

NMSA407 Linear Regression

## Lecture Notes

Michal Kulich

Last modified on September 9, 2025.



**matfyz**

Department of Probability and Mathematical Statistics  
Faculty of Mathematics and Physics, Charles University

*These lecture notes contain the whole contents of the course “NMSA407 Linear Regression”, a compulsory course of the curriculum in the Master’s program “Probability, Mathematical Statistics and Econometrics” at the Faculty of Mathematics and Physics, Charles University.*

*This document undergoes continuing development. The author will appreciate notifications by the reader of potential typos or misprints.*

*Unauthorized use of any part of this document is prohibited.*

Michal Kulich  
kulich@karlin.mff.cuni.cz

In Karlín on September 9, 2025

# Contents

<b>1. Simple Linear Regression: Technical and Historical Review</b>	<b>5</b>
<b>2. Linear Regression Model</b>	<b>12</b>
2.1. Definition and Assumptions . . . . .	12
2.2. Interpretation of Regression Coefficients . . . . .	16
2.3. Least Squares Estimation . . . . .	17
2.4. Residual Sum of Squares . . . . .	22
2.5. Equivalence of Regression Models . . . . .	22
2.6. Model for iid Response . . . . .	22
2.7. Model With Centered Covariates . . . . .	23
2.8. Relationship to Sample Covariance Matrices . . . . .	25
2.9. Decomposition of Sums of Squares . . . . .	25
2.10. Coefficient of Determination . . . . .	27
2.11. LSE Under Linear Restrictions . . . . .	28
<b>3. Properties of the Least Squares Estimator</b>	<b>31</b>
3.1. Moment Properties of the Least Squares Estimator . . . . .	31
3.2. Gauss-Markov Theorem . . . . .	33
3.3. Properties of the Least Squares Estimator Under Normality . . . . .	34
<b>4. Statistical Inference in the Linear Regression Model</b>	<b>37</b>
4.1. Exact Inference Under Normality . . . . .	37
4.1.1. Testing individual regression parameters . . . . .	38
4.1.2. Confidence intervals for individual regression parameters . . . . .	38
4.1.3. Tests and confidence intervals for linear combinations of regression parameters . . . . .	38
4.1.4. Simultaneous tests of several linear combinations of regression parameters . . . . .	38
4.1.5. Submodel testing . . . . .	40
4.1.6. Overall regression test . . . . .	41
4.1.7. One-way analysis of variance model . . . . .	42
4.1.8. Connections to maximum likelihood theory . . . . .	43
4.2. Asymptotic Inference Without Normality (Random Covariates) . . . . .	44
4.3. Asymptotic Inference Without Normality (Fixed Covariates) . . . . .	50

<b>5. Predictions</b>	<b>52</b>
5.1. Predictions and Their Pitfalls . . . . .	52
5.2. Confidence Intervals for Conditional Expectations . . . . .	54
5.3. Prediction Intervals for Future Responses . . . . .	55
<b>6. Diagnostic Methods Based on Residuals</b>	<b>57</b>
<b>7. Transformations of the Response</b>	<b>70</b>
7.1. Introduction . . . . .	70
7.2. Interpretation of Linear Model With Log-transformed Response . . . . .	71
7.3. The Box-Cox Transformation . . . . .	72
<b>8. Parametrization of a Single Covariate</b>	<b>74</b>
8.1. Parametrization of Categorical Covariates (Factors) . . . . .	74
8.1.1. Reference group parametrization . . . . .	75
8.1.2. Zero-sum parametrization . . . . .	76
8.1.3. Weighted zero-sum parametrization . . . . .	77
8.1.4. Helmert parametrization . . . . .	77
8.1.5. Orthogonal polynomials for ordered factors . . . . .	77
8.2. Parametrization of Quantitative (Numerical) Covariates . . . . .	79
8.2.1. Factorization of a numerical covariate . . . . .	79
8.2.2. Polynomial basis . . . . .	80
8.2.3. Piecewise polynomial bases, B-splines . . . . .	81
8.2.4. Sine-cosine basis . . . . .	88
<b>9. Multiple Comparisons and Simultaneous Confidence Intervals</b>	<b>89</b>
9.1. Bonferroni Procedure . . . . .	89
9.2. Tukey Method . . . . .	93
9.3. Scheffé Method . . . . .	95
9.4. Simultaneous Confidence Bounds . . . . .	96
<b>Notation</b>	<b>97</b>
<b>Bibliography</b>	<b>100</b>
<b>Index</b>	<b>101</b>
<b>A. Appendix</b>	<b>102</b>

# 1. Simple Linear Regression: Technical and Historical Review

Consider  $n$  measurements of continuous variables  $(x_i, y_i)$  for  $i = 1, \dots, n$ . Plot them as Cartesian coordinates on a scatterplot (Figure 1.1). The observations seem to be located along a line; there is a perceived linear relationship between the values of  $x$  and  $y$ , but not an exact one. The goal is to identify a line passing through the observations (see Figure 1.2) so that the line is “optimal” in some way.

**Legendre** (1805) proposed to find the line by minimizing the sum of squared vertical distances of the observed points from the fitted line (see Figure 1.3). This is called *the least squares method*.<sup>\*</sup> It can be also attributed to Gauss, who later claimed (**Gauss** 1821) that he had been using the method as early as in 1795 but had not published it.

**Adrien-Marie Legendre** (1752 – 1833) was a French mathematician who made numerous contributions to mathematics. Well-known and important concepts such as the Legendre polynomials and Legendre transformation are named after him.

Source: [https://en.wikipedia.org/wiki/Adrien-Marie\\_Legendre](https://en.wikipedia.org/wiki/Adrien-Marie_Legendre)

---

<sup>\*</sup> Český Metoda nejmenších čtverců.




Figure 1.1.: Scatterplot of two continuous variables in  $\mathbb{R}^2$ .




Figure 1.2.: Scatterplot of two continuous variables in  $\mathbb{R}^2$  with fitted line.

The least squares method is based on the presumption that the observed values of the variable  $x_i$  are measured precisely while  $y_i$  are measured with an error that shifts them away from the line that expresses the linear relationship between the two variables. This point of view justifies the minimization of vertical distances instead of e.g. perpendicular distances.

**Johann Carl Friedrich Gauss** (1777 – 1855) was a German mathematician, geodesist, and physicist who made significant contributions to many fields in mathematics and science. Gauss published the second and third complete proofs of the fundamental theorem of algebra, made important contributions to number theory and developed the theories of binary and ternary quadratic forms. He is also credited with inventing the fast Fourier transform algorithm and was instrumental in the discovery of the dwarf planet Ceres. His work on the motion of planetoids disturbed by large planets led to the introduction of the Gaussian gravitational constant and the method of least squares, which is still used in all sciences to minimize measurement error.

Source: [https://en.wikipedia.org/wiki/Carl\\_Friedrich\\_Gauss](https://en.wikipedia.org/wiki/Carl_Friedrich_Gauss)

Let us show how the idea of Legendre and Gauss works. Consider a line  $y = a + bx$  and choose  $a, b$  so that

$$SS(a, b) = \sum_{i=1}^n (y_i - a - bx_i)^2 \quad (1.1)$$

is minimized over all  $a, b \in \mathbb{R}$ . The sum in the expression (1.1) is called *the sum of squares*.\*

---

\* Český Součet čtverců.




Figure 1.3.: Zoomed subset of data from Figure 1.2 with visualized vertical distances of the points from the line (blue).

The values  $a, b$  that minimize the sum of squares are easy to find:

$$\frac{\partial SS(a, b)}{\partial a} = 2 \sum_{i=1}^n (y_i - a - bx_i)(-1),$$

$$\frac{\partial SS(a, b)}{\partial b} = 2 \sum_{i=1}^n (y_i - a - bx_i)(-x_i).$$

Thus,  $a$  and  $b$  are the solutions to the system of two equations

$$\sum_{i=1}^n y_i - na - b \sum_{i=1}^n x_i = 0,$$

$$\sum_{i=1}^n x_i y_i - a \sum_{i=1}^n x_i - b \sum_{i=1}^n x_i^2 = 0.$$

These equations are called the *normal equations*<sup>\*</sup>.

Introducing the notation  $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i$  and  $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$ , the normal equations can be solved as follows. From the first equation, we get

$$na = \sum_{i=1}^n y_i - b \sum_{i=1}^n x_i, \quad \text{hence} \quad a = \bar{y} - b\bar{x}.$$

This shows that the fitted line passes through the point  $(\bar{x}, \bar{y})$ . Next, substituting in the

---

<sup>\*</sup> Český Normální rovnice.

second equation for the optimal intercept  $a$ , we get

$$\begin{aligned} b \frac{1}{n} \sum_{i=1}^n x_i^2 &= \frac{1}{n} \sum_{i=1}^n x_i y_i - a \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i y_i - \bar{x} \bar{y} + b \bar{x}^2 \\ b \left( \frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x}^2 \right) &= \frac{1}{n} \sum_{i=1}^n x_i y_i - \bar{x} \bar{y} \\ b \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 &= \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y}) \end{aligned}$$

Finally,

$$b = \frac{\frac{1}{n} \sum_{i=1}^n x_i y_i - \bar{x} \bar{y}}{\frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x}^2} = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n (x_i - \bar{x})^2}.$$

The former version is more computationally friendly, the latter version provides an insight into the meaning of the slope  $b$ . Indeed,

$$b = \frac{\widehat{\text{cov}}(x, y)}{\widehat{\text{var}}(x)} = r_{xy} \sqrt{\frac{\widehat{\text{var}}(y)}{\widehat{\text{var}}(x)}},$$

where  $\widehat{\text{cov}}(x, y)$  is the sample covariance of the observations  $(x_i, y_i)$ ,  $\widehat{\text{var}}(x)$  is the sample variance of  $x_i$ ,  $\widehat{\text{var}}(y)$  is the sample variance of  $y_i$ , and  $r_{xy}$  is the sample correlation coefficient of the observations  $(x_i, y_i)$ .

If the observations  $x_i$  have the same sample variance as  $y_i$  then the slope of the line fitted by least squares is equal to the sample correlation coefficient  $r_{xy}$  and therefore lies in the interval  $(-1, 1)$ .

This phenomenon was noticed by sir Francis Galton (Galton 1886). He investigated the relationship of the parents' height with the height of their grown children. The recorded heights (in inches) are shown in Figure 1.4 and Galton's original visualization of the data in Figure 1.5. If we focus on the heights of sons only (to eliminate the fact that daughters are somewhat shorter) and plot them as  $y_i$  against the average height of their parents ( $x_i$ ) we obtain the scatterplot shown in Figure 1.6.

**Sir Francis Galton** (1822–1911) Darwin's cousin, prodigy child, contributor to the fields of statistics, meteorology, psychology, genetics, co-founder and proponent of eugenics.  
Source: [https://en.wikipedia.org/wiki/Francis\\_Galton](https://en.wikipedia.org/wiki/Francis_Galton)

The red line in Figure 1.6 was fitted by the method of least squares and its slope is about 0.74.\* As explained above, this value corresponds to the sample correlation between the average height of the parents and the height of their son. It means that if the average height of the parents exceeds the population mean by 10 cm the son's height is likely to be above average as well, but only by some 7.4 cm. So, tall parents tend to have tall sons, but

---

\* Galton used a different data set and estimated the slope of the fitted line to be about 0.66.



## 1. Simple Linear Regression: Technical and Historical Review

FAMILY HEIGHTS from R.F.F.  
(add 66 inches to every entry in the Table)

	Father	Mother	Sons in order of height	Daughters in order of height
1	12.5	7.0	13.2	9.2, 9.0, 9.0
2	15.5	6.5	13.5, 12.5	5.5, 5.5
3	15.0	about 4.0	11.0	8.0
4	15.0	4.0	10.5, 8.5	7.0, 4.5, 3.0
5	15.0	1.5	12.0, 9.0, 8.0	6.5, 2.5, 2.5
6	14.0	8.0		9.5
7	14.0	8.0	16.5, 14.0, 13.0, 13.0	10.5, 4.0
8	14.0	6.5		10.5, 8.0, 6.0
9	14.5	6.0		6.0
10	14.0	8.5		5.5
11	14.0	2.0	14.0, 10.0	8.0, 7.0, 7.0, 6.0, 3.5, 3.0
12	14.0	1.0		5.0
13	13.0	7.0	11.0	2.0
14	13.0	7.0	8.0, 7.0	
15	13.0	6.5	11.0, 10.5	6.7
16	13.0	about 5.0	12.0, 10.5, 10.2, 10.2, 9.2	8.7, 6.5, 4.5, 3.5
17	13.0	4.5	14.0, 13.0, 11.5, 2.5	6.5, 2.5
18	13.0	4.0		6.0, 4.5, 4.0
19	13.2	3.0		2.7
20	12.7	4.0	13.2, 13.0, 12.7	10.0, 9.0, 8.5, 8.0, 6.0
21	12.0	8.0	13.0	8.5, 8.0
22	12.0	abt 7.0	13.0, 11.0	7.0
23	12.0	5.0	14.2, 10.5, 9.5	6.0, 5.5, 5.0, 4.0
24	12.0	5.5		5.5

Figure 1.4.: Galton height data: original pen/paper records.

Source: <http://www.medicine.mcgill.ca/epidemiology/hanley/galton/>

not as tall as the parents were. Galton called this feature *regression towards the mean*. Even though the term *regression*\* originally referred to this very specific feature that appears only in certain data sets, it began to be used more generally to describe methods and techniques used for fitting lines or curves to observed data.

The least squares method can be easily extended to fit certain non-linear relationships between the two variables. For example, if the relationship is not linear but quadratic we could use the same idea with the function

$$y_i = a + bx_i + cx_i^2.$$

We could find  $a$ ,  $b$ , and  $c$  by the method of least squares by minimizing

$$SS(a, b, c) = \sum_{i=1}^n (y_i - a - bx_i - cx_i^2)^2.$$

The estimated parameters  $a$ ,  $b$ , and  $c$  are obtained by solving a system of three linear equations.

In this introductory chapter, we approached the problem of fitting a line or a curve through a cloud of bivariate data. We did not introduce any underlying probabilistic model

\* Český Regrese.

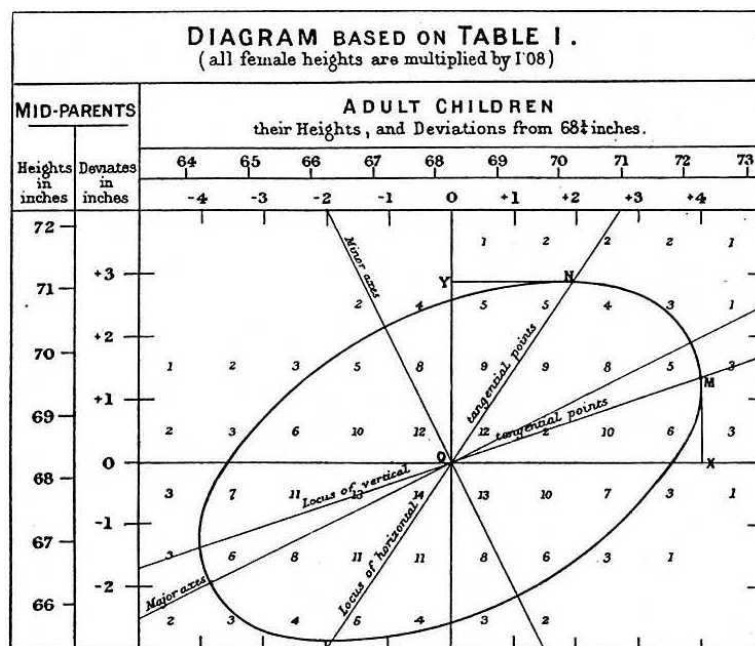


Figure 1.5.: Galton height data: original visualization by the author.

Source: <https://en.wikipedia.org>

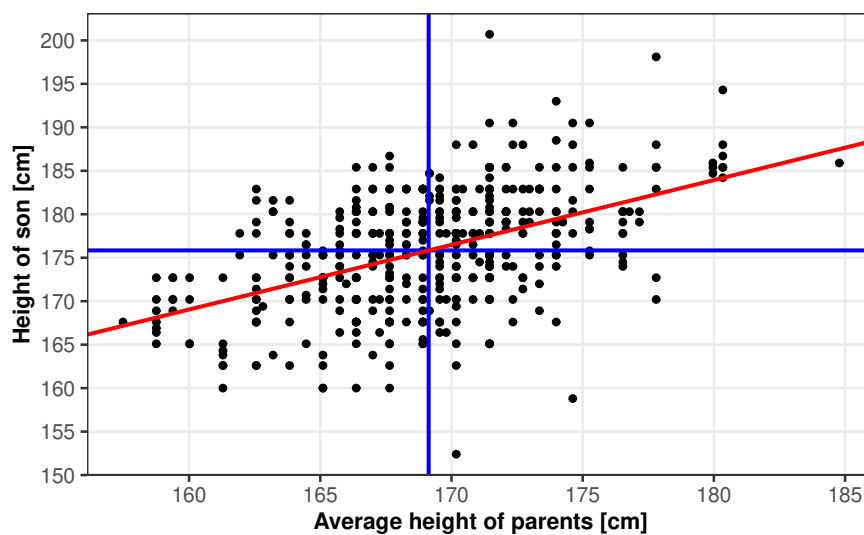


Figure 1.6.: Modified Galton data with fitted least squares line (red). The slope of the line is  $\approx 0.74$ . The means of the two variables are plotted as blue lines.

for the data, did not formulate any assumptions and were not able to find neither an interpretation for the estimates obtained by the least square method nor to investigate their theoretical properties.

## 2. Linear Regression Model

In this chapter, we formulate a general definition of the linear regression model. We explain the meaning of the regression parameters and derive a general formula for the least squares estimator. We introduce a lot of new technical terms, explain their meaning and investigate some features of linear regression models that will be important for the developments presented in subsequent chapters.

### 2.1. Definition and Assumptions

Consider a sequence of  $n$  independent random vectors  $(Y_i, \mathbf{X}_i)$ ,  $i = 1, \dots, n$ . The random variable  $Y_i$  is called *the response*<sup>\*</sup> (also *the dependent variable*<sup>†</sup>, *the outcome*). The random vector  $\mathbf{X}_i$  contains  $p < n$  components  $\mathbf{X}_i = (X_{i1}, \dots, X_{ip})^\top$  which are called *the covariates* (also explanatory variables, predictors, regressors)<sup>‡</sup>.

**Definition 2.1.** The independent observations  $(Y_i, \mathbf{X}_i)$  satisfy *the linear regression model* if the response  $Y_i$  can be written as  $Y_i = \mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i$ , that is,

$$Y_i = \beta_1 X_{i1} + \beta_2 X_{i2} + \dots + \beta_p X_{ip} + \varepsilon_i,$$

where  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^\top$  is a vector of unknown *regression parameters (coefficients)*<sup>§</sup> and *the error terms*<sup>¶</sup>  $\varepsilon_1, \dots, \varepsilon_n$  are independent random variables such that  $E[\varepsilon_i | \mathbf{X}_i] = 0$ , and  $\text{var}[\varepsilon_i | \mathbf{X}_i] = \sigma_\varepsilon^2$ .  $\nabla$

**Note.** On the covariates:

- The first covariate  $X_{i1}$  is usually taken as 1.
- The covariates  $\mathbf{X}_i$  are often created by a transformation of an originally observed random vector  $\mathbf{Z}_i$ . We suppress this in the notation.
- In certain applications, the covariates are fixed quantities rather than random variables. Because the definition of the linear model only specifies conditional moments given the observed values of the covariates it applies to fixed covariates as well. Most of the developments that follow in this course are not sensitive to differences between fixed and random covariates either. The only occasion when fixed covariates need to be treated differently than random covariates is the investigation of asymptotic properties. This will be discussed in Section ??.

---

\* Český odezva    † Český závislá proměnná    ‡ Český regresory, nezávisle proměnné, vysvětlující veličiny, prediktory, kovariáty    § Český regresní koeficienty    ¶ Český chybové členy

**Note.** On the error terms:

- The random variables  $\varepsilon_i$  are required to have zero means and equal variances. It is somewhat misleading to call them *error terms* because they include not only errors in the measurement of the response but also the effects of any factors that influence the mean of the response and are not included in the model. In econometrics, the error terms are often called *disturbances*.
- The variance  $\sigma_e^2$  of the error terms is called *the residual variance*<sup>\*</sup>.
- Sometimes, the assumptions on the error terms are strengthened to require that  $\varepsilon_i$  be independent of  $\mathbf{X}_i$ . Our definition does not require this.

The definition of the linear model can be reformulated in terms of conditional moments of the response as follows:

- $E[Y_i | \mathbf{X}_i] = \mathbf{X}_i^\top \boldsymbol{\beta} = \beta_1 X_{i1} + \beta_2 X_{i2} + \cdots + \beta_p X_{ip},$
- $\text{var}[Y_i | \mathbf{X}_i] = \sigma_e^2.$

Thus, the model makes assumptions about the first two conditional moments of the response: the conditional mean must be linear in  $\mathbf{X}_i$  through  $\boldsymbol{\beta}$  and the conditional variance must be constant.

The purpose of the linear regression model is not just to fit a line, curve or surface through a cloud of data as it was presented in Chapter 1. Instead, we aim to express how the expected value of the response  $Y_i$  changes with different values of  $\mathbf{X}_i$  and tell what influence the individual covariates have on the expectation.

**Notation.** Let

$$\mathbf{Y} = \begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{pmatrix}, \quad \mathbb{X} = \begin{pmatrix} \mathbf{X}_1^\top \\ \mathbf{X}_2^\top \\ \vdots \\ \mathbf{X}_n^\top \end{pmatrix}, \quad \boldsymbol{\varepsilon} = \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{pmatrix}.$$

The  $n$  by  $p$  matrix  $\mathbb{X}$  is called *the regression matrix*<sup>†</sup>. It includes the observed covariate vectors in the rows.

Now we can express the model for all the data together

$$\mathbf{Y} = \mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

with  $E[\boldsymbol{\varepsilon} | \mathbb{X}] = \mathbf{0}$  and  $\text{var}[\boldsymbol{\varepsilon} | \mathbb{X}] = \sigma_e^2 \mathbb{I}_n$  or

- $E[\mathbf{Y} | \mathbb{X}] = \mathbb{X}\boldsymbol{\beta},$
- $\text{var}[\mathbf{Y} | \mathbb{X}] = \sigma_e^2 \mathbb{I}_n.$

**Note.** From now on, we will often use the notation  $E, \text{var}$  for the conditional expectation and variance given the covariates. So, we will write  $E Y_i$  instead of  $E[Y_i | \mathbf{X}_i]$  and  $\text{var} Y_i$  instead of  $\text{var}[Y_i | \mathbf{X}_i]$ ; similarly for  $E \varepsilon_i, \text{var} \varepsilon_i, E \mathbf{Y}, \text{var} \mathbf{Y}$  etc.

---

<sup>\*</sup> Český residuální rozptyl    <sup>†</sup> Český regresní matice




Figure 2.1.: Two sample problem expressed as a linear regression model  $EY = \beta_1 + \beta_2 Z$ , where  $Z = \mathbb{1}(G)$ . The regression line has no interpretation except at  $Z = 0$  or  $Z = 1$ .

**Example 2.1 (Linear model for iid data).** Suppose the responses  $Y_1, \dots, Y_n$  represent a random sample of independent identically distributed random variables with  $EY_i = \mu$  and  $\text{var } Y_i = \sigma^2$ . Then

$$Y_i = \mu + \varepsilon_i,$$

where  $\varepsilon_i$ ,  $i = 1, \dots, n$  are iid with zero mean and variance  $\sigma^2$ . Thus,  $Y_i$  satisfies a linear regression model with  $X_i = 1$ ,  $\beta = \mu$  and  $\sigma_e^2 = \sigma^2$ .  $\triangle$

**Example 2.2 (Simple linear regression).** Suppose we observe a random sample of  $(Y_i, Z_i)$ , where  $Z_i$  is univariate. Define the covariate vector as  $X_i = (1, Z_i)^T$ . This leads to the regression matrix

$$\mathbb{X} = \begin{pmatrix} 1 & Z_1 \\ 1 & Z_2 \\ \vdots & \vdots \\ 1 & Z_n \end{pmatrix},$$

and the simple linear regression model (recall Chapter 1)

$$Y_i = \beta_1 + \beta_2 Z_i + \varepsilon_i,$$

with  $EY_i = \beta_1 + \beta_2 Z_i$  and  $\text{var } Y_i = \sigma_e^2$ .  $\triangle$

**Example 2.3 (Two sample problem).** In the previous example, take a special case with a binary covariate  $Z_i$ , which attains only values 0 or 1. Suppose that  $Z_i$  indicates a membership of the observation in some subgroup  $G$ , that is  $Z_i = \mathbb{1}(i \in G)$ .

## 2. Linear Regression Model




Figure 2.2.: Data following a quadratic association with a fitted quadratic curve.

The simple linear regression model has the form

$$Y_i = \beta_1 + \beta_2 \mathbb{1}(i \in G) + \varepsilon_i,$$

that is,

$$\mathbb{E} Y_i = \begin{cases} \beta_1 & \text{when } i \notin G, \\ \beta_1 + \beta_2 & \text{when } i \in G, \end{cases} \quad \text{var } Y_i = \sigma_e^2.$$

This model specifies a two-sample location problem with equal variances in both groups and possibly different expectations. The regression parameter  $\beta_2$  expresses the difference in expectations between the groups.

An illustration of the two-sample location problem is provided by Figure 2.1. The regression line is shown in red color but realize that it can only be interpreted at points that actually appear in the data, that is  $Z = 1$  (group  $G$ ) or  $Z = 0$  (group  $\neg G$ ).  $\triangle$

**Example 2.4 (Quadratic regression).** Suppose we observe a random sample of  $(Y_i, Z_i)$ , where  $Z_i$  is univariate. Define the covariate vector as  $\mathbf{X}_i = (1, Z_i, Z_i^2)^\top$ . This leads to the regression matrix

$$\mathbb{X} = \begin{pmatrix} 1 & Z_1 & Z_1^2 \\ 1 & Z_2 & Z_2^2 \\ \vdots & \vdots & \vdots \\ 1 & Z_n & Z_n^2 \end{pmatrix},$$

and the quadratic regression model (recall Chapter 1)

$$Y_i = \beta_1 + \beta_2 Z_i + \beta_3 Z_i^2 + \varepsilon_i,$$

with  $\mathbb{E} Y_i = \beta_1 + \beta_2 Z_i + \beta_3 Z_i^2$  (a quadratic function of  $Z_i$ ) and  $\text{var } Y_i = \sigma_e^2$ .

An illustration of the quadratic regression model is provided by Figure 2.2.  $\triangle$

## 2.2. Interpretation of Regression Coefficients

Recall how the regression coefficients are related to the expectation of the response:

$$\mathbb{E}[Y_i | \mathbf{X}_i = (x_{i1}, \dots, x_{ip})] = \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip}.$$

Thus, the regression coefficients capture and express the influence of  $\mathbf{X}_i$  on  $\mathbb{E} Y_i$ .

Suppose that  $X_{i1} = 1 \ \forall i \in \{1, \dots, n\}$ . Then the coefficient pertaining to this covariate is called *the intercept* (or *the absolute term*<sup>\*</sup>). Obviously,

$$\beta_1 = \mathbb{E}[Y_i | X_{i2} = 0, X_{i3} = 0, \dots, X_{ip} = 0].$$

**The intercept provides the expectation of the response for an observation with zero values of all covariates (except the first).**

Next, take an observation with any covariate vector  $\mathbf{x} = (1, x_2, \dots, x_p)$  and denote the  $j$ -th unit vector of dimension  $p$  by  $\mathbf{e}_j = (0, \dots, 0, 1, 0, \dots, 0)^T$  with 1 at the  $j$ -th position ( $j = 2, \dots, p$ ). We have

$$\mathbb{E}[Y_i | \mathbf{X}_i = \mathbf{x}] = \beta_1 + \beta_2 x_2 + \dots + \beta_p x_p$$

and

$$\mathbb{E}[Y_i | \mathbf{X}_i = \mathbf{x} + \mathbf{e}_j] = \beta_1 + \beta_2 x_2 + \dots + \beta_j(x_j + 1) + \dots + \beta_p x_p.$$

After subtracting the top equation from the bottom one, we get

$$\beta_j = \mathbb{E}[Y_i | \mathbf{X}_i = \mathbf{x} + \mathbf{e}_j] - \mathbb{E}[Y_i | \mathbf{X}_i = \mathbf{x}], \quad j = 2, \dots, p.$$

So,  $\beta_j$  expresses the increase in  $\mathbb{E} Y_i$  after the  $j$ -th covariate is increased by one unit and all other covariates stay the same.<sup>†</sup>

It is important to realize that these interpretations do not always make sense.

Obviously, the intercept cannot be interpreted if an observation with all covariates equal to zero does not exist.

In quadratic regression  $\mathbb{E}[Y_i | Z_i] = \beta_1 + \beta_2 Z_i + \beta_3 Z_i^2$ , with  $X_{i2} = Z_i$  and  $X_{i3} = Z_i^2$ , one cannot increase  $X_{i2}$  by a single unit while keeping  $X_{i3}$  the same and vice versa. So,  $\beta_2$  and  $\beta_3$  cannot be interpreted either. This is because in this model a single variable  $Z_i$  affects the values of several covariates simultaneously.

Another cautionary note applies to interpretation of the absolute value of  $\beta_j$ . It is not true that a covariate with a very large value of  $\beta_j$  affects the response more strongly than a covariate with a parameter close to zero. The strength of the influence of the covariate also depends on the units of measurement. By rescaling all values of  $X_{ij}$  to  $mX_{ij}$ , the coefficient  $\beta_j$  is made  $m$ -times smaller because  $\beta_j X_{ij} = \frac{\beta_j}{m} \cdot mX_{ij}$ . Thus, rescaling a measurement made in kilometers into meters makes the regression coefficient 1000 times smaller without changing anything about the strength of the influence of that covariate on the response.

<sup>\*</sup> Český *absolutní člen*    <sup>†</sup> Of course, if  $\beta_j < 0$ , it expresses a decrease in the expectation.



### 2.3. Least Squares Estimation

#### Definition of the least squares estimator

Consider the model

$$Y = \mathbb{X}\beta + \varepsilon$$

with  $E\varepsilon = \mathbf{0}$  and  $\text{var } \varepsilon = \sigma_e^2 \mathbb{I}_n$ . The regression matrix  $\mathbb{X}$  has  $n$  rows and  $p$  columns, with  $p < n$ , and the dimension of  $\beta$  is  $p$ .

**Definition 2.2 (Least Squares Estimator).** The *least squares estimator* (LSE)  $\hat{\beta}$  of  $\beta$  is the point in  $\mathbb{R}^p$  that minimizes the sum of squares

$$SS_e(\beta) = \sum_{i=1}^n (Y_i - \mathbf{x}_i^\top \beta)^2 = (Y - \mathbb{X}\beta)^\top (Y - \mathbb{X}\beta) = \|Y - \mathbb{X}\beta\|^2.$$

▽

In order to make the LSE unique, we will make the following assumption.

**Assumption.** Let the regression matrix  $\mathbb{X}_{n \times p}$  be of full rank, that is,  $r(\mathbb{X}) = p$ .

If the regression matrix did not have full rank there would exist at least one covariate (a column of  $\mathbb{X}$ ) that can be expressed as a linear combination of other covariates. Under such circumstances the regression coefficients are not identifiable and the LSE  $\hat{\beta}$  does not have a unique value.

**Example 2.5.** Consider the model  $EY = \beta_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4$  and suppose that  $X_4 = X_2 + X_3$ . Then there are infinitely many values of  $\beta$  that always generate the same expectation for the response:

$$\begin{aligned} EY &= \beta_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 (X_2 + X_3) = & \beta &= (\beta_1, \beta_2, \beta_3, \beta_4)^\top \\ &= \beta_1 + (\beta_2 + \beta_4) X_2 + (\beta_3 + \beta_4) X_3 = & \beta &= (\beta_1, \beta_2 + \beta_4, \beta_3 + \beta_4, 0)^\top \\ &= \beta_1 + \left(\beta_2 + \frac{\beta_4}{2}\right) X_2 + \left(\beta_3 + \frac{\beta_4}{2}\right) X_3 + \frac{\beta_4}{2} (X_2 + X_3) & \beta &= \left(\beta_1, \beta_2 + \frac{\beta_4}{2}, \beta_3 + \frac{\beta_4}{2}, \frac{\beta_4}{2}\right)^\top \end{aligned}$$

et cetera. When the regression coefficients  $\beta$  do not have a unique value the model is called *unidentifiable*\*. △

Through the entire course, **we will avoid regression matrices that are not of full rank**. It makes little sense to deal with them because such models cannot be used in practice. We can always satisfy our assumption by dropping the columns that can be expressed as linear combination of other columns and so reducing the dimension of the model and the number of parameters  $p$  until the regression matrix has a full rank.

---

\* Český neidentifikovatelný

**Note.** One could raise an objection that we consider  $\mathbb{X}$  random and hence its rank is also a random variable. The following simple example shows that it is possible to end up with a singular regression matrix by mere bad luck.

Suppose  $EY_i = \beta_1 + \beta_2 X_i$  where  $X_i \in \{0, 1\}$  is an indicator of membership of the individual in some subgroup  $\mathcal{G}$ . The rank of the regression matrix should be equal to  $p = 2$ . Let  $P[X_i = 1] \equiv \pi \in (0, 1)$ . If  $\pi = 0$  or  $\pi = 1$ , the covariate generates the same value for all observations and the regression matrix is of rank 1. But even if we exclude these cases by requiring  $\pi \in (0, 1)$ , we still get

$$\begin{aligned} P[X_i = 1 \ \forall i \in \{1, \dots, n\}] &= \pi^n > 0 \\ P[X_i = 0 \ \forall i \in \{1, \dots, n\}] &= (1 - \pi)^n > 0, \end{aligned}$$

so for any finite sample size  $n$  there is a positive probability of  $r(\mathbb{X}) = 1$ . The probability, however, converges to zero fairly quickly as  $n$  increases.

If it happens in practice, it means that either the group  $\mathcal{G}$  or the complement  $\mathcal{G}^c$  are not represented in the data at all and we cannot estimate the effect of the group on the expectation of the response. We have no choice but to drop the indicator of the group from the model and reduce the number of columns of the regression matrix.

**Note.** In the general case, express  $\mathbf{X}_i = (1, \mathbf{X}_i^M)$  (separate the intercept from the rest of the covariates). Then it holds: If  $\text{var} \mathbf{X}_i^M > 0$  then  $P[r(\mathbb{X}) = p] \rightarrow 1$  as  $n \rightarrow \infty$ .

### Derivation of the explicit form of the LSE

Let us derive the explicit form of the least squares estimator. Decompose  $SS_e(\boldsymbol{\beta})$  into several parts.

$$SS_e(\boldsymbol{\beta}) = (\mathbf{Y} - \mathbb{X}\boldsymbol{\beta})^\top (\mathbf{Y} - \mathbb{X}\boldsymbol{\beta}) = \mathbf{Y}^\top \mathbf{Y} - \boldsymbol{\beta}^\top \mathbb{X}^\top \mathbf{Y} - \mathbf{Y}^\top \mathbb{X} \boldsymbol{\beta} + \boldsymbol{\beta}^\top \mathbb{X}^\top \mathbb{X} \boldsymbol{\beta} = \mathbf{Y}^\top \mathbf{Y} - 2\boldsymbol{\beta}^\top \mathbb{X}^\top \mathbf{Y} + \boldsymbol{\beta}^\top \mathbb{X}^\top \mathbb{X} \boldsymbol{\beta}.$$

We will use rules for matrix differentiation. In particular, for any vector  $\mathbf{c}$  and any symmetric matrix  $\mathbf{A}$

$$\frac{\partial \boldsymbol{\beta}^\top \mathbf{c}}{\partial \boldsymbol{\beta}} = \mathbf{c} \quad \text{and} \quad \frac{\partial \boldsymbol{\beta}^\top \mathbf{A} \boldsymbol{\beta}}{\partial \boldsymbol{\beta}} = 2\mathbf{A}\boldsymbol{\beta}.$$

We have,

$$\frac{\partial \boldsymbol{\beta}^\top \mathbb{X}^\top \mathbf{Y}}{\partial \boldsymbol{\beta}} = \mathbb{X}^\top \mathbf{Y} \quad \text{and} \quad \frac{\partial \boldsymbol{\beta}^\top \mathbb{X}^\top \mathbb{X} \boldsymbol{\beta}}{\partial \boldsymbol{\beta}} = 2\mathbb{X}^\top \mathbb{X} \boldsymbol{\beta},$$

and hence

$$\frac{\partial SS_e(\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} = -2\mathbb{X}^\top \mathbf{Y} + 2\mathbb{X}^\top \mathbb{X} \boldsymbol{\beta}.$$

The LSE  $\hat{\boldsymbol{\beta}}$  solves the system of  $p$  linear equations

$$\mathbb{X}^\top \mathbb{X} \hat{\boldsymbol{\beta}} = \mathbb{X}^\top \mathbf{Y}, \tag{2.1}$$

which is called *the normal equations\** in this context.

When  $\mathbb{X}$  is of rank  $p$ , as we assume,  $\mathbb{X}^\top \mathbb{X}$  is a  $p \times p$  matrix of rank  $p$  and therefore its inverse exists and is unique. It follows that the normal equations have a single solution, which is

$$\hat{\beta} = (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top \mathbf{Y}. \quad (2.2)$$

This is the explicit form of the least squares estimator in linear regression.

To show that this estimator really minimizes the least squares criterion, we calculate the Hessian matrix:

$$\frac{\partial}{\partial \beta^\top} \frac{\partial SS_e(\beta)}{\partial \beta} = \frac{\partial}{\partial \beta^\top} (-2\mathbb{X}^\top \mathbf{Y} + 2\mathbb{X}^\top \mathbb{X} \beta) = 2\mathbb{X}^\top \mathbb{X},$$

which is a positive definite matrix at any argument  $\beta \in \mathbb{R}^p$ . Thus, the function  $SS_e(\beta)$  is strictly convex and we have found its global minimum.

### Alternative verification that $\hat{\beta}$ is the LSE

There is another way how to verify that the solution  $\hat{\beta}$  to the system of normal equations (2.1) is the LSE. Take any  $\beta \in \mathbb{R}^p$  and write

$$\begin{aligned} SS_e(\beta) &= \|\mathbf{Y} - \mathbb{X}\beta\|^2 = \|\mathbf{Y} - \mathbb{X}\hat{\beta} + \mathbb{X}\hat{\beta} - \mathbb{X}\beta\|^2 \\ &= \|\mathbf{Y} - \mathbb{X}\hat{\beta}\|^2 + \|\mathbb{X}(\hat{\beta} - \beta)\|^2 + 2(\mathbf{Y} - \mathbb{X}\hat{\beta})^\top \mathbb{X}(\hat{\beta} - \beta), \end{aligned}$$

where the last term is zero because

$$(\hat{\beta} - \beta)^\top \mathbb{X}^\top (\mathbf{Y} - \mathbb{X}\hat{\beta}) = (\hat{\beta} - \beta)^\top (\mathbb{X}^\top \mathbf{Y} - \mathbb{X}^\top \mathbb{X} \hat{\beta}) = \mathbf{0}$$

using the fact that  $\hat{\beta}$  solves the normal equations  $\mathbb{X}^\top \mathbb{X} \hat{\beta} = \mathbb{X}^\top \mathbf{Y}$ .

Hence, at any  $\beta \in \mathbb{R}^p$ ,

$$SS_e(\beta) = \|\mathbf{Y} - \mathbb{X}\hat{\beta}\|^2 + \|\mathbb{X}(\hat{\beta} - \beta)\|^2 \geq \|\mathbf{Y} - \mathbb{X}\hat{\beta}\|^2 = SS_e(\hat{\beta})$$

and equality is attained if and only if

$$\|\mathbb{X}(\hat{\beta} - \beta)\|^2 = (\hat{\beta} - \beta)^\top \mathbb{X}^\top \mathbb{X} (\hat{\beta} - \beta) = \mathbf{0}.$$

When  $\mathbb{X}^\top \mathbb{X}$  is of full rank, this is equivalent to  $\beta = \hat{\beta}$ . Thus,  $\hat{\beta}$  is the unique minimizer of  $SS_e(\beta)$ .

### Fitted values and residuals

#### Definition 2.3 (Fitted values, residuals).

---

\* Český *normální rovnice*

- (a)  $\hat{Y} \equiv \mathbb{X}\hat{\beta}$  are called *the fitted values*<sup>\*</sup>.  
 (b)  $u \equiv Y - \hat{Y} = Y - \mathbb{X}\hat{\beta}$  are called *the residuals*<sup>†</sup>.

Recall the definition of the linear regression model

$$Y = \mathbb{X}\beta + \varepsilon,$$

where  $\mathbb{X}\beta$  is the conditional mean of  $Y$  given the covariates and  $\varepsilon$  is random noise, and compare it with the decomposition

$$Y = \mathbb{X}\hat{\beta} + u,$$

where the fitted values  $\mathbb{X}\hat{\beta} = \hat{Y}$  represent the estimated mean of  $Y$  and the residuals  $u$  represent the estimated noise. The fitted values are the “best” approximations (or predictions) of the responses that can be calculated from the covariates alone.

We can write  $\hat{Y} = \mathbb{X}\hat{\beta} = \mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^TY = \mathbb{H}Y$ , where  $\mathbb{H} \equiv \mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^T$  is a square  $n \times n$  matrix. The matrix  $\mathbb{H}$  is called *the hat matrix*<sup>‡</sup>. It is symmetric,  $r(\mathbb{H}) = p$  because  $r(\mathbb{X}) = p$ , and it is idempotent:

$$\mathbb{H}\mathbb{H} = \mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^T\mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^T = \mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^T = \mathbb{H}.$$

Recall that any idempotent matrix satisfies  $r(\mathbb{H}) = \text{tr}(\mathbb{H})$ .

Throughout the whole course, we will frequently use the following trivial identities:

$$\mathbb{H}\mathbb{X} = \mathbb{X}, \quad (\mathbb{I} - \mathbb{H})\mathbb{X} = \mathbf{0}.$$

The main linear properties of fitted values and residuals are summarized in the following note.

**Note.**

- (a)  $\hat{Y} = \mathbb{H}Y$  where  $\mathbb{H} \equiv \mathbb{X}(\mathbb{X}^T\mathbb{X})^{-1}\mathbb{X}^T$  is a symmetric, idempotent  $n \times n$  matrix of rank  $p$ .  
 (b)  $u = (\mathbb{I} - \mathbb{H})Y$  where  $\mathbb{I} - \mathbb{H}$  is a symmetric, idempotent  $n \times n$  matrix of rank  $n - p$ . Also,  $u = (\mathbb{I} - \mathbb{H})\varepsilon$ .  
 (c)  $\hat{Y}$ ,  $u$ , and  $\hat{\beta}$  are all linear transformations of  $Y$ .  
 (d)  $\hat{Y}$  and  $u$  are always orthogonal.

Parts (a) and (c) of the note are trivial or have been proven above. As for part (b),  $(\mathbb{I} - \mathbb{H})(\mathbb{I} - \mathbb{H}) = \mathbb{I} - 2\mathbb{H} + \mathbb{H}\mathbb{H} = \mathbb{I} - \mathbb{H}$ , so  $(\mathbb{I} - \mathbb{H})$  is indeed idempotent. Its rank can be calculated using  $r(\mathbb{A}) = \text{tr}(\mathbb{A})$  for any idempotent  $\mathbb{A}$ :

$$r(\mathbb{I} - \mathbb{H}) = \text{tr}(\mathbb{I} - \mathbb{H}) = \text{tr}(\mathbb{I}) - \text{tr}(\mathbb{H}) = n - r(\mathbb{H}) = n - p. \quad (2.3)$$

---

<sup>\*</sup> Český vyrovnané hodnoty    <sup>†</sup> Český residua (sing. residuum)    <sup>‡</sup> Český nemá český ekvivalent

Finally, using the definition of the linear model and  $(\mathbb{I} - \mathbb{H})\mathbb{X} = \mathbf{0}$ ,

$$\mathbf{u} = (\mathbb{I} - \mathbb{H})\mathbf{Y} = (\mathbb{I} - \mathbb{H})(\mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}) = (\mathbb{I} - \mathbb{H})\mathbb{X}\boldsymbol{\beta} + (\mathbb{I} - \mathbb{H})\boldsymbol{\varepsilon} = (\mathbb{I} - \mathbb{H})\boldsymbol{\varepsilon}.$$

As for (d), it is easy to verify that

$$\widehat{\mathbf{Y}}^\top \mathbf{u} = \mathbf{Y}^\top \mathbb{H}(\mathbb{I} - \mathbb{H})\mathbf{Y} = \mathbf{Y}^\top (\mathbb{H} - \mathbb{H}\mathbb{H})\mathbf{Y} = \mathbf{0}.$$

**Note.** Any symmetric idempotent matrix is positive semi-definite. *Prove this yourself.*

### Geometric interpretation of the LSE

#### From Linear Algebra:

Consider a vector space  $V$  and two subspaces  $U$  and  $W$  such that  $V = U \oplus W$ .  $U$  and  $W$  are orthogonal iff  $\mathbf{u}^\top \mathbf{w} = 0$  for any  $\mathbf{u} \in U$ ,  $\mathbf{w} \in W$ . Then we denote  $W = U^\perp$ .

Any vector  $\mathbf{v} \in V$  can be uniquely decomposed as  $\mathbf{u}_v + \mathbf{w}_v$ , where  $\mathbf{u}_v \in U$  and  $\mathbf{w}_v \in U^\perp$ . This is called *orthogonal projection*. Projection is a linear transformation of the vector through a projection matrix  $\mathbb{P}$ . The columns of  $\mathbb{P}$  are the projections of basis vectors of  $V$ , and  $U$  is the image of  $\mathbb{P}$ .

A square matrix  $\mathbb{P}$  is a projection matrix if and only if it is idempotent.

Let  $\mathbb{A} = (\mathbf{a}_1, \dots, \mathbf{a}_p)$  be any basis of a subspace  $U$  of  $V$ . Then  $\mathbb{A}(\mathbb{A}^\top \mathbb{A})^{-1} \mathbb{A}^\top$  is a projection matrix of  $V$  onto  $U$ .

Let  $\mathcal{M}(\mathbb{X})$  be the linear subspace of  $\mathbb{R}^n$  generated by the columns of the regression matrix  $\mathbb{X}$  (denote them by  $\mathbf{x}_j$ ,  $j = 1, \dots, p$ ):

$$\mathcal{M}(\mathbb{X}) = \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{x} = \sum_{j=1}^p q_j \mathbf{x}_j, q_j \in \mathbb{R} \right\}.$$

Let  $\mathcal{M}(\mathbb{X})^\perp$  be the subspace orthogonal to  $\mathcal{M}(\mathbb{X})$ :

$$\mathcal{M}(\mathbb{X})^\perp = \{ \mathbf{z} \in \mathbb{R}^n : \mathbf{z}^\top \mathbf{x} = 0 \ \forall \mathbf{x} \in \mathcal{M}(\mathbb{X}) \}.$$

Then

- $\widehat{\mathbf{Y}}$  is the orthogonal projection of  $\mathbf{Y} \in \mathbb{R}^n$  to the  $p$ -dimensional subspace  $\mathcal{M}(\mathbb{X})$ , with the projection matrix  $\mathbb{H}$ ;
- $\mathbf{u}$  is the orthogonal projection of  $\mathbf{Y} \in \mathbb{R}^n$  to the  $n - p$ -dimensional subspace  $\mathcal{M}(\mathbb{X})^\perp$ , with the projection matrix  $\mathbb{I} - \mathbb{H}$ .

So,  $\mathbb{H}$  and  $\mathbb{I} - \mathbb{H}$  are projection matrices to the two orthogonal subspaces,  $\mathcal{M}(\mathbb{X})$  and  $\mathcal{M}(\mathbb{X})^\perp$ , respectively.

## 2.4. Residual Sum of Squares

The *residual sum of squares*, denoted by  $SS_e$ , is the sum of squared residuals and at the same time the minimized value of the least squares criterion  $SS_e(\boldsymbol{\beta})$ . There are several alternative ways how to express it.

$$SS_e \equiv SS_e(\hat{\boldsymbol{\beta}}) = \|\mathbf{Y} - \mathbb{X}\hat{\boldsymbol{\beta}}\|^2 = \|\mathbf{Y} - \hat{\mathbf{Y}}\|^2 = \|\mathbf{u}\|^2 = \sum_{i=1}^n u_i^2.$$

According to the note on p. 20, part (b),  $\mathbf{u} = (\mathbb{I} - \mathbb{H})\mathbf{Y} = (\mathbb{I} - \mathbb{H})\boldsymbol{\varepsilon}$ . Because  $\mathbb{I} - \mathbb{H}$  is idempotent,  $SS_e$  can be expressed as a quadratic form in two alternative ways:

$$SS_e = \mathbf{Y}^\top (\mathbb{I} - \mathbb{H}) \mathbf{Y} = \boldsymbol{\varepsilon}^\top (\mathbb{I} - \mathbb{H}) \boldsymbol{\varepsilon}.$$

Another way to express residual sum of squares is this:

$$\begin{aligned} SS_e &= (\mathbf{Y} - \mathbb{X}\hat{\boldsymbol{\beta}})^\top (\mathbf{Y} - \mathbb{X}\hat{\boldsymbol{\beta}}) = \mathbf{Y}^\top \mathbf{Y} - \mathbf{Y}^\top \mathbb{X}\hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}^\top \mathbb{X}^\top \mathbf{Y} + \hat{\boldsymbol{\beta}}^\top (\mathbb{X}^\top \mathbb{X}) \hat{\boldsymbol{\beta}} = \\ &= \mathbf{Y}^\top \mathbf{Y} - \mathbf{Y}^\top \mathbb{X}\hat{\boldsymbol{\beta}} = \mathbf{Y}^\top \mathbf{Y} - \mathbf{Y}^\top \hat{\mathbf{Y}}. \end{aligned} \quad (2.4)$$

## 2.5. Equivalence of Regression Models

Consider two different regression models for the same response  $\mathbf{Y}$ :

$$\begin{aligned} \mathbf{Y} &= \mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, & \text{where } \mathbb{X}_{n \times p} \text{ and } \boldsymbol{\beta}_{p \times 1}, \\ \text{and } \mathbf{Y} &= \mathbb{X}^*\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}^*, & \text{where } \mathbb{X}^*_{n \times q} \text{ and } \boldsymbol{\beta}^*_{q \times 1}. \end{aligned}$$

The two models are called *equivalent* if and only if  $\mathcal{M}(\mathbb{X}) = \mathcal{M}(\mathbb{X}^*)$ , that is, the linear subspaces generated by the columns of  $\mathbb{X}$  and  $\mathbb{X}^*$ , respectively, are the same. This is true if and only if there exists a  $q \times p$  matrix  $\mathbb{C}$  such that  $\mathbb{X} = \mathbb{X}^*\mathbb{C}$ . For this particular  $\mathbb{C}$ , it follows that  $\mathbb{X}\boldsymbol{\beta} = \mathbb{X}^*\mathbb{C}\boldsymbol{\beta}$  and hence  $\boldsymbol{\beta}^* = \mathbb{C}\boldsymbol{\beta}$  and  $\boldsymbol{\varepsilon}^* = \boldsymbol{\varepsilon}$ .

Because the fitted values  $\hat{\mathbf{Y}}$  in the two models are projections of the same vector  $\mathbf{Y}$  into the same linear subspace, they must be the same in both models. The same is true for the residuals  $\mathbf{u}$  and the residual sum of squares  $SS_e$ .

When  $\mathbb{X}^*_{n \times q}$  is a matrix of rank  $p < q$  then there exists a full rank matrix  $\mathbb{X}_{n \times p}$  that generates an equivalent model. This is the mechanism how to avoid ever considering non-full rank regression matrices. If a regression matrix is not of full rank we work instead with an equivalent model, which is of full rank.

## 2.6. Model for iid Response

The simplest special case of a regression model describes independent and identically distributed responses. Let  $Y_1, \dots, Y_n$  be iid random variables with  $E Y_i = \mu$  and  $\text{var } Y_i = \sigma_Y^2$ .

Write

$$Y_i = \mu + (Y_i - \mu) \equiv X_i \beta + \varepsilon_i,$$

where  $X_i = 1$  for all  $i$ ,  $\beta = \mu$ ,  $E \varepsilon_i = 0$ , and  $\text{var } \varepsilon_i = \sigma_Y^2$ . This is a linear model. We can write the vector containing all the responses in the form

$$\mathbf{Y} = \mathbb{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

where  $\mathbb{X} = (1, \dots, 1)^T \equiv \mathbf{1}_n$ ,  $\boldsymbol{\beta} = \mu$ .

**Notation.**

- Let  $\mathbf{1}_n$  be a column  $n$ -vector of ones;  $\mathbf{1}_n = (1, \dots, 1)^T$ .
- Let  $\mathbb{J}_n = \mathbf{1}_n \mathbf{1}_n^T$  be an  $n \times n$  matrix of ones.

Let us now calculate the least squares estimator and residual sum of squares. We have

$$\hat{\boldsymbol{\beta}} = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \mathbf{Y} = (\mathbf{1}_n^T \mathbf{1}_n)^{-1} (\mathbf{1}_n^T \mathbf{Y}) = \frac{1}{n} \sum_{i=1}^n Y_i \equiv \bar{Y}_n.$$

So, the least squares estimate of the common expectation is the arithmetic average. Next, the fitted values are  $\hat{\mathbf{Y}} = \bar{Y}_n \mathbf{1}_n$  and the residuals are  $\mathbf{u} = \mathbf{Y} - \bar{Y}_n \mathbf{1}_n$ . The residual sum of squares is  $SS_e = \mathbf{u}^T \mathbf{u} = \sum_{i=1}^n (Y_i - \bar{Y}_n)^2$ .

## 2.7. Model With Centered Covariates

In order to gain further insights into the meaning of the LSE procedure, we need to center the covariates. Consider the model

$$\mathbf{Y} = \mathbb{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon},$$

where the first column of  $\mathbb{X}$  is  $\mathbf{1}_n$  (the intercept column). Denote the rest of the regression matrix as  $\mathbb{X}_R$ , that is,  $\mathbb{X} = (\mathbf{1}_n | \mathbb{X}_R)$ . The vector  $\boldsymbol{\beta}$  is divided similarly into  $\boldsymbol{\beta} = \begin{pmatrix} \beta_1 \\ \boldsymbol{\beta}_R \end{pmatrix}$ .

Each observation can be expressed as

$$Y_i = \beta_1 + \beta_2 X_{i2} + \dots + \beta_p X_{ip} + \varepsilon_i.$$

Let  $\bar{X}_j = \frac{1}{n} \sum_{i=1}^n X_{ij}$  for  $j = 2, \dots, p$ . Now, subtract from the value of each covariate the respective mean (except for the intercept). We get

$$Y_i = \alpha + \beta_2 (X_{i2} - \bar{X}_2) + \dots + \beta_p (X_{ip} - \bar{X}_p) + \varepsilon_i,$$

where  $\alpha = \beta_1 + \beta_2 \bar{X}_2 + \dots + \beta_p \bar{X}_p$  to maintain the equality. This is the model with *centered covariates* (shortly, the “centered model”). It is an equivalent model (the subspaces generated

by the columns of the regression matrix have not changed) and the parameters  $\beta_2, \dots, \beta_p$  are the same. Only the intercept parameter is different. The new intercept has the interpretation  $E[Y_i | X_{i2} = \bar{X}_2, \dots, X_{ip} = \bar{X}_p]$ , the expected response for an individual with average value in all covariates.

**Message:** If any covariate is shifted by a constant (the same number is added to/subtracted from all values of the covariate)

Take  $\mathbb{J}_n = \mathbf{1}_n \mathbf{1}_n^T$ , an  $n \times n$  matrix with 1 at all positions. The centered covariates can be created by multiplication by the column centering matrix:  $\mathbb{X}_C = (\mathbb{I}_n - n^{-1} \mathbb{J}_n) \mathbb{X}_R$ . The centered model can be written as

$$\mathbf{Y} = (\mathbf{1}_n | \mathbb{X}_C) \begin{pmatrix} \alpha \\ \boldsymbol{\beta}_R \end{pmatrix} + \boldsymbol{\varepsilon}.$$

Let us find the least squares estimate of  $(\alpha, \boldsymbol{\beta}_R)$ . The original model and the centered model are equivalent, they have the same fitted values  $\hat{Y}_i$ . Let  $\hat{\boldsymbol{\beta}} = \begin{pmatrix} \hat{\beta}_1 \\ \hat{\boldsymbol{\beta}}_R \end{pmatrix}$  be the LSE in the original model,  $\hat{\boldsymbol{\beta}} = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \mathbf{Y}$ . Then for all  $i = 1, \dots, n$ ,

$$\begin{aligned} \hat{Y}_i &= \hat{\beta}_1 + \hat{\beta}_2 X_{i2} + \dots + \hat{\beta}_p X_{ip} \\ &= \hat{\alpha} + \hat{\beta}_2 (X_{i2} - \bar{X}_2) + \dots + \hat{\beta}_p (X_{ip} - \bar{X}_p), \end{aligned}$$

where  $\hat{\alpha} = \hat{\beta}_1 + \hat{\beta}_2 \bar{X}_2 + \dots + \hat{\beta}_p \bar{X}_p$ . Because  $\begin{pmatrix} \hat{\alpha} \\ \hat{\boldsymbol{\beta}}_R \end{pmatrix}$  generates the same fitted values, residuals and  $SS_e$  as the LSE of the original model, it must be the unique LSE in the centered model.

**Message:** If any covariate is shifted by a constant (the same number is added to/subtracted from all values of the covariate) there is no change in either the regression parameter for that covariate or in its LSE.

Now, apply the LSE formula to the centered model. We have

$$\begin{aligned} \begin{pmatrix} \hat{\alpha} \\ \hat{\boldsymbol{\beta}}_R \end{pmatrix} &= [(\mathbf{1}_n | \mathbb{X}_C)^T (\mathbf{1}_n | \mathbb{X}_C)]^{-1} (\mathbf{1}_n | \mathbb{X}_C)^T \mathbf{Y} \\ &= \begin{pmatrix} n & \mathbf{1}_n^T \mathbb{X}_C \\ \mathbb{X}_C^T \mathbf{1}_n & \mathbb{X}_C^T \mathbb{X}_C \end{pmatrix}^{-1} \begin{pmatrix} \sum_{i=1}^n Y_i \\ \mathbb{X}_C^T \mathbf{Y} \end{pmatrix} = \begin{pmatrix} \frac{1}{n} & \mathbf{0} \\ \mathbf{0} & (\mathbb{X}_C^T \mathbb{X}_C)^{-1} \end{pmatrix} \begin{pmatrix} \sum_{i=1}^n Y_i \\ \mathbb{X}_C^T \mathbf{Y} \end{pmatrix} = \begin{pmatrix} \bar{Y} \\ (\mathbb{X}_C^T \mathbb{X}_C)^{-1} (\mathbb{X}_C^T \mathbf{Y}) \end{pmatrix}. \end{aligned}$$

We have verified that  $\hat{\alpha} = \bar{Y}$ . The fitted values in the centered model are

$$\hat{Y}_i = \bar{Y} + \hat{\beta}_2 (X_{i2} - \bar{X}_2) + \dots + \hat{\beta}_p (X_{ip} - \bar{X}_p).$$

Because the original model has the same fitted values, we have the following conclusion.

**Conclusion:** If the model includes the intercept column, the fitted value evaluated at the average value of each of the remaining covariates is equal to the average of the responses.



We can also construct an additional way to express the residual sum of squares in a model with intercept. In the original model, we have  $SS_e = \mathbf{Y}^\top \mathbf{Y} - \mathbf{Y}^\top \mathbb{X} \hat{\boldsymbol{\beta}}$ , see (2.4). When we apply this to the centered model, we get

$$SS_e = \mathbf{Y}^\top \mathbf{Y} - \mathbf{Y}^\top (\mathbf{1}_n | \mathbb{X}_C) \begin{pmatrix} \bar{Y} \\ \hat{\boldsymbol{\beta}}_R \end{pmatrix} = \mathbf{Y}^\top \mathbf{Y} - n\bar{Y}^2 - \mathbf{Y}^\top \mathbb{X}_C \hat{\boldsymbol{\beta}}_R = \sum_{i=1}^n (Y_i - \bar{Y})^2 - \mathbf{Y}^\top \mathbb{X}_C \hat{\boldsymbol{\beta}}_R.$$

## 2.8. Relationship to Sample Covariance Matrices

In this section, we still work under the assumption that  $\mathbf{1}_n \in \mathcal{M}(\mathbb{X})$  (the intercept is included in the model). Denote by  $\mathbb{S}_{XX}$  the sample covariance matrix of the columns of  $\mathbb{X}_R$  (the remaining columns of the regression matrix after excluding the intercept). It is a  $(p-1) \times (p-1)$  matrix with diagonal elements  $\frac{1}{n-1} \sum_{i=1}^n (X_{ij} - \bar{X}_j)^2$  and off-diagonal elements  $\frac{1}{n-1} \sum_{i=1}^n (X_{ij} - \bar{X}_j)(X_{ik} - \bar{X}_k)$ . Obviously,  $\mathbb{S}_{XX} = \frac{1}{n-1} \mathbb{X}_C^\top \mathbb{X}_C$ .

Now consider the sample covariance matrix\*  $\mathbb{S}_{XY}$  of the columns of  $\mathbb{X}_R$  with the response vector  $\mathbf{Y}$ , a  $(p-1) \times 1$  matrix with elements  $\frac{1}{n-1} \sum_{i=1}^n (X_{ij} - \bar{X}_j)(Y_i - \bar{Y})$ . Because  $\sum_{i=1}^n (X_{ij} - \bar{X}_j)\bar{Y} = 0$ , we have  $\mathbb{S}_{XY} = \frac{1}{n-1} \mathbb{X}_C^\top \mathbf{Y}$ .

**Conclusion:** If the model includes the intercept column, the LSE of the non-intercept parameters can be expressed in terms of sample covariance matrices as follows:  $\hat{\boldsymbol{\beta}}_R = \mathbb{S}_{XX}^{-1} \mathbb{S}_{XY}$ .

We can also express the LSE of the intercept parameter using the results of the previous section.

$$\hat{\beta}_1 = \hat{\alpha} - \frac{1}{n} \mathbf{1}_n^\top \mathbb{X}_R \hat{\boldsymbol{\beta}}_R = \bar{Y} - \frac{1}{n} \mathbf{1}_n^\top \mathbb{X}_R \mathbb{S}_{XX}^{-1} \mathbb{S}_{XY}.$$

## 2.9. Decomposition of Sums of Squares

This can be done in two ways – for non-centered or centered response. The first decomposition is universally valid but less useful. The second is more useful but holds only if the intercept is included in the model.

### Decomposition of sums of squares with non-centered response

Start with the sum of squared responses

$$\|\mathbf{Y}\|^2 = \mathbf{Y}^\top \mathbf{Y} = \mathbf{Y}^\top \mathbb{H} \mathbf{Y} + \mathbf{Y}^\top (\mathbb{I} - \mathbb{H}) \mathbf{Y}.$$

The last term on the right-hand side can be recognized as the *residual sum of squares*  $SS_e$ . The left-hand side is called the *non-centered total sum of squares*, denoted by  $SS_T^*$ . The remaining

---

\* actually, it is a vector

term,  $\mathbf{Y}^\top \mathbb{H} \mathbf{Y}$ , is called *the non-centered regression sum of squares*, denoted by  $SS_R^*$ . We have

$$SS_R^* = \mathbf{Y}^\top \mathbb{H} \mathbf{Y} = \mathbf{Y}^\top \mathbb{H} \mathbb{H} \mathbf{Y} = \|\mathbb{H} \mathbf{Y}\|^2 = \|\widehat{\mathbf{Y}}\|^2 = \|\mathbb{X} \widehat{\boldsymbol{\beta}}\|^2 = \widehat{\boldsymbol{\beta}}^\top \mathbb{X}^\top \mathbb{X} \widehat{\boldsymbol{\beta}}$$

The non-centered decomposition is

$$\underbrace{\sum_{i=1}^n Y_i^2}_{SS_T} = \underbrace{\sum_{i=1}^n \widehat{Y}_i^2}_{SS_R^*} + \underbrace{\sum_{i=1}^n (Y_i - \widehat{Y}_i)^2}_{SS_e}.$$

### Decomposition of sums of squares with centered response

Assume that the model contains the intercept,  $\mathbf{1}_n \in \mathcal{M}(\mathbb{X})$ . Calculate the mean response  $\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i$  and subtract the mean from all responses, that is, take

$$\mathbf{Y} - \mathbf{1}_n \bar{Y} = \mathbf{Y} - \mathbf{1}_n \frac{1}{n} \mathbf{1}_n^\top \mathbf{Y} = \mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y}.$$

Now apply the decomposition of sums of squares to these centered responses.

The total (centered) sum of squares is

$$SS_T \equiv \|\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y}\|^2 = \sum_{i=1}^n (Y_i - \bar{Y})^2.$$

This can be decomposed as

$$SS_T = (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y})^\top \mathbb{H} (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y}) + (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y})^\top (\mathbb{I} - \mathbb{H}) (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y}).$$

Because the model contains the intercept,  $\mathbb{H} \mathbf{1}_n = \mathbf{1}_n$ , hence  $\mathbb{H} \mathbb{J}_n = \mathbb{J}_n$ , hence  $(\mathbb{I} - \mathbb{H}) \mathbb{J}_n = \mathbf{0}$ . Thus, the last term on the right-hand side is still the *residual sum of squares*  $SS_e$ .

The remaining term,  $(\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y})^\top \mathbb{H} (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y})$ , is *the (centered) regression sum of squares*, denoted by  $SS_R$ . We have

$$\begin{aligned} SS_R &= (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y})^\top \mathbb{H} (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y}) = \|\mathbb{H} (\mathbf{Y} - \frac{1}{n} \mathbb{J}_n \mathbf{Y})\|^2 = \|\mathbb{H} \mathbf{Y} - \frac{1}{n} \underbrace{\mathbb{H} \mathbb{J}_n}_{\mathbb{J}_n} \mathbf{Y}\|^2 \\ &= \|\widehat{\mathbf{Y}} - \frac{1}{n} \mathbb{J}_n \mathbf{Y}\|^2 = \|\widehat{\mathbf{Y}} - \mathbf{1}_n \bar{Y}\|^2 = \sum_{i=1}^n (\widehat{Y}_i - \bar{Y})^2. \end{aligned}$$

The centered decomposition of sums of squares is

$$\underbrace{\sum_{i=1}^n (Y_i - \bar{Y})^2}_{SS_T} = \underbrace{\sum_{i=1}^n (\widehat{Y}_i - \bar{Y})^2}_{SS_R} + \underbrace{\sum_{i=1}^n (Y_i - \widehat{Y}_i)^2}_{SS_e}.$$

This can be interpreted as follows. The total sum of squares  $SS_T$  captures the total variability in the response. This is decomposed into  $SS_R$ , the variability that is explained by the regression model (using the covariates), and into  $SS_e$ , which is the part of variability that could not be explained.

Notice that we have the mean of all responses in the expression for  $SS_R$  instead of the mean of the fitted values.

## 2.10. Coefficient of Determination

We continue to assume that the model contains the intercept,  $\mathbf{1}_n \in \mathcal{M}(\mathbb{X})$ , and recall the centered decomposition of sums of squares  $SS_T = SS_R + SS_e$  derived in the previous section.

**Definition 2.4 (Coefficient of determination).** The quantity

$$R^2 = \frac{SS_R}{SS_T} = 1 - \frac{SS_e}{SS_T} \quad \nabla$$

is called *the coefficient of determination*<sup>\*</sup>.

If we interpret  $SS_T$  as the total variability of the response and  $SS_R$  as the variability explained by the covariates included in the model, we can view  $R^2$  as the fraction of the variability of the response that was explained by the regression model.

### Notes on coefficient of determination

1. Obviously,  $0 \leq R^2 \leq 1$ .
2.  $\sqrt{R^2}$  is sometimes called *multiple correlation coefficient*<sup>†</sup> between the random variable  $Y$  and random vector  $\mathbf{X}$ .
3.  $R^2$  is equal to the square of the estimated correlation coefficient between  $Y$  and  $\hat{Y}$ .

**Proof.**

$$R^2 = \frac{\|\hat{Y} - \mathbf{1}_n \bar{Y}\|^2}{\|Y - \mathbf{1}_n \bar{Y}\|^2} = \frac{\|\hat{Y} - \mathbf{1}_n \bar{Y}\|^4}{\|Y - \mathbf{1}_n \bar{Y}\|^2 \|\hat{Y} - \mathbf{1}_n \bar{Y}\|^2}$$

Now express the norm in the numerator differently:

$$\begin{aligned} \|\hat{Y} - \mathbf{1}_n \bar{Y}\|^2 &= (\hat{Y} - Y + Y - \mathbf{1}_n \bar{Y})^\top (\hat{Y} - \mathbf{1}_n \bar{Y}) \\ &= \underbrace{(\hat{Y} - Y)^\top (\hat{Y} - \mathbf{1}_n \bar{Y})}_{=0} + (Y - \mathbf{1}_n \bar{Y})^\top (\hat{Y} - \mathbf{1}_n \bar{Y}) \\ &= (Y - \mathbf{1}_n \bar{Y})^\top (\hat{Y} - \mathbf{1}_n \bar{Y}) \end{aligned} \quad (2.5)$$

---

<sup>\*</sup> Český koeficient determinace    <sup>†</sup> Český koeficient mnohonásobné korelace

The first term on the second line is zero because

$$(\hat{Y} - Y)^T(\hat{Y} - \mathbf{1}_n \bar{Y}) = (\mathbb{H}Y - Y)^T(\mathbb{H}Y - \frac{1}{n}\mathbb{J}_n Y) = -Y^T(\mathbb{I} - \mathbb{H})(\mathbb{H} - \frac{1}{n}\mathbb{J}_n)Y$$

and

$$(\mathbb{I} - \mathbb{H})(\mathbb{H} - \frac{1}{n}\mathbb{J}_n) = (\mathbb{I} - \mathbb{H})\mathbb{H} - \frac{1}{n}(\mathbb{I} - \mathbb{H})\mathbb{J}_n = 0$$

because  $\mathbf{1}_n \in \mathcal{M}(\mathbb{X})$ . So,

$$R^2 = \left[ \frac{(Y - \mathbf{1}_n \bar{Y})^T(\hat{Y} - \mathbf{1}_n \bar{Y})}{\sqrt{\|Y - \mathbf{1}_n \bar{Y}\|^2 \|\hat{Y} - \mathbf{1}_n \bar{Y}\|^2}} \right]^2 = \widehat{\text{cor}}^2(Y, \hat{Y}). \quad \square$$

4. Another variant of the coefficient of determination is so called *adjusted*  $R^2$  defined as

$$R_a^2 = 1 - \frac{n-1}{n-p} \frac{SS_e}{SS_T}.$$

The motivation for this is to subtract the ratio of two unbiased estimators of  $\text{var } \varepsilon_i$  and  $\text{var } Y_i^*$ .

## 2.11. LSE Under Linear Restrictions

Consider the linear model  $Y = \mathbb{X}\beta + \varepsilon$  with  $\mathbb{X}$  of full rank. The least squares estimator  $\hat{\beta} = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T Y$  minimizes the residual sum of squares  $SS_e(\beta) = \|Y - \mathbb{X}\beta\|^2$  over all  $\beta \in \mathbb{R}^p$ .

Now we impose an additional set of linear restrictions on the parameters: let  $\mathbb{C}\beta = c$ , where  $\mathbb{C}$  is a  $q \times p$  matrix with  $\text{rank } r(\mathbb{C}) = q < p$  and  $c \in \mathbb{R}^q$ . We will minimize  $SS_e(\beta)$  over the set  $\mathcal{B} = \{\beta \in \mathbb{R}^p : \mathbb{C}\beta = c\}$ . Denote  $\hat{\beta}_C = \arg \min_{\beta \in \mathcal{B}} \|Y - \mathbb{X}\beta\|^2$ .

We can use the method of Lagrange multipliers to calculate  $\hat{\beta}_C$ . Introduce the objective function

$$S(\beta, \lambda) = SS_e(\beta) + \lambda^T(\mathbb{C}\beta - c),$$

where  $\lambda \in \mathbb{R}^q$ . Calculate

$$\frac{\partial S(\beta, \lambda)}{\partial \beta} = -2\mathbb{X}^T Y + 2\mathbb{X}^T \mathbb{X} \beta + \mathbb{C}^T \lambda$$

and set it equal to zero to find  $\hat{\beta}_C$ . We get

$$\mathbb{X}^T \mathbb{X} \hat{\beta}_C = \mathbb{X}^T Y - \frac{1}{2} \mathbb{C}^T \lambda$$

---

\* The fact that  $SS_e/(n-p)$  is an unbiased estimator of  $\sigma_e^2$  will be established in Section ??.

and hence

$$\hat{\beta}_C = (\mathbb{X}^T \mathbb{X})^{-1} [\mathbb{X}^T \mathbf{Y} - \frac{1}{2} \mathbb{C}^T \boldsymbol{\lambda}] = \hat{\beta} - \frac{1}{2} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T \boldsymbol{\lambda}. \quad (2.6)$$

The solution must satisfy the constraint  $\mathbb{C} \hat{\beta}_C = \mathbf{c}$ , i.e.,

$$\mathbb{C} \hat{\beta} - \frac{1}{2} \mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T \boldsymbol{\lambda} = \mathbf{c}.$$

Use this to identify  $\boldsymbol{\lambda}$ : it is a solution to the system of linear equations

$$\mathbb{C} \hat{\beta} - \mathbf{c} = \frac{1}{2} \mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T \boldsymbol{\lambda}.$$

Since  $r(\mathbb{X}) = p$  and  $r(\mathbb{C}) = q < p$ , the  $q \times q$  matrix  $\mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T$  is of rank  $q$ , therefore regular and invertible. Thus,

$$\boldsymbol{\lambda} = 2[\mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C} \hat{\beta} - \mathbf{c}).$$

Plug this into (2.6) to obtain the result

$$\hat{\beta}_C = \hat{\beta} - (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T [\mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C} \hat{\beta} - \mathbf{c}). \quad (2.7)$$

However, this is only a suspicious point. We still need to show that it really minimizes  $SS_e(\beta)$  over  $\beta \in \mathcal{B}$ . So, take any  $\beta \in \mathcal{B}$  and write

$$\begin{aligned} SS_e(\beta) &= \|\mathbf{Y} - \mathbb{X}\beta\|^2 = \|\mathbf{Y} - \mathbb{X}\hat{\beta}_C + \mathbb{X}\hat{\beta}_C - \mathbb{X}\beta\|^2 \\ &= \|\mathbf{Y} - \mathbb{X}\hat{\beta}_C\|^2 + \|\mathbb{X}(\hat{\beta}_C - \beta)\|^2 + 2(\mathbf{Y} - \mathbb{X}\hat{\beta}_C)^T \mathbb{X}(\hat{\beta}_C - \beta) \end{aligned}$$

Look at the last term. From (2.7), we have

$$\mathbf{Y} - \mathbb{X}\hat{\beta}_C = \mathbf{Y} - \mathbb{X}\hat{\beta} + \mathbb{X}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T [\mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C} \hat{\beta} - \mathbf{c})$$

and

$$\begin{aligned} (\mathbf{Y} - \mathbb{X}\hat{\beta}_C)^T \mathbb{X}(\hat{\beta}_C - \beta) &= \underbrace{(\mathbf{Y} - \mathbb{X}\hat{\beta})^T \mathbb{X}(\hat{\beta}_C - \beta)}_{= \mathbf{u}^T \mathbb{X} = 0} \\ &\quad + (\mathbb{C} \hat{\beta} - \mathbf{c})^T [\mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} \underbrace{\mathbb{C} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \mathbb{X}(\hat{\beta}_C - \beta)}_{= \mathbb{C}(\hat{\beta}_C - \beta) = \mathbf{c} - \mathbf{c} = 0} \\ &= 0. \end{aligned}$$

Thus, for any  $\beta \in \mathcal{B}$ ,

$$SS_e(\beta) = SS_e(\hat{\beta}_C) + (\hat{\beta}_C - \beta)^T \mathbb{X}^T \mathbb{X}(\hat{\beta}_C - \beta) \geq SS_e(\hat{\beta}_C)$$

and equality is attained if and only if  $\beta = \hat{\beta}_C$ . Thus,  $\hat{\beta}_C$  is the unique minimizer of  $SS_e(\beta)$  over  $\beta \in \mathcal{B}$  and therefore it is the restricted LSE.

## 2. Linear Regression Model

---

Now evaluate the difference between  $SS_e = SS_e(\hat{\beta})$  and  $SS_e(\hat{\beta}_C)$ . Since  $\hat{\beta}_C$  minimizes  $SS_e$  over a subspace of  $\mathbb{R}^p$ ,  $SS_e \leq SS_e(\hat{\beta}_C)$ . Write

$$\begin{aligned} SS_e(\hat{\beta}_C) &= \|Y - \mathbb{X}\hat{\beta}_C\|^2 = \|Y - \mathbb{X}\hat{\beta} + \mathbb{X}\hat{\beta} - \mathbb{X}\hat{\beta}_C\|^2 \\ &= \|Y - \mathbb{X}\hat{\beta}\|^2 + \|\mathbb{X}(\hat{\beta} - \hat{\beta}_C)\|^2 + 2(\hat{\beta} - \hat{\beta}_C)^T \underbrace{\mathbb{X}^T(Y - \mathbb{X}\hat{\beta})}_{=\mathbb{X}^T u = 0}. \end{aligned}$$

Hence

$$SS_e(\hat{\beta}_C) = SS_e + (\hat{\beta} - \hat{\beta}_C)^T \mathbb{X}^T \mathbb{X} (\hat{\beta} - \hat{\beta}_C).$$

From (2.7) we know that

$$\hat{\beta} - \hat{\beta}_C = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T [\mathbb{C}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\beta} - \mathbf{c}).$$

Plug it into the previous expression and after canceling unnecessary terms we get

$$SS_e(\hat{\beta}_C) = SS_e + (\mathbb{C}\hat{\beta} - \mathbf{c})^T [\mathbb{C}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\beta} - \mathbf{c}). \quad (2.8)$$

This result will play an important role in Chapter 4.

**Summary:** In this section, we have derived two important results for the least squares estimator  $\hat{\beta}_C$  calculated under linear restrictions  $\mathbb{C}\beta = \mathbf{c}$ :

$$\begin{aligned} \hat{\beta}_C &= \hat{\beta} - (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T [\mathbb{C}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\beta} - \mathbf{c}), \\ SS_e(\hat{\beta}_C) &= SS_e + (\mathbb{C}\hat{\beta} - \mathbf{c})^T [\mathbb{C}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\beta} - \mathbf{c}). \end{aligned}$$

### 3. Properties of the Least Squares Estimator

In this chapter, we start investigating probabilistic and statistical properties of the quantities that were introduced in the previous chapter. The first two sections apply to the general linear regression model, the third section requires the additional condition of normality of the responses (or of the error terms).

#### 3.1. Moment Properties of the Least Squares Estimator

Consider the regression model

$$Y = \mathbb{X}\beta + \varepsilon$$

with  $E\varepsilon = \mathbf{0}$  and  $\text{var } \varepsilon = \sigma_e^2 \mathbb{I}_n$  or, equivalently,  $EY = \mathbb{X}\beta$  and  $\text{var } Y = \sigma_e^2 \mathbb{I}_n$ . Let the regression matrix  $\mathbb{X}_{n \times p}$  have a full rank  $p < n$ . The least squares estimator  $\hat{\beta}$  can be expressed as

$$\hat{\beta} = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T Y.$$

The first lemma specifies the first and the second moment of  $\hat{\beta}$  (conditionally on the covariates).

**Lemma 3.1.**

- (i)  $E\hat{\beta} = \beta$ , i.e.,  $\hat{\beta}$  is an unbiased estimator of  $\beta$ .
- (ii)  $\text{var } \hat{\beta} = \sigma_e^2 (\mathbb{X}^T \mathbb{X})^{-1}$ .

◇

**Proof.** Treating  $\mathbb{X}$  as a matrix of constants and  $Y$  as a random vector, we get:

$$E\hat{\beta} = E(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T Y = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T EY = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \mathbb{X}\beta = \beta$$

and

$$\begin{aligned} \text{var } \hat{\beta} &= \text{var}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T Y = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \text{var } Y \mathbb{X} (\mathbb{X}^T \mathbb{X})^{-1} \\ &= \sigma_e^2 (\mathbb{X}^T \mathbb{X})^{-1} (\mathbb{X}^T \mathbb{X}) (\mathbb{X}^T \mathbb{X})^{-1} = \sigma_e^2 (\mathbb{X}^T \mathbb{X})^{-1}. \end{aligned}$$

□

The second lemma specifies the first and the second moments of the fitted values and residuals. Its proof is also straightforward.

**Lemma 3.2.**

- (i)  $E\hat{Y} = EY = X\beta$ ,
- (ii)  $E\mathbf{u} = \mathbf{0}$ ,
- (iii)  $\text{var}\hat{Y} = \sigma_e^2 H$ ,
- (iv)  $\text{var}\mathbf{u} = \sigma_e^2 (\mathbb{I} - H)$ .

◇

**Proof.** We have  $\hat{Y} = HY$  and  $\mathbf{u} = (\mathbb{I} - H)Y$ , where  $H = X(X^T X)^{-1} X^T$  is the projection matrix to the subspace  $\mathcal{M}(X)$ .  $H$  is symmetric, idempotent, and satisfies  $HX = X$  and  $(\mathbb{I} - H)X = \mathbf{0}$ . Hence

$$\begin{aligned} E\hat{Y} &= EHY = HEY = HX\beta = X\beta, \\ \text{var}\hat{Y} &= \text{var}HY = H\text{var}YH = \sigma_e^2 HH = \sigma_e^2 H. \end{aligned}$$

Next,

$$\begin{aligned} E\mathbf{u} &= E(\mathbb{I} - H)Y = (\mathbb{I} - H)EY = (\mathbb{I} - H)X\beta = \mathbf{0}, \\ \text{var}\mathbf{u} &= \text{var}(\mathbb{I} - H)Y = (\mathbb{I} - H)\text{var}Y(\mathbb{I} - H) = \sigma_e^2 (\mathbb{I} - H)(\mathbb{I} - H) = \sigma_e^2 (\mathbb{I} - H). \end{aligned} \quad \square$$

It is important to realize one substantial difference. We can write the responses in two different ways:

$$\begin{aligned} Y &= X\beta + \varepsilon, \\ Y &= X\hat{\beta} + \mathbf{u} \end{aligned}$$

In the first case, the error terms  $\varepsilon$  are independent and have equal variances. However, in the second case, the residuals  $\mathbf{u}$  do not share these properties: they are not independent (because the matrix  $\mathbb{I} - H$  is not diagonal) and they do not have equal variances.

Finally, we calculate the expectation of the residual sum of squares and derive an unbiased estimator for the residual variance.

**Lemma 3.3.**  $ESS_e = (n - p)\sigma_e^2$ .

◇

**Proof.** Remembering the results from Section 2.4, we can write  $SS_e = \mathbf{u}^T \mathbf{u} = \varepsilon^T (\mathbb{I} - H) \varepsilon$ . By Lemma A.1 in the Appendix and using the fact that  $\mathbb{I} - H$  is idempotent of rank  $n - p$  — see equation (2.3) — we get

$$\begin{aligned} ESS_e &= E\varepsilon^T (\mathbb{I} - H) \varepsilon = 0 + \text{tr}[(\mathbb{I} - H)\text{var}\varepsilon] \\ &= \sigma_e^2 \text{tr}(\mathbb{I} - H) = \sigma_e^2 r(\mathbb{I} - H) = \sigma_e^2 (n - p). \end{aligned} \quad \square$$

**Definition 3.1.**

$$\hat{\sigma}_e^2 = \frac{SS_e}{n - p} \equiv MS_e$$

is called *the estimated residual variance*. The symbol  $MS_e$  is just an alternative notation for the expression  $SS_e/(n - p)$ . ▽



By Lemma 3.3,  $\hat{\sigma}_e^2$  is an unbiased estimator of the residual variance.

### 3.2. Gauss-Markov Theorem

The Gauss-Markov theorem shows that the least squares estimator is in a certain sense optimal. It was originally formulated by Carl Friedrich Gauss in 1821 (Gauss 1821) under the assumption of normality. It was extended to the general case by Andrey Andreyevich Markov in 1912 (Markov 1912). Further extension to correlated errors of unequal variance was provided by Aitken (1936).\*

*Andrey Andreyevich Markov (1856 – 1922) was a Russian mathematician, who became particularly famous for his pioneering work on stochastic processes (Markov property, Markov chains, etc.).*

Source: [https://en.wikipedia.org/wiki/Andrey\\_Markov](https://en.wikipedia.org/wiki/Andrey_Markov)

Here we state the Gauss-Markov theorem in three different ways, after we introduce and explain the optimality criterion needed for all three versions.

**Definition 3.2.**  $\hat{\theta}$  is best linear unbiased estimator (BLUE) of  $\theta$  based on the data vector  $Y$  if and only if the following three conditions hold:

- (i)  $\hat{\theta}$  is linear, i.e.,  $\hat{\theta} = \mathbb{A}Y$ .
- (ii)  $\hat{\theta}$  is unbiased, i.e.,  $E\hat{\theta} = E\mathbb{A}Y = \theta$ .
- (iii) For any matrix  $\mathbb{B}$  (of the same dimension as  $\mathbb{A}$ ) that satisfies  $E\mathbb{B}Y = \theta$

$$\text{var } \mathbb{B}Y - \text{var } \hat{\theta} \geq 0,$$

that is, the matrix on the left-hand side is positive semi-definite.  $\nabla$

**Theorem 3.4 (Gauss-Markov, version I).** Let the linear regression model specified in Section 3.1 on page 31 be satisfied, let  $\hat{\beta}$  be the LSE. Then  $c^T \hat{\beta}$  is the unique best linear unbiased estimator of  $c^T \beta$  for any  $0 \neq c \in \mathbb{R}^p$ .  $\diamond$

**Proof.**

- $c^T \hat{\beta}$  is linear:  $c^T \hat{\beta} = c^T (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T Y$ .
- $c^T \hat{\beta}$  is unbiased:  $E c^T \hat{\beta} = c^T E \hat{\beta} = c^T \beta$ .
- $c^T \hat{\beta}$  has the smallest variance among all linear unbiased estimators:

Take another linear unbiased estimator  $a^T Y$  of  $c^T \beta$ , where  $a \in \mathbb{R}^n$ . We have  $E a^T Y = a^T \mathbb{X} \beta = c^T \beta$ . Hence,  $a^T \mathbb{X} = c^T$ . Now,

$$\text{var } a^T Y = \sigma_e^2 a^T a,$$

$$\text{var } c^T \hat{\beta} = \sigma_e^2 c^T (\mathbb{X}^T \mathbb{X})^{-1} c = \sigma_e^2 a^T \mathbb{X} (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T a = \sigma_e^2 a^T \mathbb{H} a.$$

\* we do not talk about that extension in this course

Finally,

$$\text{var } \mathbf{a}^\top \mathbf{Y} - \text{var } \mathbf{c}^\top \hat{\boldsymbol{\beta}} = \sigma_e^2 \mathbf{a}^\top (\mathbb{I} - \mathbb{H}) \mathbf{a} \geq 0$$

because  $\mathbb{I} - \mathbb{H}$  is positive semi-definite. The variances of both estimators are equal if and only if  $(\mathbb{I} - \mathbb{H})\mathbf{a} = \mathbf{0}$ , which is equivalent to  $\mathbf{a} = \mathbb{H}\mathbf{a}$  or  $\mathbf{a}^\top = \mathbf{a}^\top \mathbb{H}$ . It follows that the estimator  $\mathbf{a}^\top \mathbf{Y}$  can be rewritten as

$$\mathbf{a}^\top \mathbf{Y} = \mathbf{a}^\top \mathbb{H} \mathbf{Y} = \mathbf{a}^\top \mathbb{X} (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top \mathbf{Y} = \mathbf{a}^\top \mathbb{X} \hat{\boldsymbol{\beta}} = \mathbf{c}^\top \hat{\boldsymbol{\beta}}. \quad \square$$

**Theorem 3.5 (Gauss-Markov, version II).** *Let the linear regression model specified in Section 3.1 on page 31 be satisfied, let  $\hat{\boldsymbol{\beta}}$  be the LSE and  $\mathbb{C}$  any  $q \times p$  matrix. Then  $\mathbb{C}\hat{\boldsymbol{\beta}}$  is a best linear unbiased estimator of  $\mathbb{C}\boldsymbol{\beta}$ .*  $\diamond$

**Proof.** This is an easy corollary of the preceding theorem.  $\mathbb{C}\hat{\boldsymbol{\beta}}$  is obviously a linear and unbiased estimator of  $\mathbb{C}\boldsymbol{\beta}$ . Consider another linear unbiased estimator  $\mathbb{A}\mathbf{Y}$  with  $\mathbb{A}_{q \times n}$ . To be unbiased, it must satisfy  $\mathbb{E}\mathbb{A}\mathbf{Y} = \mathbb{A}\mathbb{X}\boldsymbol{\beta} = \mathbb{C}\boldsymbol{\beta}$  and hence  $\mathbb{A}\mathbb{X} = \mathbb{C}$  and  $r(\mathbb{A}) = r(\mathbb{C})$ .

Denote  $\mathbb{D} = \text{var } \mathbb{A}\mathbf{Y} - \text{var } \mathbb{C}\hat{\boldsymbol{\beta}}$  and prove that  $\mathbb{D} \geq 0$  by taking any non-zero vector  $\mathbf{d} \in \mathbb{R}^q$  and showing that  $\mathbf{d}^\top \mathbb{D} \mathbf{d} \geq 0$ . We have  $\mathbf{d}^\top \mathbb{D} \mathbf{d} = \text{var } \mathbf{d}^\top \mathbb{A}\mathbf{Y} - \text{var } \mathbf{d}^\top \mathbb{C}\hat{\boldsymbol{\beta}}$ . Also,  $\mathbb{E} \mathbf{d}^\top \mathbb{A}\mathbf{Y} = \mathbf{d}^\top \mathbb{A}\mathbb{X}\boldsymbol{\beta} = \mathbf{d}^\top \mathbb{C}\boldsymbol{\beta}$ .

Hence,  $\mathbf{d}^\top \mathbb{A}\mathbf{Y}$  is a linear unbiased estimator of  $\mathbf{d}^\top \mathbb{C}\boldsymbol{\beta}$  and it follows from Theorem 3.4 that  $\mathbf{d}^\top \mathbb{D} \mathbf{d} \geq 0$ .  $\square$

**Note.** It follows from Theorem 3.5 that  $\hat{\boldsymbol{\beta}}$  is the BLUE of  $\boldsymbol{\beta}$ . Just take a special case with  $\mathbb{C} = \mathbb{I}$ .

Another special case of Theorem 3.5, with  $\mathbb{C} = \mathbb{X}$ , shows that  $\hat{\mathbf{Y}}$  is the BLUE of  $\mathbb{E}\mathbf{Y}$ . This produces the third version of the Gauss-Markov theorem.

**Theorem 3.6 (Gauss-Markov, version III).** *Let the linear regression model specified in Section 3.1 on page 31 be satisfied, let  $\hat{\boldsymbol{\beta}}$  be the LSE. Then  $\hat{\mathbf{Y}}$  is the best linear unbiased estimator of  $\mathbb{E}\mathbf{Y}$ .*  $\diamond$

### 3.3. Properties of the Least Squares Estimator Under Normality

In this section, we consider the linear regression model with the normality assumption. In particular,

$$\mathbf{Y} = \mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

with

$$\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}, \sigma_e^2 \mathbb{I}_n) \quad \text{or, equivalently,} \quad \mathbf{Y} \sim N_n(\mathbb{X}\boldsymbol{\beta}, \sigma_e^2 \mathbb{I}_n).$$

The regression matrix  $\mathbb{X}_{n \times p}$  still has a full rank  $p < n$ . All the results of the previous two sections are still valid. Under normality, we can derive additional results about distributions of various quantities, which are summarized in the following lemma.

**Lemma 3.7.** *Under the assumptions of the current section,*

- (i)  $\hat{\boldsymbol{\beta}} \sim N_p(\boldsymbol{\beta}, \sigma_e^2(\mathbb{X}^T \mathbb{X})^{-1});$
- (ii)  $\hat{Y} \sim N_n(\mathbb{X}\boldsymbol{\beta}, \sigma_e^2 \mathbb{H});$
- (iii)  $\mathbf{u} \sim N_n(\mathbf{0}, \sigma_e^2(\mathbb{I} - \mathbb{H}));$
- (iv)  $\frac{SS_e}{\sigma_e^2} \sim \chi_{n-p}^2;$
- (v)  $\hat{\boldsymbol{\beta}}$  and  $SS_e$  are independent. ◇

**Proof.**

- (i)–(iii) This is obvious:  $\hat{\boldsymbol{\beta}}$ ,  $\hat{Y}$ , and  $\mathbf{u}$  are just linear transformations of  $\mathbf{Y}$ . The first and second moments have been provided by Lemmas 3.1 and 3.2.
- (iv) As shown in Section 2.4,  $SS_e = \boldsymbol{\varepsilon}^T(\mathbb{I} - \mathbb{H})\boldsymbol{\varepsilon}$ , where  $\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}, \sigma_e^2 \mathbb{I}_n)$ . Because  $\mathbb{I} - \mathbb{H}$  is idempotent of rank  $n - p$  it follows from Lemma A.2 in the Appendix that  $SS_e/\sigma_e^2 \sim \chi_{n-p}^2$ .
- (v) We have  $\hat{\boldsymbol{\beta}} = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T \mathbf{Y} \equiv \mathbb{B}\mathbf{Y}$  and  $\sigma_e^2 = \mathbf{Y}^T(\mathbb{I} - \mathbb{H})\mathbf{Y} \equiv \mathbf{Y}\mathbb{A}\mathbf{Y}$ . By Lemma A.3 in the Appendix it suffices to show that  $\mathbb{B}\mathbb{A} = \mathbf{0}$ . But

$$\mathbb{B}\mathbb{A} = (\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T(\mathbb{I} - \mathbb{H}) = \mathbf{0}$$

because  $(\mathbb{I} - \mathbb{H})\mathbb{X} = \mathbf{0}$ . □

The linear regression model with normally distributed responses is a parametric model. Let us derive the maximum likelihood estimators (MLE) of  $\boldsymbol{\beta}$  and  $\sigma_e^2$ .

We have  $\mathbf{Y} \sim N_n(\mathbb{X}\boldsymbol{\beta}, \sigma_e^2 \mathbb{I}_n)$  with unknown parameters  $\boldsymbol{\theta} = (\boldsymbol{\beta}^T, \sigma_e^2)^T$ . The likelihood is

$$L(\boldsymbol{\theta} | \mathbf{Y}) = \frac{1}{(2\pi)^{n/2}(\sigma_e^2)^{n/2}} e^{-\frac{1}{2\sigma_e^2}(\mathbf{Y} - \mathbb{X}\boldsymbol{\beta})^T(\mathbf{Y} - \mathbb{X}\boldsymbol{\beta})}$$

and the log-likelihood

$$\ell(\boldsymbol{\beta}, \sigma_e^2 | \mathbf{Y}) = -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log \sigma_e^2 - \frac{1}{2\sigma_e^2}(\mathbf{Y} - \mathbb{X}\boldsymbol{\beta})^T(\mathbf{Y} - \mathbb{X}\boldsymbol{\beta}).$$

Regardless of  $\sigma_e^2$ , to maximize this over  $\boldsymbol{\beta}$  it is enough to minimize  $\|\mathbf{Y} - \mathbb{X}\boldsymbol{\beta}\|^2 = SS_e(\boldsymbol{\beta})$ . So, the least squares estimator  $\hat{\boldsymbol{\beta}}$  is the maximum likelihood estimator of  $\boldsymbol{\beta}$  in the normal linear regression model. Plug this into the log-likelihood to find the MLE of  $\sigma_e^2$ :

$$\ell(\hat{\boldsymbol{\beta}}, \sigma_e^2 | \mathbf{Y}) = -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log \sigma_e^2 - \frac{1}{2\sigma_e^2} SS_e. \quad (3.1)$$

Now,

$$\frac{\partial \ell(\hat{\boldsymbol{\beta}}, \sigma_e^2 | \mathbf{Y})}{\partial \sigma_e^2} = -\frac{n}{2} \cdot \frac{1}{\sigma_e^2} + \frac{1}{2} \frac{SS_e}{(\sigma_e^2)^2}$$

The MLE solves the equation

$$\frac{n}{\sigma_e^2} = \frac{SS_e}{(\sigma_e^2)^2}$$

and the solution is  $SS_e/n$ . We have proven the following lemma.

**Lemma 3.8.** *In the normal linear regression model, the maximum likelihood estimator of  $\boldsymbol{\beta}$  is the LSE  $\hat{\boldsymbol{\beta}} = (\mathbb{X}^T \mathbb{X})^{-1}(\mathbb{X}^T \mathbf{Y})$  and the maximum likelihood estimator of  $\sigma_e^2$  is  $SS_e/n$ .  $\diamond$*

**Note.** The MLE of  $\sigma_e^2$  differs from the unbiased estimator  $\hat{\sigma}_e^2$  of Definition 3.1 by dividing  $SS_e$  with  $n$  instead of  $n - 1$ . This difference becomes negligible as  $n$  increases.

## 4. Statistical Inference in the Linear Regression Model

### 4.1. Exact Inference Under Normality

In this section, we work under the assumption of normality, when the regression model can be formulated as

$$Y \sim N_n(\mathbb{X}\beta, \sigma_e^2 \mathbb{I}_n)$$

and we use the results of Section 3.3, in particular, Lemma 3.7. First, we formulate the exact distribution of the normalized linear combination of estimated regression coefficients.

**Lemma 4.1.** *Under the assumptions of the current section, for any  $\mathbf{c} \neq \mathbf{0}$ ,*

$$\frac{\mathbf{c}^\top \hat{\boldsymbol{\beta}} - \mathbf{c}^\top \boldsymbol{\beta}}{\sqrt{\hat{\sigma}_e^2 \mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}}} \sim t_{n-p}.$$

$t_{n-p}$  is the Student's  $t$ -distribution with  $n - p$  degrees of freedom. ◇

**Proof.** By Lemma 3.7, part (i),

$$\mathbf{c}^\top \hat{\boldsymbol{\beta}} \sim N(\mathbf{c}^\top \boldsymbol{\beta}, \sigma_e^2 \mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c})$$

and hence

$$\frac{\mathbf{c}^\top \hat{\boldsymbol{\beta}} - \mathbf{c}^\top \boldsymbol{\beta}}{\sigma_e \sqrt{\mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}}} \sim N(0, 1).$$

By Lemma 3.7, parts (iv) and (v),  $\frac{SS_e}{\sigma_e^2} \sim \chi_{n-p}^2$  and  $\hat{\boldsymbol{\beta}}$  and  $SS_e$  are independent. It follows that

$$\frac{\frac{\mathbf{c}^\top \hat{\boldsymbol{\beta}} - \mathbf{c}^\top \boldsymbol{\beta}}{\sigma_e \sqrt{\mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}}}}{\sqrt{\frac{SS_e}{\sigma_e^2 (n-p)}}} = \frac{\mathbf{c}^\top \hat{\boldsymbol{\beta}} - \mathbf{c}^\top \boldsymbol{\beta}}{\sqrt{\hat{\sigma}_e^2 \mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}}} \sim t_{n-p}. \quad \square$$

We can use this lemma to perform tests and construct confidence intervals for any linear combinations of regression coefficients. For example, if we take  $\mathbf{c} = \mathbf{e}_j$ , we get  $\mathbf{c}^\top \hat{\boldsymbol{\beta}} = \hat{\beta}_j$  and  $\mathbf{c}^\top \boldsymbol{\beta} = \beta_j$ . Next,  $\mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c} \equiv v_{jj}$ , where  $v_{jj}$  denotes the  $j$ -th diagonal element of  $(\mathbb{X}^\top \mathbb{X})^{-1}$ .

#### 4.1.1. Testing individual regression parameters

Consider the hypothesis  $H_0 : \beta_j = a$  against the two-sided alternative  $H_0 : \beta_j \neq a$  (usually, we take  $a = 0$ ). Based on Lemma 4.1, we reject  $H_0$  if

$$\frac{|\hat{\beta}_j - a|}{\sqrt{\hat{\sigma}_e^2 v_{jj}}} \geq t_{n-p}\left(1 - \frac{\alpha}{2}\right),$$

where  $t_{n-p}(1 - \alpha/2)$  is the  $(1 - \alpha/2)$ -quantile of  $t$  distribution with  $n - p$  degrees of freedom. Note that  $\hat{\sigma}_e^2 v_{jj}$  is the estimated variance of  $\hat{\beta}_j$ . This test has the exact level  $\alpha$ .

#### 4.1.2. Confidence intervals for individual regression parameters

Let  $\beta_j$  be the true value of the  $j$ -th regression parameter and  $\hat{\beta}_j$  be the LSE of  $\beta_j$ . By Lemma 4.1,

$$P\left[-t_{n-p}\left(1 - \frac{\alpha}{2}\right) \leq \frac{\hat{\beta}_j - \beta_j}{\sqrt{\hat{\sigma}_e^2 v_{jj}}} \leq t_{n-p}\left(1 - \frac{\alpha}{2}\right)\right] = 1 - \alpha.$$

By a simple manipulation, we get

$$P\left[\hat{\beta}_j - t_{n-p}\left(1 - \frac{\alpha}{2}\right)\sqrt{\hat{\sigma}_e^2 v_{jj}} \leq \beta_j \leq \hat{\beta}_j + t_{n-p}\left(1 - \frac{\alpha}{2}\right)\sqrt{\hat{\sigma}_e^2 v_{jj}}\right] = 1 - \alpha.$$

Thus,

$$\hat{\beta}_j \mp t_{n-p}\left(1 - \frac{\alpha}{2}\right)\sqrt{\hat{\sigma}_e^2 v_{jj}}$$

are the boundary points of a confidence interval for  $\beta_j$  with coverage probability exactly  $1 - \alpha$ .

#### 4.1.3. Tests and confidence intervals for linear combinations of regression parameters

Choose the desired  $\mathbf{c}$  and use Lemma 4.1 in the same way as in Sections 4.1.1 and 4.1.2.

#### 4.1.4. Simultaneous tests of several linear combinations of regression parameters

Consider a matrix of constants  $\mathbb{C}_{q \times p}$  with  $q \leq p$  and  $r(\mathbb{C}) = q$ . By Theorem 3.5 (Gauss-Markov, ver. II),  $\mathbb{C}\hat{\boldsymbol{\beta}}$  is a best linear unbiased estimator of  $\mathbb{C}\boldsymbol{\beta}$  even the data are not normal. The next lemma provides the exact distribution of  $\mathbb{C}\hat{\boldsymbol{\beta}}$  under normality.

**Lemma 4.2.** Under the assumptions of the current section, for any  $\mathbb{C}_{q \times p}$  with  $q \leq p$  and  $r(\mathbb{C}) = q$ ,

$$\frac{1}{q\hat{\sigma}_e^2}(\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta})^T [\mathbb{C}(\mathbf{X}^T \mathbf{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta}) \sim F_{q, n-p}. \quad (4.1)$$

$F_{q, n-p}$  is the Fisher's F-distribution with  $q$  and  $n - p$  degrees of freedom.  $\diamond$

**Proof.** By Lemma 3.7, part (i),

$$\mathbb{C}\hat{\boldsymbol{\beta}} \sim \mathcal{N}(\mathbb{C}\boldsymbol{\beta}, \sigma_e^2 \mathbb{C}(\mathbf{X}^T \mathbf{X})^{-1} \mathbb{C}^T)$$

and hence

$$\frac{1}{\sigma_e^2}(\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta})^T [\mathbb{C}(\mathbf{X}^T \mathbf{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta}) \sim \chi_q^2.$$

By Lemma 3.7, parts (iv) and (v),

$$\frac{SS_e}{\sigma_e^2} \sim \chi_{n-p}^2$$

and  $\hat{\boldsymbol{\beta}}$  and  $SS_e$  are independent. Hence

$$\frac{\frac{1}{q\hat{\sigma}_e^2}(\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta})^T [\mathbb{C}(\mathbf{X}^T \mathbf{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta})}{\frac{SS_e}{(n-p)\sigma_e^2}} \sim F_{q, n-p}.$$

This proves the lemma.  $\square$

This lemma can be used to perform simultaneous tests of several linear combinations of regression parameters. Put the desired coefficients into the  $q$  rows of the matrix  $\mathbb{C}$  and consider testing

$$H_0 : \mathbb{C}\boldsymbol{\beta} = \mathbf{c} \quad \text{against} \quad H_1 : \mathbb{C}\boldsymbol{\beta} \neq \mathbf{c},$$

where  $\mathbf{c}$  is a  $q$ -vector of constants (often zeros). The hypothesis is rejected when

$$\frac{1}{q\hat{\sigma}_e^2}(\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbf{c})^T [\mathbb{C}(\mathbf{X}^T \mathbf{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbf{c}) \geq F_{q, n-p}(1 - \alpha),$$

where  $F_{q, n-p}(1 - \alpha)$  is the  $(1 - \alpha)$ -quantile of the F distribution with  $q$  and  $n - p$  degrees of freedom. This test has the exact level  $\alpha$ .

The test statistic can be expressed in a more convenient form if we realize that the hypothesis  $H_0$  specifies a set of linear constraints on the regression parameters. We know from section 2.11, equation (2.8) that the numerator of the test statistic can be written as

$$(\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbf{c})^T [\mathbb{C}(\mathbf{X}^T \mathbf{X})^{-1} \mathbb{C}^T]^{-1} (\mathbb{C}\hat{\boldsymbol{\beta}} - \mathbf{c}) = SS_e(\hat{\boldsymbol{\beta}}_C) - SS_e,$$

where  $SS_e(\hat{\boldsymbol{\beta}}_C)$  is the residual sum of squares under linear constraints, that is, under  $H_0$ , and  $SS_e$  is the residual sum of squares without restrictions, i.e., when  $H_0$  is not assumed to hold.

This consideration allows us to summarize the previous results in the following way.

**Lemma 4.3.** When the hypothesis  $H_0 : \mathbb{C}\boldsymbol{\beta} = \mathbf{c}$  is true, where  $\mathbb{C}$  is a  $q \times p$  matrix of constants with  $q \leq p$  and  $r(\mathbb{C}) = q$ , and  $\mathbf{c}$  is a  $q$ -vector of constants, then

$$\frac{n-p}{q} \frac{SS_e^0 - SS_e}{SS_e} \sim F_{q, n-p},$$

where  $SS_e^0$  is the residual sum of squares calculated under  $H_0$  and  $SS_e$  is the residual sum of squares calculated without any restrictions.

The hypothesis is rejected if

$$\frac{n-p}{q} \frac{SS_e^0 - SS_e}{SS_e} \geq F_{q, n-p}(1 - \alpha).$$

This test has the exact level  $\alpha$ . ◇

The lemma provides a much easier formulation of the test, which can be further extended to any submodel testing (see the next section). The numerator of the test statistic expresses how much  $SS_e$  increased under  $H_0$  relative to what it was when  $H_0$  was not assumed to hold.

From Section 2.9, the numerator can be rewritten in terms of regression sums of squares as well. In the model  $\mathbf{Y} = \mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$ , we decompose the total sum of squares as  $SS_T = SS_R + SS_e$ . Under the null hypothesis, the same total sum of squares is decomposed as  $SS_T = SS_R^0 + SS_e^0$ . The numerator

$$SS_e^0 - SS_e = (SS_T - SS_R^0) - (SS_T - SS_R) = SS_R - SS_R^0$$

shows how much the regression sum of squares improved after removing the restriction imposed by the null hypothesis.

#### 4.1.5. Submodel testing

Consider two different regression models for the same response  $\mathbf{Y}$  (with  $q \geq 1$ ):

$$\begin{array}{ll} \mathbf{Y} = \mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \text{ with } \mathbb{X}_{n \times p}, \boldsymbol{\beta}_{p \times 1} & \text{Model } (M) \\ \text{and } \mathbf{Y} = \mathbb{Z}\boldsymbol{\gamma} + \boldsymbol{\varepsilon}^*, \text{ with } \mathbb{Z}_{n \times (p-q)}, \boldsymbol{\gamma}_{(p-q) \times 1} & \text{Model } (M_1) \end{array}$$

Model  $M_1$  is called a *submodel* of Model  $M$  if and only if  $\mathcal{M}(\mathbb{Z}) \subset \mathcal{M}(\mathbb{X})$ , that is, the linear space generated by the columns of  $\mathbb{Z}$  is a subspace of the linear space generated by the columns of  $\mathbb{X}$ . The submodel explains the response through fewer covariates and fewer parameters. Whenever the submodel  $M_1$  is true the larger model  $M$  must also be true.

Because  $\mathcal{M}(\mathbb{Z}) \subset \mathcal{M}(\mathbb{X})$ , each column of  $\mathbb{Z}$  can be expressed as a linear combination of the columns of  $\mathbb{X}$ . Thus, there exists a  $p \times (p - q)$  matrix  $\mathbb{A}$  such that  $\mathbb{Z} = \mathbb{X}\mathbb{A}$ . Take



any matrix  $\mathbb{B}_{p \times q}$  such that  $(\mathbb{A}|\mathbb{B})$  is of full rank. Then  $(\mathbb{X}\mathbb{A}|\mathbb{X}\mathbb{B})$  is another basis of  $\mathcal{M}(\mathbb{X})$ . Therefore,  $\mathbb{X}\boldsymbol{\beta}$  can be written as

$$\mathbb{X}\boldsymbol{\beta} = \mathbb{X}\mathbb{A}\boldsymbol{\gamma} + \mathbb{X}\mathbb{B}\boldsymbol{\delta} = \mathbb{Z}\boldsymbol{\gamma} + \mathbb{X}\mathbb{B}\boldsymbol{\delta}$$

Model  $M$  is equivalent to the model

$$Y = \mathbb{Z}\boldsymbol{\gamma} + \mathbb{X}\mathbb{B}\boldsymbol{\delta} + \boldsymbol{\varepsilon}$$

and the submodel  $M_1$  is true if and only if  $\boldsymbol{\delta} = \mathbf{0}$ . Thus, any submodel can be obtained by linear restriction testing applied on a model that is equivalent to the larger model  $M$ .

This proves that the test specified in Lemma 4.3 can be applied to any submodel testing problem. To be specific, let

- $SS_e^0$  be the residual sum of squares under the submodel;
- $SS_e$  be the residual sum of squares under the larger model;
- $q$  be the difference in the number of parameters between the larger model and the submodel.

If the submodel holds then

$$\frac{n-p}{q} \frac{SS_e^0 - SS_e}{SS_e} \sim F_{q, n-p}.$$

The submodel is rejected in favor of the larger model if

$$\frac{n-p}{q} \frac{SS_e^0 - SS_e}{SS_e} \geq F_{q, n-p}(1 - \alpha).$$

Submodel testing is the most important tool for building regression models, that is, for deciding which covariates should be included in the model and in what functional form.

#### 4.1.6. Overall regression test

Let us consider a special case of submodel (or linear restriction) testing. Take  $\mathbb{C} = (\mathbf{0}|\mathbb{I}_{p-1})$  and test  $H_0 : \mathbb{C}\boldsymbol{\beta} = \mathbf{0}$ . This is equivalent to  $\beta_2 = \cdots = \beta_p = 0$ , that is, all regression coefficients except the intercept are zero. The number of tested parameters is  $q = p - 1$ . The submodel contains only the intercept. When this hypothesis is true, the responses have the same expectation, which does not (linearly) depend on any of the covariates.

Recall the decomposition of centered sums of squares derived in Section 2.9. We have  $SS_T = SS_R + SS_e$ , or explicitly

$$\sum_{i=1}^n (Y_i - \bar{Y})^2 = \sum_{i=1}^n (\hat{Y}_i - \bar{Y})^2 + \sum_{i=1}^n (Y_i - \hat{Y}_i)^2.$$

Under  $H_0$  (intercept only), we have  $\hat{Y}_i^0 = \bar{Y}$  for all  $i$ ,  $SS_R^0 = 0$ ,  $SS_e^0 = SS_T$ . The test statistic for testing this hypothesis is

$$\frac{n-p}{p-1} \frac{SS_e^0 - SS_e}{SS_e} = \frac{n-p}{p-1} \frac{SS_R}{SS_e}.$$

This can be also expressed as

$$\frac{n-p}{p-1} \frac{SS_R}{SS_T - SS_R} = \frac{n-p}{p-1} \frac{R^2}{1-R^2},$$

where  $R^2$  is the coefficient of determination (see Section 2.10). The hypothesis that no covariates affect the expectation can be rejected if this test statistic exceeds  $F_{p-1, n-p}(1-\alpha)$ .

The calculation of the overall regression test is traditionally visualized in the form of so called analysis-of-variance (ANOVA) table, see Table 4.1.

#### 4.1.7. One-way analysis of variance model

Let the observed data be  $(Y_i, Z_i)$  independent pairs,  $i = 1, \dots, n$ , where  $Y_i$  is the response and  $Z_i \in \{1, 2, \dots, m\}$  classifies the subjects into one of  $m$  disjoint groups. Let the expectation of the response in the  $j$ -th group be  $\mu_j = E[Y_i | Z_i = j]$ ,  $j = 1, \dots, m$ . Let the conditional distribution of the response in the  $j$ -th group be  $N(\mu_j, \sigma_e^2)$ , that is, all observations have normal distributions with potentially different means in the  $m$  groups and equal variances. This is the classical one-way analysis of variance (ANOVA) model.

The one-way ANOVA model can be formulated as a linear regression model by

$$\begin{aligned} Y_i &= \sum_{j=1}^m \mu_j \mathbb{1}(Z_i = j) + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma_e^2) \\ &= \mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i, \end{aligned}$$

where  $\mathbf{X}_i = \mathbf{e}_j$  when  $Z_i = j$  and  $\boldsymbol{\beta} = (\mu_1, \dots, \mu_m)^\top$ .

Denote the group sizes by  $n_j = \sum_{i=1}^n \mathbb{1}(Z_i = j)$ . Sort the observations so that the  $n_1$  observations coming from group 1 are listed first, followed by the  $n_2$  observations belonging

Source of variation	SS	d.f.	MS	F
All covariates	$SS_R$	$p-1$	$MS_R = \frac{SS_R}{p-1}$	$F = \frac{MS_R}{MS_e}$
Error	$SS_e$	$n-p$	$MS_e = \frac{SS_e}{n-p}$	
Total	$SS_T$	$n-1$		

Table 4.1.: Analysis of variance table for the overall regression test.

to group 2 and so on. Consider the corresponding regression matrix  $\mathbb{X}$ . The least squares estimator of  $\beta$  is

$$\hat{\beta} = (\mathbb{X}^T \mathbb{X})^{-1} (\mathbb{X}^T \mathbf{Y}).$$

We have

$$\mathbb{X}^T \mathbb{X} = \text{diag}(n_1, \dots, n_m) \quad \text{and} \quad \mathbb{X}^T \mathbf{Y} = \left( \sum_{i=1}^n Y_i \mathbb{1}(Z_i = 1), \dots, \sum_{i=1}^n Y_i \mathbb{1}(Z_i = m) \right).$$

Hence  $\hat{\beta} = (\bar{Y}_1, \dots, \bar{Y}_m)$ , where  $\bar{Y}_j$  is the arithmetic average of the observations belonging to the  $j$ -th group. This is also the least squares estimator of the expectation  $\mu_j$  in the  $j$ -th group.

The fitted values in the one-way ANOVA model are

$$\hat{Y}_i = \sum_{j=1}^m \bar{Y}_j \mathbb{1}(Z_i = j),$$

the regression sum of squares is

$$SS_R = \sum_{i=1}^n (\hat{Y}_i - \bar{Y})^2 = \sum_{i=1}^n \sum_{j=1}^m (\bar{Y}_j - \bar{Y})^2 \mathbb{1}(Z_i = j) = \sum_{j=1}^m n_j (\bar{Y}_j - \bar{Y})^2.$$

In classical ANOVA, this is also denoted by  $SS_A$ . The residual sum of squares is

$$SS_e = \sum_{i=1}^n (Y_i - \hat{Y}_i)^2 = \sum_{i=1}^n \sum_{j=1}^m (Y_i - \bar{Y}_j)^2 \mathbb{1}(Z_i = j).$$

Consider the hypothesis  $H_0 : \mu_1 = \mu_2 = \dots = \mu_m$  that all the groups have the same mean. This is equivalent to  $H_0 : \beta_1 = \beta_2 = \dots = \beta_m$ . Under this hypothesis, the data can be described by an intercept-only model and the test of this hypothesis is the overall regression test constructed in the previous section (with  $p = m$ ). The test statistic of the overall regression test is

$$F = \frac{n-m}{m-1} \frac{SS_R}{SS_e} = \frac{SS_R/(m-1)}{SS_e/(n-m)}$$

The hypothesis is rejected if  $F \geq F_{m-1, n-m}(1-\alpha)$ . This is the classical one-way ANOVA F-test. We have derived it as a special case of an overall regression test in a linear model.

#### 4.1.8. Connections to maximum likelihood theory

Separate the regression parameter into  $\beta = (\beta_1^T, \beta_2^T)^T$ , where  $\beta_1$  has  $p-q$  elements and  $\beta_2$   $q$  elements. Consider the hypothesis  $H_0 : \beta_2 = \mathbf{0}$ . This corresponds to a submodel test and also to the test of the hypothesis  $H_0 : \mathbb{C}\beta = \mathbf{0}$  with  $\mathbb{C} = (\mathbf{0} | \mathbb{I}_q)_{q \times p}$ .

By Lemma 4.2 and Lemma 4.3, the test statistic for this test can be expressed as

$$\frac{n-p}{q} \frac{SS_e^0 - SS_e}{SS_e} = \frac{1}{q\hat{\sigma}_e^2} (\mathbf{C}\hat{\boldsymbol{\beta}})^\top [\mathbf{C}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{C}^\top]^{-1} (\mathbf{C}\hat{\boldsymbol{\beta}}) \stackrel{H_0}{\sim} F_{q, n-p}.$$

Let us now consider the likelihood ratio test of the same hypothesis. First, consider  $\sigma_e^2$  known. The likelihood ratio statistic is

$$LR = 2[\ell(\hat{\boldsymbol{\beta}}) - \ell(\tilde{\boldsymbol{\beta}})],$$

where  $\tilde{\boldsymbol{\beta}}$  is the MLE (which is the same as the LSE) of  $\boldsymbol{\beta}$  calculated under the submodel. Let  $\tilde{\mathbf{Y}}$  be the fitted values in the submodel. The maximum likelihood theory stipulates that  $LR \xrightarrow{D} \chi_q^2$  when the submodel is true.

From (3.1), the log-likelihoods of the larger model and the submodel are

$$\begin{aligned} \ell(\hat{\boldsymbol{\beta}}) &= -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log \sigma_e^2 - \frac{1}{2\sigma_e^2} SS_e, \\ \ell(\tilde{\boldsymbol{\beta}}) &= -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log \sigma_e^2 - \frac{1}{2\sigma_e^2} SS_e^0. \end{aligned}$$

Thus,

$$LR = \frac{1}{\sigma_e^2} (SS_e^0 - SS_e) = \frac{1}{\sigma_e^2} (\mathbf{C}\hat{\boldsymbol{\beta}})^\top [\mathbf{C}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{C}^\top]^{-1} (\mathbf{C}\hat{\boldsymbol{\beta}}) \stackrel{H_0}{\sim} \chi_q^2$$

and its distribution under  $H_0$  is exact, not only asymptotic.

With unknown  $\sigma_e^2$ , we modify the LR test as follows:

$$F = \frac{\frac{1}{q} LR}{\frac{(n-p)\hat{\sigma}_e^2}{\sigma_e^2} / (n-p)} \stackrel{H_0}{\sim} F_{q, n-p}.$$

So, the F-test for submodel testing is equivalent to the likelihood ratio test. With some more effort we could prove that also the Wald test and Rao score test yield the same test statistic.

## 4.2. Asymptotic Inference Without Normality (Random Covariates)

In this section, we show that all the results derived in Section 4.1 under the assumption of normality can be extended to the general case as long as certain moment conditions are fulfilled. The results become asymptotic, though.

We assume that covariates are random and  $(Y_i, \mathbf{X}_i)$  are a random sample of independent identically distributed vectors drawn from some  $(p+1)$ -variate distribution.

Let the data satisfy the linear model

$$Y = \mathbb{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$$

with  $E \boldsymbol{\varepsilon} = \mathbf{0}$  and  $\text{var } \boldsymbol{\varepsilon} = \sigma_e^2 \mathbb{I}_n$  or  $E[Y_i | X_i] = X_i^\top \boldsymbol{\beta}$  and  $\text{var}[Y_i | X_i] = \sigma_e^2$ .

We assume finite second moments of the response and the covariates and linear independence of the components of the covariate vector. We still assume that  $r(\mathbb{X}) = p^*$ . No additional assumptions are imposed on the distributions of  $Y_i$  or  $\varepsilon_i$ .

**Assumption.**

(AS1)  $\text{var}[Y_i | X_i] = \sigma_e^2 < \infty$ ;

(AS2)  $E_X X_i X_i^\top \equiv \mathbb{V}_X < \infty$ ;

(AS3)  $\mathbb{V}_X > 0$  (full rank regular invertible matrix).

$E_X$  denotes the expectation over the marginal distribution of the covariates.

In this section, we need to distinguish three kinds of expectations. The notation we will use is as follows.

$$E_{(Y,X)} h(Y, X) = \int h(y, \mathbf{x}) f(y, \mathbf{x}) d\mu(y, \mathbf{x})$$

is the expectation with respect to the joint distribution of  $(Y, X)$ . The joint density with respect to the measure  $\mu$  is denoted by  $f(y, \mathbf{x})$ .

$$E h(Y, X) = E[h(Y, X) | X]$$

is the conditional expectation given the covariates.

$$E_X h(X) = \int h(\mathbf{x}) f(\mathbf{x}) d\nu(\mathbf{x})$$

is the expectation with respect to the marginal distribution of  $X$ . The marginal density with respect to the measure  $\nu$  is denoted by  $f(\mathbf{x})$ .

With this notation, we have

$$\begin{aligned} E_{(Y,X)} h(Y, X) &= E_X[E h(Y, X)], \\ \text{var}_{(Y,X)} h(Y, X) &= E_X \text{var} h(Y, X) + \text{var}_X E h(Y, X). \end{aligned}$$

We know from Lemma 3.1 that the least squares estimator  $\hat{\boldsymbol{\beta}}$  is unbiased under these circumstances. The LSE

$$\hat{\boldsymbol{\beta}} = (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top \mathbf{Y}.$$

is the unique solution to the system of normal equations  $\mathbb{X}^\top \mathbb{X} \hat{\boldsymbol{\beta}} = \mathbb{X}^\top \mathbf{Y}$  stated in (2.1).

---

\* See the note on page 18.

We can write

$$\mathbb{X}^T \mathbb{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n) \begin{pmatrix} \mathbf{X}_1^T \\ \vdots \\ \mathbf{X}_n^T \end{pmatrix} = \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^T$$

and  $\mathbb{X}^T \mathbf{Y} = \sum_{i=1}^n \mathbf{X}_i Y_i$ . Notice that by the weak law of large numbers,  $\frac{1}{n} \mathbb{X}^T \mathbb{X} \xrightarrow{P} \mathbb{V}_X$ . The normal equations can be rewritten as  $\mathbb{X}^T \mathbf{Y} - \mathbb{X}^T \mathbb{X} \hat{\boldsymbol{\beta}} = \mathbf{0}$ . For any  $\boldsymbol{\beta} \in \mathbb{R}^p$ , define

$$\mathbf{U}(\boldsymbol{\beta}) \equiv \mathbb{X}^T \mathbf{Y} - \mathbb{X}^T \mathbb{X} \boldsymbol{\beta} = \sum_{i=1}^n (\mathbf{X}_i Y_i - \mathbf{X}_i \mathbf{X}_i^T \boldsymbol{\beta}) = \sum_{i=1}^n \mathbf{X}_i (Y_i - \mathbf{X}_i^T \boldsymbol{\beta}).$$

Take a single term from the sum and denote it by

$$\mathbf{U}_i(\boldsymbol{\beta}) \equiv \mathbf{X}_i (Y_i - \mathbf{X}_i^T \boldsymbol{\beta}).$$

The LSE  $\hat{\boldsymbol{\beta}}$  is the single solution to the system of equations  $\mathbf{U}(\boldsymbol{\beta}) = \sum_{i=1}^n \mathbf{U}_i(\boldsymbol{\beta}) = \mathbf{0}$ . Thus,  $\mathbf{U}(\boldsymbol{\beta})$  plays the role of the score statistic, except that we do not have a parametric model and so cannot derive the score statistic from the likelihood. This kind of an ad-hoc score statistic is sometimes called a *pseudoscore*.

The next theorem shows that under the current assumptions the LSE  $\hat{\boldsymbol{\beta}}$  defined by this particular pseudoscore is a consistent and asymptotically normal estimator of the true  $\boldsymbol{\beta}$ .

**Theorem 4.4 (Asymptotic properties of the LSE without normality).** *Under the assumptions of the current section, when  $\boldsymbol{\beta}$  denotes the true regression parameters,*

- (i)  $\hat{\boldsymbol{\beta}} \xrightarrow{P} \boldsymbol{\beta}$  ( $\hat{\boldsymbol{\beta}}$  is consistent);
- (ii)  $\frac{1}{\sqrt{n}} \mathbf{U}(\boldsymbol{\beta}) \xrightarrow{D} N_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X)$ ;
- (iii)  $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{D} N_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X^{-1})$ . ◇

**Note.** Rewrite point (iii) to see the approximate distribution of  $\hat{\boldsymbol{\beta}}$ :

$$\begin{aligned} \sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) &\dot{\sim} N_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X^{-1}) \\ \sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) &\dot{\sim} N_p(\mathbf{0}, \sigma_e^2 (\frac{1}{n} \mathbb{X}^T \mathbb{X})^{-1}) \\ \hat{\boldsymbol{\beta}} - \boldsymbol{\beta} &\dot{\sim} N_p(\mathbf{0}, \sigma_e^2 (\mathbb{X}^T \mathbb{X})^{-1}) \\ \hat{\boldsymbol{\beta}} &\dot{\sim} N_p(\boldsymbol{\beta}, \sigma_e^2 (\mathbb{X}^T \mathbb{X})^{-1}) \end{aligned}$$

The exact distribution of  $\hat{\boldsymbol{\beta}}$  under normality is **exactly the same** as the approximate distribution of  $\hat{\boldsymbol{\beta}}$  without assuming normality.

**Proof (of Theorem 4.4).**

(i)

$$\hat{\beta} = (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top \mathbf{Y} = \left( \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top \right)^{-1} \left( \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i Y_i \right)$$

By the weak law of large numbers and continuous transformation theorem,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top &\xrightarrow{P} \mathbb{E}_X \mathbf{X}_i \mathbf{X}_i^\top = \mathbb{V}_X, \quad \left( \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top \right)^{-1} \xrightarrow{P} \mathbb{V}_X^{-1} \\ \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i Y_i &\xrightarrow{P} \mathbb{E}_{(Y,X)} \mathbf{X}_i Y_i = \mathbb{E}_X (\mathbf{X}_i \mathbb{E} Y_i) = \mathbb{E}_X (\mathbf{X}_i \mathbf{X}_i^\top \beta) = \mathbb{V}_X \beta. \end{aligned}$$

Hence,  $\hat{\beta} \xrightarrow{P} \mathbb{V}_X^{-1} \mathbb{V}_X \beta = \beta$ .

(ii)  $\mathbf{U}_i(\beta) = \mathbf{X}_i(Y_i - \mathbf{X}_i^\top \beta)$  are independent identically distributed random vectors. Calculate their first and second moments at the true  $\beta$ , noticing that  $\mathbf{U}_i(\beta) = \mathbf{X}_i \varepsilon_i$ .

$$\begin{aligned} \mathbb{E}_{(Y,X)} \mathbf{U}_i(\beta) &= \mathbb{E}_X \mathbb{E} \mathbf{X}_i \varepsilon_i = \mathbf{0}, \\ \text{var}_{(Y,X)} \mathbf{U}_i(\beta) &= \mathbb{E}_X \text{var} \mathbf{U}_i(\beta) + \underbrace{\text{var}_X \mathbb{E} \mathbf{U}_i(\beta)}_{=0} = \mathbb{E}_X \text{var} (\mathbf{X}_i \varepsilon_i) = \mathbb{E}_X \mathbf{X}_i \mathbf{X}_i^\top \sigma_e^2 = \mathbb{V}_X \sigma_e^2. \end{aligned}$$

By the central limit theorem for iid random vectors,

$$\frac{1}{\sqrt{n}} \mathbf{U}(\beta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{U}_i(\beta) \xrightarrow{D} \mathbf{N}_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X).$$

(iii) Consider the difference between  $\mathbf{U}(\beta)$  evaluated at the true  $\beta$  and at the LSE  $\hat{\beta}$ .

$$\mathbf{U}(\beta) - \underbrace{\mathbf{U}(\hat{\beta})}_{=0} = \sum_{i=1}^n [\mathbf{X}_i(Y_i - \mathbf{X}_i^\top \beta) - \mathbf{X}_i(Y_i - \mathbf{X}_i^\top \hat{\beta})] = \sum_{i=1}^n [\mathbf{X}_i \mathbf{X}_i^\top (\hat{\beta} - \beta)].$$

Next,

$$\frac{1}{\sqrt{n}} \mathbf{U}(\beta) = \left( \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top \right) \sqrt{n}(\hat{\beta} - \beta)$$

and

$$\sqrt{n}(\hat{\beta} - \beta) = \underbrace{\left( \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top \right)^{-1}}_{\xrightarrow{P} \mathbb{V}_X^{-1}} \underbrace{\frac{1}{\sqrt{n}} \mathbf{U}(\beta)}_{\xrightarrow{D} \mathbf{N}_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X)}.$$

By Slutsky's Theorem,

$$\sqrt{n}(\hat{\beta} - \beta) \xrightarrow{D} \mathbf{N}_p(\mathbf{0}, \sigma_e^2 \underbrace{\mathbb{V}_X^{-1} \mathbb{V}_X \mathbb{V}_X^{-1}}_{=\mathbb{V}_X^{-1}}).$$

□

**Note.** Consistence of  $\hat{\beta}$  can be proven directly, as shown in the proof of Theorem 4.4(i) but it also follows as a direct consequence of Theorem 4.4(iii). It is easy to see that if  $\hat{\beta}$  did not converge in probability to  $\beta$  then  $\sqrt{n}(\hat{\beta} - \beta)$  cannot converge in distribution.

By Lemma 3.3,  $\hat{\sigma}_e^2 = \frac{SS_e}{n-p}$  is an unbiased estimator of the residual variance  $\sigma_e^2$ . Now we need to show that this estimator is consistent.

**Lemma 4.5.** Under the assumptions of the current section,  $\hat{\sigma}_e^2 \xrightarrow{P} \sigma_e^2$ . ◇

**Proof.**

$$\hat{\sigma}_e^2 = \underbrace{\frac{n}{n-p}}_{\rightarrow 1} \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \hat{\beta})^2.$$

Further,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \hat{\beta})^2 &= \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \beta + \mathbf{X}_i^\top \beta - \mathbf{X}_i^\top \hat{\beta})^2 \\ &= \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \beta)^2 + \frac{1}{n} \sum_{i=1}^n [\mathbf{X}_i^\top (\beta - \hat{\beta})]^2 + \left[ \frac{2}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \beta) \mathbf{X}_i^\top \right] (\beta - \hat{\beta}) \end{aligned}$$

The first term converges in probability to  $E \varepsilon_i^2 = \text{var } \varepsilon_i = \sigma_e^2$ . The second term can be written as

$$\underbrace{(\beta - \hat{\beta})^\top}_{\xrightarrow{P} 0} \underbrace{\left( \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top \right)}_{\xrightarrow{P} \mathbb{V}_X} \underbrace{(\beta - \hat{\beta})}_{\xrightarrow{P} 0};$$

therefore, it converges to zero in probability. The third term also converges to zero in probability, because  $\hat{\beta} \xrightarrow{P} \beta$  and

$$\frac{2}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i^\top \beta) \mathbf{X}_i^\top \xrightarrow{P} 2E \varepsilon_i \mathbf{X}_i^\top = \mathbf{0}.$$

This completes the proof. □

Now we are ready to restate the key results from Section 4.1 in their asymptotic versions. Let us start with the asymptotic version of Lemma 4.1.

**Lemma 4.6.** Under the assumptions of the current section, for any  $\mathbf{c} \neq \mathbf{0}$ ,

$$\frac{\mathbf{c}^\top \hat{\beta} - \mathbf{c}^\top \beta}{\sqrt{\hat{\sigma}_e^2 \mathbf{c}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}}} \xrightarrow{D} N(0, 1). \quad \diamond$$



**Proof.** By Theorem 4.4(iii),

$$\begin{aligned}\sqrt{n}(\hat{\beta} - \beta) &\xrightarrow{D} N_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X^{-1}) \\ \sqrt{n}(\mathbf{c}^\top \hat{\beta} - \mathbf{c}^\top \beta) &\xrightarrow{D} N(0, \sigma_e^2 \mathbf{c}^\top \mathbb{V}_X^{-1} \mathbf{c}) \\ \frac{\mathbf{c}^\top \hat{\beta} - \mathbf{c}^\top \beta}{\sqrt{\sigma_e^2 \mathbf{c}^\top (n \mathbb{V}_X)^{-1} \mathbf{c}}} &\xrightarrow{D} N(0, 1)\end{aligned}$$

Replace  $\sigma_e^2$  by  $\hat{\sigma}_e^2$  and  $\mathbb{V}_X$  by  $\frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top$ . By Lemma 4.5 and Slutsky's Theorem, this does not change the limiting distribution.  $\square$

**Note.** Lemma 4.1 states that the distribution of the left-hand side is exactly  $t_{n-p}$  under normality. The limiting distribution in Lemma 4.6 is standard normal. Because the distribution function of  $t_{n-p}$  converges to the distribution function of the standard normal distribution as  $n \rightarrow \infty$ , Lemma 4.1 can be used as an asymptotic approximation when the responses are not normally distributed. Thus, the methods for testing and for constructing confidence intervals introduced in Sections 4.1.1 and 4.1.2 can be used with non-normal responses if the sample size is large enough.

Next, we formulate an asymptotic version of Lemma 4.2.

**Lemma 4.7.** Under the assumptions of the current section, for any  $\mathbb{C}_{q \times p}$  with  $q \leq p$  and  $r(\mathbb{C}) = q$ ,

$$\frac{1}{\hat{\sigma}_e^2} (\mathbb{C} \hat{\beta} - \mathbb{C} \beta)^\top [\mathbb{C} (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{C}^\top]^{-1} (\mathbb{C} \hat{\beta} - \mathbb{C} \beta) \xrightarrow{D} \chi_q^2 \quad \text{as } n \rightarrow \infty.$$

$\diamond$

**Proof.** By Theorem 4.4(iii),

$$\begin{aligned}\sqrt{n}(\hat{\beta} - \beta) &\xrightarrow{D} N_p(\mathbf{0}, \sigma_e^2 \mathbb{V}_X^{-1}) \\ \sqrt{n}(\mathbb{C} \hat{\beta} - \mathbb{C} \beta) &\xrightarrow{D} N_q(\mathbf{0}, \sigma_e^2 \mathbb{C} \mathbb{V}_X^{-1} \mathbb{C}^\top) \\ \frac{1}{\hat{\sigma}_e^2} (\mathbb{C} \hat{\beta} - \mathbb{C} \beta)^\top [\mathbb{C} (n \mathbb{V}_X)^{-1} \mathbb{C}^\top]^{-1} (\mathbb{C} \hat{\beta} - \mathbb{C} \beta) &\xrightarrow{D} \chi_q^2.\end{aligned}$$

Replace  $\sigma_e^2$  by  $\hat{\sigma}_e^2$  and  $\mathbb{V}_X$  by  $\frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i^\top$ . By Lemma 4.5 and Slutsky's Theorem, this does not change the limiting distribution.  $\square$

**Note.** Denote the left-hand side of the expression (4.1) in Lemma 4.2 by  $F$ . Lemma 4.2 claims that  $F \sim F_{q, n-p}$  under normality. It follows that  $qF \xrightarrow{D} \chi_q^2$  (which is exactly the claim of Lemma 4.7). Thus the claim of Lemma 4.2 can be considered as an asymptotic approximation when the responses are not normal. The same is true for Lemma 4.3.

When we test a submodel against a larger model and reject the submodel if

$$\frac{n-p}{q} \frac{SS_e^0 - SS_e}{SS_e} \geq F_{q, n-p}(1 - \alpha),$$

the test has the exact level  $\alpha$  when the responses are normal and a level that converges to  $\alpha$  if the responses are not normal and  $n \rightarrow \infty$ .

All the results derived in Sections 4.1.1–4.1.7 under the assumption of normality hold asymptotically even if normality is violated, as long as the assumptions (AS1)–(AS3) given on page 45 hold. Therefore, normality of  $Y_i$  or  $\varepsilon_i$  should not be considered a necessary condition for the validity of results obtained by linear regression analyses. Normality only makes asymptotic results exact.

**Note.** When the number of observations is large enough, we do not care whether the responses are normal or not. How large is “large enough”, though? The answer depends on the complexity of the model and the degree of violation of normality. In a simple linear regression model, 25 observations may be enough for the asymptotic results to provide acceptable approximation even if the responses are strongly non-normal. In complex models with many covariates and high-order interactions, we may need hundreds or thousands of observations.

**Note.** Trusting regression results obtained on small datasets is dangerous in both situations. If the data are not normal, asymptotics cannot be relied upon. The only possibility to proceed with the analysis is to *assume* that the data are normal. However, on a small dataset we cannot verify this assumption and therefore we cannot trust the results either.

### 4.3. Asymptotic Inference Without Normality (Fixed Covariates)

In the previous section, we have shown that the results derived under the assumption of normality can be used as asymptotic approximations when the covariates are random and satisfy certain moment conditions. In that setup, it is particularly easy to prove the asymptotic results because the data form a sequence of independent and identically distributed random vectors.

However, in certain applications such as industrial experiments the covariates cannot be considered random because their values are pre-determined by the experimenter and set to the desired values.\* In this section, we will state conditions under which the claims of the previous sections remain valid even if the covariates are constant. The proofs will be omitted, though.

---

\* Imagine evaluating the effect of temperature on the performance of a certain product. The temperature is set by the investigator to equidistant values such as 5°, 10°, 15°, 20° etc. These values are definitely not random.

#### 4. Statistical Inference in the Linear Regression Model

---

We assume that covariates  $\mathbf{x}_i$  are fixed vectors of constants. The observations are  $(Y_i, \mathbf{x}_i)$ ,  $i = 1, \dots, n$ , where  $Y_i$  are independent random variables.

Suppose the data follow the linear model

$$Y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \varepsilon_i$$

where  $E\varepsilon_i = 0$  and  $\text{var}\varepsilon_i = \sigma_\varepsilon^2 < \infty$ . Let the regression matrix be of full rank. The distribution of  $Y_i$  is otherwise arbitrary.

Redefine  $\mathbb{V}_X = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$ . [Huber \(1973\)](#) formulated the following condition for the validity of asymptotic results.

**Condition (Huber's condition).** The largest diagonal element of  $\mathbb{H} = \mathbb{X}(\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top$  converges to zero as  $n \rightarrow \infty$ .

**Proposition 4.8 ([Huber 1973](#)).** Under the assumptions of the current section, when  $\boldsymbol{\beta}$  denotes the true regression parameters,

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{D} N_p(\mathbf{0}, \sigma_\varepsilon^2 \mathbb{V}_X^{-1})$$

if and only if Huber's condition holds. ◇

Huber's condition is sufficient to prove consistency of  $\hat{\boldsymbol{\beta}}$  and necessary and sufficient to prove normality. The proof relies on Feller-Lindeberg central limit theorem and is omitted.

[Arnold \(1980\)](#) showed that Lemmas [4.6](#) and [4.7](#) also hold if Huber's condition is fulfilled. Thus, analysis of the linear regression model with fixed covariates proceeds exactly in the same way as with random covariates.

## 5. Predictions

### 5.1. Predictions and Their Pitfalls

Obtaining predictions for an individual observation based on the observed values of the covariates is one of the common goals of regression analysis. Actually, regression analysis may have a number of possible objectives — each of them requires somewhat different approach to development and evaluation of the model. Among the possible objectives are, e.g.,

1. Separate the signal  $X_i^\top \beta$  from the noise  $\varepsilon_i$ .
2. Predict the *expectation* of a future observation with known covariates.
3. Predict the *value* of a future observation with known covariates.
4. Determine the functional shape of  $m(x) = E[Y | X = x]$ .
5. Find out which covariates affect the expectation of the response and evaluate their influence.
6. Evaluate the influence of a single specific covariate on the expectation of the response.
7. *Et cetera*.

Objectives 1–3 are related to making predictions of the response. A point prediction can be easily obtained by taking the fitted value  $\hat{Y}$ , which is the best linear unbiased estimator of the true expectation. However, one must be extremely careful not to make predictions for covariate values that are outside the scope of the covariates that were used to build the model. The validity of the model can be verified only for covariate values that were present in the data. Predicting responses outside the scope of the covariates is called *extrapolation*. Extrapolation represents one of the most frequent abuses of regression models in practice.

**Example 5.1.** Consider the situation illustrated in Figure 5.1. There is one covariate  $X$  and the true conditional expectation is  $E[Y | X = x] = m(x) = x + 2$  for  $x \leq 3$  and

$$m(x) = \frac{\exp\{2.5x - 7.5\}}{1 + \exp\{2.5x + 7.5\}} + 4.5 \quad \text{for } x \geq 3.$$

The conditional expectation is linear up to  $x = 3$ , then the speed of increase slows down and the expectation approaches the limit  $m(x) = 5.5$  as  $x$  increases to infinity.\* The observed covariates range from 0 to 3. The regression model captures well the linear part of  $m(x)$  over the interval  $(0, 3)$  but it cannot recognize that the relationship changes for  $x > 3$ . If the linear model is extrapolated to obtain predictions for  $x > 3$ , the predictions will be seriously biased upwards.  $\triangle$

---

\* We can call this function a logistic pipe curve.




Figure 5.1.: Incorrect extrapolation of a logistic pipe curve beyond the range of data.




Figure 5.2.: Incorrect extrapolation of a sinus curve beyond the range of data.

**Example 5.2.** Figure 5.2 shows the results of extrapolation when the true expectation follows the function  $m(x) = \sin x$  and the regression model is linear, with the covariate observed within the interval  $(-0.5, 0.5)$ . Inside that interval, the fitted regression line approximates the sinus function relatively well. Outside that interval, however, the predictions obtained from the regression line are worthless.  $\triangle$

The only case when extrapolation is allowed is the situation when the true shape of the function  $m(x)$  is known (there is some physical law that determines that shape without any uncertainty) and we are certain that the estimate of  $m(x)$  applies to covariate values that are beyond the scope of the data.

**Example 5.3.** The problem of extrapolation may be difficult to spot when multiple covari-




Figure 5.3.: Extrapolation beyond the scope of data: two covariates  $X_1$  and  $X_2$ , the prediction is made within the range of both but outside the area where observations are available.

ates are present in the model. Figure 5.3 shows observed values of two covariates  $X_1$  and  $X_2$  (the response is not shown). The first covariate has values in the interval  $(-4.2, 4)$ , the other covariate lies within  $(-3, 3)$ . We intend to make a prediction at  $x_1 = -2$  and  $x_2 = 2$ . Even though both covariates are within the ranges represented in the data, this particular combination is out of the scope of the observed pairs and thus the prediction at this point suffers from the extrapolation issue.  $\triangle$

**Note.** When presenting the results of regression models it is important to include detailed description of the covariate values that were used to fit the model. Otherwise the estimated regression parameters could be misused to obtain extrapolated predictions at covariate values which are far from the observations represented in the data.

## 5.2. Confidence Intervals for Conditional Expectations

Consider a vector of covariates  $\mathbf{X} = \mathbf{x}$ , which is in the scope of data (so that we avoid the extrapolation mistake). We want to estimate  $E[Y | \mathbf{X} = \mathbf{x}]$  for this particular covariate vector. The point estimate is  $\hat{Y} = \mathbf{x}^\top \hat{\boldsymbol{\beta}}$ , and it is the BLUE by Gauss-Markov Theorem 3.6.

Let us construct a confidence interval for the unknown  $E[Y | \mathbf{X} = \mathbf{x}] = \mathbf{x}^\top \boldsymbol{\beta}$  to capture the uncertainty in the prediction appropriately. By Lemma 4.1,

$$U = \frac{\mathbf{x}^\top \hat{\boldsymbol{\beta}} - \mathbf{x}^\top \boldsymbol{\beta}}{\sqrt{\hat{\sigma}_e^2 \mathbf{x}^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{x}}} \sim t_{n-p}$$

under normality. By Lemma 4.6, the same statement holds approximately for non-normal

data. Now,

$$P[-t_{n-p}(1 - \alpha/2) < U < t_{n-p}(1 - \alpha/2)] = 1 - \alpha$$

and, after an easy manipulation,

$$P\left[\mathbf{x}^\top \hat{\boldsymbol{\beta}} - t_{n-p}(1 - \alpha/2)\hat{\sigma}_e \sqrt{\mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x}} < \mathbf{x}^\top \boldsymbol{\beta} < \mathbf{x}^\top \hat{\boldsymbol{\beta}} + t_{n-p}(1 - \alpha/2)\hat{\sigma}_e \sqrt{\mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x}}\right] = 1 - \alpha$$

For non-normal data, this claim holds asymptotically, as  $n \rightarrow \infty$ .

Thus, the confidence interval for  $E[Y | X = \mathbf{x}] = \mathbf{x}^\top \boldsymbol{\beta}$  with coverage probability  $1 - \alpha$  (exact or asymptotic) is

$$\mathbf{x}^\top \hat{\boldsymbol{\beta}} \mp t_{n-p}(1 - \alpha/2)\hat{\sigma}_e \sqrt{\mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x}}.$$

Notice that for an observation that was used to fit the model, when  $\mathbf{x} = \mathbf{X}_i$  for some  $i$ ,  $\mathbf{X}_i^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{X}_i = h_{ii}$ , the  $i$ -th diagonal element of the projection matrix  $\mathbb{H}$ .

### 5.3. Prediction Intervals for Future Responses

Consider a future observation with a known vector of covariates  $\mathbf{X} = \mathbf{x}$ , which is in the scope of data. We want to construct an interval that includes the future observed response of such an observation with a desired probability, i.e., find  $C_L$  and  $C_U$  such that

$$P[C_L < Y < C_U] = 1 - \alpha.$$

Because we want the interval to cover a realization of a random variable rather than an unknown fixed quantity, we call the interval a *prediction interval* rather than a confidence interval.

The construction of the interval must be based on the distribution of a single observation  $Y$  and that cannot be approximated by asymptotic results. There fore we *must assume* in this section that the observation follows the normal distribution, in particular

$$Y \sim N(\mathbf{x}^\top \boldsymbol{\beta}, \sigma_e^2).$$

Write  $Y$  as  $Y = \mathbf{x}^\top \boldsymbol{\beta} + \varepsilon$  where  $\varepsilon \sim N(0, \sigma_e^2)$ . We assume that the new observation  $Y$  is independent of the observations  $Y_1, \dots, Y_n$  contained in the dataset used to estimate the parameters. Calculate the fitted value  $\hat{Y} = \mathbf{x}^\top \hat{\boldsymbol{\beta}}$ . By Lemma 3.7, part (i),

$$\hat{Y} \sim N(\mathbf{x}^\top \boldsymbol{\beta}, \sigma_e^2 \mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x})$$

and, by independence and normality of  $Y$  and  $\hat{Y}$ ,

$$\hat{Y} - Y \sim N(0, \sigma_e^2 + \sigma_e^2 \mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x}).$$

Hence,

$$\frac{\mathbf{x}^\top \hat{\boldsymbol{\beta}} - Y}{\hat{\sigma}_e \sqrt{1 + \mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x}}} \sim t_{n-p}$$

and the resulting prediction interval for  $Y$  with coverage probability  $1 - \alpha$  is

$$\mathbf{x}^\top \hat{\boldsymbol{\beta}} \mp t_{n-p}(1 - \alpha/2) \hat{\sigma}_e \sqrt{1 + \mathbf{x}^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{x}}. \quad (5.1)$$

It differs from the interval derived in the previous section by adding 1 under the square root and by the necessity to assume normality.



## 6. Diagnostic Methods Based on Residuals

In this section we introduce and illustrate residual-based methods for checking model assumptions and assessing the validity of the model.

Lemma 3.2 assures that the  $i$ -th residual  $u_i = Y_i - \hat{Y}_i = Y_i - \mathbf{X}_i^T \hat{\boldsymbol{\beta}}$  has zero mean and variance  $\text{var } u_i = \sigma_e^2(1 - h_{ii})$ , where  $h_{ii}$  is the  $i$ -th diagonal element of the projection matrix  $\mathbb{H} = \mathbb{X}(\mathbb{X}^T \mathbb{X})^{-1} \mathbb{X}^T$ . Because raw residuals do not have the same variance, we will use so called *standardized residuals*.

**Definition 6.1.**

$$u_i^* = \frac{u_i}{\hat{\sigma}_e \sqrt{1 - h_{ii}}}$$

are called *standardized residuals*. ▽

If the model is valid and all assumptions are fulfilled, standardized residuals have approximately zero mean and unit variance. This can be verified by plotting standardized residuals in various ways. Common examples are:

- Scatterplots of residuals against various continuous variables, smoothed by some non-parametric smoother to facilitate recognition of patterns.
- Boxplots of residuals for specific subgroups of observations.
- Histograms and Q-Q plots of residuals.

### Scatterplot of residuals against order of observation

This type of plot puts standardized residuals on the vertical axis and observation order  $i$  on the horizontal axis. It is useful when the ordering of observations has some real meaning, for example, if it captures the time sequence in which the observations are recorded. If the assumptions are satisfied, the plot shows just a random cloud of points centered around the line  $y = 0$ . If the plot suggests some effect of order on the residuals there is a suspicion that the observation order has some unaccounted effect on the response and the data are not truly independent.

The top panel of Figure 6.1 shows the situation when the assumptions are fulfilled. The data are generated from the model  $Y_i = 1 + X_i + \varepsilon_i$  with  $\varepsilon_i \sim N(0, 0.25)$ ,  $i = 1, \dots, 100$ . The fitted model was  $E Y_i = \beta_1 + \beta_2 X_i$ . The model is correct and there is no recognizable pattern

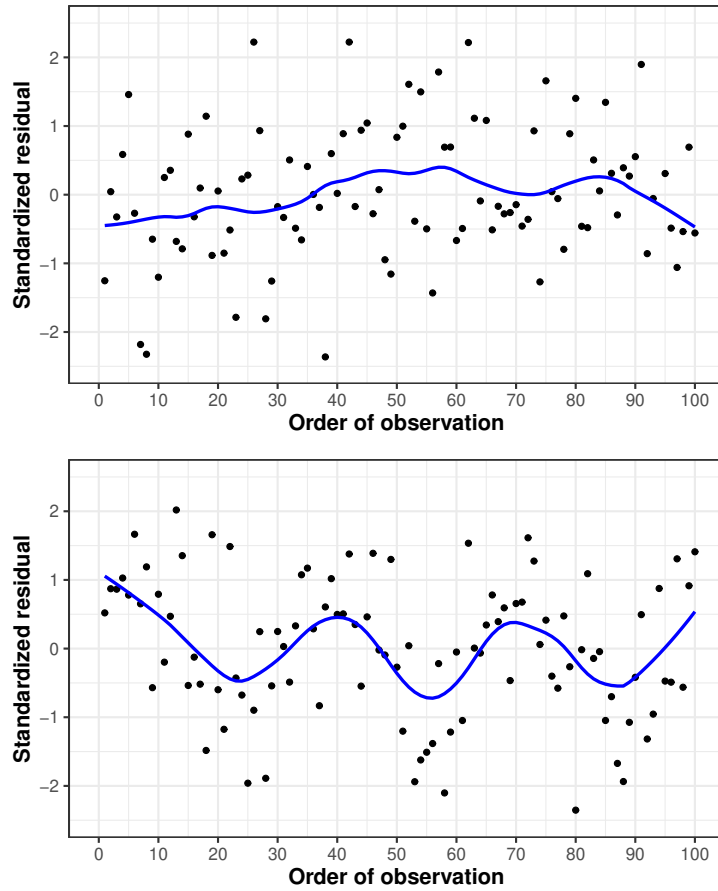


Figure 6.1.: Standardized residuals against observation order. Top panel: all assumptions are satisfied. Bottom panel: an unaptured periodic effect. Both plots were smoothed by lowess smoother with window over 1/4 of the data range (blue).

in the top panel. The bottom panel shows residuals from the same linear model for data generated from the model  $Y_i = 1 + X_i + \sin(i/5) + \varepsilon_i$ . A periodic effect of the observation order was added to the responses but the analysis ignored that and proceeded in the same way as above. The omitted periodic effect is clearly demonstrated by the smoothed scatterplot in the bottom panel.

### Scatterplot of residuals against fitted values

This version plots fitted values  $\hat{Y}_i$  on the horizontal axis. It provides kind of general assessment of the validity of the model. If the assumptions are satisfied, the plot shows just a random cloud of points centered around the line  $y = 0$ . If the plot suggests some pattern across the fitted values, there is a suspicion that the effect of some covariate on the response is modeled inappropriately.




Figure 6.2.: Standardized residuals against fitted values. Top panel: all assumptions are satisfied. Bottom panel: omitted quadratic effect. Both plots were smoothed by lowess smoother with window over 1/2 of the data range (blue).

The top panel of Figure 6.2 shows the situation when the assumptions are fulfilled. The data are again generated from the model  $Y_i = 1 + X_i + \varepsilon_i$  with  $\varepsilon_i \sim N(0, 0.25)$ ,  $i = 1, \dots, 100$ . The fitted model was  $EY_i = \beta_1 + \beta_2 X_i$ . The model is correct and there is no recognizable pattern in the top panel. The bottom panel shows residuals from the same linear model for data generated from the model  $Y_i = 1 + X_i + (X_i - 1.5)^2 + \varepsilon_i$ . A quadratic effect of the covariate was added to the responses but the analysis ignored that and proceeded in the same way as above. The omitted quadratic effect is clearly demonstrated by the smoothed scatterplot in the bottom panel.

When the model includes several covariates, the plot of residuals against fitted values may not reflect clearly that one of the covariates was modeled in an inappropriate way or to determine which covariate it was. For checking more complex models it is much better to plot the residuals against individual covariates rather than against fitted values.

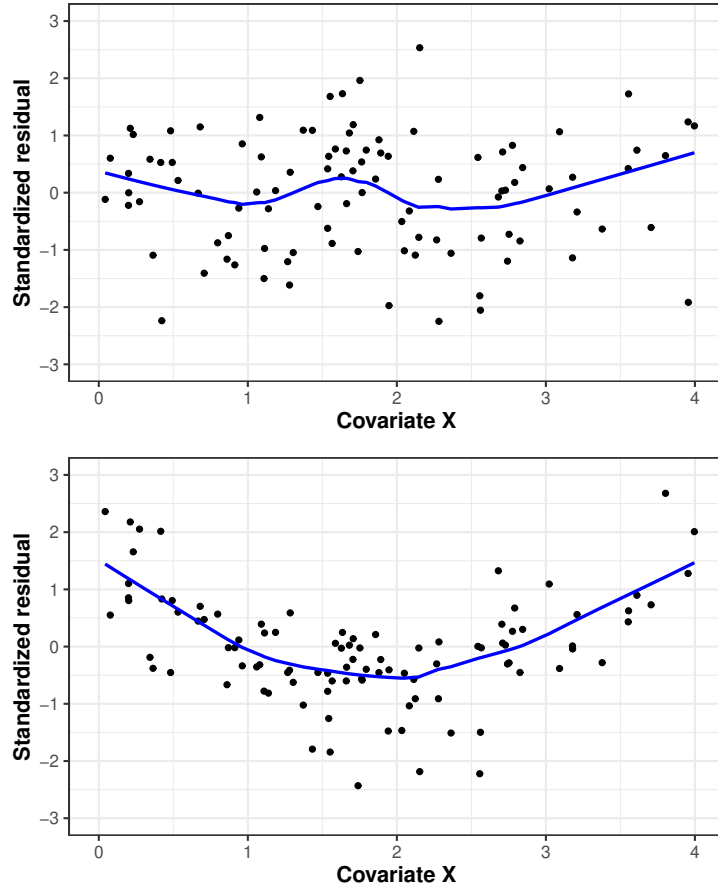


Figure 6.3.: Standardized residuals against a covariate. Top panel: all assumptions are satisfied. Bottom panel: omitted quadratic effect. Smoothed by lowess smoother with window over 1/2 of the data range (blue).

### Scatterplot of residuals against continuous covariates

This plot uses values of a particular covariate on the horizontal axis. It checks whether the covariate was included in the model in an appropriate form. If it was and other assumptions are also satisfied, the plot shows just a random cloud of points centered around the line  $y = 0$ . If the plot suggests some pattern depending on the covariate, there is a suspicion that the effect of this covariate on the response is modeled inappropriately.

The top panel of Figure 6.3 shows the situation when the assumptions are fulfilled. The data are generated from the model  $Y_i = 1 + Z_i + 0.5X_i + \varepsilon_i$  with  $\varepsilon_i \sim N(0, 1)$ ,  $i = 1, \dots, 100$ . The covariate  $Z_i$  is binary, the covariate  $X_i$  is continuous and correlated with  $Z_i$ . The fitted model was correct:  $EY_i = \beta_1 + \beta_2 Z_i + \beta_3 X_i$ . There is no recognizable pattern in the top panel, which plots the standardized residuals against the values of  $X_i$ . The bottom panel shows residuals from the same model but for data generated from the model  $Y_i = 1 + Z_i +$




Figure 6.4.: Standardized residuals against a covariate. Top panel: all assumptions are satisfied. Bottom panel: mild increase of residual variance with covariate  $X$ .

$0.5X_i + (X_i - 1.5)^2 + \varepsilon_i$ . A quadratic effect of the covariate  $X_i$  was added to the response but the analysis ignored that and proceeded with an incorrect model. The omitted quadratic effect is clearly demonstrated on the smoothed scatterplot in the bottom panel.

The plot of the residuals against a covariate is able to capture an inappropriately modeled effect of that covariate even if the model contains other covariates. Even better way to determine a suitable functional format for a continuous covariate provides the plot of partial residuals, which will be discussed later in this section.

This plot can be also used to check whether the variance is constant or whether it changes with the values of the covariate. If the variance is constant and other assumptions are also satisfied, the plot shows points that are about equally spread across the range of values of the covariate. If the residuals seem more variable in some regions than in others, there is a suspicion that the assumption of equal variance is not true.

The top panel of Figure 6.4 shows the situation when the assumptions are fulfilled. The




Figure 6.5.: Square root of absolute standardized residuals against a covariate. Top panel: all assumptions are satisfied. Bottom panel: mild increase of residual variance with covariate  $X$ .

data are generated from the model  $Y_i = 1 + Z_i + 0.5X_i + (X_i - 1.5)^2 + \varepsilon_i$  with  $\varepsilon_i \sim N(0, 1)$ ,  $i = 1, \dots, 100$  (as before). This time, the correct model with quadratic term was used to perform the analysis and no assumptions are violated. We can see in the top panel, that the points are about equally spread in the vertical direction at all values of  $X_i$ . The bottom panel shows residuals from the same model but with unequal variance depending on  $X_i$  —  $\text{var } \varepsilon_i = (X_i + 1)^2/9$ . The variance increases with  $X_i$  and the spread of the residuals seems to slightly increase with the value of  $X_i$ .

From this type of plot, it is relatively difficult to see whether the variance changes or not, and it is not possible to seek assistance from a smoother.




Figure 6.6.: Boxplots of standardized residuals by a factor covariate. Top panel: all assumptions are satisfied. Bottom panel: unequal variances of error terms and an omitted effect of another covariate.

### Scatterplots of square root of absolute standardized residuals against continuous covariates

A better way to check the assumption of constant variability is to plot square root of absolute standardized residuals against a covariate. Such a plot can be smoothed to facilitate the interpretation. Figure 6.5 shows this type of plot for the same situations as in Figure 6.4. In the top panel, where the assumptions are satisfied, we see some increase in a certain range of covariate values but there is no overall trend. In the bottom panel, where there is an actual increase in variability, we clearly see increasing trend in the absolute standardized residuals.

### Boxplots of standardized residuals by a categorical covariate

For categorical (factor) covariates that classify the observations into disjoint subgroups, boxplots of standardized residuals can be used to assess model assumptions. If the assumptions are satisfied, the medians of residuals in each group (the bars inside the boxes) will be close to zero and the height of the boxes (interquartile range) will be similar in all the subgroups (Figure 6.6, top panel). In the bottom panel of the same figure, the heights of the boxes visibly differ because the residual variance in Group 1 was larger than that in Group 0 and the box for Group 1 is not centered around zero because an important covariate correlated with group membership was omitted from the model.

### Boxplots of standardized residuals by a factorized continuous covariate

Continuous covariates can be factorized in order to assess model assumptions using boxplots of standardized residuals. In the top panel of Figure 6.7, all the model assumptions were satisfied. In the bottom panel, the heights of the boxes somewhat increase from the left to the right because the residual variance increased with the covariate  $X$  and the boxes are located around a quadratic curve because the model failed to include an important quadratic effect of this covariate.

### Histograms of standardized residuals

Histograms are helpful to visualize the distribution of residuals and to detect the presence of observations with unusually large absolute residuals. Figure 6.8 shows results of such a visualization when the residuals are normally distributed (top panel), when the residuals have a relatively heavy tailed distribution (middle panel) and when the distribution of residuals is skewed to the right (bottom panel).

### Quantile plots of standardized residuals

Quantile plots (Q-Q plots) are generally more helpful tools to assess normality than histograms. These plots contain ordered standardized residuals on the vertical axis and plot them against theoretical expected values of corresponding order statistics calculated under normality (on the horizontal axis). More precisely, the Q-Q plot is a scatterplot of pairs  $(u_{(i)}^*, z_i)$ , where  $u_{(i)}^*$  is the  $i$ -th smallest standardized residual and

$$z_i = \Phi^{-1}\left(\frac{i}{n+1}\right)$$

approximates  $E Z_{(i)}$ , the expectation of the  $i$ -th order statistic in a random sample from  $N(0, 1)$  of the size  $n$ . If the distribution of error terms is normal, the points displayed on the Q-Q plot approximately follow a line (the top panel of Figure 6.9). If the distribution of errors has heavier tails than the normal distribution, the Q-Q plot displays an S-shaped curve as shown in the middle panel of Figure 6.9. When the errors come from a skewed distribution, the Q-Q plot shows a bow-shaped curve (bottom panel of Figure 6.9).






Figure 6.7.: Boxplots of standardized residuals by a factorized continuous covariate. Top panel: all assumptions are satisfied. Bottom panel: omitted quadratic effect and mildly increasing variance with  $X$ .

### Assessing the correct functional form of a covariate

In linear regression, it is important to consider whether the effect of a particular continuous covariate can be modeled in a linear way or whether a more complex functional form (e.g., quadratic) is required. Scatterplots of standardized residuals against the values of that covariate can help to assess whether the covariate was included in an appropriate way (see Figure 6.3). A more convenient way to find the correct transformation of the covariate is to plot so called *partial residuals*.

Partial residuals are differences between the observed response  $Y_i$  and the fitted value  $\hat{Y}_i$  from which the estimated effect of the covariate was entirely removed. E.g., if a continuous covariate  $X_{i4}$  is included in the current model in the linear form, we obtain the partial residuals for that covariate by

$$Y_i - (\hat{Y}_i - \hat{\beta}_4 X_{i4}) = u_i + \hat{\beta}_4 X_{i4}.$$

A smoothed scatterplot of partial residuals plotted against the values of the covariate of interest directly suggests an appropriate functional form for that covariate. If the linear form is sufficient the partial residuals seem to follow a line (see the top panel of Figure 6.10). In the bottom panel of the same figure, the true effect of the covariate is quadratic and partial residuals testify to that by following a parabolic function.

### Testing model assumptions

A number of formal statistical tests have been developed to verify the validity of various model assumptions: equal variances, independence of error terms, normality of error terms. These tests always take as the null hypothesis the situation when the assumptions are fulfilled and reject the null hypothesis when the data contain sufficient evidence that a particular assumption is violated. The larger the data set is the more likely it is that the test will discover even minute and entirely unimportant violations. With small data sets, the test will often fail to find even very large and harmful violations. The impact of violated model assumptions on the results of the model depends on the extent of the violation – the test results depend on both the extent of violation and the sample size.

Therefore we recommend to assess model assumptions by various descriptive methods (residual plots, descriptive statistics) rather than by formal tests. Tests of normality in linear regression are particularly discouraged: with small sample sizes, normality is important for the validity of statistical inference, however the test of normality is not able to confirm that the data are truly normal. Even substantial departures from normality will go unnoticed in small data sets. On the other hand, with large sample sizes, the test will reject normality more frequently, however this assumption is entirely unimportant because asymptotic results provide reliable approximations even if the data are strongly non-normal. Formal tests of normality in linear regression are thus entirely worthless.




Figure 6.8.: Histograms of standardized residuals. Top panel: normal distribution of errors. Middle panel: heavy-tailed distribution of errors ( $t_4$ ). Bottom panel: right-skewed distribution of errors (negative Gumbel).

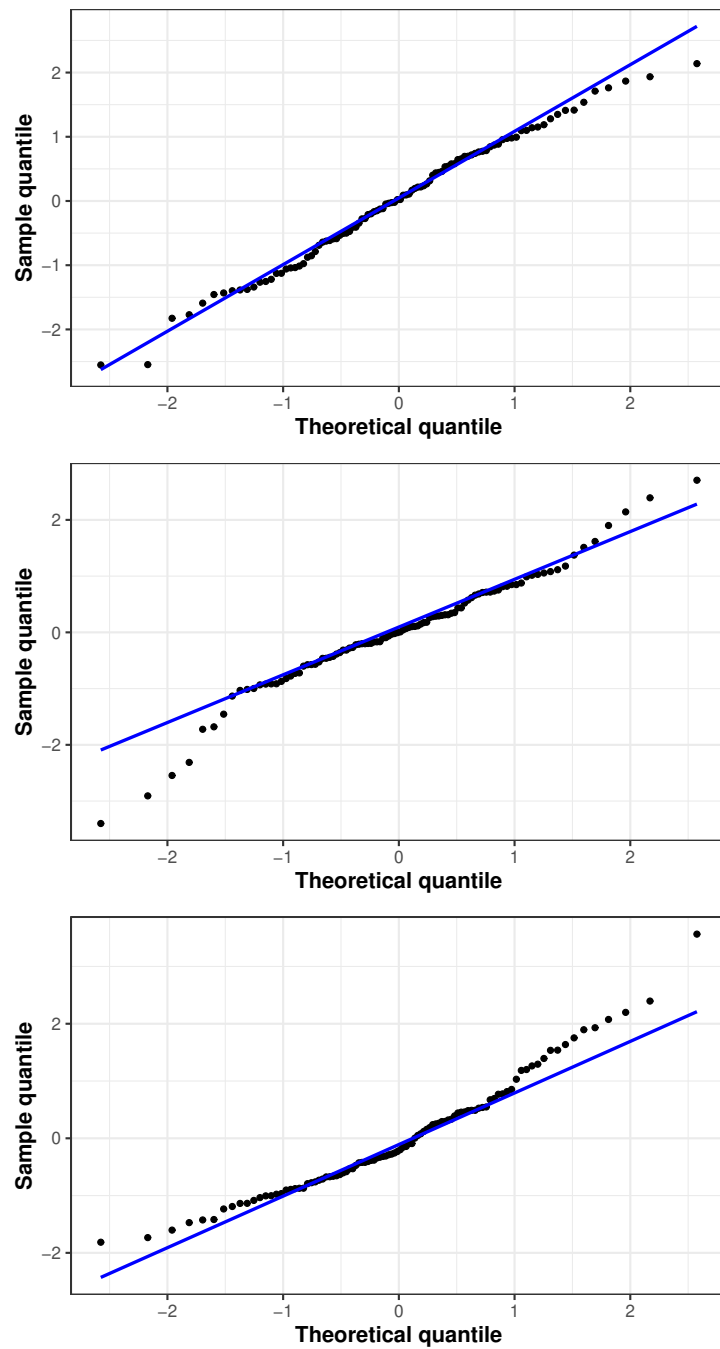


Figure 6.9.: Q-Q plots of standardized residuals. Top panel: normal distribution of errors. Middle panel: heavy-tailed distribution of errors ( $t_4$ ). Bottom panel: right-skewed distribution of errors (negative Gumbel).

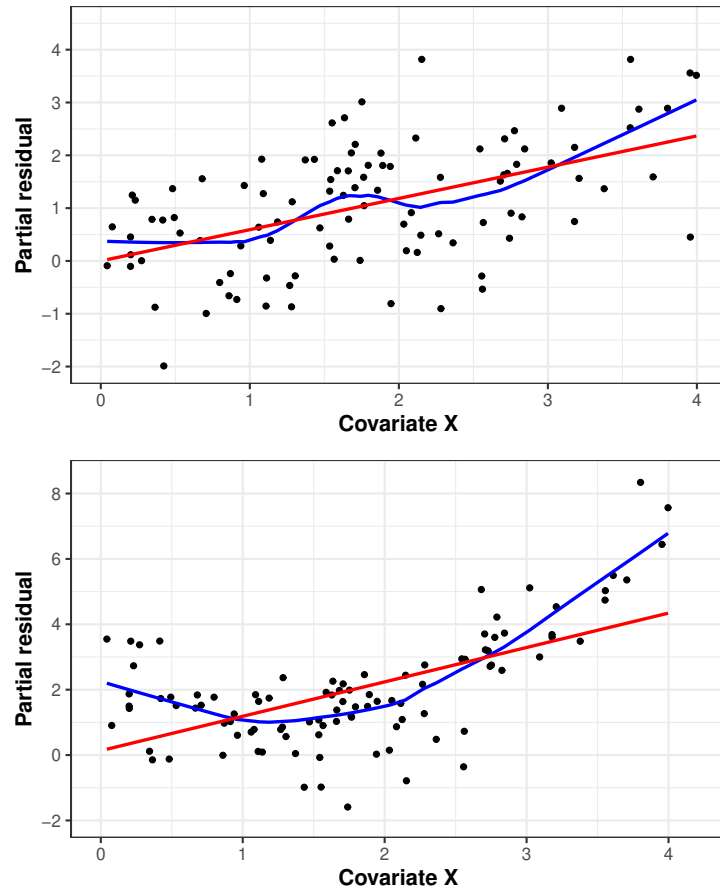


Figure 6.10.: Scatterplots of partial residuals against a covariate. Top panel: all assumptions are satisfied. Bottom panel: omitted quadratic effect. Smoothed by lowess smoother with window over 1/2 of the data range (blue).

## 7. Transformations of the Response

### 7.1. Introduction

Responses that do not follow the normal distribution typically violate the assumption of equal variances as well. For example, consider a response  $Y$  that can be close to zero but cannot be negative. If we use a linear model without intercept to express the dependence of  $EY$  on a covariate  $X > 0$

$$Y = \beta X + \varepsilon$$

it is clear that for very small  $X$  the variability in  $\varepsilon$  cannot be large because  $Y$  is positive and  $\varepsilon$  can only take small negative values. On the other hand, when  $X$  is large, the variability in  $\varepsilon$  can be much larger. The variance of  $\varepsilon$  thus tends to increase with increasing  $X$ . This is what frequently happens with this type of data.

The fact that the response is not normal is not of much concern if the sample size is large enough. However, non-constant variance is always of concern, no matter how large the data set is.

In these situations, a transformation of the response may help to satisfy the model assumptions. Suppose there exists a strictly monotone smooth function  $g$  such that

$$g(Y_i) = \mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i,$$

where  $E\varepsilon_i = 0$  and  $\text{var } \varepsilon_i = \sigma_\varepsilon^2$ . Thus, we impose the linear model on  $g(Y_i)$  rather than on  $Y_i$  itself.

We have  $E g(Y_i) = \mathbf{X}_i^\top \boldsymbol{\beta}$ . The parameters of this model express the effect of the covariate vector on the expectation of  $g(Y_i)$ . In general, the model does not allow us to estimate the effects of the covariates on  $EY_i$ . The implied model for  $Y_i$  is actually  $Y_i = g^{-1}(\mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i)$  and hence  $EY_i = E g^{-1}(\mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i)$ . We only know that  $E\varepsilon_i = 0$  and evaluation of  $EY_i$  from this is not possible – unless  $g$  is a logarithm (see the next section).

A backward transformation of  $E g(Y_i)$  does not help. If  $g$  is convex and strictly increasing, we have by Jensen inequality  $E g(Y_i) > g(EY_i)$  and by application of the inverse on both sides,  $EY_i < g^{-1}(E g(Y_i))$ .

If the goal of the transformed model is to make predictions about the response, the situation is less dire. Suppose that the density of  $\varepsilon_i$  in the transformed model is symmetric around 0. Then the mean and the median of  $g(Y_i)$  are both equal to  $\mathbf{X}_i^\top \boldsymbol{\beta}$ . Unlike the mean,

the median can be back-transformed and we get

$$\text{med}(Y_i) = g^{-1}(\text{med}(g(Y_i))) = g^{-1}(X_i^T \beta). \quad (7.1)$$

Thus, if we manage to transform the response so that the errors have a symmetric distribution we can predict the median of  $Y_i$  by  $g^{-1}(X_i^T \hat{\beta})$  for any strictly monotone transformation  $g$ .

## 7.2. Interpretation of Linear Model With Log-transformed Response

Let  $Y_i > 0$  almost surely and consider the model

$$\log Y_i = X_i^T \beta + \varepsilon_i$$

where  $E \varepsilon_i = 0$  and  $\text{var } \varepsilon_i = \sigma_\varepsilon^2$ . Thus, the linear regression model holds after the logarithmic transformation of the response. Then

$$Y_i = e^{X_i^T \beta} e^{\varepsilon_i} = e^{\beta_1} \cdot e^{\beta_2 X_{i2}} \dots e^{\beta_p X_{ip}} \cdot e^{\varepsilon_i}.$$

Assume that  $e^{\varepsilon_i}$  has finite variance. We obtained a model where covariate effects multiply each other and the error also acts multiplicatively. The first two moments of the untransformed response can be calculated as follows:

$$\begin{aligned} E Y_i &= e^{X_i^T \beta} E e^{\varepsilon_i}, \\ \text{var } Y_i &= (e^{X_i^T \beta})^2 \text{var } e^{\varepsilon_i}. \end{aligned}$$

If  $\varepsilon_i$  are independent and identically distributed, then  $\text{var } e^{\varepsilon_i}$  is a constant and the variance of the response is an increasing function of  $X_i^T \beta$ .

Denote  $\gamma = E e^{\varepsilon_i}$ . We know that  $E \varepsilon_i = 0$  and from Jensen's inequality ( $e^x$  is convex),  $\gamma \equiv E e^{\varepsilon_i} > e^{E \varepsilon_i} = 1$ .

Let us determine the interpretation of  $\beta$  in the untransformed model. We have

$$\begin{aligned} E[Y | X = x] &= e^{x^T \beta} \gamma \quad \text{and} \\ E[Y | X = x + e_j] &= e^{x^T \beta + \beta_j} \gamma \quad j = 2, \dots, p. \end{aligned}$$

Taking the ratio of the two equations, the unknown  $\gamma$  cancels and we get

$$\frac{E[Y | X = x + e_j]}{E[Y | X = x]} = \frac{e^{x^T \beta + \beta_j} \gamma}{e^{x^T \beta} \gamma} = e^{\beta_j}.$$

In the linear model with a log-transformed response,  $e^{\beta_j}$  expresses the proportional increase in  $E Y$  after increasing the  $j$ -th covariate by one unit while keeping all other covariates fixed.

The log-transformed linear model estimates the multiplicative effects of the covariates on the expectation of the response and these are expressed by the exponentiated regression coefficients.

What about the intercept term? Set  $\mathbf{X} = (1, 0, \dots, 0)^\top$ . Then  $E[Y | \mathbf{X} = \mathbf{e}_1] = \gamma e^{\beta_1} = e^{\beta_1 + \log \gamma}$ . But  $\gamma$  is not specified, we only know that  $\gamma > 1$ . Thus, the intercept cannot tell anything about the expectation of the response when all covariates are zero unless we assume a specific distribution of the response. For example, when  $\varepsilon \sim N(0, \sigma_\varepsilon^2)$  (that is,  $Y$  is log-normal) then  $\gamma = E e^\varepsilon = e^{\sigma_\varepsilon^2/2}$  and  $E[Y | \mathbf{X} = \mathbf{e}_1] = e^{\beta_1 + \sigma_\varepsilon^2/2}$ .

The logarithmic transformation of the response is a valuable and frequently employed tool when the responses are strictly positive, attain values that are close to zero and have long right tails with increasing variability as the mean increases. The logarithmic transformation of such a response improves equality of variances and brings the distribution of the transformed response closer to normality.

**Note.**

- (i) When the response attains non-negative values with  $P[Y = 0] > 0$ , the log transform can be used after shifting the responses by a constant  $c > 0$ . Thus, the transformed response is  $\log(Y + c)$ . Question for thought: What is the interpretation of  $\beta$  after such transformation?
- (ii) If the response is log-transformed, the covariates usually require a transformation as well.

### 7.3. The Box-Cox Transformation

Let  $Y_i > 0$  almost surely. Consider the following family of transformations of the response depending on a parameter  $\lambda \in \mathbb{R}$ :

$$\begin{aligned} g_\lambda(y) &= \frac{y^\lambda - 1}{\lambda} && \text{for } \lambda \neq 0, \\ g_\lambda(y) &= \log y && \text{for } \lambda = 0. \end{aligned}$$

The transformations are defined in this way so that  $\lim_{\lambda \rightarrow 0} g_\lambda(y) = g_0(y)$ , that is, they are continuous in  $\lambda$ . Otherwise, subtracting a constant and dividing the response by a constant is irrelevant. Among the transformations included in the Box-Cox family are all power function transformations such as  $Y^p$ ,  $\sqrt[p]{Y}$ ,  $\frac{1}{Y^p}$ ,  $\frac{1}{\sqrt[p]{Y}}$ ,  $p = 2, 3, \dots$

Assuming that there exists  $\lambda \in \mathbb{R}$  such that  $g_\lambda(Y)$  is normally distributed, such  $\lambda$  can be estimated via maximum likelihood methods based on the normal density. This is done numerically, the estimated  $\lambda$  does not have a closed-form expression.

However, there need not exist any  $\lambda$  that makes the response normal after the Box-Cox transformation and, even if it existed, the model does not have interpretable parameters.



However, this approach can be valuable when the goal of the regression analysis is to make predictions, we are happy with predicting the median instead of the expectation (see (7.1)) and the Box-Cox transformation is successful in making the responses normal.

When  $g_\lambda(Y)$  is normal, we can use the methods of Section 5.3 to make predictions of individual future responses. By (5.1) we calculate a prediction interval  $(C_L, C_U)$  satisfying

$$P[C_L < g_\lambda(Y) < C_U] = 1 - \alpha$$

and because  $g_\lambda(Y)$  is a monotone transformation of  $Y$ , we get

$$P[g_\lambda^{-1}(C_L) < Y < g_\lambda^{-1}(C_U)] = 1 - \alpha$$

(this is for an increasing  $g_\lambda(Y)$ , otherwise we switch the bounds).

Calculating a prediction interval for  $EY$  of a future observation via Box-Cox transformation is not possible, however.

## 8. Parametrization of a Single Covariate

### 8.1. Parametrization of Categorical Covariates (Factors)

*Categorical covariates* (factors, qualitative variables)\* are variables that indicate group membership. Their values, typically integers, do not have any numerical meaning.

Let the studied population be divided into  $m$  disjoint groups. Consider a covariate  $Z_i \in \{1, \dots, m\}$  where  $Z_i = j$  means that the  $i$ -th observation belongs to group  $j$ . The numerical coding of the groups is frequently arbitrary, in fact any permutation of integers  $\{1, \dots, m\}$  could be used to distinguish the groups from each other. Therefore the linear model  $E[Y_i | Z_i] = \beta_1 + \beta_2 Z_i$  does not make sense for a categorical covariate. Instead, we need to allow each group to have its own expectation which is not related to the expectations of other groups. So, let  $\mu_j = E[Y_i | Z_i = j]$  denote the expectation of the response in the  $j$ -th group.

This situation has been already discussed in the context of one-way ANOVA models, see Section 4.1.7. There, the model was specified as

$$E[Y_i | Z_i] = \beta_1 \mathbb{1}(Z_i = 1) + \dots + \beta_m \mathbb{1}(Z_i = m) \quad (8.1)$$

The regression matrix  $\mathbb{X}$  includes the indicators  $\mathbb{1}(Z_i = j)$  in the  $j$ -th column. The columns of the regression matrix sum to one, so  $\mathbf{1}_n \in \mathcal{M}(\mathbb{X})$  and the model includes an intercept. The parameters correspond to the group means:  $\beta_j = \mu_j$  for  $j = 1, \dots, m$ . This is called the *cell-means parametrization* of a factor covariate. This parametrization is suitable for calculation the LSE of group expectations, fitted values and sums of squares. However, we cannot use it for parametrization of multiple factors included in the same model because the columns sum to one for each of the factors and the model does not have full rank.

We need a model that includes the intercept explicitly ( $X_{i1} = 1$  for all  $i$ ) and models the effects of the factor by  $m - 1$  additional parameters. We start from

$$E[Y_i | Z_i] = \beta_1 + \gamma_1 \mathbb{1}(Z_i = 1) + \dots + \gamma_m \mathbb{1}(Z_i = m).$$

Here, we only recoded the parameters of the previous model as  $\gamma_1, \dots, \gamma_m$  and added an intercept denoted by  $\beta_1$ . Of course, this model does not have full rank. To deal with this, we introduce a single linear constraint on the parameters  $\gamma_1, \dots, \gamma_m$ . We can choose different constraints and each leads to a different parametrization of the model and different interpretation of the regression coefficients. In the rest of this section, we give examples of the most common parametrizations of factor variables.

---

\* Český faktor; kategoriální veličina

All the parametrizations generate models that are equivalent to the cell-means model (8.1) in the sense of Section 2.5; all have the same fitted values, residuals, residual sums of squares and regression sums of squares. They differ in the meaning and interpretation of regression parameters and their LSE's.

### 8.1.1. Reference group parametrization

Take the restriction  $\gamma_1 = 0$ . This leads to the model

$$\begin{aligned} E[Y_i | Z_i] &= \beta_1 + \gamma_2 \mathbb{1}(Z_2 = 1) + \cdots + \gamma_m \mathbb{1}(Z_i = m) \\ &= \beta_1 + \beta_2 \mathbb{1}(Z_2 = 1) + \cdots + \beta_m \mathbb{1}(Z_i = m) = \mathbf{X}_i^\top \boldsymbol{\beta}. \end{aligned}$$

The following table displays the group expectations and the rows of the regression matrix used for each group.

Group	$E Y_i$	$\mathbf{X}_i^\top$					
1	$\mu_1 = \beta_1$	1	0	0	$\cdots$	0	
2	$\mu_2 = \beta_1 + \beta_2$	1	1	0	$\cdots$	0	
3	$\mu_3 = \beta_1 + \beta_3$	1	0	1	$\cdots$	0	
$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\ddots$		
$m$	$\mu_m = \beta_1 + \beta_m$	1	0	0	$\cdots$	1	

The second up to the last column of the matrix shown on the right of the above table defines an  $m \times (m - 1)$  matrix of so called *dummy variables* used to implement this parametrization in a regression matrix (the first column represents the intercept).

Let us determine the interpretation of the components of  $\boldsymbol{\beta}$  using the above table. We have

$$\begin{aligned} \beta_1 &= \mu_1 = E[Y_i | Z_i = 1] \\ \beta_2 &= \mu_2 - \mu_1 = E[Y_i | Z_i = 2] - E[Y_i | Z_i = 1] \\ &\vdots \\ \beta_m &= \mu_m - \mu_1 = E[Y_i | Z_i = m] - E[Y_i | Z_i = 1] \end{aligned}$$

Thus, the first group serves as a reference group and the expectations in the other groups are all compared to the first group.

**Note.**

- (i) This parametrization is used by R as the default coding of factors (there, it is called “treatment contrasts”).
- (ii) Any group can be selected as the reference group. It is undesirable, however, to choose a very small group as the reference group.

### 8.1.2. Zero-sum parametrization

This parametrization arises by the restriction  $\sum_{j=1}^m \gamma_j = 0$ .

With this restriction, we can express  $\gamma_m$  as  $\gamma_m = -\gamma_1 - \dots - \gamma_{m-1}$  and plug it into the model.

$$\begin{aligned} E[Y_i | Z_i] &= \beta_1 + \gamma_1 \mathbb{1}(Z_i = 1) + \dots + \gamma_{m-1} \mathbb{1}(Z_i = m-1) - (\gamma_1 + \dots + \gamma_{m-1}) \mathbb{1}(Z_i = m) \\ &= \beta_1 + \underbrace{\gamma_1}_{\beta_2} \underbrace{[\mathbb{1}(Z_i = 1) - \mathbb{1}(Z_i = m)]}_{X_{i2}} + \dots + \underbrace{\gamma_{m-1}}_{\beta_m} \underbrace{[\mathbb{1}(Z_i = m-1) - \mathbb{1}(Z_i = m)]}_{X_{im}} = \mathbf{X}_i^T \boldsymbol{\beta}. \end{aligned}$$

Now we can determine the group expectations and the rows of the regression matrix used for each group and display them in a table.

Group	$E Y_i$	$\mathbf{X}_i^T$				
1	$\mu_1 = \beta_1 + \beta_2$	1	1	0	$\dots$	0
2	$\mu_2 = \beta_1 + \beta_3$	1	0	1	$\dots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\ddots$	
$m-1$	$\mu_{m-1} = \beta_1 + \beta_m$	1	0	0	$\dots$	1
$m$	$\mu_m = \beta_1 - \sum_{j=2}^m \beta_j$	1	-1	-1	$\dots$	-1

Again, the second up to the last column of the matrix shown on the right of the above table defines an  $m \times (m-1)$  matrix of so called *dummy variables* used to implement the zero-mean parametrization in a regression matrix (the first column represents the intercept).

From the above table, we have  $\sum_{j=1}^m \mu_j = m\beta_1$ , hence  $\beta_1 = \frac{1}{m} \sum_{j=1}^m \mu_j \equiv \bar{\mu}$ . Thus, the intercept corresponds to the arithmetic average of the expectations of the  $m$  groups. The interpretation of the other components of  $\boldsymbol{\beta}$  is as follows.

$$\begin{aligned} \beta_2 &= \gamma_1 = \mu_1 - \beta_1 = \mu_1 - \bar{\mu} \\ \beta_3 &= \gamma_2 = \mu_2 - \beta_1 = \mu_2 - \bar{\mu} \\ &\vdots \\ \beta_m &= \gamma_{m-1} = \mu_{m-1} - \beta_1 = \mu_{m-1} - \bar{\mu} \end{aligned}$$

Finally, we can calculate

$$\gamma_m = -\sum_{j=1}^{m-1} \gamma_j = (m-1)\bar{\mu} - \sum_{j=1}^{m-1} \mu_j = \sum_{j=1}^m \mu_j - \bar{\mu} - \sum_{j=1}^{m-1} \mu_j = \mu_m - \bar{\mu}.$$

Thus, the parameters  $\gamma_j$  compare the expectation in the  $j$ -th group to the arithmetic average of the expectations of all groups.

**Note.** Zero-sum parametrization is specified in R by assigning this type of contrasts (dummy variables) to the factor  $\mathbf{z}$  through the function `C(z, contr.sum)` and using this in the regression model instead of the original factor  $\mathbf{z}$ .

### 8.1.3. Weighted zero-sum parametrization

This parametrization arises by the restriction  $\sum_{j=1}^m n_j \gamma_j = 0$  where  $n_j = \sum_{i=1}^n \mathbb{1}(Z_i = j)$  is the size of the  $j$ -th group. This restriction leads to  $\beta_1 = \frac{1}{n} \sum_{j=1}^m n_j \mu_j$ , which can be estimated by the overall mean of all observations and is itself an estimate of the unconditional mean of the response. The other parameters have the interpretation  $\gamma_j = \beta_{j+1} = \mu_j - \beta_1$  and represent the differences between the conditional mean of the response in the  $j$ -th group and the unconditional mean of the response.

This parametrization is not recommended despite its seemingly attractive interpretation. The reason is that the meaning of its parameters depends on the data (through  $n_j$ ) and thus it changes when the data set is modified. This is undesirable. We strongly prefer parametrizations that have the same interpretation in any data set.

### 8.1.4. Helmert parametrization

This is another parametrization offered by R. It leads to parameters that can be interpreted as follows:  $\beta_1 = \mu_1$ ,  $\beta_2 = \mu_2 - \mu_1$ ,  $\beta_3 = \mu_3 - \frac{\mu_1 + \mu_2}{2}$ ,  $\beta_4 = \mu_4 - \frac{\mu_1 + \mu_2 + \mu_3}{3}$  etc. The parameter  $\beta_j$  for  $j > 1$  expresses the difference between the conditional mean in the  $j$ -th group and the arithmetic average of the conditional means in the preceding groups.

This parametrization is rarely useful, it is suitable only for certain specific industrial applications.

### 8.1.5. Orthogonal polynomials for ordered factors

We say that  $Z \in \{1, \dots, m\}$  is an *ordered factor* (or ordered categorical variable) if it indicates membership in one of  $m$  groups and the groups are in some sense ordered, that is, group 1 precedes group 2, which precedes group 3, etc.

Examples of ordered factors are school grades (1 = excellent, 2 = very good, ..., 5 = very poor), opinion scales (1 = strongly agree, 2 = somewhat agree, ..., 5 = strongly disagree), or grouped continuous variables such as age groups. The values of ordered factors have some numerical interpretation and it is desirable to take this into account in the analysis. In particular, it makes sense to consider a polynomial model expressing the mean in the  $j$ -th group as follows:

$$\mathbb{E}[Y_i | Z_i] = \beta_1 + \gamma_1 Z_i + \gamma_2 Z_i^2 + \dots + \gamma_{m-1} Z_i^{m-1}.$$

This model is equivalent to the cell-means model (8.1) because any constellation of values of the  $m$  means can be exactly fitted by a polynomial of degree  $m-1$ . However, the interpretation of the regression coefficients is difficult and their least squares estimators are strongly correlated because certain columns of the resulting regression matrix are correlated.

Instead, we consider an equivalent polynomial model

$$\mathbb{E}[Y_i | Z_i] = \beta_1 + \gamma_1 g_1(Z_i) + \gamma_2 g_2(Z_i) + \cdots + \gamma_{m-1} g_{m-1}(Z_i),$$

where  $g_j(Z_i)$  are mutually orthogonal polynomials in  $Z_i$  of degree  $j$ ,  $j = 1, \dots, m-1$ . The form of the regression matrix can be described by the following table.

Group		$\mathbf{X}_i^\top$			
1	1	$g_1(1)$	$g_2(1)$	$\cdots$	$g_{m-1}(1)$
2	1	$g_1(2)$	$g_2(2)$	$\cdots$	$g_{m-1}(2)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$		$\vdots$
$m-1$	1	$g_1(m-1)$	$g_2(m-1)$	$\cdots$	$g_{m-1}(m-1)$
$m$	1	$g_1(m)$	$g_2(m)$	$\cdots$	$g_{m-1}(m)$
		$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	
		$\mathbf{g}_1$	$\mathbf{g}_2$	$\mathbf{g}_{m-1}$	

The vectors  $\mathbf{g}_j$  are polynomials of degree  $j$  evaluated at the group codes  $1, 2, \dots, m$  satisfying  $\mathbf{1}^\top \mathbf{g}_j = 0$  for  $j \in \{1, \dots, m-1\}$  and  $\mathbf{g}_j^\top \mathbf{g}_k = 0$  for  $j \neq k \in \{1, \dots, m-1\}$ . Such polynomials are not uniquely determined. For the most frequent situations when  $m = 3$  or  $m = 4$ , we can use

$$m = 3: \quad (\mathbf{g}_1, \mathbf{g}_2) = \begin{pmatrix} -1 & 1 \\ 0 & -2 \\ 1 & 1 \end{pmatrix}, \quad m = 4: \quad (\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3) = \begin{pmatrix} -3 & 1 & -1 \\ -1 & -1 & 3 \\ 1 & -1 & -3 \\ 3 & 1 & 1 \end{pmatrix}.$$

The functional dependence of  $\mathbb{E}[Y_i | Z_i]$  on  $Z_i$  can be determined by testing the significance of the individual parameters  $\gamma_1, \dots, \gamma_{m-1}$ . If there is interest in the linear part, the coefficient  $\gamma_1$  describes the slope of the linear relationship and a test of linearity can be performed by testing  $H_0: \gamma_2 = \cdots = \gamma_{m-1} = 0$ . The LSE of  $\gamma_1$  has exactly the same value and standard error as if the model only contained the linear term without any higher order terms (see Chapter ?? for a justification of this claim).

**Note.** This is the default parametrization in R when a factor is declared as ordered by calling the function `factor(z, ordered=TRUE)` or `ordered(z)`. For any factor, this parametrization can be also requested by calling the function `C(z, contr.poly)` and using this in the regression model instead of the original factor `z`. In R, it is called “polynomial contrasts”.

**Note.** If the ordered factor arose by grouping a numerical covariate into intervals (for example age groups defined as 0–10 years, 11–20 years, etc.), it is better to code the groups by midpoints of those intervals rather than by integers  $1, 2, \dots, m$ , especially if the intervals do not have the same width. Such covariates can be also treated by the methods described in the next section, though.

## 8.2. Parametrization of Quantitative (Numerical) Covariates

Quantitative (numerical) covariates are measurements of some numerical quantity that often has some associated physical units (age, height, price, temperature, percentage, number of children). Thus, their values have numerical interpretation. Quantity covariates may be realizations of a continuous random variable with many distinct values and infrequent ties due to rounding, may attain a limited number of pre-specified values (temperature in industrial experiments) or may be purely discrete. A discrete quantitative variable can be also treated as a factor or an ordered factor by the methods discussed in the previous section.

The simplest way how to include a quantitative covariate in the model is in a linear form  $m(Z) \equiv E[Y|Z] = \beta_1 + \beta_2 Z$ . (In this section, we will use the notation  $m(Z)$  to express the functional form of the dependence of  $EY$  on  $Z$ .) However, the linear form of  $m(Z)$  may not be appropriate.

To come up with more complex functional forms, we can transform  $Z$  by a pre-selected set of basis functions  $g_1, \dots, g_m$  and fit the model

$$m(Z) = \beta_1 + \gamma_1 g_1(Z_i) + \gamma_2 g_2(Z_i) + \dots + \gamma_m g_m(Z_i).$$

This form can capture any function that can be expressed as a linear combination of the selected basis functions.

How to choose the basis functions? We can learn from the results of descriptive analysis of residuals (see Chapter 6) or we can use submodel testing to compare more complex models with simpler models and remove the unnecessary terms. In the rest of this section, we will describe various possibilities for the choice of the basis functions.

### 8.2.1. Factorization of a numerical covariate

We can transform a numerical variable  $Z$  into a factor variable by specifying groups determined by ranges of values of  $Z$ . Take the points

$$t_0 = -\infty < t_1 < \dots < t_{m-1} < t_m = \infty,$$

define intervals  $A_j = (t_{j-1}, t_j)$  for  $j = 1, \dots, m$  and set

$$Z_i^* = j \Leftrightarrow Z_i \in A_j.$$

Then treat  $Z_i^*$  as a factor variable (or an ordered factor variable) with  $m$  levels using the methods of Section 8.1. This is equivalent to approximating  $m(Z)$  by a piecewise constant function on the prespecified intervals. We can visualize the fitted piecewise constant function and investigate it to choose a suitable continuous functional form for further investigation or keep the piecewise constant approximation in the final model.

**Note.** We can test whether the true effect of the covariate is linear by fitting the model that includes the linear form as well as the factorized form

$$m(Z) = \beta_1 + \beta_2 Z + \gamma_2 \mathbb{1}(Z \in A_2) + \cdots + \gamma_m \mathbb{1}(Z \in A_m)$$

and testing the hypothesis  $H_0 : \gamma_2 = \cdots = \gamma_m = 0$  by comparing this model with  $m(Z) = \beta_1 + \beta_2 Z$ .

**Note.** When choosing the cutoff points  $t_1, \dots, t_{m-1}$ , we should make sure that there is enough observations in each interval  $A_j$  so that we can estimate the expectation on that interval with sufficient precision.

**Advantages:** Simple flexible method, easy to implement, with straightforward interpretation of model parameters.

**Disadvantages:** The results depend on the choice of the cutoff points  $t_j$ , the intervals must be wide enough to include enough observations, the model may include too many parameters, factorization entails a partial loss of information contained in the original form of the covariate and is inefficient.

### 8.2.2. Polynomial basis

Choose  $g_j$  as a polynomials of degree  $j$  for  $j = 1, \dots, m$ . Then  $m(Z)$  has the form of a polynomial of degree  $m$ . For example,

$$m(Z) = \beta_1 + \gamma_1 Z + \gamma_2 Z^2 + \cdots + \gamma_m Z_m.$$

Because this choice of basis functions creates a regression matrix with highly correlated columns (see the discussion in Section 8.1.5) one could create the basis from orthogonal polynomials instead, e.g. using Legendre polynomials after transforming all the values of the covariate into the closed interval  $\langle -1, 1 \rangle$ .

**Note.** In R, an orthogonal polynomial basis of degree  $m$  can be calculated by calling the function `poly(z, degree=m)`, which can be used directly on the right-hand side of the model formula.

According to the Weierstrass Theorem, any continuous function defined on a finite closed interval can be approximated by a polynomial with an arbitrary precision. In practice, it is not really so. Fitting a polynomial of a very high degree requires a lot of parameters and their estimates are imprecise. The resulting polynomial tends to be wiggly and may fit the data poorly at the edges as well as over flat areas.\* A visualization of the fitted polynomial may give a misleading impression about the true shape of the function  $m(Z)$ .

**Advantages:** Simple flexible method, easy to implement, theoretically may provide a good approximation to  $m(Z)$ .

---

\* This is called *Runge's phenomenon*



**Disadvantages:** Difficult interpretation of model parameters, poor fit at the edges of the data and over flat areas.

### 8.2.3. Piecewise polynomial bases, B-splines

Splines represent another convenient and flexible method for modeling complex relationships between a continuous covariate and the mean of the response.

#### Definition 8.1.

- A spline of degree  $k$  is a piecewise polynomial function of degree  $k$  with  $k - 1$  continuous derivatives everywhere.
- A spline basis is a set of functions  $B_{j,k}(x)$  such that any spline of degree  $k$  can be expressed as a linear combination of  $B_{j,k}(x)$ .  $\nabla$

The most convenient spline basis is so called B-spline basis that satisfies the conditions

$$B_{j,k}(x) \geq 0 \quad \forall x \quad \text{and} \quad \sum_j B_{j,k}(x) = 1 \quad \forall x.$$

Consider a continuous variable  $Z$  with observed values  $Z_1, \dots, Z_n$ . Choose points  $t_1 < t_2 < \dots < t_m < t_{m+1}$  such that  $t_1 < Z_i < t_{m+1} \quad \forall i$ . These points are called *knots*<sup>\*</sup>. The knots divide the range of data values  $Z_1, \dots, Z_n$  into  $m$  disjoint intervals  $I_j \equiv \langle t_j, t_{j+1} \rangle$ . They should be chosen so that each of these intervals contains a sufficient number of observations. The knots  $t_1$  and  $t_{m+1}$  are called *boundary knots*, the other  $m - 1$  knots are called *inner knots*.

Given the set of knots  $t_j$  and intervals  $I_j$ , define the B-spline basis functions of degree  $k = 0$  as

$$B_{j,0}(x) = \mathbb{1}(x \in I_j) = \begin{cases} 1 & x \in \langle t_j, t_{j+1} \rangle = I_j, \\ 0 & \text{otherwise,} \end{cases}$$

for  $j = 1, \dots, m$ . Linear combinations of these zero-degree basis functions are discontinuous piecewise constant functions. The basis functions are non-negative and sum to one, as required. They can be included in the linear model in the form

$$m(Z) = \gamma_1 B_{1,0}(Z) + \gamma_2 B_{2,0}(Z) + \dots + \gamma_m B_{m,0}(Z)$$

or, preferably, with an explicit intercept term after dropping the first basis function

$$m(Z) = \beta_1 + \gamma_2 B_{2,0}(Z) + \dots + \gamma_m B_{m,0}(Z)$$

These basis function create exactly the same approximation by a piecewise constant function as the factorization approach discussed in Section 8.2.1.

---

<sup>\*</sup> Český uzly




Figure 8.1.: Recurrent procedure for generation of B-spline bases.

The higher degree B-spline basis functions are constructed by a recurrent relationship called *the Cox-de Boor recursion formula*:

$$B_{j,k}(x) = w_{j-1,k}(x)B_{j-1,k-1}(x) + [1 - w_{j,k}(x)]B_{j,k-1}(x)$$

for  $j = 1, \dots, m+k$  and  $k = 1, 2, \dots$ , where

$$w_{j,k}(x) = \begin{cases} \frac{x-t_j}{t_{j+k}-t_j} & \text{for } x \in \langle t_j, t_{j+k} \rangle, j \neq 1 \text{ and } j \neq m+k, \\ 0 & \text{for } x \notin \langle t_j, t_{j+k} \rangle \text{ or } j = 1, \\ 1 & \text{for } j = m+k. \end{cases}$$

This creates higher degree basis functions by linear combinations of two neighboring lower-degree basis functions, with coefficients that linearly depend on the argument. A graphical illustration of the recurrent procedure is provided by Figure 8.1.

For degree zero, each basis function spans only one interval. For degree one, each basis functions spans two neighboring intervals (except at the boundary intervals). For degree  $k$ , each basis functions spans  $k+1$  neighboring intervals (except for the first  $k$  and the last  $k$  basis functions). With  $m$  intervals, we have  $m-1$  inner knots between them,  $m+k$  basis functions and  $m+k-1$  model parameters (the basis functions are linearly dependent, so the first of them is dropped).

Using the Cox-de Boor formula, the B-spline basis functions of degree 1 can be expressed as follows:

$$B_{1,1}(x) = 0 + \left[ 1 - \frac{x - t_1}{t_2 - t_1} \right] \mathbb{1}(x \in \langle t_1, t_2 \rangle) = \frac{t_2 - x}{t_2 - t_1} \mathbb{1}(x \in \langle t_1, t_2 \rangle)$$

$$B_{2,1}(x) = \frac{x - t_1}{t_2 - t_1} \mathbb{1}(x \in \langle t_1, t_2 \rangle) + \frac{t_3 - x}{t_3 - t_2} \mathbb{1}(x \in \langle t_2, t_3 \rangle)$$

et cetera. Figures 8.2, 8.3, and 8.4 illustrate the linear, quadratic and cubic B-spline bases.

When selecting the knot locations, it is important to make sure that there are enough observations in the interval between two neighboring knots. One could place the inner knots equidistantly, or place them at certain empirical quantiles of the observed values of the covariate, or make them denser at regions where one expects more abrupt changes of the fitted function and sparser elsewhere. Outer knots should be placed independently of the data (not at the minimum and maximum of the observed covariate values).

One can use a submodel F-test to verify the degree of the spline as well as the significance of one or more selected knots: lowering a degree of the spline or removing knots creates a submodel.

**Advantages:**

- very flexible tool, many functions can be well approximated by splines with suitably selected knots
- each basis spans at most  $k + 1$  intervals; individual observations have only local influence on the fitted curve, not a global influence over the whole range of the data
- splines are easy to use in R using the function `bs` in the library `spline`

**Disadvantages:**

- the estimated parameters are hard to interpret (except in case of linear splines)
- the fitted function must be plotted in order to understand the meaning of the results
- it may be difficult to calculate predictions correctly in an external data set (one must make sure that the spline bases are calculated in exactly the same way as in the original data)

Figures 8.5 and 8.6 show examples of nonlinear mean functions fitted by linear regression lines, high-order polynomials and cubic splines.




Figure 8.2.: B-spline bases of degree 1 over the interval  $\langle 0, 10 \rangle$ .




Figure 8.3.: B-spline bases of degree 2 over the interval  $\langle 0, 10 \rangle$ .




Figure 8.4.: B-spline bases of degree 3 over the interval  $\langle 0, 10 \rangle$ .

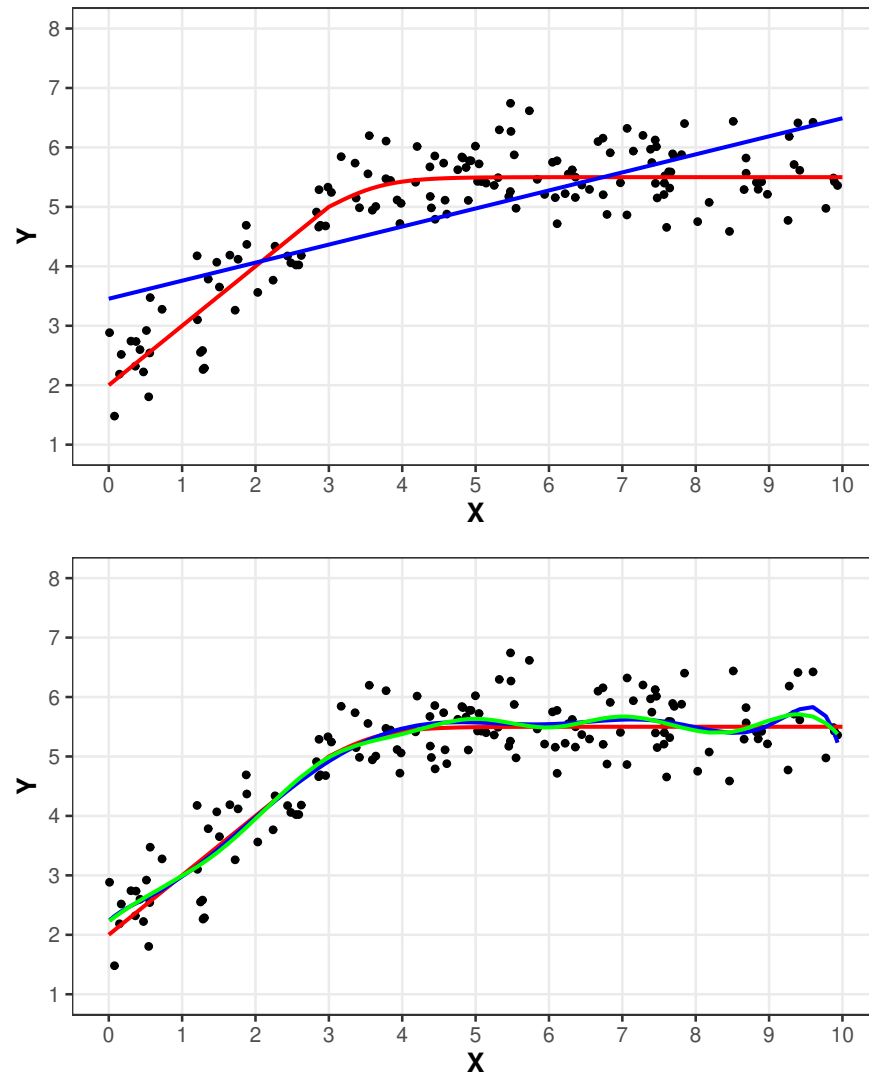


Figure 8.5.: Logistic pipe function (red curve = true mean) fitted by a regression line (top panel), by a 10-th degree polynomial (blue curve in bottom panel), and by a 3rd degree spline with 7 inner knots (green curve in bottom panel).

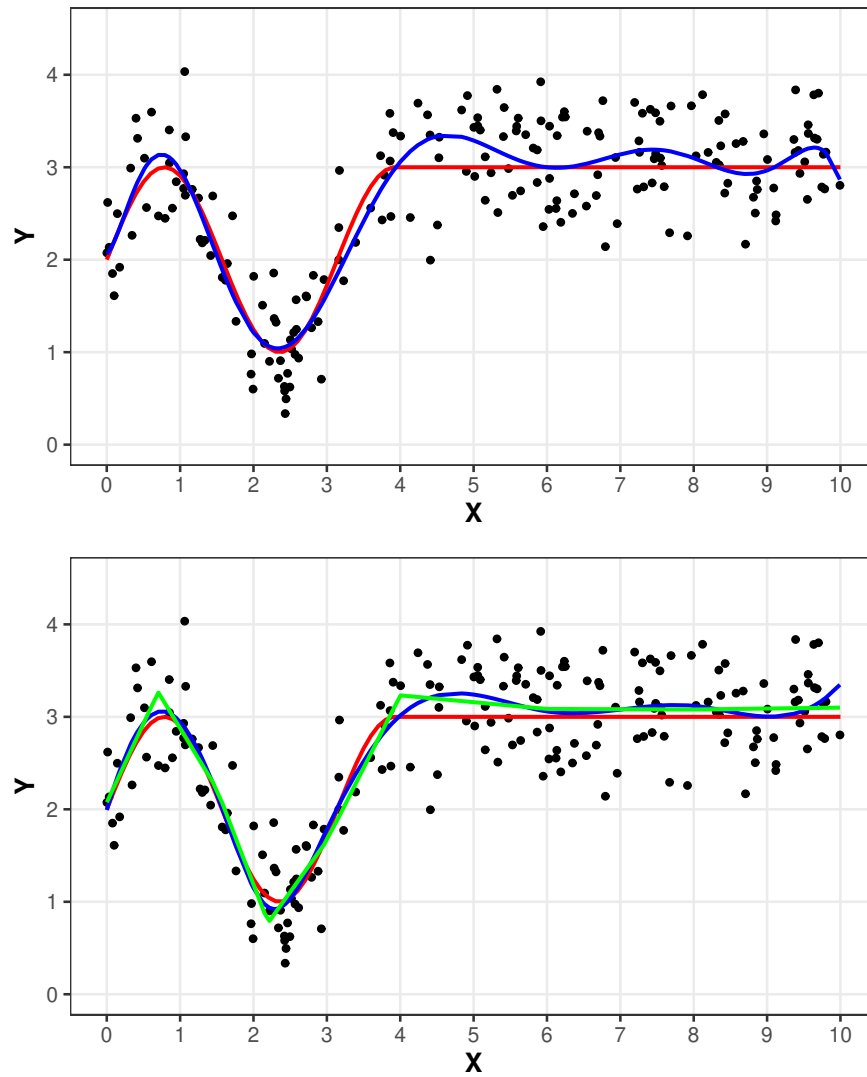


Figure 8.6.: Sine-constant function (red curve = true mean) fitted by 10th degree polynomial (blue curve in top panel) and by 1st degree (green curve in bottom panel) and 3rd degree (blue curve in bottom panel) spline with 7 inner knots.

#### 8.2.4. Sine-cosine basis

Suppose  $m(z) \equiv \mathbb{E}[Y | Z = z]$  is a periodic function of a bounded variation in the argument  $z$ , with a period  $d$ . Such a function can be expressed using Fourier series in a sine-cosine form as follows

$$m(z) = a_0 + \sum_{j=1}^M a_j \cos \varphi_j \cos\left(\frac{2\pi j}{d}z\right) + \sum_{j=1}^M a_j \sin \varphi_j \sin\left(\frac{2\pi j}{d}z\right),$$

where  $M$  may be infinite,  $a_j$  represents the amplitude of the  $j$ -th term and  $\varphi_j$  represents the phase shift of the  $j$ -th term.

This expression can be approximates in a linear regression model by choosing the maximum period  $d$  and a finite number of terms  $M$  and transforming the covariate  $Z$  in the following way:

$$m(Z) = \mathbb{E}[Y | Z] = \beta_1 + \sum_{j=1}^M \gamma_{1j} \cos\left(\frac{2\pi j}{d}Z\right) + \sum_{j=1}^M \gamma_{2j} \sin\left(\frac{2\pi j}{d}Z\right).$$

The parameters  $\gamma_{1j}$  and  $\gamma_{2j}$  are estimated as regression parameters using least squares.



## 9. Multiple Comparisons and Simultaneous Confidence Intervals

Recall Section 4.1.4, where we assumed normality and considered testing the hypothesis

$$H_0 : \mathbb{C}\boldsymbol{\beta} = \mathbf{0} \quad \text{against} \quad H_1 : \mathbb{C}\boldsymbol{\beta} \neq \mathbf{0},$$

where  $\mathbb{C}_{q \times p}$  with  $q \leq p$  and  $r(\mathbb{C}) = q$  is a matrix of constants. This is tested by a submodel F-test. By Lemma 4.7, the same test can be applied asymptotically when the responses are not normal. The hypothesis  $H_0$  is equivalent to the set of  $q$  partial hypotheses  $H_{0j} : \mathbf{c}_j^T \boldsymbol{\beta} = 0$  for  $j = 1, \dots, q$ , where  $\mathbf{c}_j$  is the  $j$ -th row of the matrix  $\mathbb{C}$ .

Now imagine the composite hypothesis  $H_0$  is not true. This means that at least one of the partial hypotheses is violated. We want to test the partial hypotheses in order to find out which of them can be rejected. As we will see in the following section, this cannot be done by performing the  $q$  partial hypotheses with the usual level because the overall rate of type I error would be much larger than desired.

### 9.1. Bonferroni Procedure

Let us state the problem quite generally. Suppose  $\Theta$  is the parameter space and consider subsets  $\Theta_k \subset \Theta$  for  $k = 1, \dots, M$  such that  $\bigcap_{k=1}^M \Theta_k \neq \emptyset$ . Formulate hypotheses about the true parameter  $\boldsymbol{\theta}$  falling into those subsets:

$$H_{0k} : \boldsymbol{\theta} \in \Theta_k, \quad k = 1, \dots, M.$$

These hypotheses are to be tested with tests specified by test statistics  $T_k$  and critical regions  $C_k$ . Hence, the hypothesis  $H_{0k}$  is rejected if and only if  $T_k \in C_k$ . Suppose the critical regions are selected so that the level of each test is  $\alpha^*$ , that is

$$P(T_k \in C_k \mid H_{0k} \text{ holds}) = \alpha^*$$

(replace the equality by a limit if the test is asymptotic). In this way, we can control the probability of type I error for each test separately. But suppose that we need to control these errors *simultaneously* for all the tests. We need a procedure that assures

$$P(\{T_1 \in C_1\} \cup \{T_2 \in C_2\} \cup \dots \cup \{T_M \in C_M\} \mid H_{01}, \dots, H_{0M} \text{ hold}) \leq \alpha$$

so that the probability of committing at least one type I error is bounded from above by  $\alpha$ .

For any random events  $A_1, \dots, A_M$ ,

$$P[A_1 \cup A_2 \cup \dots \cup A_M] \leq \sum_{k=1}^M P[A_k]$$

and equality holds if and only if  $A_j$  and  $A_k$  are disjoint for any  $j \neq k$ . This is called *Bonferroni inequality* (Bonferroni 1936). Applying this, we get

$$\begin{aligned} P\left(\bigcup_{k=1}^M \{T_k \in C_k\} \mid H_{01}, \dots, H_{0M} \text{ hold}\right) \\ \leq \sum_{k=1}^M P(T_k \in C_k \mid H_{01}, \dots, H_{0M} \text{ hold}) = \sum_{k=1}^M P(T_k \in C_k \mid H_{0k} \text{ holds}) = M\alpha^*. \end{aligned}$$

**Carlo Emilio Bonferroni** (1892 – 1960) was an Italian mathematician who worked on probability theory. Bonferroni is best known for the Bonferroni inequalities (a generalization of the union bound), and for the Bonferroni correction in statistics, which he did not invent, but which is related to the Bonferroni inequalities.

Source: [https://en.wikipedia.org/wiki/Carlo\\_Emilio\\_Bonferroni](https://en.wikipedia.org/wiki/Carlo_Emilio_Bonferroni)

If we conduct  $M$  tests on the level  $\alpha^*$ , the overall level (probability of committing at least one type I error) may be as large as  $M\alpha^*$ . *The more tests are performed the more likely it is that at least one valid hypothesis is rejected.* To protect against this, it is sufficient to take  $\alpha^* = \alpha/M$  for the level of each individual test. The overall level of all tests is then guaranteed to be at most the desired  $\alpha$ . This is called *the Bonferroni method*.

The same principle applies to confidence intervals. Suppose  $U_1, \dots, U_M$  are confidence intervals for parameters  $\theta_1, \dots, \theta_M$ , each of them has coverage probability  $1 - \alpha^*$  (or converging to  $1 - \alpha^*$  if the interval is asymptotic). This is equivalent to  $P[U_k \not\supset \theta_k] = \alpha^*$ .

We want to control the errors simultaneously requiring that the probability that at least one interval does not cover the true parameter is at most  $\alpha$ , i.e.,

$$P\left[\bigcup_{k=1}^M \{U_k \not\supset \theta_k\}\right] \leq \alpha$$

and, equivalently,

$$P\left[\bigcap_{k=1}^M \{U_k \supset \theta_k\}\right] \geq 1 - \alpha.$$

By Bonferroni inequality, we can achieve this by setting  $\alpha^* = \alpha/M$ , that is, by setting the coverage probabilities of individual intervals to  $1 - \alpha/M$ .

### Bonferroni method in linear regression

Let us return to the problem of testing

$$H_0 : \mathbb{C}\boldsymbol{\beta} = \mathbf{0} \quad \text{against} \quad H_1 : \mathbb{C}\boldsymbol{\beta} \neq \mathbf{0},$$

and the partial hypotheses  $H_{0j} : \mathbf{c}_j^\top \boldsymbol{\beta} = 0$  in linear regression.

By Lemma 4.2 and Lemma 4.7,  $H_0$  is rejected when

$$\frac{1}{q\hat{\sigma}_e^2} \widehat{\boldsymbol{\beta}}^\top \mathbb{C}^\top [\mathbb{C}(\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{C}^\top]^{-1} \mathbb{C} \widehat{\boldsymbol{\beta}} \geq F_{q, n-p}(1-\alpha).$$

This test has the level  $\alpha$  (exact or asymptotic) but if it rejects we do not know which of the hypotheses  $H_{0j}$  should be rejected and which should not.

Instead, we can test each  $H_{0j} : \mathbf{c}_j^\top \boldsymbol{\beta} = 0$  separately by a t-test at the level  $\alpha/q$  (recall Lemma 4.1 and Lemma 4.6). We reject  $H_{0j}$  when

$$\frac{|\mathbf{c}_j^\top \widehat{\boldsymbol{\beta}}|}{\sqrt{\hat{\sigma}_e^2 \mathbf{c}_j^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}_j}} \geq t_{n-p} \left( 1 - \frac{\alpha}{2q} \right).$$

Of course, the Bonferroni method has the true overall level bounded from above by  $\alpha$ , and sometimes the true overall level is much less than  $\alpha$ . It is quite possible that  $H_0$  is rejected by the overall F-test but none of the  $H_{0j}$  is rejected with the adjusted level  $\alpha/M$ .

Now consider individual confidence intervals for the true values of  $\mathbf{c}_j^\top \boldsymbol{\beta}$ . By Lemma 4.1 and Lemma 4.6, these confidence intervals have the form

$$\mathbf{c}_j^\top \widehat{\boldsymbol{\beta}} \mp t_{n-p} \left( 1 - \frac{\alpha}{2} \right) \hat{\sigma}_e \sqrt{\mathbf{c}_j^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}_j}$$

and their coverage probability is  $1 - \alpha$  (exactly or asymptotically).

From Lemma 4.2 and Lemma 4.7, we can get a confidence region for  $\mathbb{C}\boldsymbol{\beta}$  with coverage probability  $1 - \alpha$ . It includes all values  $\mathbb{C}\boldsymbol{\beta} \in \mathbb{R}^q$  such that

$$\frac{1}{q\hat{\sigma}_e^2} (\mathbb{C}\widehat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta})^\top [\mathbb{C}(\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{C}^\top]^{-1} (\mathbb{C}\widehat{\boldsymbol{\beta}} - \mathbb{C}\boldsymbol{\beta}) \leq F_{q, n-p}(1-\alpha).$$

However, this is a hyper-ellipsoid in  $\mathbb{R}^q$ , which is difficult to describe or evaluate. We would much prefer to have confidence intervals for the individual linear combinations  $\mathbf{c}_j^\top \boldsymbol{\beta}$ , keeping their simultaneous coverage at  $\geq 1 - \alpha$ . Such intervals can be obtained using Bonferroni method. They have the form

$$\mathbf{c}_j^\top \widehat{\boldsymbol{\beta}} \mp t_{n-p} \left( 1 - \frac{\alpha}{2q} \right) \hat{\sigma}_e \sqrt{\mathbf{c}_j^\top (\mathbb{X}^\top \mathbb{X})^{-1} \mathbf{c}_j}$$

and their simultaneous coverage probability is  $\geq 1 - \alpha$ .

### Pairwise comparison of means in one-way ANOVA using Bonferroni method

Consider a factor variable  $G_i \in \{1, \dots, m\}$  specifying membership in one of  $m$  groups for each observation. Let the expectations of the response in the groups be  $\mu_j \equiv E[Y_i | G_i = j]$ . The question is: which of the groups differ in the expectations from each other and which are the same? This is one of the most common questions in regression analyses of factor covariates. Suppose for simplicity that there is no other predictor in the model.

We use the one-way ANOVA model with reference group parametrization:

$$E[Y_i | G_i] = \beta_1 + \gamma_2 \mathbb{1}(G_i = 2) + \dots + \gamma_m \mathbb{1}(G_i = m).$$

The group expectations can be expressed as  $\mu_1 = \beta_1, \mu_2 = \beta_1 + \gamma_2, \dots, \mu_m = \beta_1 + \gamma_m$ .

Consider the overall hypothesis  $H_0 : \mu_1 = \dots = \mu_m$  and the  $\binom{m}{2}$  partial hypotheses comparing the individual expectations with each other:  $H_{jk} : \mu_j = \mu_k$  for  $j < k$ . Denote by  $n_j$  the number of observations in the  $j$ -th group. In one-way ANOVA model, we have

$$\begin{aligned} \hat{\mu}_j &= \frac{1}{n_j} \sum_{i=1}^n Y_i \mathbb{1}(G_i = j) \\ \text{var } \hat{\mu}_j &= \frac{\sigma_e^2}{n_j} \end{aligned}$$

Also,  $\hat{\mu}_j$  and  $\hat{\mu}_k$  are independent for  $j \neq k$ .

The hypothesis  $H_0$  is violated if and only if any of the hypotheses  $H_{jk}$  is violated. If we reject  $H_0$  by the one-way ANOVA F-test, we want to know which of the  $H_{jk}$  can be rejected. Using the Bonferroni method, we test each of these partial hypotheses by a t-test at the level  $\alpha^* = \alpha / \binom{m}{2} = \frac{2\alpha}{m(m-1)}$  and we reject  $H_{jk} : \mu_j = \mu_k$  when

$$\frac{|\hat{\mu}_j - \hat{\mu}_k|}{\hat{\sigma}_e \left( \frac{1}{n_j} + \frac{1}{n_k} \right)} \geq t_{n-p} \left( 1 - \frac{\alpha}{m(m-1)} \right).$$

It is not difficult to see that for larger  $m$  this procedure is very inefficient: it will be rarely able to detect differences between the individual group means unless the differences are very large.

### Bonferroni method: summary

The Bonferroni method is extremely general and has no restrictive assumptions. It can be applied to any testing problem and any set of confidence intervals no matter what the model is and no matter how the tests and confidence intervals are constructed.

On the other hand, the Bonferroni method is very conservative. The resulting tests have low power and the confidence intervals are unnecessarily wide. The larger is the number of partial tests or confidence intervals  $M$  the poorer is the performance of this procedure.

In the next two sections we introduce two alternative methods that are specifically designed for linear regression.

## 9.2. Tukey Method

The Tukey method [Tukey \(1949\)](#) is specifically designed to address the problem of pairwise comparisons of the expectations in one-way ANOVA.\*

**John Tukey (1915 – 2000)** was an American mathematician and statistician, best known for the development of the fast Fourier Transform (FFT) algorithm and the box plot. The Tukey HSD test, the Tukey lambda distribution, the Tukey test of additivity, and the Teichmüller–Tukey lemma all bear his name. He is also credited with coining the term bit and the first published use of the word software.

Source: [https://en.wikipedia.org/wiki/John\\_Tukey](https://en.wikipedia.org/wiki/John_Tukey)

### Tukey method: introduction

Consider  $M$  independent variables  $T_j \sim N(\mu_j, \sigma^2)$ ,  $j = 1, \dots, M$ . Let  $\hat{\sigma}^2$  be an estimator of  $\sigma^2$  such that

- (i)  $\frac{f\hat{\sigma}^2}{\sigma^2} \sim \chi_f^2$  for some natural number  $f$  and
- (ii)  $\hat{\sigma}^2$  is independent of  $(T_1, \dots, T_M)$ .

Normalize the observations and look at their range. Define

$$Q = \frac{\max_{j=1, \dots, M} \left( \frac{T_j - \mu_j}{\sigma} \right) - \min_{j=1, \dots, M} \left( \frac{T_j - \mu_j}{\sigma} \right)}{\hat{\sigma}/\sigma} = \frac{\max_{j=1, \dots, M} (T_j - \mu_j) - \min_{j=1, \dots, M} (T_j - \mu_j)}{\hat{\sigma}}$$

Because  $\frac{T_j - \mu_j}{\sigma}$  are independent variables distributed as  $N(0, 1)$  and  $\frac{\hat{\sigma}}{\sigma}$  is a transformation of a random variable with a  $\chi_f^2$  distribution, which is independent of the variables in the numerator, the following lemma holds.

**Lemma 9.1 (Tukey).** *The statistic  $Q$  has a distribution that depends on  $M$  and  $f$  but not on  $\mu_1, \dots, \mu_n$  and  $\sigma^2$ .*  $\diamond$

The distribution function of the statistic  $Q$  can be calculated. The distribution is called *Tukey distribution*. Denote its quantile function by  $q_{f,M}(u)$ .

From this, we have

$$P \left[ \max_j (T_j - \mu_j) - \min_j (T_j - \mu_j) < \hat{\sigma} q_{f,M}(1 - \alpha) \right] = 1 - \alpha.$$

The difference between the largest and the smallest number from a finite set is smaller than a constant if and only if each of the distances between any pair of them is smaller than the same constant. Hence,

$$P \left[ |(T_j - \mu_j) - (T_k - \mu_k)| < \hat{\sigma} q_{f,M}(1 - \alpha) \quad \forall j, k \right] = 1 - \alpha$$

\* This method is also called Tukey's honest significant differences or Tukey HSD

and

$$P\left[T_j - T_k - \hat{\sigma}q_{f,M}(1 - \alpha) < \mu_j - \mu_k < T_j - T_k + \hat{\sigma}q_{f,M}(1 - \alpha) \quad \forall j, k\right] = 1 - \alpha$$

We have obtained simultaneous confidence intervals for the differences in the expectations and the probability that *all of them* contain the true values is jointly  $1 - \alpha$ .

The confidence intervals can be easily converted into tests. The hypothesis  $H_{jk} : \mu_j = \mu_k$  is rejected when

$$\frac{|T_j - T_k|}{\hat{\sigma}} > q_{f,M}(1 - \alpha).$$

The overall level of all  $\binom{M}{2}$  tests is  $\alpha$ .

#### Tukey method: extension

The Tukey method assumes that the observations have all the same variance, which is very restrictive. So, relax the assumption to  $T_j \sim N(\mu_j, v_j\sigma^2)$ , where  $v_j$  is a known positive constant. We can repeat the derivation of the Tukey method until we get the equality

$$P\left[\left|\frac{T_j - \mu_j}{\sqrt{v_j}} - \frac{T_k - \mu_k}{\sqrt{v_k}}\right| < \hat{\sigma}q_{f,M}(1 - \alpha) \quad \forall j, k\right] = 1 - \alpha.$$

[Hayter \(1984\)](#) has proven that this equality under the current assumptions implies

$$P\left[\left|(T_j - \mu_j) - (T_k - \mu_k)\right| < \hat{\sigma}\sqrt{\frac{v_j + v_k}{2}}q_{f,M}(1 - \alpha) \quad \forall j, k\right] \geq 1 - \alpha.$$

What appears inside is actually the square root of the average estimated variance in the two groups:  $\sqrt{\frac{v_j\hat{\sigma}^2 + v_k\hat{\sigma}^2}{2}}$ . Thus, we can use the Tukey method even with unequal variances replacing the equal estimated variance by the average of estimated variances in the two groups, otherwise the procedure is unchanged. However, the Tukey method becomes conservative under such circumstances.

#### Tukey method: application in linear regression

In linear regression, we could take  $T_j = \mathbf{c}_j^\top \hat{\boldsymbol{\beta}}$  and try to apply the Tukey method on this. We have  $T_j \sim N(\mathbf{c}_j^\top \boldsymbol{\beta}, \sigma_e^2 \mathbf{c}_j^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{c}_j)$ . So we take  $\mu_j \equiv \mathbf{c}_j^\top \boldsymbol{\beta}$  and  $v_j \equiv \mathbf{c}_j^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{c}_j$ . However, the Tukey method also requires independence between the individual  $T_j$ 's. This is very difficult to achieve because the  $T_j$ 's are different linear combinations of the same random vector  $\hat{\boldsymbol{\beta}}$ . The only case when this assumption clearly holds is the case of simple one-way ANOVA with  $\mu_j$  equal to the expectation of the  $j$ -th group and  $T_j = \hat{\mu}_j$  is the average of the observations in the  $j$ -th group. The constants  $v_j$  are defined as  $v_j = 1/n_j$ .

To conclude, in a simple one-way ANOVA with  $m$  groups and no additional covariates, we can use the Tukey method to test the hypotheses  $H_{jk} : \mu_j = \mu_k$ . We reject  $H_{jk}$  when

$$\frac{|\hat{\mu}_j - \hat{\mu}_k|}{\hat{\sigma}_e \sqrt{\frac{1}{2} \left( \frac{1}{n_j} + \frac{1}{n_k} \right)}} \geq q_{n-m, m}(1 - \alpha).$$

These test have an overall level at most  $\alpha$ . We can also construct simultaneous confidence intervals for the differences between the group means. The interval for  $\mu_j - \mu_k$  is

$$\hat{\mu}_j - \hat{\mu}_k \mp \hat{\sigma}_e \sqrt{\frac{1}{2} \left( \frac{1}{n_j} + \frac{1}{n_k} \right) q_{n-m, m}(1 - \alpha)},$$

and the overall coverage of these intervals is  $\geq 1 - \alpha$ .

The Tukey method is theoretically appealing but it has so strong assumptions that it can rarely be used in linear regression models. The only case when this method can be applied is for comparison of group means in one-way ANOVA.

### 9.3. Scheffé Method

*Henry Scheffé (1907 – 1977) was an American statistician. He is known for the Lehmann–Scheffé theorem and Scheffé’s method. Scheffé was president of the Institute of Mathematical Statistics in 1954, and also served as vice president of the American Statistical Association from 1954 to 1956.*

Source: [https://en.wikipedia.org/wiki/Henry\\_Scheffé](https://en.wikipedia.org/wiki/Henry_Scheffé)

Scheffé (1953) proposed another method for simultaneous tests and confidence intervals in linear regression. It is based on the following result from matrix algebra.

**Lemma 9.2.** Let  $\mathbb{A}$  be a symmetric positive definite matrix. Then for any column vector  $\mathbf{x}$  of a suitable dimension,

$$\mathbf{x}^T \mathbb{A} \mathbf{x} \leq 1 \quad \text{if and only if} \quad (\mathbf{c}^T \mathbf{x})^2 \leq \mathbf{c}^T \mathbb{A}^{-1} \mathbf{c} \quad \forall \mathbf{c}. \quad \diamond$$

**Proof.** (i) Let  $\mathbf{x}^T \mathbb{A} \mathbf{x} \leq 1$ . Consider the Cauchy-Schwartz inequality

$$(\mathbf{a}^T \mathbf{b})^2 \leq \|\mathbf{a}\|^2 \|\mathbf{b}\|^2.$$

Apply it with  $\mathbf{a} = \mathbb{A}^{1/2} \mathbf{x}$  and  $\mathbf{b} = \mathbb{A}^{-1/2} \mathbf{c}$ . We get

$$(\mathbf{x}^T \mathbf{c})^2 \leq (\mathbf{x}^T \mathbb{A} \mathbf{x}) (\mathbf{c}^T \mathbb{A}^{-1} \mathbf{c}) \leq \mathbf{c}^T \mathbb{A}^{-1} \mathbf{c}$$

and this inequality holds for any vector  $\mathbf{c}$ .

- (i) Let  $(\mathbf{c}^\top \mathbf{x})^2 \leq \mathbf{c}^\top \mathbb{A}^{-1} \mathbf{c} \quad \forall \mathbf{c}$ . Choose  $\mathbf{c} = \mathbb{A} \mathbf{x}$ . We get  $(\mathbf{x}^\top \mathbb{A} \mathbf{x})^2 \leq \mathbf{x}^\top \mathbb{A} \mathbb{A}^{-1} \mathbb{A} \mathbf{x} = \mathbf{x}^\top \mathbb{A} \mathbf{x}$ .  
When  $\mathbf{x} = \mathbf{0}$  then indeed  $0 = \mathbf{x}^\top \mathbb{A} \mathbf{x} \leq 1$ . Otherwise, divide the inequality by  $\mathbf{x}^\top \mathbb{A} \mathbf{x} > 0$ .  
 $\square$

**Theorem 9.3 (Scheffé).** Let  $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \sigma^2 \mathbb{V})$ , where  $\mathbb{V} > 0$  is known. Let  $\hat{\sigma}^2$  be an estimator of  $\sigma^2$  satisfying  $\frac{f \hat{\sigma}^2}{\sigma^2} \sim \chi_f^2$  for some natural number  $f$ . Let  $\hat{\sigma}^2$  be independent of  $\mathbf{X}$ . Take  $\mathcal{B}$  an  $m$ -dimensional linear subspace of  $\mathbb{R}^n$  with  $m \leq n$ . Then

$$\mathbb{P} \left[ |\mathbf{a}^\top \mathbf{X} - \mathbf{a}^\top \boldsymbol{\mu}| \leq \sqrt{m \hat{\sigma}^2 F_{m,f}(1-\alpha) \mathbf{a}^\top \mathbb{V} \mathbf{a}} \quad \forall \mathbf{a} \in \mathcal{B} \right] = 1 - \alpha. \quad \diamond$$

## 9.4. Simultaneous Confidence Bounds

CITE HUBER AND ARNOLD in CHAP 4!!!!



## Notation

Here we list symbols that are consistently used in the same meaning throughout the whole text (perhaps with a few exceptions). Symbols that are introduced and used locally (e.g., in one section) are usually not listed here.

$\boldsymbol{\varepsilon}$	column vector of error terms
$\mathbb{H}$	hat matrix
$h_{ii}$	the $i$ -th diagonal element of the hat matrix $\mathbb{H}$
$\mathbf{1}_n$	column vector of ones of length $n$
$\mathbb{J}_n$	$\mathbf{1}_n \mathbf{1}_n^T$ , $n \times n$ matrix of ones
$\mathcal{M}(\mathbb{X})$	subspace generated by the columns of $\mathbb{X}$
$\mathcal{M}(\mathbb{X})^\perp$	subspace orthogonal to the columns of $\mathbb{X}$
$R^2$	coefficient of determination
$SS_e(\boldsymbol{\beta})$	sum of squares taken as a function of $\boldsymbol{\beta}$
$SS_e$	residual sum of squares (minimized over $\boldsymbol{\beta}$ )
$SS_R$	regression sum of squares (centered)
$SS_T$	total sum of squares (centered)
$\mathbf{u}$	column vector of residuals
$u_i^*$	the $i$ -th standardized residual
$\mathbb{X}$	regression matrix containing covariate vectors in rows
$\mathbf{Y}$	column vector of responses
$\hat{\mathbf{Y}}$	column vector of fitted values

## List of Figures

1.1. Scatterplot of two continuous variables in $\mathbb{R}^2$ . . . . .	5
1.2. Scatterplot of two continuous variables in $\mathbb{R}^2$ with fitted line. . . . .	6
1.3. Zoomed subset of data from Figure 1.2 with visualized vertical distances of the points from the line (blue). . . . .	7
1.4. Galton height data: original pen/paper records. . . . .	9
1.5. Galton height data: original visualization by the author. . . . .	10
1.6. Modified Galton data with fitted least squares line. . . . .	10
2.1. Two sample problem expressed as a linear regression model $EY = \beta_1 + \beta_2 Z$ , where $Z = \mathbb{1}(G)$ . The regression line has no interpretation except at $Z = 0$ or $Z = 1$ . . . . .	14
2.2. Data following a quadratic association with a fitted quadratic curve. . . . .	15
5.1. Incorrect extrapolation of a logistic pipe curve beyond the range of data. . . . .	53
5.2. Incorrect extrapolation of a sinus curve beyond the range of data. . . . .	53
5.3. Extrapolation beyond the scope of data: two covariates. . . . .	54
6.1. Scatterplots of standardized residuals against observation order. . . . .	58
6.2. Scatterplots of standardized residuals against fitted values. . . . .	59
6.3. Scatterplots of standardized residuals against a covariate. . . . .	60
6.4. Scatterplots of standardized residuals against a covariate (unequal variance). . . . .	61
6.5. Scatterplots of square root of absolute standardized residuals against a covariate. . . . .	62
6.6. Boxplots of standardized residuals by a factor covariate. . . . .	63
6.7. Boxplots of standardized residuals by a factorized continuous covariate. . . . .	65
6.8. Histograms of standardized residuals. . . . .	67
6.9. Q-Q plots of standardized residuals. . . . .	68
6.10. Scatterplots of partial residuals against a covariate. . . . .	69
8.1. Recurrent procedure for generation of B-spline bases. . . . .	82
8.2. B-spline bases of degree 1 over the interval $\langle 0, 10 \rangle$ . . . . .	84
8.3. B-spline bases of degree 2 over the interval $\langle 0, 10 \rangle$ . . . . .	84
8.4. B-spline bases of degree 3 over the interval $\langle 0, 10 \rangle$ . . . . .	85
8.5. Fitting a logistic pipe function by polynomials and splines . . . . .	86
8.6. Fitting a sine-constant function by polynomials and splines . . . . .	87

## List of Tables

4.1. Analysis of variance table for the overall regression test. . . . .	42
--	----

## Bibliography

- Aitken, A. C. (1936). IV.—on least squares and linear combination of observations, *Proceedings of the Royal Society of Edinburgh* **55**: 42–48.
- Arnold, S. F. (1980). Asymptotic validity of  $f$  tests for the ordinary linear model and the multiple correlation model, *Journal of the American Statistical Association* **75**(372): 890–894.
- Bonferroni, C. E. (1936). Teoria statistica delle classi e calcolo delle probabilita, Pubbl. d. R. Ist. Super. di Sci. Econom. e Commerciali di Firenze. Firenze: Libr. Internaz. Seeber.
- Galton, F. (1886). Regression towards mediocrity in hereditary stature, *The Journal of the Anthropological Institute of Great Britain and Ireland* **15**: 246–263.
- Gauss, C. F. (1821). *Theoria combinationis observationum erroribus minimis obnoxiae*, Heinrich Dieterich, Göttingen.
- Hayter, A. J. (1984). A proof of the conjecture that the Tukey-Kramer multiple comparisons procedure is conservative, *The Annals of Statistics* **12**(1): 61–75.
- Huber, P. J. (1973). Robust regression: Asymptotics, conjectures and monte carlo, *The Annals of Statistics* **1**(5): 799–821.
- Legendre, A.-M. (1805). *Nouvelles méthodes pour la détermination des orbites des comètes*, F. Didot, Paris.
- Markov, A. A. (1912). *Wahrscheinlichkeitsrechnung*, 2nd edn, Leipzig and Berlin.
- Scheffé, H. (1953). A method for judging all contrasts in the analysis of variance, *Biometrika* **40**(1/2): 87–104.
- Tukey, J. W. (1949). Comparing individual means in the analysis of variance, *Biometrics* **5**(2): 99–114.

# Index

- best linear unbiased estimator (BLUE), [33](#),  
[52](#), [54](#)
- Bonferroni inequality, [90](#)
- Bonferroni method
  - for confidence intervals, [90](#)
  - for tests, [90](#)
- categorical covariate (factor), [74](#)
  - ordered, [77](#)
- coefficient of determination, [27](#), [42](#)
  - adjusted, [28](#)
- equivalent regression models, [22](#)
- factor (categorical covariate), [74](#)
  - ordered, [77](#)
- fitted values, [20](#), [32](#), [52](#)
- hat matrix, [20](#)
- intercept, [16](#)
- least squares estimator, [17](#), [19](#), [31](#)
  - with linear restrictions, [28](#)
- least squares method, [5](#)
- linear regression model, [12](#)
- multiple correlation coefficient, [27](#)
- normal equations, [7](#), [19](#), [45](#)
- one-way ANOVA model, [42](#)
- overall regression test, [42](#)
- projection matrix, [21](#)
- regression matrix, [13](#)
- regression sum of squares
  - centered, [26](#)
  - non-centered, [26](#)
- residual sum of squares, [22](#), [25](#), [26](#), [30](#), [32](#)
- residual variance, [13](#)
  - estimator, [32](#)
- residuals, [20](#), [32](#), [57](#)
  - partial, [65](#)
  - standardized, [57](#)
- submodel, [40](#)
- submodel testing, [41](#)
- sum of squares, [7](#)
- total sum of squares
  - centered, [26](#)
  - non-centered, [25](#)
- Tukey distribution, [93](#)
- Tukey method, [93](#)
- Tukey's honest significant differences, [93](#)

## A. Appendix

The Appendix presents some useful results that are used in this course.

**Lemma A.1.** Let  $X$  be any random vector of dimension  $n$  with mean  $\mu$  and finite variance matrix  $\Sigma$ . Let  $\mathbb{A}$  be any  $n \times n$  matrix. Then

$$E X^T \mathbb{A} X = \mu^T \mathbb{A} \mu + \text{tr}(\mathbb{A} \Sigma). \quad \diamond$$

**Proof.**

$$\begin{aligned} E X^T \mathbb{A} X &= E (X - \mu + \mu)^T \mathbb{A} (X - \mu + \mu) \\ &= E \text{tr}[(X - \mu)^T \mathbb{A} (X - \mu)] + E (X - \mu)^T \mathbb{A} \mu + E \mu^T \mathbb{A} (X - \mu) + E \mu^T \mathbb{A} \mu \\ &= \text{tr}[E (X - \mu)(X - \mu)^T \mathbb{A}] + 0 + 0 + \mu^T \mathbb{A} \mu \\ &= \text{tr}[(\text{var} X) \mathbb{A}] + \mu^T \mathbb{A} \mu = \mu^T \mathbb{A} \mu + \text{tr}(\mathbb{A} \Sigma). \quad \square \end{aligned}$$

**Lemma A.2.** Let  $X \sim N_n(\mathbf{0}, \Sigma)$ . Let  $\mathbb{A}$  be an  $n \times n$  matrix such that  $\mathbb{A} \Sigma$  is idempotent. Then

$$X^T \mathbb{A} X = \chi_{\text{tr}(\mathbb{A} \Sigma)}^2. \quad \diamond$$

**Lemma A.3.** Let  $X \sim N_n(\mu, \Sigma)$ . Then  $X^T \mathbb{A} X$  and  $\mathbb{B} X$  are independent if and only if

$$\mathbb{B} \Sigma \mathbb{A} = \mathbf{0}. \quad \diamond$$