Chi-Square Tests

AP Statistics ▪ Unit 8

AP Statistics

What we will cover

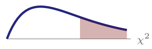

1 Idea

The chi-square idea

The **chi-square** 卡方 tests compare counts:

- **observed** 观测频数 – what happened
- **expected** 期望频数 – what the null predicts

> Big gaps ⇒ doubt the claim. **Goodness-of-fit** 拟合优度 checks one variable; a **two-way table** 双向表 checks two.

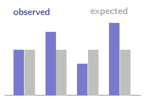

Compare bar by bar.

2 GOF

Goodness-of-fit test

Does one variable follow a claimed distribution?

$$\chi^2 = \sum \frac{(\text{obs} - \text{exp})^2}{\text{exp}}$$

- expected $=$ total $\times$ claimed proportion
- condition: every expected ≥ 5
- $df =$ categories $- 1$

> Large χ^2 ⇒ big gaps; the p-value is the right-tail area.

Goodness-of-fit test, in detail

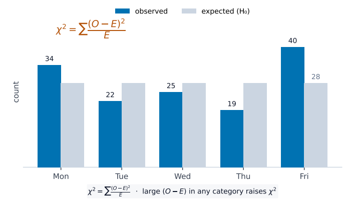

obs vs exp under H_0

> Does one variable follow a claimed distribution?

Setting Up a Goodness-of-Fit Test

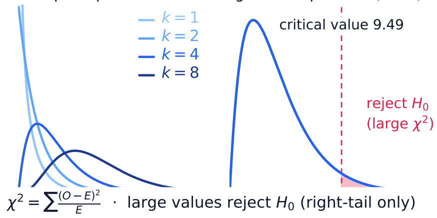

> The chi-square distribution is right-skewed. A large statistic lands in the shaded right tail past the critical value – that is where you reject the model.

3 Two-way

Two-way tables

For a two-way table, compute each **expected count** 期望频数 from the margins (under "no relationship"):

$$\text{exp} = \frac{(\text{row total})(\text{column total})}{\text{grand total}}$$

- every expected ≥ 5
- $df = (\text{rows} - 1)(\text{cols} - 1)$

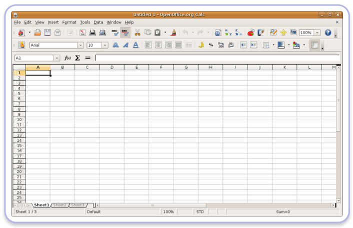

Organise categorical counts before the test

4 Homogeneity

Homogeneity vs independence

Same χ^2 mechanics – the *design* sets the name:

Homogeneity 齐性	**Independence** 独立性
several separate samples, one variable	*one* sample, two variables recorded
H_0: same distribution across groups	H_0: the two variables are unrelated

> Separately survey men and women = homogeneity; one sample recording gender & preference = independence.

Three chi-square tests, one calculation

the shared arithmetic	$\chi^2 = \sum \frac{(O-E)^2}{E}$ in every case – only the question and the *df* change
goodness-of-fit 拟合优度	one categorical variable against a claimed distribution. $df = k-1$
test for homogeneity 同质性检验	several populations, one variable – are the distributions the same? Data from separate samples
test for independence 独立性检验	one population, two variables – are they associated? Data from one sample
the two-way ***df***	$(\text{rows}-1)(\text{columns}-1)$ for both
the condition	every expected count at least 5

> Homogeneity and independence use identical arithmetic and answer different questions. Which you have is decided by how the data were collected, not by the table.

The three chi-square tests

Goodness-of-fit	The **goodness-of-fit (GOF)** test: one categorical variable against a claimed distribution. Expected counts come from the claim.
Homogeneity	Several populations, one variable: are the distributions the same across the groups?
Independence	One population, two variables: are they associated? Expected counts are $\frac{\text{row total} \times \text{column total}}{n}$.
Choosing	Count the samples and the variables. One sample, one variable is GOF; several samples is homogeneity; one sample, two variables is independence.

> All three use the same statistic and the same conditions – expected counts at least 5. Only the degrees of freedom and the conclusion differ.

5 Choose

Choosing a procedure

Count the samples and variables:

- **One proportion**: one-sample z
- **Two proportions**: two-sample z
- **Goodness-of-fit**: one variable vs a claim
- **Homogeneity**: several samples, one variable
- **Independence**: one sample, two variables

> **Number of samples** 样本数 + **number of variables** 变量数 point straight to the test.

Exam tips

- Use $\chi^2 = \sum \frac{(O-E)^2}{E}$ for **categorical** data; always divide by the **expected** count.
- Pick the right test: goodness-of-fit (one variable), independence, or homogeneity (two-way table).
- Compute expected counts as $\frac{\text{row total} \times \text{column total}}{\text{grand total}}$ and check each is ≥ 5.
- A large χ^2 (small p-value) means observed counts differ from expected by more than chance.
- State degrees of freedom correctly (categories -1, or $(r-1)(c-1)$).

Are My Results Unexpected?

When data are **counts** spread across several categories, we test whether the observed counts differ from what a claim predicts.

> The tool is the **chi-square 卡方** (χ^2) statistic, which adds up the standardized gaps between observed and expected counts:
>
> $$\chi^2 = \sum \frac{(O-E)^2}{E}.$$

6 Recap

Worked example – chi-square goodness of fit

A die is rolled 100 times: four faces give $30, 20, 25, 25$ over four categories. Test whether it is fair.

- H_0: all four categories are equally likely $\Rightarrow$ expected $= 100/4 = 25$ each.
- $\chi^2 = \sum \frac{(O-E)^2}{E} = \frac{25}{25} + \frac{25}{25} + 0 + 0 = 2.0$
- $\text{df} =$ categories $- 1 = 3$; critical value at $\alpha = 0.05$ is 7.81.
- $2.0 < 7.81 \Rightarrow$ fail to reject H_0: no evidence the die is unfair.

Use counts, never percentages, and check every expected count is ≥ 5. "Fail to reject" is not "proved fair" – you just lack evidence against it.

Unit 8 – key takeaways

1 Idea
compare observed to expected counts; right-tail p-value

2 Statistic
$\chi^2 = \sum(\text{O} - \text{E})^2/\text{E}$; every expected ≥ 5

3 Tables
expected $=$ row$\times$col/grand; $df = (r-1)(c-1)$

4 Which test
fit, homogeneity, independence by design

Check yourself

1. Write the chi-square statistic formula.
2. What is the *df* for a goodness-of-fit with 6 categories?
3. How do you find an expected count in a two-way table?
4. Several separate samples, one variable: which test?

Answers **1.** $\chi^2 = \sum(\text{O} - \text{E})^2/\text{E}$ **2.** $6 - 1 = 5$ **3.** row$\times$col/grand total **4.** homogeneity