Analysis of Covariance

Statistics II (LIX002X05)

University of Groningen, Faculty of Arts, Information Science
Wilbert Heeringa

Introduction

- Analysis of covariance (ANCOVA) blends ANOVA and regression analysis.
- Evaluates whether population means of a dependent variable are equal across levels of a categorical independent variable, while controlling for the effects of other continuous variables.
- Continuous variables are called **covariates**: variables that have the potential to be related to the dependent variable; a covariate is not of primary interest.
- ANCOVA tests whether the independent variable still influences the dependent variable after the influence of the covariate(s) has been removed.
- Intuitively, ANCOVA can be thought of as 'adjusting' the dependent variable by the group means of the covariate(s).

Why including covariates?

- To reduce within-group error variance (residual variance):
  the larger the amount of variability that is explained in terms of other variables (covariates), the more the error variance is reduced, the more accurately the effect of the independent variable is assessed.

- To eliminate confound:
  if any variables are known to influence the dependent variable being measured, then ANCOVA will remove the bias of that variable when the confounding variable is entered as a covariate in the analysis.

Hypotheses

- Are there differences in level between groups given the covariate(s)?

H_0: $\mu_1 = \mu_2 = \ldots = \mu_I$
after controlling for the covariatie

H_a: not all of the μ_i are equal
after controlling for the covariatie

Assumptions

- 1. Independence:
  observations are independent of each other.
- 2. Interval scale:
  the dependent variable is measured on at least an interval scale.
- 3. Normality:
  the residuals are normally distributed. Use normal quantile plots and the Shapiro-Wilk test.
- 4. Homogeneity of variance:
  the groups have the same variance. Use Levene's test and Hartley's test.
- 5. Dissociation:
  the covariate and the factor variable are independent of each other.
- The covariate does not overlap with the effect of the factor variable(s). They each explain a different part of the variance in the dependent variable.

Assumptions

- In order to test this assumption, run an ANOVA with the factor variable(s) as independent variable(s), and the potential covariate as dependent variable.
- When the result of the ANOVA test is not significant, there is no dependency between the independent variable and the covariate.
- 6. Homogeneity of regression slopes:
  The dependent variable and any covariate(s) have the same slopes across all levels of the categorical grouping variable (factor).
- In SPSS run the ANCOVA with a model which includes the interaction between the grouping variable and the covariate.
- A scatter plot of the covariate (x-axis) and the dependent variable (y-axis) by factor group should show that all lines have a similar slope.

Variables

- Experiment Van Bezooijen & Heeringa (2006): measure intuitions of non-linguists about dialects in the Netherlands and Flanders.
- Task: rate the dialect distance compared to standard Dutch per province in a map: 0=no distance, 100=maximal distance.
- 140 Dutch subjects were involved in the experiment.

Average intuitive linguistic distances per province compared to standard Dutch.

Example

- Some dialects are recognized as minority languages by the European Union.
- Frisian (province of Friesland), Low Saxon (province of Groningen, Drenthe, Overijssel), Limburgish (province of Limburg).
- Will recognized minority languages have larger intuitive linguistic distances (compared to standard Dutch) then dialects that are not recognized?
- We need to control for objective linguistic distance, in our example measured as pronunciation distance.

Example

- We use ANCOVA with:
  - Factor:
    recognized
  - Covariate:
    pronunciation difference
  - Dependent variable:
    intuitive linguistic distance

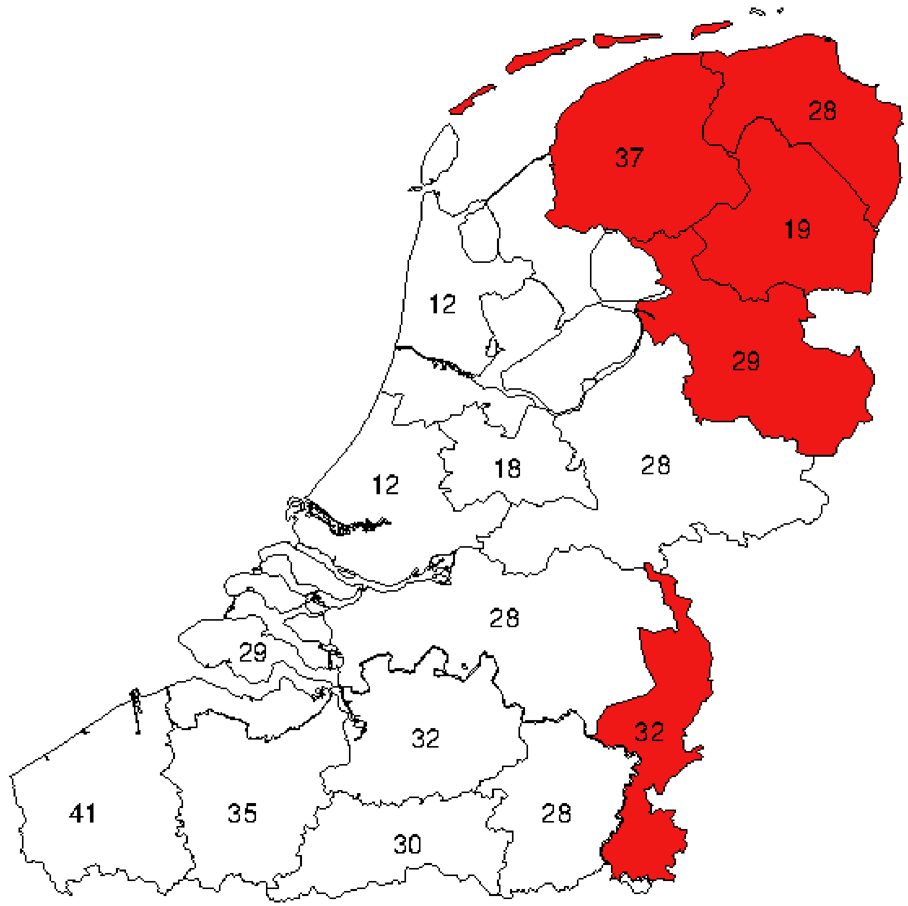

Dialect areas that are recognized as minority languages are shown in red. Average pronunciation differences are shown in the provinces.

Hypotheses

- Are there differences in level between groups given the covariate(s)?

H_0: $\mu_{recognized} = \mu_{not\ recognized}$
after controlling for pronunciation distance

H_a: $\mu_{recognized} \neq \mu_{not\ recognized}$
after controlling for pronunciation distance

1. Independence

- All the values of the dependent variable *intuitive linguistic distance* are independent of each other.

2. Interval scale

- The values of the dependent variable *intuitive linguistic distance* are measured on at least the interval scale (namely the ratio scale).

3. Normality

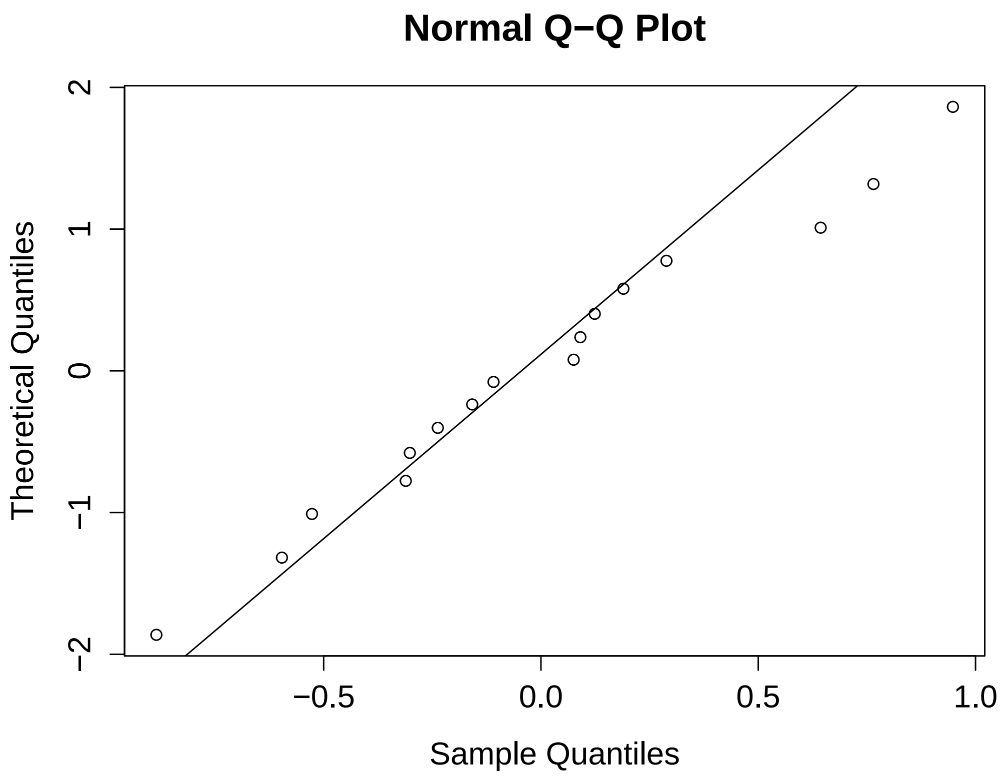

Normal quantile plot of the residuals.

3. Normality

- Normality of residues tested with the Shapiro-Wilk test:

```
Shapiro-Wilk normality test
W = 0.9761, p-value = 0.9254
```

- Distribution of residuals does not significantly differ from a normal distribution, therefore we may assume that the residuals are normally distributed.

4. Homogeneity of variance

- Results of Levene's test (generated when doing the ANCOVA test):

```
Levene's Test for Homogeneity of Variance (center = median)
      Df F value Pr(>F)
group  1  2.0802 0.1712
      14
```

- Variance of groups do not significantly differ from each other.

4. Homogeneity of variance

- Hartley's test: we look at the variances:

	variance	n
non-recognized	270.8727	11
recognized	126.3	5

- Largest / smallest = 270.8727/126.3 = 2.144677
- There are $k = 2$ groups, the smallest group has $n = 5$ observations. Given k, n and $\alpha = 0.05$ the critical value is 9.6 (see the table at `http://www.csulb.edu/~acarter3/course-biostats/tables/table-Fmax-values.pdf`).
- Since $2.144677 < 9.6$ Hartley's test confirms that the variances are the same across the levels of *recognized*.

5. Dissociation

- We perform a one-factor ANOVA with *pronunciation difference* as dependent variable and *recognized* as factor:

```
Anova Table (Type III tests)

Response: pronunciation_difference
             Sum Sq Df F value Pr(>F)
(Intercept)  0.0879  1  0.0836 0.7767
recognized   0.2814  1  0.2676 0.6130
Residuals   14.7186 14
```

- $p >> 0.05$, *pronunciation difference* and *recognized* are independent.

6. Homogeneity of regression slopes

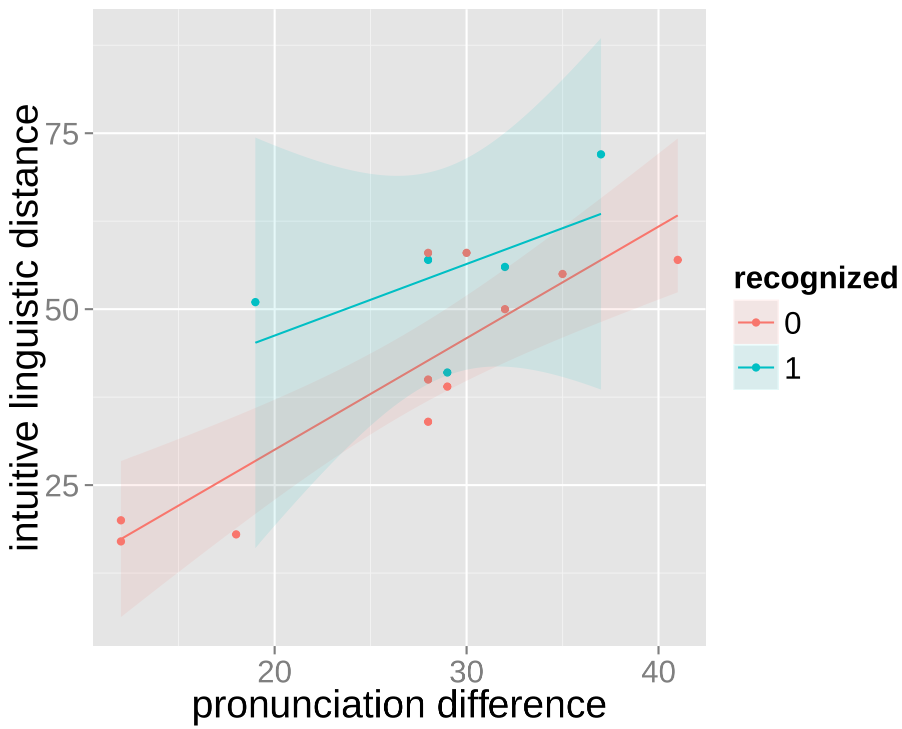

The dependent variable and any covariate(s) should have the same slopes across all levels of the categorical grouping variable (factors).

6. Homogeneity of regression slopes

- Run the ANCOVA with a model which includes the interaction between the factor *recognized* and the covariate *pronunciation difference*.

```
Anova Table (Type III tests)

Response: intuitive_linguistic_distance
                                     Sum Sq Df F value    Pr(>F)
(Intercept)                          0.4967  1  1.6706 0.2205159
recognized                           1.7816  1  5.9917 0.0307170 *
pronunciation_difference             7.8879  1 26.5284 0.0002408 ***
recognized:pronunciation_difference 0.1757  1  0.5908 0.4569581
Residuals                            3.5681 12
```

- We find $p = 0.4569581$ which is larger than $\alpha = 0.05$. We may assume that the regression slopes are the same across all levels of *recognized*.

Running a one-way ANOVA

- What happens when running an ANOVA (i.e. the covariate is not considered)?
- Results:

```
Anova Table (Type III tests)

Response: intuitive_linguistic_distance
             Sum Sq Df F value  Pr(>F)
(Intercept)  0.8950  1  1.0325 0.32680
recognized   2.8641  1  3.3041 0.09057 .
Residuals   12.1359 14
```

Running a one-way ANOVA

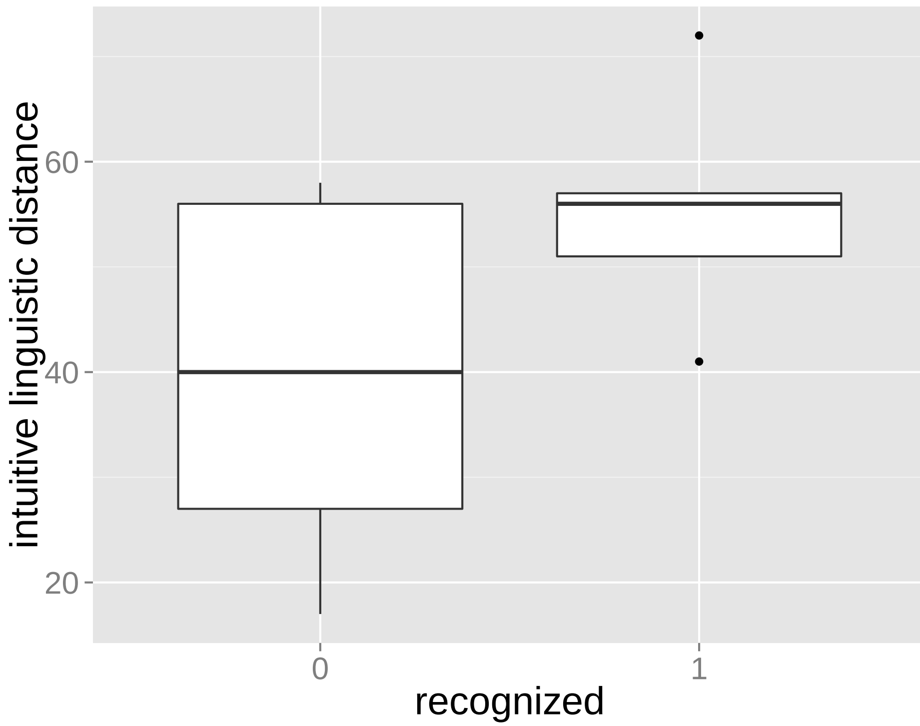

Running a simple linear regression analysis

- What happens when running a simple linear regression analysis (i.e. the factor is not considered)?
- Results:

```
Coefficients:
                          Estimate Std. Error t value Pr(>|t|)
(Intercept)              7.273e-18  1.550e-01   0.000 1.000000
pronunciation_difference 8.008e-01  1.601e-01   5.002 0.000194 ***
```

Running an ANCOVA

- Results of the ANCOVA test:

```
Anova Table (Type III tests)

Response: intuitive_linguistic_distance
                         Sum Sq Df F value    Pr(>F)
(Intercept)              0.5184  1  1.8002 0.2026587
recognized               1.6375  1  5.6863 0.0330235 *
pronunciation_difference 8.3921  1 29.1414 0.0001215 ***
Residuals                3.7437 13
```

Effect size

- In a one-factor ANOVA the **determination coefficient** is defined as:

$$R^2 = \frac{SSG}{SST}$$

- R^2 is also refered to as η^2 (eta squared)
- In ANCOVA we want to calculate the η^2 per factor.
- However, we cannot divide by SST, since SST represents the variance of all predictors together, while we want to keep the predictors separated.
- Solution: we calculate the **partial** eta squared per factor. Partial squared eta for a factor A:

$$\eta_p^2 = \frac{SS_A}{SS_A + SSE}$$

Effect size

- In our example:
  - Model:

$$R^2 = \eta^2 = \frac{SSG}{SST} = \frac{8.3921 + 1.6375}{8.3921 + 1.6375 + 3.7437} = 0.73$$

  - Recognized:

$$\eta^2_{recognized} = \frac{SS_{recognized}}{SS_{recognized} + SSE} = \frac{8.3921}{8.3921 + 3.7437} = 0.69$$

  - Pronunciation difference:

$$\eta^2_{pronunciation} = \frac{SS_{pronunciation}}{SS_{pronunciation} + SSE} = \frac{1.6375}{1.6375 + 3.7437} = 0.30$$

Effect size

- Rule of thumb:

$$\begin{array}{rcll} 0.01 \leq & \eta_p^2 & < 0.06 & \text{small effect} \\ 0.06 \leq & \eta_p^2 & < 0.14 & \text{medium effect} \\ & \eta_p^2 & \geq 0.14 & \text{large effect} \end{array}$$

- Large effects are found for both predictors. The distinction according to *recognition* explains 69% of the variation of the measurements, and the variation in *pronunciation difference* explains 30% of the variation in the measurements.

Contrasts and multiple comparisons

- In SPSS it is not possible to specify contrasts in the way we did for one-way ANOVA's. However, **multiple comparisons** can be carried out.
- When doing multiple comparisons, choose the Bonferroni or Sidak correction. Bonferroni is more conservative than Sidak.