skew, right-tail tables.

Questions

Short-Answer Questions

1 Define sampling distribution (one–two lines).

2 Write the t equation; name each symbol.

3 What are degrees of freedom? State df for one-sample t.

4 Why is the t curve heavier-tailed than normal at small df?

5 Write the F equation; which variance goes on top, and why?

6 Write the chi-square formula; define O and E.

7 df for: (a) goodness of fit, 5 categories; (b) 3 × 2 table.

8 How do you read the t table? (row, column, tails.)

9 Name one pharmacy application each of t, F, chi-square.

Long-Answer Questions

1 Explain sampling distributions with a pharmacy illustration; how do they differ from the distribution of individual observations?

2 Explain $t = (\bar{x} - \mu_0)/(s/\sqrt{n})$ part by part; df and its effect on the curve; how to read the t table, with an example.

3 Explain F: equation, larger-on-top rule, df_1/df_2, curve shape, table reading; two pharmacy applications.

4 Explain chi-square: equation, why counts are required, curve shape, df rules, table reading — illustrate with the capsule-colour example, every step.

Multiple Choice Questions

1 Sampling distribution = distribution of: a) individuals b) a statistic over repeated samples c) parameters d) errors

2 t statistic = a) $(\bar{x}-\mu_0)/\sigma$ b) $(\bar{x}-\mu_0)/(s/\sqrt{n})$ c) s_1^2/s_2^2 d) $\Sigma(O-E)^2/E$

3 One-sample t, n = 16 → df: a) 16 b) 15 c) 8 d) 4

4 As df grows, t: a) skews more b) approaches standard normal c) shifts right d) turns bimodal

5 F numerator = a) smaller variance b) larger variance c) pooled mean d) sample size

6 F takes values: a) $-\infty$ to $+\infty$ b) 0 to $+\infty$ c) -1 to 1 d) integers only

7 F table columns = a) df_2 b) df_1 (numerator) c) α only d) n

8 Chi-square uses: a) means and SDs b) observed and expected counts c) variances d) Z scores

9 Goodness of fit, 4 categories → df: a) 4 b) 3 c) 2 d) 1

10 2 × 2 independence test → df: a) 3 b) 2 c) 1 d) 4

Answer key: 1-b, 2-b, 3-b, 4-b, 5-b, 6-b, 7-b, 8-b, 9-b, 10-c

Viva Questions

1 Statistic vs parameter?

2 t instead of Z — when?

3 t = 2.5 in plain words?

4 Why larger variance on top in F?

5 Why can't chi-square be negative?

6 df for independence?

Application-Based Questions

1 n = 9 tablets: $\bar{x}$ = 498 mg, s = 6 mg (illustrative); claim 500 mg. Calculate t, df; compare with 2.306 (df = 8, two-tail 5%, standard table); conclude in words.

2 Two analysts, 7 assays each (illustrative): s^2 = 9.0 vs 2.0. Calculate F, df_1, df_2; compare with 4.28 (df 6, 6, 5%, standard table). Conclude on precision.

3 Defect counts: capping 12, lamination 18, chipping 15, discolouration 15 (total 60, illustrative); expected 15 each. Calculate χ^2, df; compare with 7.815 (df = 3, 5%); conclude.

References

- Bolton (Pharmaceutical Statistics)
- Daniel (Biostatistics)
- Montgomery (Design and Analysis of Experiments)