Introduction to Quantitative Research Methods: Week 7

Chi-Square, t-Tests, and Seminar Exercises

James Rice

Thursday, March 5, 2026

Chi-Square: what it measures

- **Chi-square** tests whether observed frequencies differ from what we would expect under the null hypothesis.
- It is used mainly with **categorical** variables.
- It is a **measure of difference**, not a measure of strength or direction.

Chi-square statistic

$$\chi^2 = \sum \frac{(f_O - f_E)^2}{f_E}$$

- f_O = observed frequency
- f_E = expected frequency
- Larger χ^2 values indicate greater departure from randomness

Chi-Square: one-way example

Subject	**Obs.**	**Exp.**	**Residual**	**Residual^2**	$\frac{(f_O - f_E)^2}{f_E}$
Chemical	120	133.33	-13.33	177.78	1.33
Electrical	130	133.33	-3.33	11.11	0.08
Mechanical	150	133.33	16.67	277.78	2.08
				Total	**3.49**

$$\chi^2 \approx 3.5 \qquad df = k - 1 = 3 - 1 = 2$$

- At the 5% level, the critical value is 5.99.
- Our test statistic is $\chi^2 = 3.49$.
- Since $3.49 < 5.99$, we **fail to reject** H_0.

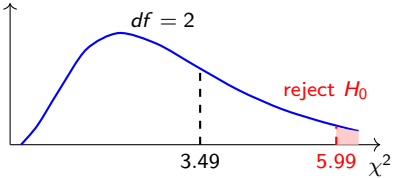

Chi-Square: two-way example in R

Create the contingency table

```
my.table <- table(pension, class)
```

Observed counts

Pension	Mgr	Int	Manual
Member	132	37	42
Not member	71	44	75

```
chisq.test(my.table)

## Pearson's Chi-squared test
## data: my.table
## X-squared = 27.2176, df = 2, p-value = 1.23e-06
```

Interpretation:

- $\chi^2 = 27.22$ with $p < 0.001$, so we reject the null that pension status and class are independent.
- Pension membership varies systematically across class groups.

t-Tests: the basic idea

What a t-test does

A t-test asks whether an observed difference in means is **large relative to the amount of random variation** we would normally expect in samples.

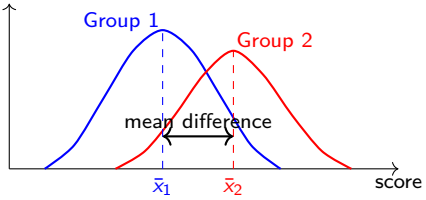

$$H_0: \ \mu_1 - \mu_2 = 0 \qquad H_1: \ \mu_1 - \mu_2 \neq 0$$

$$t = \frac{\bar{x}_1 - \bar{x}_2}{SE(\bar{x}_1 - \bar{x}_2)}$$

- $\bar{x}_1 - \bar{x}_2$: the **observed gap** between group means
- $SE(\bar{x}_1 - \bar{x}_2)$: the **typical sampling variation** in that gap
- So t is a signal-to-noise ratio:

bigger separation $\Rightarrow$ bigger $|t|$

more overlap $\Rightarrow$ smaller $|t|$

t-Tests: from statistic to decision

How we decide

Once we calculate t, we compare it with a **critical value** from the t-distribution (or use the p-value).

$$\text{If } |t| > t_{\text{critical}}, \text{ reject } H_0$$

$$\text{If } |t| \leq t_{\text{critical}}, \text{ fail to reject } H_0$$

- The centre of the distribution contains values that are plausible under H_0.
- The tails contain values that are unlikely if H_0 is true.
- A very large positive or negative t-value gives stronger evidence against H_0.

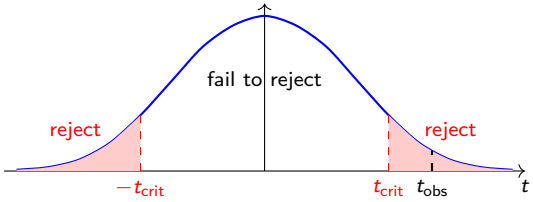

$$|t_{\text{obs}}| > t_{\text{crit}} \Rightarrow \text{reject } H_0$$

Seminar task: what you are asked to do

Using the `ONS.WB.csv` dataset:

1. Examine summary statistics for:

   `MCZ_1, MCZ_2, MCZ_3, MCZ_4, rsex, agex, Ethnicity2r`

2. Present descriptive statistics using:
   - tables,
   - histograms,
   - bar charts,
   - and/or density plots
3. Recode non-response values (e.g. `Refusal`) to missing
4. Compare mean well-being scores for **White** and **Other** respondents using `describeBy()` to compare group means

Seminar task: inferential part

The seminar then asks you to move from description to inference:

5. State hypotheses for `MCZ_1` (life satisfaction)
6. Run an **independent samples t-test**:

```
t.test(MCZ_1 ~ Ethnicity2r, data = ONS)
```

7. Report:
   - the test statistic,
   - the p-value,
   - the 95% confidence interval,
   - and a short interpretation

Seminar answers: Q1 loading the data

Step 1: read in the dataset

```
ONS <- read.csv("ONS.WB.csv")
```

Step 2: inspect the main well-being variables

```
summary(ONS$MCZ_1)
summary(ONS$MCZ_2)
summary(ONS$MCZ_3)
summary(ONS$MCZ_4)
```

Output for life satisfaction (`MCZ_1`)

```
## Min. 1st Qu. Median Mean 3rd Qu. Max.
## 0.00 7.00 8.00 7.38 9.00 10.00
```

Interpretation: life satisfaction is generally quite high in the sample, with most responses clustered toward the upper end of the 0-10 scale.

Seminar answers: Q1 tables and plots

Frequency tables for key categorical variables

```
table(ONS$Ethnicity2r)
table(ONS$rsex)
```

Example output

```
## round(prop.table(table(ONS$Ethnicity2r)),3)
##
## Other Refusal White
## 0.092 0.001 0.907

## round(prop.table(table(ONS$rsex)), 3)
##
## Female Male
## 0.542 0.458
```

Histograms

Histograms for the well-being measures

```
plot1 <- ggplot(ONS, aes(x = MCZ_1)) + geom_histogram()
plot2 <- ggplot(ONS, aes(x = MCZ_2)) + geom_histogram()
plot3 <- ggplot(ONS, aes(x = MCZ_3)) + geom_histogram()
plot4 <- ggplot(ONS, aes(x = MCZ_4)) + geom_histogram()
grid.arrange(plot1, plot2, plot3, plot4, ncol=2)
```

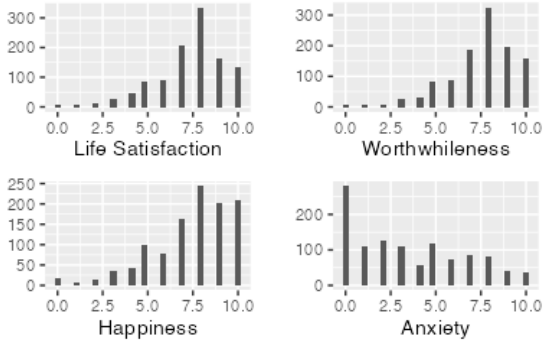

Seminar answers: Q3 cleaning the ethnicity variable

Step 1: recode refusal as missing

```
ONS$Ethnicity2r[ONS$Ethnicity2r == "Refusal"] <- NA
```

Step 2: make sure the variable is treated as a factor

```
ONS$Ethnicity2r <- factor(ONS$Ethnicity2r)
```

Why do this?

- `Refusal` is not a meaningful ethnic category for comparison.
- Coding it as `NA` removes it from the analysis.
- Converting to a factor ensures R treats the variable as categorical.

Seminar answers: Q4 descriptives by ethnicity

Compare group means for well-being

```
describeBy(ONS$MCZ_1, ONS$Ethnicity2r)
describeBy(ONS$MCZ_4, ONS$Ethnicity2r)
```

Output for `MCZ_1` (life satisfaction)

```
Descriptive statistics for life satisfaction by group
group: Other
   vars n mean sd median trimmed mad
X1 1 102 6.52 2.18 7 6.67 1.48
   min max range skew kurtosis se
X1 1 10 9 -0.64 -0.22 0.22
--------------------------------
group: White
   vars n mean sd median trimmed mad
X1 1 1010 7.47 1.88 8 7.65 1.48
   min max range skew kurtosis se
X1 0 10 10 -1.08 1.62 0.06
```

Interpretation: White respondents score higher on life satisfaction.

Seminar answers: Q5 hypotheses

For `MCZ_1` (life satisfaction), the model answer should clearly state:

Null hypothesis

$$H_0 : \mu_{\text{White}} - \mu_{\text{Other}} = 0$$

There is no population difference in mean life satisfaction between White and Other ethnic groups.

Alternative hypothesis

$$H_1 : \mu_{\text{White}} - \mu_{\text{Other}} \neq 0$$

There is a population difference in mean life satisfaction between White and Other ethnic groups.

This is a **two-tailed independent samples t-test**.

Seminar answers: Q6-Q7 t-test output + interpretation

```
t.test(MCZ_1 ~ Ethnicity2r, data = ONS)

    Welch Two Sample t-test

data: MCZ_1 by Ethnicity2r
t = -4.242, df = 116.6, p-value =
4.464e-05
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
 -1.3946 -0.5068
sample estimates:
mean in group Other mean in group White
            6.5196 7.4703
```

Interpretation:

- Since $p = 4.46 \times 10^{-5} < 0.05$, we reject H_0.
- The confidence interval does not include 0, which supports the same conclusion.
- White respondents have a higher mean life satisfaction score in this sample.
- The estimated mean gap is about 0.95 points on the 0–10 scale.