

|||| Chapter 5

Simple Linear regression

Contents

5	Simple Linear regression	
5.1	Linear regression and least squares	1
5.2	Parameter estimates and estimators	4
5.2.1	Estimators are central	10
5.3	Variance of estimators	11
5.4	Distribution and testing of parameters	19
5.4.1	Confidence and prediction intervals for the line	23
5.5	Matrix formulation of simple linear regression	29
5.6	Correlation	32
5.6.1	Inference on the sample correlation coefficient	32
5.6.2	Correlation and regression	33
5.7	Model validation	35
	Glossaries	40
	Acronyms	41

5.1 Linear regression and least squares

In engineering applications we are often faced with the problem of determining the best model of some outcome given a known input

$$y = f(x), \quad (5-1)$$

hence x is the input and the function f is the model. The task is now to find the best model given the input variables (x) and the outcome (y). The simplest model, besides just a mean value (covered in Chapters 3 and 4), would be a model where f is a linear function of x

$$y = \beta_0 + \beta_1 x. \quad (5-2)$$

When the outcome y is the result of some experiment, the model will not be perfect, and we need to add an error term

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad i = \{1, \dots, n\}, \quad (5-3)$$

where ε_i is called the *error* and is a sequence of independent random variables with expectation equal zero (i.e. the mean $E(\varepsilon_i) = 0$ and some variance ($V(\varepsilon_i) = \sigma^2$). The statistical interpretation of (5-2) is therefore that it expresses the expected value of the outcome

$$E(Y_i) = \beta_0 + \beta_1 x_i, \quad (5-4)$$

also called the *model prediction*.

It is of course a very unusual situation that we actually know the values of β_0 and β_1 and we will have to rely on estimates based on some observations (y_1, \dots, y_n). As usual we express this by putting a “hat” on the parameters

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_i, \quad (5-5)$$

meaning that we expect or predict \hat{y}_i (in mean or average) under the conditions given by x_i .

|||| Example 5.1

A car manufacturer wants to find the relation between speed and fuel consumption, to do so she sets up the following model

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad (5-6)$$

here $E(Y_i)$ is the expected fuel consumption at the speed x_i . Further, there will be uncontrollable variations, e.g. due to differences in weather condition, but also non-linear effects not included in the model might be present. These variations are captured by the ε_i 's. We see that speed is something we control here, and we then observe the outcome (here fuel consumption), at different experimental conditions (speeds).

In this chapter we will deal with estimation and inference of β_0, β_1 , and prediction of Y_i given x_i . At some point we will have realizations (or observations) of the outcome, in this case we write

$$y_i = \beta_0 + \beta_1 x_i + e_i, \quad i = \{1, \dots, n\}. \quad (5-7)$$

Now y_i is a realization and e_i is the deviation between the model prediction and the actual observation: a realization of the error ε_i , it is called a *residual*. Clearly, we want the residuals to be small in some sense, the usual choice (and the one treated in this chapter) is in the Residual Sum of Squares (RSS) sense, i.e. we want to minimize the residual sum of squares

$$RSS(\beta_0, \beta_1) = \sum_{i=1}^n \varepsilon_i^2 = \sum_{i=1}^n (Y_i - (\beta_0 + \beta_1 x_i))^2, \quad (5-8)$$

where we have emphasized that the residual sum of squares is a function of the parameters (β_0, β_1) . The parameter estimates $(\hat{\beta}_0, \hat{\beta}_1)$ are the values of β_0 and β_1 which minimize RSS. Note, that we use Y_i and ε_i rather than the observed values (y_i and e_i), this is to emphasize that the estimators are random variables, in actual calculations after the experiments are carried out we will just replace Y_i with y_i and ε_i with e_i . Figure 5.1 sketches the linear regression problem.




Figure 5.1: Conceptual diagram for the simple linear regression problem.

|||| Remark 5.2 Estimates and estimators

In (5-8) the RSS is a function of the random variables (Y_i) , thus making RSS a random variable. If we replace Y_i with the realizations y_i then RSS is also a realization.

In this chapter the result of optimizing RSS with respect to β_0 and β_1 will be denoted $\hat{\beta}_0$ and $\hat{\beta}_1$. Sometimes $\hat{\beta}_0$ and $\hat{\beta}_1$ will be functions of Y_i and sometimes they will be functions of the realizations y_i , they are referred to as:

1. **Estimators:** before the experiment has been carried out, then $\hat{\beta}_0$ and $\hat{\beta}_1$ are functions of Y_i and they are also random variables, and we call them *estimators*.
2. **Estimates:** after the experiment had been carried out, then $\hat{\beta}_0$ and $\hat{\beta}_1$ are functions of y_i and they are also realizations of random variables, and we call them *estimates*.

|||| **Remark 5.3 Two types of examples**

In this chapter we will use two types of examples, one is labelled “Simulation”, which are simulation studies intended to illustrate the consequences of theorems and general results. While the other type of examples (not labelled “Simulation”), are intended to illustrate the use of the theorems on practical examples.

5.2 Parameter estimates and estimators

When $\hat{\beta}_0$ and $\hat{\beta}_1$ is a result of minimizing the function in Equation (5-8), we refer to the estimators as *least squares estimators*. The least squares estimators are given in the following theorem:

|||| **Theorem 5.4 Least squares estimators**

The least squares estimators of β_0 and β_1 are given by

$$\hat{\beta}_1 = \frac{\sum_{i=1}^n (Y_i - \bar{Y})(x_i - \bar{x})}{S_{xx}}, \quad (5-9)$$

$$\hat{\beta}_0 = \bar{Y} - \hat{\beta}_1 \bar{x}, \quad (5-10)$$

where $S_{xx} = \sum_{i=1}^n (x_i - \bar{x})^2$.

As we can see above the estimators ($\hat{\beta}_1$ and $\hat{\beta}_2$) are functions of random variables (Y_i and \bar{Y}), and thus the estimators are themselves random variables. We can therefore talk about the expectation, variance and distribution of the estimators. In analyses with data we will of course only see realizations of Y_i and we just replace Y_i and \bar{Y} with their realizations y_i and \bar{y} . In this case we speak about *estimates* of β_0 and β_1 .

Before we go on with the proof of Theorem 5.4, the application of the theorem is illustrated in the following example:

||| Example 5.5 Student height and weight

Consider the student height and weight data presented in Chapter 1,

Heights (x_i)	168	161	167	179	184	166	198	187	191	179
Weights (y_i)	65.5	58.3	68.1	85.7	80.5	63.4	102.6	91.4	86.7	78.9

We want to find the best least squares regression line for these points, this is equivalent to calculating the least squares estimates of $\hat{\beta}_0$ and $\hat{\beta}_1$.

We start by finding the two sample means

$$\bar{x} = \frac{1}{10} (168 + 161 + \dots + 179) = 178,$$

$$\bar{y} = \frac{1}{10} (65.5 + 58.3 + \dots + 78.9) = 78.11.$$

The value of S_{xx} is calculated by

$$S_{xx} = (168 - 178)^2 + \dots + (179 - 178)^2 = 1342.$$

We can now calculate $\hat{\beta}_1$ as

$$\hat{\beta}_1 = \frac{1}{1342} ((65.5 - 78.11)(168 - 179) + \dots + (78.9 - 78.11)(179 - 178)) = 1.11,$$

and finally, we can calculate $\hat{\beta}_0$ as

$$\hat{\beta}_0 = 78.11 - 1.11 \cdot 178 = -120.$$

The calculations can be implemented by

```
# Read data
x = np.array([168, 161, 167, 179, 184, 166, 198, 187, 191, 179])
y = np.array([65.5, 58.3, 68.1, 85.7, 80.5, 63.4, 102.6, 91.4, 86.7, 78.9])

# Calculate averages
xbar = x.mean()
ybar = y.mean()

# Parameters estimates
Sxx = sum((x - xbar)**2)
betalhat = sum((x - xbar)*(y - ybar)) / Sxx
beta0hat = ybar - betalhat * xbar
print(round(beta0hat.mean(),2), round(betalhat.mean(),2))

-119.96 1.11
```

Rather than using “manual” calculations, we can use the builtin functions

```
student = pd.DataFrame({'x': x, 'y': y})
fitStudents = smf.ols(formula = 'y ~ x', data=student).fit()
print(fitStudents.summary(slim=True))
```

OLS Regression Results						
Dep. Variable:	y	R-squared:	0.932			
Model:	OLS	Adj. R-squared:	0.924			
No. Observations:	10	F-statistic:	110.3			
Covariance Type:	nonrobust	Prob (F-statistic):	5.87e-06			
	coef	std err	t	P> t	[0.025	0.975]
Intercept	-119.9581	18.897	-6.348	0.000	-163.535	-76.381
x	1.1127	0.106	10.504	0.000	0.868	1.357

As we can see the two calculations give the same results regarding the parameter estimates. We can also see that the high level calculation gives some more information. How to interpret and calculate these numbers will be treated in the following pages.

Before we go on with the analysis of the results, the proof of Theorem 5.4 is presented:

|||| Proof

Of Theorem 5.4: In order to find the minimum of the function RSS we differentiate the residual sum of squares with respect to the parameters

$$\frac{\partial RSS}{\partial \hat{\beta}_0} = -2 \sum_{i=1}^n (y_i - (\hat{\beta}_0 + \hat{\beta}_1 x_i)), \quad (5-11)$$

now equating with zero we get

$$\begin{aligned} 0 &= -2 \sum_{i=1}^n (y_i - (\hat{\beta}_0 + \hat{\beta}_1 x_i)) \\ &= -2n\bar{y} + 2n\hat{\beta}_0 + 2n\hat{\beta}_1\bar{x}, \end{aligned} \quad (5-12)$$

solving for $\hat{\beta}_0$ gives

$$\hat{\beta}_0 = \bar{y} - \hat{\beta}_1\bar{x}, \quad (5-13)$$

and by similar calculations we get

$$\begin{aligned}
 \frac{\partial RSS}{\partial \hat{\beta}_1} &= \frac{\partial}{\partial \hat{\beta}_1} \left(\sum_{i=1}^n (y_i - (\bar{y} - \hat{\beta}_1 \bar{x} + \hat{\beta}_1 x_i))^2 \right) \\
 &= \frac{\partial}{\partial \hat{\beta}_1} \left(\sum_{i=1}^n ((y_i - \bar{y}) - \hat{\beta}_1 (x_i - \bar{x}))^2 \right) \\
 &= -2 \sum_{i=1}^n ((y_i - \bar{y}) - \hat{\beta}_1 (x_i - \bar{x}))(x_i - \bar{x}) \\
 &= -2 \left[\sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x}) - \hat{\beta}_1 \sum_{i=1}^n (x_i - \bar{x})^2 \right],
 \end{aligned} \tag{5-14}$$

equating with zero and solving for $\hat{\beta}_1$ gives

$$\begin{aligned}
 \hat{\beta}_1 &= \frac{\sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x})}{\sum_{i=1}^n (x_i - \bar{x})^2} \\
 &= \frac{\sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x})}{S_{xx}}.
 \end{aligned} \tag{5-15}$$

The estimates $\hat{\beta}_0$ and $\hat{\beta}_1$ are called **least squares estimates**, because they minimize the sum of squared residuals (i.e. RSS). Replacing y_i with Y_i give the estimators in the theorem.

■

When we have obtained parameter estimates in the linear regression model above, we would like to make quantitative statements about the uncertainty of the parameters, and in order to design tests we will also need the probability distribution of the parameter estimators. The usual assumption is that the errors are normal random variables

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad \text{where } \varepsilon_i \sim N(0, \sigma^2), \tag{5-16}$$

or in other words the errors are independent identically distributed (i.i.d.) normal random variables with zero mean and variance σ^2 . When random variables are involved we know that repeating the experiment will result in different values of the response (Y_i), and therefore in different values of the parameter estimates. To illustrate this we can make simulation experiments to analyse the behaviour of the parameter estimates. Recall that the role of simulation examples are to illustrate probabilistic behaviour of e.g. estimators, not how actual data is analysed.

|||| Remark 5.6 How to write a statistical model

In Remark 3.2 it was explained how to write the model behind the t -tests, i.e.

$$X_i \sim N(\mu, \sigma^2) \text{ and i.i.d.} \quad (5-17)$$

Remember, that i.i.d. is short for independently and identically distributed, which essentially means that the observations are selected randomly from population, see the text after Example 1.2.

Using this notation the linear regression model could be written

$$Y_i \sim N(\beta_0 + \beta_1 x_i, \sigma^2) \text{ and independent,} \quad (5-18)$$

however we will write models as above in Equation (5-16).
Note, if $\beta_1 = 0$ the model is

$$Y_i = \beta_0 + \varepsilon_i, \quad \text{where } \varepsilon_i \sim N(0, \sigma^2) \text{ and i.i.d.,} \quad (5-19)$$

which is exactly the model above in Equation (5-17), and the estimate of the mean of the population, from which the sample (i.e. (y_1, \dots, y_n)) was taken, is then

$$\hat{\mu} = \hat{\beta}_0. \quad (5-20)$$

|||| Example 5.7 Simulation of parameter estimation

Consider the linear model

$$Y_i = 10 + 3x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, 5^2). \quad (5-21)$$

We can make repetitions of this experiment by


```
np.random.seed(124)
n = 10
k = 500
beta0 = 10
beta1 = 3
sigma = 5
x = np.linspace(-2, 5, num=n)
y0 = np.zeros((k,n))
y = y0 + beta0 + beta1*x + np.random.normal(0, sigma, size=(k,n))
```

The variable `y` now contains `n` rows and `k` columns, representing `k` experiments, for each of the experiment we can calculate the parameter estimates:

```
b0 = np.zeros(k)
b1 = np.zeros(k)
for i in range(k):
    D = pd.DataFrame({'x': x, 'y': y[i,:]})
    result = smf.ols(formula = 'y ~ x', data=D).fit()
    b0[i] = result.params["Intercept"]
    b1[i] = result.params["x"]
print(round(b0.mean(),2), round(b1.mean(),2))

10.09 2.97
```

As we can see the average of the parameter estimates (`b0.mean` and `b1.mean`) are very close to the true parameter values ($\beta_0 = 10$ and $\beta_1 = 3$). We can of course also look at the empirical density (the normalized histogram, see Section 1.6.1) of the parameter estimates:



The estimates seem to be rather symmetrically distributed around the true parameter values. It is also clear that there is some variation in the estimates: the estimates of β_0 range from about 4 to about 16 and the estimates of β_1 range from about 1 to 5.

Try changing the code (see the accompanying chapter script):



What happens to the mean value of the estimates if you change the number of data points (n)?



What happens to the empirical density and the scatter plot of the parameter estimates if you change:

- The number of data points (n)?
- The range of x -values?
- The residual variance (σ^2)?
- The values of β_0 and β_1 ?

In the example above we saw that the average of the parameter estimates were very close to the true values, this is of course a nice property of an estimator. When this is the case in general, i.e. when $E[\hat{\beta}_i] = \beta_i$ we say that the estimator is central or unbiased. The estimators $\hat{\beta}_0$ and $\hat{\beta}_1$ are in fact central, and we show this in Section 5.2.1 below.

In order to test hypothesis about β_0 and β_1 we will also need to give exact statements about the distribution of the parameters. We saw in Example 5.7 above that the distributions seem to be symmetric around the true values, but we will need more precise statements about the distributions and their variances. This important part will be dealt with in the Sections 5.3 and 5.4.

5.2.1 Estimators are central

In the linear regression model we assume that the observed values of Y_i can be split into two parts: the prediction (the part explained by the regression line $(\beta_0 + \beta_1 x_i)$) and the error (a random part (ε_i)). As usual we view our estimators as functions of random variables (the ε_i 's), so it makes sense to calculate the expectation of the estimators. The assumption $E(\varepsilon_i) = 0$ is central for the presented arguments, and will be used repeatedly.

In order to find the expectation of the parameter estimators we rewrite our estimators as functions of the true parameters (β_0 and β_1)

$$\hat{\beta}_1 = \frac{\sum_{i=1}^n (Y_i - \bar{Y})(x_i - \bar{x})}{S_{xx}}, \quad (5-22)$$

inserting $Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i$ and $\bar{Y} = \frac{1}{n} \sum_{i=1}^n (\beta_0 + \beta_1 x_i + \varepsilon_i) = \beta_0 + \beta_1 \bar{x} + \bar{\varepsilon}$ gives

$$\hat{\beta}_1 = \frac{\sum_{i=1}^n [(\beta_0 + \beta_1 x_i + \varepsilon_i - (\beta_0 + \beta_1 \bar{x} + \bar{\varepsilon})) (x_i - \bar{x})]}{S_{xx}}, \quad (5-23)$$

now the sum is divided into a part which depends on ε_i (the random part) and a part which is independent of ε_i

$$\begin{aligned}\hat{\beta}_1 &= \frac{\sum_{i=1}^n \beta_1 (x_i - \bar{x})^2}{S_{xx}} + \frac{\sum_{i=1}^n (\varepsilon_i - \bar{\varepsilon})(x_i - \bar{x})}{S_{xx}} \\ &= \beta_1 + \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}} - \frac{\bar{\varepsilon} \sum_{i=1}^n (x_i - \bar{x})}{S_{xx}},\end{aligned}\tag{5-24}$$

now observe that $\sum_{i=1}^n (x_i - \bar{x}) = 0$ to get

$$\hat{\beta}_1 = \beta_1 + \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}},\tag{5-25}$$

for $\hat{\beta}_0$ we get

$$\begin{aligned}\hat{\beta}_0 &= \bar{y} - \hat{\beta}_1 \bar{x} \\ &= \frac{1}{n} \sum_{i=1}^n (\beta_0 + \beta_1 x_i + \varepsilon_i) - \left(\beta_1 + \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}} \right) \bar{x} \\ &= \beta_0 + \beta_1 \bar{x} + \frac{1}{n} \sum_{i=1}^n \varepsilon_i - \left(\beta_1 + \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}} \right) \bar{x} \\ &= \beta_0 + \frac{1}{n} \sum_{i=1}^n \varepsilon_i - \left(\frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}} \right) \bar{x}.\end{aligned}\tag{5-26}$$

Since expectation is a linear operation (see Chapter 2) and the expectation of ε_i is zero we find that $E[\hat{\beta}_0] = \beta_0$ and $E[\hat{\beta}_1] = \beta_1$, and we say that $\hat{\beta}_0, \hat{\beta}_1$ are central estimators.

5.3 Variance of estimators

In order for us to be able to construct confidence intervals for parameter estimates, talk about uncertainty of predictions and test hypothesis, we will need the variance of the parameter estimates as well as an estimator of the error variance (σ^2).

Parameter variance and covariance of estimators are given in the following theorem:

|||| Theorem 5.8 Variance of estimators

The variance and covariance of the estimators in Theorem 5.4 are given by

$$V[\hat{\beta}_0] = \frac{\sigma^2}{n} + \frac{\bar{x}^2 \sigma^2}{S_{xx}}, \quad (5-27)$$

$$V[\hat{\beta}_1] = \frac{\sigma^2}{S_{xx}}, \quad (5-28)$$

$$\text{Cov}[\hat{\beta}_0, \hat{\beta}_1] = -\frac{\bar{x} \sigma^2}{S_{xx}}, \quad (5-29)$$

where σ^2 is usually replaced by its estimate ($\hat{\sigma}^2$). The central estimator for σ^2 is

$$\hat{\sigma}^2 = \frac{RSS(\hat{\beta}_0, \hat{\beta}_1)}{n - 2}. \quad (5-30)$$

When the estimate of σ^2 is used the variances also become estimates and we'll refer to them as $\hat{\sigma}_{\hat{\beta}_0}^2$ and $\hat{\sigma}_{\hat{\beta}_1}^2$.

The variance of $\hat{\beta}_1$ is a function of the true error variance (σ^2) and S_{xx} . For most (all reasonable) choices of the regressors (x), S_{xx} will be an increasing function of n , and the variance of $\hat{\beta}_1$ will therefore decrease as n increases. This expresses that we will be more certain about the estimates as we increase the number of points in our sample. The same is true for the variance of $\hat{\beta}_0$, and the covariance between $\hat{\beta}_1$ and $\hat{\beta}_0$. The error variance estimate ($\hat{\sigma}^2$) is the residual sum of squares divided by $n - 2$, the intuitive explanation for the $n - 2$ (rather than n or $n - 1$) is that if we only have two ($n = 2$) pairs (x_i, y_i) , it will not be possible to say anything about the variation (the residuals will be zero). Or another phrasing is that; we have used 2 degrees of freedom to estimate $\hat{\beta}_0$ and $\hat{\beta}_1$.

Before we turn to the proof of Theorem 5.8, we will take a look at a couple of examples.

|||| Example 5.9 (Example 5.5 cont.)

In Example 5.5 we found the parameter estimates

$$\hat{\beta}_0 = -120, \quad \hat{\beta}_1 = 1.11,$$

we can now find predicted values of the dependent variable by

$$\hat{y}_i = -120 + 1.11 \cdot x_i,$$

and the values of the residuals

$$e_i = y_i - \hat{y}_i,$$

and finally the error variance estimate is

$$\hat{\sigma}^2 = \frac{1}{10-2} \sum_{i=1}^{10} e_i^2.$$

We can implement the results by:

```
beta0 = fitStudents.params["Intercept"]
beta1 = fitStudents.params["x"]
e = student["y"] - (beta0 + beta1 * student["x"])
n = len(e)
sigma = math.sqrt(np.sum(e**2) / (n - 2))
sigma_beta0 = math.sqrt(sigma**2 * (1 / n + xbar**2 / Sxx))
sigma_beta1 = math.sqrt(sigma**2 / Sxx)
print(round(sigma,2), round(sigma_beta0,2), round(sigma_beta1,2))

3.88 18.9 0.11
```

As usual we use standard deviations rather than variances, this also means that we can compare with the results from `smf.ols` (see Example 5.5). Again we can find our estimates in the Python-output, the parameter standard deviations are given in the second column of the coefficient matrix. The estimated standard deviation is not reported in the summary from `smf.ols`, but can be extracted by

```
round(np.sqrt(fitStudents.scale),2)

np.float64(3.88)
```

The simulation example (Example 5.7) can also be extended to check the equations of Theorem 5.8:

|||| Example 5.10 Simulation continued

In Example 5.7 we looked at simulation from the model

$$Y_i = 10 + 3x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, 5^2)$$

In order to calculate the variance estimates we need to calculate \bar{x} and S_{xx} :

```

np.random.seed(1241)
Sxx = (n-1)*x.var()
print(x.mean(), round(Sxx,2))

1.5 44.92

k = 500
y0 = np.zeros((k,n))
y = y0 + beta0 + beta1*x + np.random.normal(0, 5, size=(k,n))

```

and we would expect to obtain the variance estimates close to

$$V[\hat{\beta}_0] = 5^2 \left(\frac{1}{10} + \frac{1.50^2}{49.91} \right) = 3.63$$

$$V[\hat{\beta}_1] = \frac{5^2}{49.91} = 0.501.$$

With simulations we find:

```

b0 = np.zeros(k)
b1 = np.zeros(k)
sigma = np.zeros(k)
for i in range(k):
    D = pd.DataFrame({'x': x, 'y': y[i,:]}
    result = smf.ols(formula = 'y ~ x', data=D).fit()
    b0[i] = result.params["Intercept"]
    b1[i] = result.params["x"]
    sigma[i] = math.sqrt(result.scale)

print(round(b0.var(),2), round(b1.var(),2), round(sigma.mean(),2))

3.58 0.46 4.81

```

We can see that the simulated values are close to the theoretical values. You are invited to play around with different settings for the simulation, in particular increasing k will increase the accuracy of the estimates of the variances.

The example above shows how Theorem 5.8 can be illustrated by simulation, a formal proof is given by:

||| **Proof**

Of Theorem 5.8. Using (5-26) we can write the variance of $\hat{\beta}_0$ as

$$V(\hat{\beta}_0) = V\left[\beta_0 + \frac{1}{n} \sum_{i=1}^n \varepsilon_i - \left(\frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}}\right) \bar{x}\right], \quad (5-31)$$

using the definition of the variance ($V(X) = E[(X - E[X])^2]$) and $E(\varepsilon) = 0$ we get

$$\begin{aligned} V(\hat{\beta}_0) &= V\left[\frac{1}{n} \sum_{i=1}^n \varepsilon_i\right] + V\left[\left(\frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}}\right) \bar{x}\right] - \\ &\quad 2E\left[\frac{1}{n} \sum_{i=1}^n \varepsilon_i \left(\frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}}\right) \bar{x}\right], \end{aligned} \quad (5-32)$$

now use independence between ε_i and ε_j ($i \neq j$) to get

$$\begin{aligned} V(\hat{\beta}_0) &= \frac{\sigma^2}{n} + \frac{\bar{x}^2 \sigma^2 \sum_{i=1}^n (x_i - \bar{x})^2}{(S_{xx})^2} + \frac{\bar{x} \sigma^2}{n S_{xx}} \sum_{i=1}^n (x_i - \bar{x}) \\ &= \frac{\sigma^2}{n} + \frac{\bar{x}^2 \sigma^2}{S_{xx}}. \end{aligned} \quad (5-33)$$

Finally, the variance of $\hat{\beta}_1$ is (again using the definition of variance and independence of the ε 's)

$$\begin{aligned} V(\hat{\beta}_1) &= V\left[\beta_1 + \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}}\right] \\ &= \frac{\sum_{i=1}^n (x_i - \bar{x})^2 V(\varepsilon_i)}{(S_{xx})^2} \\ &= \frac{\sigma^2}{S_{xx}}, \end{aligned} \quad (5-34)$$

and the covariance between the parameters estimates becomes

$$\begin{aligned} \text{Cov}(\hat{\beta}_0, \hat{\beta}_1) &= E[(\hat{\beta}_0 - \beta_0)(\hat{\beta}_1 - \beta_1)] \\ &= E\left[\left(\frac{1}{n} \sum_{i=1}^n \varepsilon_i - \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}}\right) \bar{x} \frac{\sum_{i=1}^n \varepsilon_i (x_i - \bar{x})}{S_{xx}}\right] \\ &= \frac{\bar{x}}{n S_{xx}} E\left[\sum_{i=1}^n \varepsilon_i \sum_{j=1}^n \varepsilon_j (x_j - \bar{x})\right] - \frac{\bar{x}}{(S_{xx})^2} E\left[\sum_{i=1}^n \varepsilon_i^2 (x_i - \bar{x})^2\right] \\ &= \frac{\bar{x} \sigma^2 (n \bar{x} - n \bar{x})}{n S_{xx}} - \frac{\bar{x}}{(S_{xx})^2} \sigma^2 \sum_{i=1}^n (x_i - \bar{x})^2 \\ &= -\frac{\bar{x} \sigma^2}{S_{xx}}. \end{aligned} \quad (5-35)$$

To get an estimate of the residual variance we calculate the expected value of the residual sum of squares

$$E(RSS) = E \left[\sum_{i=1}^n (Y_i - (\hat{\beta}_0 + \hat{\beta}_1 x_i))^2 \right], \quad (5-36)$$

inserting $Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i$ and rearranging gives

$$\begin{aligned} E(RSS) &= \sum_{i=1}^n E[(-(\hat{\beta}_0 - \beta_0) - (\hat{\beta}_1 - \beta_1)x_i + \varepsilon_i)^2] \\ &= \sum_{i=1}^n \{E[(\hat{\beta}_0 - \beta_0)^2] + E[(\hat{\beta}_1 - \beta_1)^2] x_i^2 + E[\varepsilon_i^2] + \\ &\quad 2E[(\hat{\beta}_0 - \beta_0)(\hat{\beta}_1 - \beta_1)] x_i - 2E[(\hat{\beta}_0 - \beta_0)\varepsilon_i] - 2E[(\hat{\beta}_1 - \beta_1)\varepsilon_i] x_i\}, \end{aligned} \quad (5-37)$$

now observe that $E[(\hat{\beta}_0 - \beta_0)^2] = V[\hat{\beta}_0]$, $E[(\hat{\beta}_1 - \beta_1)^2] = V[\hat{\beta}_1]$, $E[\varepsilon_i^2] = \sigma^2$, and $E[(\hat{\beta}_0 - \beta_0)(\hat{\beta}_1 - \beta_1)] = \text{Cov}(\hat{\beta}_0, \hat{\beta}_1)$, and insert $\hat{\beta}_0$ and $\hat{\beta}_1$ in the last two terms

$$\begin{aligned} E(RSS) &= n V(\hat{\beta}_0) + V(\hat{\beta}_1) \sum_{i=1}^n x_i^2 + n\sigma^2 + 2 \sum_{i=1}^n \text{Cov}(\hat{\beta}_0, \hat{\beta}_1) x_i - \\ &\quad 2 \sum_{i=1}^n \left\{ E \left[\left(\frac{1}{n} \sum_{j=1}^n \varepsilon_j - \frac{\sum_{j=1}^n \varepsilon_j (x_j - \bar{x})}{S_{xx}} \right) \varepsilon_i \right] - E \left[\frac{\sum_{j=1}^n \varepsilon_j (x_j - \bar{x})}{S_{xx}} \varepsilon_i \right] x_i \right\} \\ &= \sigma^2 + \frac{n\bar{x}^2\sigma^2}{S_{xx}} + \frac{\sigma^2 \sum_{i=1}^n x_i^2}{S_{xx}} + n\sigma^2 - 2 \sum_{i=1}^n \frac{\bar{x}\sigma^2}{S_{xx}} x_i - \\ &\quad 2 \sum_{i=1}^n \left(\frac{\sigma^2}{n} - \frac{\sigma^2(x_i - \bar{x})}{S_{xx}} \right) - 2 \sum_{i=1}^n \frac{\sigma^2(x_i - \bar{x})x_i}{S_{xx}}, \end{aligned} \quad (5-38)$$

now collect terms and observe that $\sum x_i = n\bar{x}$

$$\begin{aligned} E(RSS) &= \sigma^2(n+1) + \frac{\sigma^2}{S_{xx}} \sum_{i=1}^n (x_i^2 + \bar{x}^2) - 2 \frac{n\bar{x}^2\sigma^2}{S_{xx}} - 2\sigma^2 - 2 \frac{\sigma^2 \sum_{i=1}^n (x_i^2 - x_i\bar{x})}{S_{xx}} \\ &= \sigma^2(n-1) + \frac{\sigma^2}{S_{xx}} \sum_{i=1}^n (-x_i^2 - \bar{x}^2 + 2x_i\bar{x}) \\ &= \sigma^2(n-1) - \frac{\sigma^2}{S_{xx}} S_{xx} \\ &= \sigma^2(n-2), \end{aligned} \quad (5-39)$$

and thus a central estimator for σ^2 is $\hat{\sigma}^2 = \frac{RSS}{n-2}$.

■

Before we continue with parameter distributions and hypothesis testing, the next example illustrates the behaviour of the parameter variance estimates:

||| Example 5.11 Simulation continued

Consider the following model

$$Y_i = 1 + x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, 1), \quad (5-40)$$

also assume that $x_i = \frac{i-1}{n-1}$, $i = 1, \dots, n$ where n is the number of pairs (x_i, y_i) . We want to make a simulation experiment for increasing number of pairs, and extract the parameter variance, parameter covariance and residual variance estimates. In order to do so we need to extract these numbers from a linear model in Python. This can be done by:

```
np.random.seed(134)
x = np.linspace(0, 1, num=10)
y = 1 + x + np.random.normal(0, 1, size=10)
# Fit the model (estimate parameter)
fit = smf.ols(formula = 'y ~ x', data=pd.DataFrame({'x': x, 'y': y})).fit()
# Print summary of model fit
print(fit.summary(slim=True))
```

OLS Regression Results


Dep. Variable:	y	R-squared:	0.003			
Model:	OLS	Adj. R-squared:	-0.122			
No. Observations:	10	F-statistic:	0.02230			
Covariance Type:	nonrobust	Prob (F-statistic):	0.885			
	coef	std err	t	P> t	[0.025	0.975]
Intercept	1.8541	0.445	4.162	0.003	0.827	2.881
x	0.1122	0.751	0.149	0.885	-1.620	1.844

```
# Residual standard deviation
sigma = math.sqrt(fit.scale)
# Estimated standard deviation of parameters
print(round(fit.bse,2))
```

```
Intercept    0.45
x            0.75
dtype: float64
```

Now let's return to the simulation example, the number of independent variables (x) is increased and we draw the residual from the standard normal distribution, in this particular case we can find S_{xx} as a function of n , and compare the expected values (fix $\sigma^2 = 1$) with the simulation results

```
np.random.seed(134)
k = 500
sigma_beta = np.zeros(shape=(k,2))
sigma = np.zeros(k)
n = np.arange(3, k+3)
for i in range(k):
    x = np.linspace(0, 1, num=n[i])
    y = 1 + x + np.random.normal(0, 1, size=n[i])
    D = pd.DataFrame({'x': x, 'y': y})
    fit = smf.ols(formula = 'y ~ x', data=D).fit()
    sigma_beta[i,:] = fit.bse
    sigma[i] = math.sqrt(fit.scale)
```



We see that the residual variance converge to the true value with smaller and smaller variation, while the parameter variances converge to zero. In a plot like this we can therefore see the gain from obtaining more observations of the model.

Again you are encouraged to change some of the specifications of the simulation set up and see what happens.

5.4 Distribution and testing of parameters

The regression model is given by

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad (5-41)$$

where the estimators of the parameters and their variances are given by Theorems 5.4 and 5.8. Since the estimators are linear functions of normal random variables (ε_i) they will also be normal random variables. To give the full stochastic model we need to use the estimate of the residual variance, and take the uncertainty of this estimator into account when constructing tests.

As we already saw in Example 5.7 we cannot expect to get the true value of the parameter, but there will be some deviations from the true value due to the stochastic nature of the model/real world application. The purpose of this section is to give the precise description of the parameter distributions. We aim at testing hypothesis of the type

$$H_{0,i}: \beta_i = \beta_{0,i}, \quad (5-42)$$

against some alternatives. The general remarks on hypothesis testing from Chapter 3 still apply, but we will go through the specific construction for linear regression here.

The central estimator of σ^2 (Equation (5-30)) is χ^2 -distributed with $n - 2$ degrees of freedom. In order to test the hypothesis in Equation (5-42) we need the normalized distance to a null hypothesis (i.e the distance from the observed estimate $\hat{\beta}_{0,i}$ to the value under the null hypothesis $\beta_{0,i}$). From Theorem 5.8 the standard deviations of the parameter estimates are found to

$$\hat{\sigma}_{\beta_0} = \sqrt{\frac{\hat{\sigma}^2}{n} + \frac{\bar{x}^2 \hat{\sigma}^2}{S_{xx}}} = \hat{\sigma} \sqrt{\frac{1}{n} + \frac{\bar{x}^2}{\sum_{i=1}^n (x_i - \bar{x})^2}}, \quad (5-43)$$

$$\hat{\sigma}_{\beta_1} = \sqrt{\frac{\hat{\sigma}^2}{S_{xx}}} = \hat{\sigma} \sqrt{\frac{1}{\sum_{i=1}^n (x_i - \bar{x})^2}}, \quad (5-44)$$

under the null hypothesis the normalized (with standard deviations) distance between the estimators and the true values are both t -distributed with $n - 2$ degrees of freedom, and hypothesis testing and confidence intervals are based on this t -distribution:

|||| Theorem 5.12 Test statistics

Under the null hypothesis ($\beta_0 = \beta_{0,0}$ and $\beta_1 = \beta_{0,1}$) the statistics

$$T_{\beta_0} = \frac{\hat{\beta}_0 - \beta_{0,0}}{\hat{\sigma}_{\beta_0}}, \quad (5-45)$$

$$T_{\beta_1} = \frac{\hat{\beta}_1 - \beta_{0,1}}{\hat{\sigma}_{\beta_1}}, \quad (5-46)$$

are t -distributed with $n - 2$ degrees of freedom, and inference should be based on this distribution.

|||| Proof

The proof is omitted, but rely on the fact that $\hat{\beta}_j$ is normally distributed, $\hat{\sigma}_{\beta_j}^2$ is related to the χ^2 -distribution, and (provided independence) a standard normal random variable divided by the square root of a χ^2 distributed random variable is t -distributed. ■

In this material we only test two-sided hypothesis. The hypothesis can be concluded using p -values or critical values, in the same way as we saw for hypothesis regarding mean values in Chapter 3 Section 3.1.7.

|||| Example 5.13 Example 5.9 cont.

We continue with the data from Examples 5.5 and 5.9, where we found the parameter estimates and the variance estimates. We want to test the hypotheses

$$H_{00} : \beta_0 = 0 \quad \text{vs.} \quad H_{10} : \beta_0 \neq 0, \quad (5-47)$$

$$H_{01} : \beta_1 = 1 \quad \text{vs.} \quad H_{11} : \beta_1 \neq 1, \quad (5-48)$$

on confidence level $\alpha = 0.05$. With reference to Examples 5.5 and 5.9, and Theorem 5.12, we can calculate the t -statistics as

$$t_{\text{obs},\beta_0} = \frac{-119.96}{18.897} = -6.35, \quad (5-49)$$

$$t_{\text{obs},\beta_1} = \frac{1.113 - 1}{0.1059} = 1.07. \quad (5-50)$$

H_{00} is rejected if $|t_{\text{obs},\beta_0}| > t_{1-\alpha/2}$, and H_{01} is rejected if $|t_{\text{obs},\beta_1}| > t_{1-\alpha/2}$, as usual we can find the critical values by:

```
round(stats.t.ppf(0.975, df=10-2), 2)

np.float64(2.31)
```

and we see that with significance level $\alpha = 0.05$, then H_{00} is rejected and H_{01} isn't. If we prefer p -values rather than critical values, these can be calculated by:

```
pv0 = 2 * (1 - stats.t.cdf(6.35, df=10-2))
pv1 = 2 * (1 - stats.t.cdf(1.07, df=10-2))
print(round(pv0,5), round(pv1,2))

0.00022 0.32
```

The p -value for the intercept (β_0) is less than 0.05, while the p -value for β_1 is greater than 0.05, hence we conclude that $\beta_0 \neq 0$, but we cannot reject that $\beta_1 = 1$. The summary of linear model buildin function, also give t -statistics and p -values (see Example 5.5). The test statistic and the p -value for H_{01} is different from the one we obtained above. The reason for this is that `smf.ols` tests the default hypothesis $H_{0i} : \beta_i = 0$ against the alternative $H_{1i} : \beta_i \neq 0$. Even though this choice is reasonable in many situations it does not cover all situation, and we need to calculate p -values from the summary statistics ourselves if the hypotheses are different from the default ones.

|||| Method 5.14 Level α t -tests for parameter

1. Formulate the *null hypothesis*: $H_{0,i} : \beta_i = \beta_{0,i}$, and the alternative hypothesis $H_{1,i} : \beta_i \neq \beta_{0,i}$
2. Compute the test statistic $t_{\text{obs},\beta_i} = \frac{\hat{\beta}_i - \beta_{0,i}}{\hat{\sigma}_{\beta_i}}$
3. Compute the evidence against the *null hypothesis*

$$p\text{-value}_i = 2 \cdot P(T > |t_{\text{obs},\beta_i}|) \quad (5-51)$$

4. If $p\text{-value}_i < \alpha$ reject $H_{0,i}$, otherwise accept $H_{0,i}$

In many situations we will be more interested in quantifying the uncertainty of

the parameter estimates rather than testing a specific hypothesis. This is usually given in the form of confidence intervals for the parameters:

|||| **Method 5.15** **Parameter confidence intervals**

$(1 - \alpha)$ confidence intervals for β_0 and β_1 are given by

$$\hat{\beta}_0 \pm t_{1-\alpha/2} \cdot \hat{\sigma}_{\beta_0}, \quad (5-52)$$

$$\hat{\beta}_1 \pm t_{1-\alpha/2} \cdot \hat{\sigma}_{\beta_1}, \quad (5-53)$$

where $t_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -quantile of a t -distribution with $n - 2$ degrees of freedom. Where $\hat{\sigma}_{\beta_0}$ and $\hat{\sigma}_{\beta_1}$ are calculated from the results in Theorem 5.8, and Equations (5-43) and (5-44).

|||| **Remark 5.16**

We will not show (prove) the results in Method 5.15, but see Remark 3.34.

|||| **Example 5.17** **Example 5.13 cont.**

Based on Method 5.15 we immediately find the 95% confidence intervals for the parameters

$$I_{\beta_0} = -119.96 \pm t_{0.975} \cdot 18.9 = [-163.5, -76.4],$$

$$I_{\beta_1} = 1.113 \pm t_{0.975} \cdot 0.1059 = [0.869, 1.36],$$

with the degrees of freedom for the t -distribution equal 8, and we say with high confidence that the intervals contain the true parameter values. Of course we can get these directly from the result returned by `ols.smf`:

```
print(round(fitStudents.conf_int(alpha=0.05),2))
```

```

              0      1
Intercept -163.53 -76.38
x           0.87   1.36
```


5.4.1 Confidence and prediction intervals for the line

It is clearly of interest to predict outcomes of future experiments. Here we need to distinguish between prediction intervals, where we predict the outcome of one single experiment, and confidence intervals, where we predict the mean value of future outcomes. In the latter case we only need to account for the uncertainty in the parameter estimates while in the first case we will also need to account for the uncertainty of the error (the random part ε_i).

If we conduct a new experiment with $x_i = x_{\text{new}}$ the *expected outcome* is

$$\hat{y}_{\text{new}} = \hat{\beta}_0 + \hat{\beta}_1 x_{\text{new}} \quad (5-54)$$

where the only source of variation comes from the variance of the parameter estimates, and we can calculate the variance of \hat{Y}_{new}

$$\begin{aligned} V(\hat{Y}_{\text{new}}) &= V(\hat{\beta}_0 + \hat{\beta}_1 x_{\text{new}}) \\ &= V(\hat{\beta}_0) + V(\hat{\beta}_1 x_{\text{new}}) + 2 \text{Cov}(\hat{\beta}_0, \hat{\beta}_1 x_{\text{new}}), \end{aligned} \quad (5-55)$$

now use the calculation rules for variances and covariances (Section 2.7), and insert the variances and the covariance from Theorem 5.8

$$\begin{aligned} V(\hat{Y}_{\text{new}}) &= \frac{\sigma^2}{n} + \frac{\sigma^2 \bar{x}^2}{S_{xx}} + \frac{\sigma^2 x_{\text{new}}^2}{S_{xx}} - 2 \frac{\sigma^2 \bar{x} x_{\text{new}}}{S_{xx}} \\ &= \sigma^2 \left(\frac{1}{n} + \frac{(x_{\text{new}} - \bar{x})^2}{S_{xx}} \right), \end{aligned} \quad (5-56)$$

to find the variance of a single new point, we are using

$$Y_{\text{new}} = \hat{\beta}_0 + \hat{\beta}_1 x_{\text{new}} + \varepsilon_{\text{new}}, \quad (5-57)$$

and therefore need to add the variance of the residuals (ε_{new} is independent from $\hat{\beta}_0$ and $\hat{\beta}_1$)

$$V(Y_{\text{new}}) = \sigma^2 \left(1 + \frac{1}{n} + \frac{(x_{\text{new}} - \bar{x})^2}{S_{xx}} \right). \quad (5-58)$$

When we construct confidence and prediction intervals we need to account for the fact that σ^2 is estimated from data and thus use the t -distribution:

|||| **Method 5.18 Intervals for the line**

The $(1-\alpha)$ **confidence interval** for the line $\hat{\beta}_0 + \hat{\beta}_1 x_{\text{new}}$ is

$$\hat{\beta}_0 + \hat{\beta}_1 x_{\text{new}} \pm t_{1-\alpha/2} \hat{\sigma} \sqrt{\frac{1}{n} + \frac{(x_{\text{new}} - \bar{x})^2}{S_{xx}}}, \quad (5-59)$$

and the $(1-\alpha)$ **prediction interval** is

$$\hat{\beta}_0 + \hat{\beta}_1 x_{\text{new}} \pm t_{1-\alpha/2} \hat{\sigma} \sqrt{1 + \frac{1}{n} + \frac{(x_{\text{new}} - \bar{x})^2}{S_{xx}}}, \quad (5-60)$$

where $t_{1-\alpha/2}$ is the $(1 - \alpha/2)$ -quantile of the t -distribution with $n - 2$ degrees of freedom.

|||| **Remark 5.19**

We will not show the results in Method 5.18, but use Equations (5-54)–(5-58) and Remark 3.34.

As illustrated in Figure 5.2 the confidence interval width will approach zero for an increasing number of data points (n) increase or as S_{xx} increase (actually, in most situations S_{xx} will also increase as n increase). Note also, that the confidence and prediction interval widths are smallest when $x_{\text{new}} = \bar{x}$. The prediction interval width will approach $2z_{1-\alpha/2} \cdot \sigma$ as $n \rightarrow \infty$. The difference between the intervals are that the prediction interval covers a new observation in $(1 - \alpha) \cdot 100\%$ of the times, while the confidence interval is expected to cover the true regression line $(1 - \alpha) \cdot 100\%$ of the times. One important point is: when we have calculated the prediction interval based on some particular sample, then we actually don't know the probability of this interval covering new observations. What we know is: if we repeat the experiment, then in $(1 - \alpha) \cdot 100\%$ of the times a new observation will be covered (we make a new observation each time). Same goes for the confidence interval: we don't know if the true regression line is covered by a particular interval, we only know that if we repeat the experiment, then in $(1 - \alpha) \cdot 100\%$ of the times the true regression line will be covered.

In the following: first an example on calculating confidence and prediction intervals, second an example on the width of the intervals, and finally Example 5.22 on the prediction interval coverage, are given.




Figure 5.2: Best linear fit (red line), truth (blue line), 95% prediction interval for the points (light grey area), 95 % CI for the line (dark grey area), and observed values (black dots), for simulated data (see Example 5.21).

|||| Example 5.20 Student height and weight Example (5.17 cont.)

With reference to Example 5.17 suppose we want to calculate prediction and confidence intervals for the line for a new student with $x_{\text{new}} = 200$ cm, the prediction is $\hat{y}_{\text{new}} = 102.6$ kg and the 95% confidence and prediction intervals become

$$I_{\text{pred}} = -120 + 1.113 \cdot 200 \pm t_{0.975}(8) \cdot 3.88 \sqrt{1 + \frac{1}{10} + \frac{(178 - 200)^2}{1342}} = [91.8, 113], \quad (5-61)$$

$$I_{\text{conf}} = -120 + 1.113 \cdot 200 \pm t_{0.975}(8) \cdot 3.88 \sqrt{\frac{1}{10} + \frac{(178 - 200)^2}{1342}} = [96.5, 109], \quad (5-62)$$

where $t_{0.975}$ is the 0.975-quantile of a t -distribution with $n - 2$ degrees of freedom.

The intervals can be calculated directly by:

```
new_data = pd.DataFrame({'x': [200]})

# Get prediction and confidence intervals
pd.set_option("display.float_format", None) ## unset option
pred = fitStudents.get_prediction(new_data).summary_frame(alpha=0.05)
print(round(pred,2))
```

	mean	mean_se	mean_ci_lower	mean_ci_upper	obs_ci_lower	obs_ci_upper
0	102.59	2.63	96.52	108.66	91.77	113.41

in the output mean refer to the predicted expected value, mean_se refer to the standard error of the predicted expected value, mean_ci_.. refer to confidence intervals while obs_ci_.. refer to prediction intervals.

|||| Example 5.21 Simulation

Figure 5.2 illustrates the difference between the confidence and prediction intervals for simulated data, with different numbers of observations. The simulated model is

$$y_i = 10 + 2x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, 5^2). \quad (5-63)$$

When n increases the width of the confidence interval for the line narrows and approaches 0, while the prediction interval width does not approach 0, but rather $2z_{1-\alpha/2}\sigma$. Further, the width of the prediction interval will always be larger than the width of the confidence interval.

|||| Example 5.22 Prediction interval coverage

In this example it is illustrated that we actually don't know the probability that a prediction interval covers new observations, when it is calculated using a sample (i.e. we have a realization of the prediction interval). First, a prediction interval is calculated using a single sample and it is investigated how many of k new observations falls inside it:

```

np.random.seed(129)
n = 30
beta0 = 10
beta1 = 3
sigma = 0.5
# Generate some input values
x = np.random.uniform(-10, 10, n)
# Simulate output values
y = beta0 + beta1*x + np.random.normal(0, sigma, size=n)
# Fit a simple linear regression model to the sample
D = pd.DataFrame({'x': x, 'y': y})
fit = smf.ols(formula = 'y ~ x', data=D).fit()

# The number of new observations
k = 10000
# Generate k new input values
xnew = pd.DataFrame({'x': np.random.uniform(-10, 10, k)})
# Calculate the prediction intervals for the new input values
PI = fit.get_prediction(xnew).conf_int(obs=True)
# Simulate new output observations
ynew = beta0 + beta1*xnew['x'] + np.random.normal(0, sigma, size=k)
# Calculate the fraction of times the prediction interval covered the
# new observation
y_within_ci = (lambda y, ci_low, ci_high: (y > ci_low) & (y < ci_high))(ynew, PI[:,0],
print(sum(y_within_ci)/k)

0.8162

```

We see that the interval covered only 81.6% of the new observations, which is quite a bit below 95% (per default `.get_prediction` use $\alpha = 5\%$).

Now, let's repeat the sampling, so we make a new sample k times and each time calculate a new fit and prediction interval, and each time check if a new observation falls inside it:

```

# The number of simulated samples
k = 10000
# Define function to make model and check new y observation
def check_newobs_within_ci():
    # The number of observations and the parameters
    n = 30
    beta0 = 10
    beta1 = 3
    sigma = 0.5
    # Generate some input values
    x = np.random.uniform(-10, 10, n)
    # Simulate output values
    y = beta0 + beta1*x + np.random.normal(0, sigma, size=n)
    # Fit a simple linear regression model to the sample
    D = pd.DataFrame({'x': x, 'y': y})
    fit = smf.ols(formula = 'y ~ x', data=D).fit()
    # Generate ONE new input value
    xnew = pd.DataFrame({'x': np.random.uniform(-10, 10, 1)})
    # The prediction interval for the new value
    PI = fit.get_prediction(xnew).conf_int(obs=True)
    # Simulate a single new observations
    ynew = beta0 + beta1*xnew['x'] + np.random.normal(0, sigma, size=1)
    # Check if the new observation was inside the interval
    y_within_ci = (lambda y, ci_low, ci_high: (y > ci_low) & (y < ci_high))(ynew, PI[:,0])
    return y_within_ci

# Replicate the calculation 10000 times
replicated_results = [check_newobs_within_ci() for _ in range(k)]

# The fraction of covered new observations
print(sum(replicated_results)/k)

0      0.9536
Name: x, dtype: float64

```

It is found that coverage is now very close to the expected 95% and this is indeed the way the coverage probability should be interpreted: with repeated sampling the probability is $1 - \alpha$ that a prediction interval will cover a randomly chosen new observation. Same goes for confidence intervals (of any kind): with repeated sampling the probability is $1 - \alpha$ that a confidence interval will cover the true value.

5.5 Matrix formulation of simple linear regression

The simple linear regression problem can be formulated in vector-matrix notation as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (5-64)$$

or

$$\begin{bmatrix} Y_1 \\ \vdots \\ Y_n \end{bmatrix} = \begin{bmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \vdots \\ \varepsilon_n \end{bmatrix}, \quad \varepsilon_i \sim N(0, \sigma^2). \quad (5-65)$$

One of the advantages of the matrix formulation is that the analysis generalize to higher dimensions in a straight forward way (i.e. more x s and parameters as in the following chapter). The residual sum of squares is given by

$$RSS = \boldsymbol{\varepsilon}^T \boldsymbol{\varepsilon} = (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})^T (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}), \quad (5-66)$$

and the parameter estimators are given by:

|||| Theorem 5.23

The estimators of the parameters in the simple linear regression model are given by

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}, \quad (5-67)$$

and the covariance matrix of the estimates is

$$V[\hat{\boldsymbol{\beta}}] = \sigma^2 (\mathbf{X}^T \mathbf{X})^{-1}, \quad (5-68)$$

and central estimate for the error variance is

$$\hat{\sigma}^2 = \frac{RSS}{n-2}. \quad (5-69)$$

Here $V[\hat{\boldsymbol{\beta}}]$ is a matrix with elements $(V[\hat{\boldsymbol{\beta}}])_{11} = V[\hat{\beta}_0]$, $(V[\hat{\boldsymbol{\beta}}])_{22} = V[\hat{\beta}_1]$, and $(V[\hat{\boldsymbol{\beta}}])_{12} = (V[\hat{\boldsymbol{\beta}}])_{21} = \text{Cov}[\hat{\beta}_0, \hat{\beta}_1]$.

When we want to find the minimum of RSS , we again need to differentiate RSS

with respect to the parameters

$$\begin{aligned}\frac{\partial RSS}{\partial \beta} &= -2\mathbf{X}^T(\mathbf{Y} - \mathbf{X}\beta) \\ &= -2(\mathbf{X}^T\mathbf{Y} - \mathbf{X}^T\mathbf{X}\beta).\end{aligned}\tag{5-70}$$

Solving for β gives

$$\hat{\beta} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y},\tag{5-71}$$

taking the expectation of $\hat{\beta}$ we get

$$\begin{aligned}\mathbb{E}[\hat{\beta}] &= \mathbb{E}[(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y}] \\ &= (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbb{E}[\mathbf{X}\beta + \varepsilon] \\ &= (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{X}\beta \\ &= \beta.\end{aligned}\tag{5-72}$$

The variance of the parameters estimates are given by

$$\begin{aligned}\mathbb{V}[\hat{\beta}] &= \mathbb{V}[(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y}] \\ &= (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbb{V}[\mathbf{X}\beta + \varepsilon]\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1} \\ &= (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T(\mathbb{V}[\mathbf{X}\beta] + \mathbb{V}[\varepsilon])\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1} \\ &= (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\sigma^2\mathbf{I}\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1} \\ &= \sigma^2(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1} \\ &= \sigma^2(\mathbf{X}^T\mathbf{X})^{-1}.\end{aligned}\tag{5-73}$$

Again a central estimate for σ^2 is

$$\hat{\sigma}^2 = \frac{RSS(\hat{\beta})}{n-2},\tag{5-74}$$

and the estimate of the parameter covariance matrix is

$$\hat{\Sigma}_{\beta} = \hat{\sigma}^2(\mathbf{X}^T\mathbf{X})^{-1}.\tag{5-75}$$

Marginal tests ($H_0 : \beta_i = \beta_{i,0}$) are constructed by observing that

$$\frac{\hat{\beta}_i - \beta_{i,0}}{\sqrt{(\hat{\Sigma}_{\beta})_{ii}}} \sim t(n-2).\tag{5-76}$$

The matrix calculations are illustrated in the next example.

|||| Example 5.24 Student height and weight

To illustrate how the matrix formulation works in student height and weight data is worked through below:

```
# Data
y = student['y']
n = len(student['y'])
X = np.array([np.repeat(1,n),student['x']]).T

# Parameter estimates and variance
beta = np.linalg.inv(X.T@X)@X.T@y
e      = y - X@beta
s      = math.sqrt(np.sum(e**2) / (n - 2))
Vbeta  = s**2 * np.linalg.inv(X.T@X)
sbeta  = np.sqrt(Vbeta.diagonal())
Tstat  = beta / sbeta
pval   = 2 * (1 - stats.t.cdf(np.abs(Tstat), df = n-2))

# Print the results
coefMat = np.array([beta, sbeta, Tstat, pval]).T
col_names = ["Estimates", "Std.Error", "t.value", "p.value"]
row_names = ["beta0", "beta1"]
coefMat = pd.DataFrame(coefMat, columns = col_names, index = row_names)
pd.set_eng_float_format(accuracy=3)
print(coefMat)

           Estimates      Std.Error      t.value      p.value
beta0 -119.958E+00    18.897E+00   -6.348E+00   221.088E-06
beta1   1.113E+00    105.939E-03   10.504E+00    5.875E-06

print(round(s,2))

3.88

pd.set_option("display.float_format", None) ## unset option
```

```
# Prediction and confidence interval
xnew = np.array([1, 200]).T
ynew = xnew@beta
Vconf = xnew@Vbeta@xnew.T
Vpred = Vconf + s**2
se_pred = np.sqrt(np.array([Vconf, Vpred]))
np.round(se_pred, 2)

array([2.630, 4.690])
```

5.6 Correlation

In the analysis above we focus on situations where we are interested in one variable (y) as a function of another variable (x). In other situations we might be more interested in how x and y vary together. Examples could be ecosystems, where the number of predators is a function of the number of preys, but the reverse relation is also true, further both of these numbers are affected by random variations and knowledge of one only gives partial knowledge of the other. Another example is individual student grade in 2 different courses, before any grade has been given we will expect that a high grade in one course will imply a high grade in the other course, but none of them is controlled or known in advance.

In the cases above we talk about correlation analysis and to this end we will need the sample correlation coefficient, as defined in Section 1.4.3

$$\hat{\rho} = \frac{1}{n-1} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s_x} \right) \left(\frac{y_i - \bar{y}}{s_y} \right). \quad (5-77)$$

In Section 1.4.3 we notated sample correlation with r , but here we use $\hat{\rho}$, since it is an estimate for the correlation ρ (see Section 2.8), and imply that there is a meaningful interpretation of the ρ .

5.6.1 Inference on the sample correlation coefficient

In order to answer the question: are X and Y correlated? We will be interested in constructing a test of the type

$$H_0 : \rho = 0, \quad H_1 : \rho \neq 0. \quad (5-78)$$

Consider the model

$$Y_i = \beta_0 + \beta_1 X_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad (5-79)$$

in this case we can rewrite the sample correlation as

$$\begin{aligned} \hat{\rho} &= \frac{1}{n-1} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s_x} \right) \left(\frac{y_i - \bar{y}}{s_y} \right) \\ &= \frac{S_{xx}}{n-1} \frac{1}{S_{xx}} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s_x} \right) \left(\frac{y_i - \bar{y}}{s_y} \right) \\ &= \frac{S_{xx}}{n-1} \frac{1}{s_x s_y} \hat{\beta}_1 \\ &= \frac{s_x}{s_y} \hat{\beta}_1, \end{aligned} \quad (5-80)$$

implying that the hypothesis (5-78) can be tested by testing the hypothesis

$$H_0 : \beta_1 = 0; \quad H_1 : \beta_1 \neq 0. \quad (5-81)$$

since clearly the relationship in Equation (5-79) can be reversed. It should be noted that we cannot use the test to construct a confidence interval for ρ .

It should be stressed that correlation does not imply causality, it just implies that the variables x and y vary together. As an example consider the number of beers sold at the university bar and the number of students attending the introductory course in statistics. Let's say that both numbers have increased and therefore have a high correlation coefficient, but it does not seem reasonable to conclude that students are more interested in statistics when drinking beers. A closer look might reveal that the number of enrolled students have actually increased and this can indeed explain the increase in both numbers.

5.6.2 Correlation and regression

In the linear regression models we would like to measure how much of the variation in the outcome (Y) is explained by the input (x). A commonly used measure for this is the coefficient of determination (explanation) or R^2 -value (see also the Python summary in Example 5.5).

|||| **Definition 5.25** **Coefficient of determination R^2**

The coefficient of determination expresses the proportion of variation in the outcome (Y) explained by the regression line

$$r^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}. \quad (5-82)$$

In order to find this we will split the variance of y into a component due to the regression line and a component due to the residual variation

$$\begin{aligned} s_y^2 &= \frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})^2 \\ &= \frac{1}{n-1} \sum_{i=1}^n (\hat{\beta}_0 + \hat{\beta}_1 x_i + e_i - \frac{1}{n} \sum_{i=1}^n (\hat{\beta}_0 + \hat{\beta}_1 x_i + e_i))^2 \\ &= \frac{1}{n-1} \sum_{i=1}^n (\hat{\beta}_1 (x_i - \bar{x}) + e_i)^2 \\ &= \hat{\beta}_1^2 s_x^2 + \frac{n-2}{n-1} \hat{\sigma}^2, \end{aligned} \quad (5-83)$$

where the first term on the right hand side is the variability explained by the regression line and the second term is the residual variation. Dividing with the variance of Y gives a splitting in the relative variation from each of the terms. If we write out the variation explained by the regression line we get

$$\begin{aligned} \frac{\hat{\beta}_1^2 s_x^2}{s_y^2} &= \left(\frac{\sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x})}{\sum_{i=1}^n (x_i - \bar{x})^2} \right)^2 \frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n-1} \frac{n-1}{\sum_{i=1}^n (y_i - \bar{y})^2} \\ &= \left(\frac{1}{n-1} \sum_{i=1}^n (y_i - \bar{y})(x_i - \bar{x}) \right)^2 \frac{n-1}{\sum_{i=1}^n (x_i - \bar{x})^2} \frac{n-1}{\sum_{i=1}^n (y_i - \bar{y})^2} \\ &= \left(\frac{1}{n-1} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s_x} \right) \left(\frac{y_i - \bar{y}}{s_y} \right) \right)^2 \\ &= \hat{\rho}^2. \end{aligned} \quad (5-84)$$

We can therefore conclude that the proportion of variability (R^2) in Y explained by the regression line is equal to the squared sample correlation coefficient ($\hat{\rho}^2$).

|||| Example 5.26 Student weight and height (Example 5.20 cont.)

With reference to Example 5.20 above we can calculate the correlation coefficient by:

```
np.round(np.corrcoef(student['x'], student['y'])[0, 1]**2,3)

np.float64(0.932)
```

or we can base our calculations on the estimated slope:

```
round(fitStudents.params["x"]**2 * student['x'].var() / student['y'].var(),3)

np.float64(0.932)
```

or we can find it directly in the summary of the regression model (see Example 5.5): where the number is called R-squared.

5.7 Model validation

So far we have discussed how to estimate parameters, predict future values, make inference etc. in the model

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2). \quad (5-85)$$

In all we have done so far the basic assumption is that the residuals are normally distributed with zero mean and constant variance, and further the residuals are mutually independent. These are assumptions which should be checked and if the assumptions are not fulfilled some actions should be taken in order to fix this. This is called *model validation* or *residual analysis* and is exactly the same idea behind the validation needed for the mean model used for *t*-tests in Section 3.1.8, though here including a few more steps.

The normality assumption can be checked by a normal q-q plot, and the constant variance assumption may be checked by plotting the residuals as a function of the fitted values. The normal q-q plot have been treated in Section 3.1.8 and should be applied equivalently. Plotting the residuals as a function of the fitted values should not show a systematic behaviour, this means that the range should be constant and the mean value should be constant, as illustrated in the following example:

|||| Example 5.27 Residuals analysis

We consider data generated from the following three models

$$Y_{1,i} = \beta_0 + \beta_1 x_{1,i} + \varepsilon_i, \quad \varepsilon_i \sim N(0, 1), \quad (5-86)$$

$$Y_{2,i} = \beta_0 + \beta_1 x_{1,i} + \beta_2 x_{2,i} + \varepsilon_i, \quad \varepsilon_i \sim N(0, 1), \quad (5-87)$$

$$Y_{3,i} = e^{\beta_0 + \beta_1 x_{1,i} + \varepsilon_i}, \quad \varepsilon_i \sim N(0, 1) \quad (5-88)$$


In all cases we fit the model

$$Y_i = \beta_0 + \beta_1 x_{1,i} + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad (5-89)$$

to the data: from the first model we would expect that the residual analysis do not show any problems, for the second model we have a linear dependence which is not included in the model and we should see this in the residual analysis, and the third is a non-linear function of the residuals as well as the regressors and one way to handle this will be discussed.

The first model is simulated, estimated and analysed by ($\beta_0 = 0$, $\beta_1 = 1$, and $\sigma^2 = 1$):


```
np.random.seed(134)
n = 100
x1 = np.linspace(1, 10, n)
y = x1 + np.random.normal(0, 1, size=n)
D = pd.DataFrame({'x': x1, 'y': y})
fit = smf.ols(formula = 'y ~ x', data=D).fit()
# Get the predictions (fitted values)
ypred = fit.predict(D) # or fit.fittedvalues
# Plot the Q-Q plot and the residuals vs. fitted values
fig, ax = plt.subplots(1,2)
# Q-Q plot of the residuals
sm.qqplot(y - ypred, line="q", a=1/2, ax=ax[0])
ax[0].set_title("Q-Q plot")
# Scatter plot
ax[1].scatter(ypred, y - ypred)
ax[1].set_xlabel("Fitted values")
ax[1].set_ylabel("Residuals")
ax[1].set_title("Residuals vs Fitted values")
plt.tight_layout()
plt.show()
```



As we can see there is no serious departure from normality and there are no patterns in the residuals as a function of the fitted values.

The second model (with $\beta_0 = 0$, $\beta_1 = 1$, $\beta_2 = 0.5$ and $\sigma^2 = 1$) is simulated, estimated and analysed by (plot functions omitted):


```
np.random.seed(134)
n = 100
x1 = np.linspace(1, 10, n)
x2 = x1**2
y = x1 + 0.5 * x2 + stats.norm.rvs(loc = 0, scale = 1, size=n)
D = pd.DataFrame({'x': x1, 'y': y})
fit = smf.ols(formula = 'y ~ x', data=D).fit()
```



We see some departure from normality, but also that the residuals are related to the fitted values with a clear pattern. In the next chapter we will learn that we should find the hidden dependence (x_2) and include it in the model.


The third model (with $\beta_0 = 0$, $\beta_1 = 1$, $\beta_2 = 0.5$ and $\sigma^2 = 1$) is simulated, estimated and analysed by (plot function omitted):

```
np.random.seed(134)
n = 100
x1 = np.linspace(4, 10, n)
y = np.exp( 0.2 * x1 + np.random.normal(0, 0.15, size=n))
D = pd.DataFrame({'x': x1, 'y': y})
fit = smf.ols(formula = 'y ~ x', data=D).fit()
```



We see some departure from normality, and also that the variance increases as a function of the fitted values. When the variance is clearly related with the fitted values one should try to transform the dependent variable. The following do the analysis based in log-transformed data:

```
np.random.seed(134)
n = 100
x1 = np.linspace(4, 10, n)
y = np.exp( 0.2 * x1 + np.random.normal(0, 0.15, size=n))
y = np.log(y)
D = pd.DataFrame({'x': x1, 'y': y})
fit = smf.ols(formula = 'y ~ x', data=D).fit()
```

From the q-q plot it is found that the distribution is now very close to a normal distribution compared to previous q-q plot. And, as we can see the residuals are no longer related clearly to the fitted values.

|||| **Method 5.28 Model validation (or residual analysis)**

1. Check the normality assumption with a q-q plot of the residuals
2. Check the systematic behaviour by plotting the residuals e_i as a function of fitted values \hat{y}_i

|||| **Remark 5.29 Independence**

In general independence should also be checked, while there are ways to do this we will not discuss them here.

Glossaries

Correlation [Korrelation] The sample correlation coefficient are a summary statistic that can be calculated for two (related) sets of observations. It quantifies the (linear) strength of the relation between the two. See also: Covariance [32–34](#)

Covariance [Kovarians] The sample covariance coefficient are a summary statistic that can be calculated for two (related) sets of observations. It quantifies the (linear) strength of the relation between the two. See also: Correlation [11, 12, 15, 17, 23, 29, 30](#)

Critical value *Kritisk værdi* As an alternative to the p -value one can use the so-called critical values, that is the values of the test-statistic which matches exactly the significance level [20, 21](#)

Degrees of freedom [Frihedsgrader] The number of "observations" in the data that are free to vary when estimating statistical parameters often defined as $n - 1$ [12, 19, 20, 22, 24, 25](#)

Expectation [Forventningsværdi] A function for calculating the mean. The value we expect for a random variable (or function of random variables), hence of the population [1, 4, 10, 11, 30](#)

Histogram [Histogram] The default histogram uses the same width for all classes and depicts the raw frequencies/counts in each class. By dividing the raw counts by n times the class width the density histogram is found where the area of all bars sum to 1 [9](#)

Independence [Uafhængighed] [15, 39](#)

(Statistical) Inference [Statistisk inferens (følgeslutninger baseret på data)] [2, 20, 35](#)

Least squares [Mindste kvadraters (metode)] [4, 5, 7](#)

Linear regression [Lineær regression (-sanalyse)] [2](#), [7](#), [10](#), [19](#), [29](#), [33](#)

Normal distribution [Normal fordeling] [18](#)

Null hypothesis [Nulhypotese (H_0)] [19–21](#)

P-value [p -værdi (for faktisk udfald af en teststørrelse)] [20](#), [21](#)

Two-sided (test) [Tosidet test (test med tosidet alternativ)] Is also called non-directional (test) [20](#)

Acronyms

ANOVA Analysis of Variance *Glossary:* [Analysis of Variance](#)

cdf cumulated distribution function *Glossary:* [cumulated distribution function](#)

CI confidence interval [11](#), [19](#), [22–24](#), [26](#), [33](#), *Glossary:* [confidence interval](#)

CLT Central Limit Theorem *Glossary:* [Central Limit Theorem](#)

IQR Inter Quartile Range *Glossary:* [Inter Quartile Range](#)

LSD Least Significant Difference *Glossary:* [Least Significant Difference](#)

pdf probability density function *Glossary:* [probability density function](#)