

Linear Statistical Models

Math 6010-1; Fall 2016

Davar Khoshnevisan

155 South 1400 East JWB 233, Department of Mathematics, University
of Utah, Salt Lake City UT 84112-0090

E-mail address: `davar@math.utah.edu`

URL: `http://www.math.utah.edu/~davar`

Contents

<i>Lecture 1.</i>	Linear statistical models	1
§1.	Introduction	1
§2.	The method of least squares	2
§3.	Simple linear regression	4
<i>Lecture 2.</i>	Random vectors	7
§1.	Expectation	7
§2.	Covariance	8
§3.	Mathematical properties of variance and covariance	9
§4.	A relation to positive-semidefinite matrices	10
<i>Lecture 3.</i>	Some linear algebra	11
§1.	Symmetric matrices	11
§2.	Positive-semidefinite matrices	13
§3.	The rank of a matrix	15
§4.	Projection matrices	17
<i>Lecture 4.</i>	Quadratic forms	21
§1.	Random quadratic forms	21
§2.	Examples of quadratic forms	22
§3.	The variance of a random quadratic form	24
<i>Lecture 5.</i>	Moment-generating functions and independence	27
<i>Lecture 6.</i>	Gaussian Random Vectors	31
§1.	The multivariate normal distribution	31

§2.	The nondegenerate case	32
§3.	The bivariate normal distribution	34
§4.	A few important properties of multivariate normal distributions	35
§5.	Quadratic forms	37
<i>Lecture 7.</i>	Linear Models	39
§1.	The basic model	39
§2.	The least-squares estimator of $\theta := X\beta$	41
§3.	The least-squares estimator of β	41
§4.	Optimality	43
§5.	Regression and prediction	43
§6.	Estimation of σ^2	44
§7.	The normal model	44
§8.	Some examples	45
<i>Lecture 8.</i>	Assessing Normality	49
§1.	Visual data exploration	49
§2.	General remarks	49
§3.	Histograms	50
§4.	QQ-Plots	51
§5.	The Correlation Coefficient of the QQ-Plot	51
§6.	Some benchmarks	52
<i>Lecture 9.</i>	Hypothesis Testing	59
§1.	A test for one parameter	59
§2.	Least-squares estimates for contrasts	60
§3.	The normal model	62
§4.	Examples	64
<i>Lecture 10.</i>	Confidence Intervals and Sets	71
§1.	Confidence intervals for one parameter	71
§2.	Confidence ellipsoids	72
§3.	Bonferonni bounds	73
§4.	Scheffé's simultaneous conservative confidence bounds	74
§5.	Confidence bounds for the regression surface	76
§6.	Prediction intervals	77
<i>Lecture 11.</i>	Polynomial Regression	79
§1.	Introduction	79

§2. Orthogonal polynomials	80
§3. The Forsythe–Hayes 2-step method	82
<i>Lecture 12.</i> Splines	85
§1. Introduction	85
§2. Linear splines	85
§3. Quadratic splines	86
§4. Cubic splines	87

Linear statistical models

1. Introduction

The goal of this course is, in rough terms, to predict a variable y , given that we have the opportunity to observe variables x_1, \dots, x_{p-1} . This is a very important statistical problem. Therefore, let us spend a bit of time and examine a simple example:

Given the various vital statistics of a newborn baby, you wish to predict his height y at maturity. Examples of those “vital statistics” might be $x_1 :=$ present height, and $x_2 :=$ present weight. Or perhaps there are still more predictive variables $x_3 :=$ your height and $x_4 :=$ your spouse’s height, etc.

In actual fact, a visit to your pediatricist might make it appear that this prediction problem is trivial. But that is not so [though it is nowadays fairly well understood in this particular case]. A reason the problem is nontrivial is that there is no *a priori* way to know “how” y depends on x_1, \dots, x_4 [say, if we plan to use all 4 predictive variables]. In such a situation, one resorts to writing down a reasonable model for this dependence structure, and then analyzing that model. [Finally, there might be need for model verification as well.]

In this course, we study the general theory of “linear statistical models.” That theory deals with the simplest possible nontrivial setting where such problems arises in various natural ways. Namely, in that theory we *posit* that y is a linear function of (x_1, \dots, x_4) , possibly give or take some “noise.” In other words, the theory of linear statistical models posits that there exist unknown parameters $\beta_0, \dots, \beta_{p-1}$ [here, $p = 5$] such that

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_{p-1} x_{p-1} + \varepsilon, \quad (1)$$

where ε is a random variable. The problem is still not fully well defined [for instance, what should be the distribution of ε , etc.?). But this is roughly

the starting point of the theory of linear statistical models. And one begins asking natural questions such as, “how can we estimate β_0, \dots, β_4 ?” or “can we perform inference for these parameters?” [for instance, can we test to see if y does not depend on x_1 in this model; i.e., test for $H_0 : \beta_1 = 0$?]. And so on.

We will also see, at some point, how the model can be used to improve itself. For instance, suppose we have only one predictive variable x , but believe to have a nonlinear dependence between y and x . Then we could begin by thinking about *polynomial regression*; i.e., a linear statistical model of the form

$$y = \beta_0 + \beta_1 x + \beta_2 x^2 + \dots + \beta_{p-1} x^{p-1} + \varepsilon.$$

Such a model fits in the general form (1) of linear statistical models, as well: We simply define new predictive variables $x_j := x^j$ for all $1 \leq j < p$. One of the conclusions of this discussion is that we are studying models that are linear functions of unknown parameters $\beta_0, \dots, \beta_{p-1}$ and not x_1, \dots, x_{p-1} . This course studies statistical models with such properties. And as it turns out, not only these models are found in a great number of diverse applications, but also they have a rich mathematical structure.

2. The method of least squares

Suppose we have observed n data points in pairs: $(x_1, y_1), \dots, (x_n, y_n)$. The basic problem here is, what is the best straight line that fits this data? There is of course no unique sensible answer, because “best” might mean different things.

We will use the *method of least squares*, introduced by C.-F. Gauss. Here is how the method works: If we used the line $y = \beta_0 + \beta_1 x$ to describe how the x_i 's affect the y_i 's, then the error of approximation, at $x = x_i$, is $y_i - (\beta_0 + \beta_1 x_i)$; this is called the *i th residual error*. The sum of the squared residual errors is $\text{SSE} := \sum_{i=1}^n (y_i - \beta_0 - \beta_1 x_i)^2$, and the method of least squares is to find the line with the smallest SSE. That is, we need to find the optimal β_0 and β_1 —written $\hat{\beta}_0$ and $\hat{\beta}_1$ —that solve the following optimization problem:

$$\min_{\beta_0, \beta_1} \sum_{i=1}^n (y_i - \beta_0 - \beta_1 x_i)^2. \quad (2)$$

Theorem 1 (Gauss). *The least-squares solution to (2) is given by*

$$\hat{\beta}_1 := \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^n (x_i - \bar{x})^2} \quad \text{and} \quad \hat{\beta}_0 := \bar{y} - \hat{\beta}_1 \bar{x}.$$

Proof. Define

$$L(\beta_0, \beta_1) := \sum_{i=1}^n (y_i - \beta_0 - \beta_1 x_i)^2.$$

Our goal is to minimize the function L . An inspection of the graph of L shows that L has a unique minimum; multivariable calculus then tells us that it suffices to set $\partial L / \partial \beta_j = 0$ for $j = 1, 2$ and solve. Because

$$\begin{aligned}\frac{\partial}{\partial \beta_0} L(\beta_0, \beta_1) &= -2 \sum_{i=1}^n (y_i - \beta_0 - \beta_1 x_i), \\ \frac{\partial}{\partial \beta_1} L(\beta_0, \beta_1) &= -2 \sum_{i=1}^n x_i (y_i - \beta_0 - \beta_1 x_i),\end{aligned}$$

a few lines of simple arithmetic finish the derivation. \square

The preceding implies that, given the points $(x_1, y_1), \dots, (x_n, y_n)$, the best line of fit through these points—in the sense of least squares—is

$$y = \hat{\beta}_0 + \hat{\beta}_1 x. \quad (3)$$

For all real numbers x and y , define x_{SU} and y_{SU} to be their respective “standardizations.” That is,

$$x_{\text{SU}} := \frac{x - \bar{x}}{\sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2}}, \quad y_{\text{SU}} := \frac{y - \bar{y}}{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2}}.$$

Then, (3) can be re-written in the following equivalent form:

$$\begin{aligned}y_{\text{SU}} &= \frac{\hat{\beta}_0 - \bar{y}}{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2}} + \frac{\hat{\beta}_1 x}{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2}} \\ &= \frac{\hat{\beta}_1}{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2}} (x - \bar{x}),\end{aligned}$$

the last line following from the identity $\hat{\beta}_0 = \bar{y} - \hat{\beta}_1 \bar{x}$. We re-write the preceding again:

$$y_{\text{SU}} = \frac{\hat{\beta}_1 \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2}}{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2}} x_{\text{SU}}.$$

Because

$$\hat{\beta}_1 = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2},$$

we can re-write the best line of fit, yet another time, this time as the following easy-to-remember formula:

$$y_{\text{SU}} = r x_{\text{SU}},$$

where r [a kind of “correlation coefficient”] is defined as

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

3. Simple linear regression

Suppose Y_1, \dots, Y_n are observations from a distribution, and they satisfy

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i \quad (1 \leq i \leq n), \quad (4)$$

where $\varepsilon_1, \dots, \varepsilon_n$ are [unobserved] i.i.d. $N(0, \sigma^2)$ for a fixed [possibly unknown] $\sigma > 0$. We assume that x_1, \dots, x_n are known, and seek to find the “best” β_0 and β_1 .¹

In other words, we believe that we are observing a certain linear function of the variable x at the x_i 's, but our measurement [and/or modeling] contains noise sources $[\varepsilon_i$'s] which we cannot observe.

Theorem 2 (Gauss). *Suppose $\varepsilon_1, \dots, \varepsilon_n$ are i.i.d. with common distribution $N(0, \sigma^2)$, where $\sigma > 0$ is fixed. Then the maximum likelihood estimators of β_1 and β_0 are, respectively,*

$$\hat{\beta}_1 = \frac{\sum_{j=1}^n (x_j - \bar{x})(Y_j - \bar{Y})}{\sum_{j=1}^n (x_j - \bar{x})^2} \quad \text{and} \quad \hat{\beta}_0 = \bar{Y} - \hat{\beta}_1 \bar{x}.$$

Therefore, based on the data $(x_1, Y_1), \dots, (x_n, Y_n)$, we predict the y -value at $x = x_*$ to be $y_* := \hat{\beta}_0 + \hat{\beta}_1 x_*$.

Proof. Note that Y_1, \dots, Y_n are independent [though not i.i.d.], and the distribution of Y_j is $N(\beta_0 + \beta_1 x_j, \sigma^2)$. Therefore, the joint probability density function of (Y_1, \dots, Y_n) is

$$f(y_1, \dots, y_n) := \frac{1}{(2\pi\sigma^2)^{n/2}} \exp \left[-\frac{1}{2\sigma^2} \sum_{j=1}^n (y_j - \beta_0 - \beta_1 x_j)^2 \right].$$

According to the MLE principle we should maximize $f(Y_1, \dots, Y_n)$ over all choices of β_0 and β_1 . But this is equivalent to minimizing $L(\beta_0, \beta_1) := \sum_{i=1}^n (Y_i - \beta_0 - \beta_1 x_i)^2$. But this was exactly what we did in Theorem 1 [except for real variables y_1, \dots, y_n in place of the random variables Y_1, \dots, Y_n]. \square

In this way we have the following “regression equation,” which uses the observed data $(x_1, Y_1), \dots, (x_n, Y_n)$ in order to predict a y -value corresponding to $x = x_*$:

$$Y(x) = \hat{\beta}_0 + \hat{\beta}_1 x.$$

But as we shall see, this method is good not only for prediction, but also for inference. Perhaps a first important question is, “do the y 's depend linearly on the x 's”? Mathematically speaking, we are asking to test the hypothesis that $\beta_1 = 0$. If we could compute the distribution of $\hat{\beta}_1$, then standard methods can be used to accomplish this. We will see later on how this can be accomplished.

¹In actual applications, the x_i 's are often random. In such a case, we assume that the model holds after conditioning on the x_i 's.

But before we develop the theory of linear inference we need to know a few things about linear algebra and some of its probabilistic consequences.

Random vectors

It will be extremely helpful to us if we worked directly with random vectors and not a group of individual random variables. Throughout, all vectors are written columnwise; and so are random ones. Thus, for instance, a random vector $\mathbf{X} \in \mathbf{R}^n$ is written columnwise as

$$\mathbf{X} = \begin{pmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{pmatrix} = (X_1, \dots, X_n)'.$$

And even more generally, we might sometimes be interested in random matrices. For instance, a random $m \times n$ matrix is written coordinatewise as

$$\mathbf{X} = \begin{pmatrix} X_{1,1} & \cdots & X_{1,n} \\ \vdots & & \vdots \\ X_{m,1} & \cdots & X_{m,n} \end{pmatrix}.$$

1. Expectation

If \mathbf{X} is a random $m \times n$ matrix, then we define its *expectation* in the most natural possible way as

$$E\mathbf{X} := \begin{pmatrix} EX_{1,1} & \cdots & EX_{1,n} \\ \vdots & & \vdots \\ EX_{m,1} & \cdots & EX_{m,n} \end{pmatrix}.$$

Many of the properties of expectations continue to hold in the setting of random vectors and/or matrices. The following summarizes some of those properties.

Proposition 1. Suppose \mathbf{A} , \mathbf{B} , \mathbf{C} , and \mathbf{D} are nonrandom matrices, and \mathbf{X} and \mathbf{Y} are random matrices. Then,

$$\mathbf{E}(\mathbf{AXB} + \mathbf{CYD}) = \mathbf{A}(\mathbf{EX})\mathbf{B} + \mathbf{C}(\mathbf{EY})\mathbf{D},$$

provided that the matrix dimensions are sensible.

Proof. Because

$$(\mathbf{AXB} + \mathbf{CYD})_{i,j} = (\mathbf{AXB})_{i,j} + (\mathbf{CYD})_{i,j},$$

it suffices to prove that $\mathbf{E}(\mathbf{AXB}) = \mathbf{AE}(\mathbf{X})\mathbf{B}$. But then, we can work coordinatewise as follows:

$$\begin{aligned} \mathbf{E}[(\mathbf{AXB})_{i,j}] &= \mathbf{E}\left[\sum_{k=1}^n \sum_{\ell=1}^n A_{i,k} X_{k,\ell} B_{\ell,j}\right] = \sum_{k=1}^n \sum_{\ell=1}^n A_{i,k} \mathbf{E}(X_{k,\ell}) B_{\ell,j} \\ &= \sum_{k=1}^n \sum_{\ell=1}^n A_{i,k} [\mathbf{EX}]_{k,\ell} B_{\ell,j} = [\mathbf{AE}(\mathbf{X})\mathbf{B}]_{i,j}. \end{aligned}$$

That is, $\mathbf{E}(\mathbf{AXB}) = \mathbf{AE}(\mathbf{X})\mathbf{B}$ coordinatewise. This proves the result. \square

2. Covariance

Suppose $\mathbf{X} = (X_1, \dots, X_m)'$ and $\mathbf{Y} = (Y_1, \dots, Y_n)'$ are two jointly distributed random vectors. We define their *covariance* as

$$\text{Cov}(\mathbf{X}, \mathbf{Y}) := \begin{pmatrix} \text{Cov}(X_1, Y_1) & \cdots & \text{Cov}(X_1, Y_n) \\ \vdots & & \vdots \\ \text{Cov}(X_m, Y_1) & \cdots & \text{Cov}(X_m, Y_n) \end{pmatrix}.$$

Proposition 2. We always have

$$\text{Cov}(\mathbf{X}, \mathbf{Y}) = \mathbf{E}[(\mathbf{X} - \mathbf{EX})(\mathbf{Y} - \mathbf{EY})'].$$

Proof. The (i, j) th entry of the matrix $(\mathbf{X} - \mathbf{EX})(\mathbf{Y} - \mathbf{EY})'$ is $(X_i - \mathbf{EX}_i)(Y_j - \mathbf{EY}_j)$, whose expectation is $\text{Cov}(X_i, Y_j)$. Because this is true for all (i, j) , the result holds coordinatewise. \square

Warning. Note where the transpose is: Except in the case that n and m are the same integer, $(\mathbf{X} - \mathbf{EX})'(\mathbf{Y} - \mathbf{EY})$ does not even make sense, whereas $(\mathbf{X} - \mathbf{EX})(\mathbf{Y} - \mathbf{EY})'$ is always a random $m \times n$ matrix. \square

An important special case occurs when we have $\mathbf{X} = \mathbf{Y}$. In that case we write

$$\text{Var}(\mathbf{X}) := \text{Cov}(\mathbf{X}, \mathbf{X}).$$

We call $\text{Var}(\mathbf{X})$ the *variance-covariance* matrix of \mathbf{X} . The terminology is motivated by the fact that

$$\text{Var}(\mathbf{X}) = \begin{pmatrix} \text{Var}(X_1) & \text{Cov}(X_1, X_2) & \text{Cov}(X_1, X_3) & \cdots & \text{Cov}(X_1, X_m) \\ \text{Cov}(X_2, X_1) & \text{Var}(X_2) & \text{Cov}(X_2, X_3) & \cdots & \text{Cov}(X_2, X_m) \\ \text{Cov}(X_3, X_1) & \text{Cov}(X_3, X_2) & \text{Var}(X_3) & \cdots & \text{Cov}(X_3, X_m) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \text{Cov}(X_m, X_1) & \text{Cov}(X_m, X_2) & \text{Cov}(X_m, X_3) & \cdots & \text{Var}(X_m) \end{pmatrix}.$$

Note that $\text{Var}(\mathbf{X})$ is *always* a square and symmetric matrix; its dimension is $m \times m$ when \mathbf{X} is m -dimensional. On-diagonal entries of $\text{Var}(\mathbf{X})$ are always nonnegative; off-diagonal entries can be arbitrary real numbers.

3. Mathematical properties of variance and covariance

- Because $(\mathbf{X} - \mathbf{EX})(\mathbf{X} - \mathbf{EX})' = \mathbf{XX}' - \mathbf{X}(\mathbf{EX})' - (\mathbf{EX})\mathbf{X}' + (\mathbf{EX})(\mathbf{EX})'$, it follows that

$$\begin{aligned} \text{Var}(\mathbf{X}) &= \mathbf{E}(\mathbf{XX}') - 2(\mathbf{EX})(\mathbf{EX})' + (\mathbf{EX})(\mathbf{EX})' \\ &= \mathbf{E}(\mathbf{XX}') - (\mathbf{EX})(\mathbf{EX})', \end{aligned}$$

after expansion. This is a multidimensional extension of the formula $\text{Var}(Z) = \mathbf{E}(Z^2) - (\mathbf{EZ})^2$, valid for every [univariate] random variable Z .

- If $\mathbf{a} \in \mathbf{R}^n$ is nonrandom, then $(\mathbf{X} - \mathbf{a}) - \mathbf{E}(\mathbf{X} - \mathbf{a}) = \mathbf{X} - \mathbf{EX}$. Therefore,

$$\text{Var}(\mathbf{X} - \mathbf{a}) = \mathbf{E}[(\mathbf{X} - \mathbf{EX})(\mathbf{X} - \mathbf{EX})'] = \text{Var}(\mathbf{X}).$$

This should be a familiar property in the one-dimensional case.

- If \mathbf{X} , \mathbf{Y} , and \mathbf{Z} are three jointly-distributed random vectors [with the same dimensions], then $\mathbf{X}((\mathbf{Y} + \mathbf{Z}) - \mathbf{E}(\mathbf{Y} + \mathbf{Z}))' = \mathbf{X}(\mathbf{Y} - \mathbf{EY})' + \mathbf{X}(\mathbf{Z} - \mathbf{EZ})'$. Therefore,

$$\text{Cov}(\mathbf{X}, \mathbf{Y} + \mathbf{Z}) = \text{Cov}(\mathbf{X}, \mathbf{Y}) + \text{Cov}(\mathbf{X}, \mathbf{Z}).$$

- Suppose \mathbf{A} , \mathbf{B} are nonrandom matrices. Then, $(\mathbf{AX} - \mathbf{E}(\mathbf{AX}))(\mathbf{BY} - \mathbf{E}(\mathbf{BY}))' = \mathbf{A}(\mathbf{X} - \mathbf{EX})(\mathbf{Y} - \mathbf{EY})'\mathbf{B}'$. Therefore,

$$\text{Cov}(\mathbf{AX}, \mathbf{BY}) = \mathbf{A}\text{Cov}(\mathbf{X}, \mathbf{Y})\mathbf{B}'.$$

The special case that $\mathbf{X} = \mathbf{Y}$ is worth pointing out: In that case we obtain the identity,

$$\text{Var}(\mathbf{AX}) = \mathbf{A}\text{Var}(\mathbf{X})\mathbf{A}'.$$

4. A relation to positive-semidefinite matrices

Let $\mathbf{a} \in \mathbf{R}^n$ be a nonrandom vector and \mathbf{X} be an n -dimensional random vector. Then, the properties of variance-covariance matrices ensure that

$$\text{Var}(\mathbf{a}'\mathbf{X}) = \mathbf{a}'\text{Var}(\mathbf{X})\mathbf{a}.$$

Because $\mathbf{a}'\mathbf{X} = \sum_{j=1}^n a_j X_j$ is univariate, $\text{Var}(\mathbf{a}'\mathbf{X}) \geq 0$, and hence

$$\mathbf{a}'\text{Var}(\mathbf{X})\mathbf{a} \geq 0 \quad \text{for all } \mathbf{a} \in \mathbf{R}^n. \quad (1)$$

A real and symmetric $n \times n$ matrix \mathbf{A} is said to be *positive semidefinite* if $\mathbf{x}'\mathbf{A}\mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbf{R}^n$. And \mathbf{A} is *positive definite* if $\mathbf{x}'\mathbf{A}\mathbf{x} > 0$ for every nonzero $\mathbf{x} \in \mathbf{R}^n$.

Proposition 3. *If \mathbf{X} is an n -dimensional random vector, then $\text{Var}(\mathbf{X})$ is positive semidefinite. If $P\{\mathbf{a}'\mathbf{X} = b\} = 0$ for every $\mathbf{a} \in \mathbf{R}^n$ and $b \in \mathbf{R}$, then $\text{Var}(\mathbf{X})$ is positive definite.*

Proof. We have seen already in (1) that $\text{Var}(\mathbf{X})$ is positive semidefinite. Now suppose that $P\{\mathbf{a}'\mathbf{X} = b\} = 0$, as indicated. Then, $\mathbf{a}'\mathbf{X}$ is a genuine random variable and hence $\mathbf{a}'\text{Var}(\mathbf{X})\mathbf{a} = \text{Var}(\mathbf{a}'\mathbf{X}) > 0$ for all $\mathbf{a} \in \mathbf{R}^n$. \square

Remark 4. The very same argument can be used to prove the following improvement: Suppose $P\{\mathbf{a}'\mathbf{X} = b\} < 1$ for all $b \in \mathbf{R}$ and $\mathbf{a} \in \mathbf{R}^n$. Then $\text{Var}(\mathbf{X})$ is positive definite. The proof is the same because $P\{\mathbf{a}'\mathbf{X} = E(\mathbf{a}'\mathbf{X})\} < 1$ implies that the variance of the random variable $\mathbf{a}'\mathbf{X}$ cannot be zero when $\mathbf{a} \neq \mathbf{0}$. \square

Some linear algebra

Recall the convention that, for us, all vectors are column vectors.

1. Symmetric matrices

Let \mathbf{A} be a real $n \times n$ matrix. Recall that a complex number λ is an *eigenvalue* of \mathbf{A} if there exists a real and nonzero vector \mathbf{x} —called an *eigenvector* for λ —such that $\mathbf{Ax} = \lambda\mathbf{x}$. Whenever \mathbf{x} is an eigenvector for λ , so is $a\mathbf{x}$ for every real number a .

The *characteristic polynomial* $\chi_{\mathbf{A}}$ of matrix \mathbf{A} is the function

$$\chi_{\mathbf{A}}(\lambda) := \det(\lambda\mathbf{I} - \mathbf{A}),$$

defined for all complex numbers λ , where \mathbf{I} denotes the $n \times n$ identity matrix. It is not hard to see that a complex number λ is an eigenvalue of \mathbf{A} if and only if $\chi_{\mathbf{A}}(\lambda) = 0$. We see by direct computation that $\chi_{\mathbf{A}}$ is an n th-order polynomial. Therefore, \mathbf{A} has precisely n eigenvalues, thanks to the fundamental theorem of algebra. We can write them as $\lambda_1, \dots, \lambda_n$, or sometimes more precisely as $\lambda_1(\mathbf{A}), \dots, \lambda_n(\mathbf{A})$.

1. The spectral theorem. The following important theorem is the starting point of our discussion. It might help to recall that vectors $\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbf{R}^n$ are *orthonormal* if $\mathbf{x}_i' \mathbf{x}_j = 0$ when $i \neq j$ and $\mathbf{x}_i' \mathbf{x}_i = \|\mathbf{x}_i\|^2 = 1$.

Theorem 1. *If \mathbf{A} is a real and symmetric $n \times n$ matrix, then $\lambda_1, \dots, \lambda_n$ are real numbers. Moreover, there exist n orthonormal eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ that correspond respectively to $\lambda_1, \dots, \lambda_n$.*

Proof. Let λ be an eigenvalue of $\mathbf{A} := (a_{ij})$, with corresponding eigenvector \mathbf{v} and observe that $\bar{\lambda}$ is an eigenvalue of $\bar{\mathbf{A}} := (\bar{a}_{ij})$, with corresponding

eigenvector $\bar{\mathbf{v}}$. Therefore,

$$\lambda \|\mathbf{v}\|^2 = \lambda \mathbf{v}' \bar{\mathbf{v}} = (\lambda \mathbf{v})' \bar{\mathbf{v}} = (\mathbf{A}\mathbf{v})' \bar{\mathbf{v}} = \mathbf{v}' \mathbf{A}' \bar{\mathbf{v}}.$$

Because \mathbf{A} is symmetric [and real], $\mathbf{A}'\mathbf{v} = \overline{\mathbf{A}\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$. This proves that

$$\lambda \|\mathbf{v}\|^2 = \bar{\lambda} \mathbf{v}' \bar{\mathbf{v}} = \bar{\lambda} \|\mathbf{v}\|^2.$$

Divide by $\|\mathbf{v}\|^2 \neq 0$ to deduce that $\lambda = \bar{\lambda}$, which means that the eigenvalue λ is real valued. Since λ is arbitrary, this proves that all eigenvalues of \mathbf{A} are real. \square

Theorem 2 (The spectral theorem). *Let \mathbf{A} denote a symmetric $n \times n$ matrix with real eigenvalues $\lambda_1, \dots, \lambda_n$ and corresponding orthonormal eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$. Define $\mathbf{D} := \text{diag}(\lambda_1, \dots, \lambda_n)$ to be the diagonal matrix of the λ_i 's and \mathbf{P} to be the matrix whose columns are \mathbf{v}_1 through \mathbf{v}_n respectively; that is,*

$$\mathbf{D} := \begin{pmatrix} \lambda_1 & 0 & 0 & \cdots & 0 \\ 0 & \lambda_2 & 0 & \cdots & 0 \\ 0 & 0 & \lambda_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_n \end{pmatrix}, \quad \mathbf{P} := (\mathbf{v}_1, \dots, \mathbf{v}_n).$$

Then \mathbf{P} is orthogonal [$\mathbf{P}' = \mathbf{P}^{-1}$] and $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1} = \mathbf{P}\mathbf{D}\mathbf{P}'$.

Proof. \mathbf{P} is orthogonal because the orthonormality of the \mathbf{v}_i 's implies that

$$\mathbf{P}'\mathbf{P} = \begin{pmatrix} \mathbf{v}_1 \\ \vdots \\ \mathbf{v}_n \end{pmatrix}' (\mathbf{v}_1, \dots, \mathbf{v}_n) = \mathbf{I}.$$

Furthermore, because $\mathbf{A}\mathbf{v}_j = \lambda_j \mathbf{v}_j$, it follows that $\mathbf{A}\mathbf{P} = \mathbf{P}\mathbf{D}$, which is another way to say that $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$. \square

Recall that the *trace* of an $n \times n$ matrix \mathbf{A} is the sum $A_{1,1} + \cdots + A_{n,n}$ of its diagonal entries.

Corollary 3. *If \mathbf{A} is a real and symmetric $n \times n$ matrix with real eigenvalues $\lambda_1, \dots, \lambda_n$, then*

$$\text{tr}(\mathbf{A}) = \lambda_1 + \cdots + \lambda_n \quad \text{and} \quad \det(\mathbf{A}) = \lambda_1 \times \cdots \times \lambda_n.$$

Proof. Write \mathbf{A} , in spectral form, as $\mathbf{P}\mathbf{D}\mathbf{P}^{-1}$. Since the determinant of \mathbf{P}^{-1} is the reciprocal of that of \mathbf{A} , it follows that $\det(\mathbf{A}) = \det(\mathbf{D})$, which is clearly

$\lambda_1 \times \cdots \times \lambda_n$. In order to compute the trace of \mathbf{A} we compute directly also:

$$\begin{aligned} \operatorname{tr}(\mathbf{A}) &= \sum_{i=1}^n \sum_{j=1}^n P_{ij} (\mathbf{D}\mathbf{P}^{-1})_{ij} = \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n P_{ij} D_{i,k} P_{j,k}^{-1} \\ &= \sum_{i=1}^n \sum_{k=1}^n (\mathbf{P}\mathbf{P}^{-1})_{i,k} D_{i,k} = \sum_{i=1}^n D_{i,i} = \operatorname{tr}(\mathbf{D}), \end{aligned}$$

which is $\lambda_1 + \cdots + \lambda_n$. \square

2. The square-root matrix. Let \mathbf{A} continue to denote a real and symmetric $n \times n$ matrix.

Proposition 4. *There exists a complex and symmetric $n \times n$ matrix \mathbf{B} —called the square root of \mathbf{A} and written as $\mathbf{A}^{1/2}$ or even sometimes as $\sqrt{\mathbf{A}}$ —such that $\mathbf{A} = \mathbf{B}^2 := \mathbf{B}\mathbf{B}$.*

The proof of Proposition 4 is more important than its statement. So let us prove this result.

Proof. Apply the spectral theorem and write $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$. Since \mathbf{D} is a diagonal matrix, its square root can be defined unambiguously as the following complex-valued $n \times n$ diagonal matrix:

$$\mathbf{D}^{1/2} := \begin{pmatrix} \lambda_1^{1/2} & 0 & 0 & \cdots & 0 \\ 0 & \lambda_2^{1/2} & 0 & \cdots & 0 \\ 0 & 0 & \lambda_3^{1/2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda_n^{1/2} \end{pmatrix}.$$

Define $\mathbf{B} := \mathbf{P}\mathbf{D}^{1/2}\mathbf{P}^{-1}$, and note that

$$\mathbf{B}^2 = \mathbf{P}\mathbf{D}^{1/2}\mathbf{P}^{-1}\mathbf{P}\mathbf{D}^{1/2}\mathbf{P}^{-1} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1} = \mathbf{A},$$

since $\mathbf{P}^{-1}\mathbf{P} = \mathbf{I}$ and $(\mathbf{D}^{1/2})^2 = \mathbf{D}$. \square

2. Positive-semidefinite matrices

Recall that an $n \times n$ matrix \mathbf{A} is *positive semidefinite* if it is symmetric and

$$\mathbf{x}'\mathbf{A}\mathbf{x} \geq 0 \quad \text{for all } \mathbf{x} \in \mathbf{R}^n.$$

Recall that \mathbf{A} is *positive definite* if it is symmetric and

$$\mathbf{x}'\mathbf{A}\mathbf{x} > 0 \quad \text{for all nonzero } \mathbf{x} \in \mathbf{R}^n.$$

Theorem 5. *A symmetric matrix \mathbf{A} is positive semidefinite if and only if all of its eigenvalues are ≥ 0 . \mathbf{A} is positive definite if and only if all of its eigenvalues are > 0 . In the latter case, \mathbf{A} is also nonsingular.*

The following is a ready consequence.

Corollary 6. *All of the eigenvalues of a variance-covariance matrix are always ≥ 0 .*

Now let us establish the theorem.

Proof of Theorem 5. Suppose \mathbf{A} is positive semidefinite, and let λ denote one of its eigenvalues, together with corresponding eigenvector \mathbf{x} . Since $0 \leq \mathbf{x}'\mathbf{A}\mathbf{x} = \lambda\|\mathbf{x}\|^2$ and $\|\mathbf{x}\| > 0$, it follows that $\lambda \geq 0$. This proves that all of the eigenvalues of \mathbf{A} are nonnegative. If \mathbf{A} is positive definite, then the same argument shows that all of its eigenvalues are > 0 . Because $\det(\mathbf{A})$ is the product of all n eigenvalues of \mathbf{A} (Corollary 3), it follows that $\det(\mathbf{A}) > 0$, whence \mathbf{A} is nonsingular.

This proves slightly more than half of the proposition. Now let us suppose that all eigenvalues of \mathbf{A} are ≥ 0 . We write \mathbf{A} in spectral form $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}'$, and observe that \mathbf{D} is a diagonal matrix of nonnegative numbers. By virtue of its construction, $\mathbf{A}^{1/2} = \mathbf{P}\mathbf{D}^{1/2}\mathbf{P}'$, and hence for all $\mathbf{x} \in \mathbf{R}^n$,

$$\mathbf{x}'\mathbf{A}\mathbf{x} = (\mathbf{D}^{1/2}\mathbf{P}\mathbf{x})'(\mathbf{D}^{1/2}\mathbf{P}\mathbf{x}) = \|\mathbf{D}^{1/2}\mathbf{P}\mathbf{x}\|^2, \quad (1)$$

which is ≥ 0 . Therefore, \mathbf{A} is positive semidefinite.

If all of the eigenvalues of \mathbf{A} are > 0 , then (1) tells us that

$$\mathbf{x}'\mathbf{A}\mathbf{x} = \|\mathbf{D}^{1/2}\mathbf{P}\mathbf{x}\|^2 = \sum_{j=1}^n \left([\mathbf{D}^{1/2}\mathbf{P}\mathbf{x}]_j \right)^2 = \sum_{j=1}^n \lambda_j ([\mathbf{P}\mathbf{x}]_j)^2, \quad (2)$$

where $\lambda_j > 0$ for all j . Therefore,

$$\mathbf{x}'\mathbf{A}\mathbf{x} \geq \min_{1 \leq j \leq n} \lambda_j \cdot \sum_{j=1}^n ([\mathbf{P}\mathbf{x}]_j)^2 = \min_{1 \leq j \leq n} \lambda_j \cdot \mathbf{x}'\mathbf{P}'\mathbf{P}\mathbf{x} = \min_{1 \leq j \leq n} \lambda_j \cdot \|\mathbf{x}\|^2.$$

Since $\min_{1 \leq j \leq n} \lambda_j > 0$, it follows that $\mathbf{x}'\mathbf{A}\mathbf{x} > 0$ for all nonzero \mathbf{x} . This completes the proof. \square

Let us pause and point out a consequence of the proof of this last result.

Corollary 7. *If \mathbf{A} is positive semidefinite, then its extremal eigenvalues satisfy*

$$\min_{1 \leq j \leq n} \lambda_j = \min_{\|\mathbf{x}\| > 0} \frac{\mathbf{x}'\mathbf{A}\mathbf{x}}{\|\mathbf{x}\|^2}, \quad \max_{1 \leq j \leq n} \lambda_j = \max_{\|\mathbf{x}\| > 0} \frac{\mathbf{x}'\mathbf{A}\mathbf{x}}{\|\mathbf{x}\|^2}.$$

Proof. We saw, during the course of the previous proof, that

$$\min_{1 \leq j \leq n} \lambda_j \cdot \|\mathbf{x}\|^2 \leq \mathbf{x}'\mathbf{A}\mathbf{x} \quad \text{for all } \mathbf{x} \in \mathbf{R}^n. \quad (3)$$

Optimize over all \mathbf{x} to see that

$$\min_{1 \leq j \leq n} \lambda_j \leq \min_{\|\mathbf{x}\| > 0} \frac{\mathbf{x}'\mathbf{A}\mathbf{x}}{\|\mathbf{x}\|^2}. \quad (4)$$

But $\min_{1 \leq j \leq n} \lambda_j$ is an eigenvalue for \mathbf{A} ; let \mathbf{z} denote a corresponding eigenvector in order to see that

$$\min_{1 \leq j \leq n} \lambda_j \leq \min_{\|\mathbf{x}\| > 0} \frac{\mathbf{x}' \mathbf{A} \mathbf{x}}{\|\mathbf{x}\|^2} \leq \frac{\mathbf{z}' \mathbf{A} \mathbf{z}}{\|\mathbf{z}\|^2} = \min_{1 \leq j \leq n} \lambda_j.$$

So both inequalities are in fact equalities, and hence follows the formula for the minimum eigenvalue. The one for the maximum eigenvalue is proved similarly. \square

Finally, a word about the square root of positive semidefinite matrices:

Proposition 8. *If \mathbf{A} is positive semidefinite, then so is $\mathbf{A}^{1/2}$. If \mathbf{A} is positive definite, then so is $\mathbf{A}^{1/2}$.*

Proof. We write, in spectral form, $\mathbf{A} = \mathbf{P} \mathbf{D} \mathbf{P}'$ and observe [by squaring it] that $\mathbf{A}^{1/2} = \mathbf{P} \mathbf{D}^{1/2} \mathbf{P}'$. Note that $\mathbf{D}^{1/2}$ is a real diagonal matrix since the eigenvalues of \mathbf{A} are ≥ 0 . Therefore, we may apply (1) to $\mathbf{A}^{1/2}$ [in place of \mathbf{A}] to see that $\mathbf{x}' \mathbf{A}^{1/2} \mathbf{x} = \|\mathbf{D}^{1/4} \mathbf{P} \mathbf{x}\|^2 \geq 0$ where $\mathbf{D}^{1/4}$ denotes the [real] square root of $\mathbf{D}^{1/2}$. This proves that if \mathbf{A} is positive semidefinite, then so is $\mathbf{A}^{1/2}$. Now suppose there exists a positive definite \mathbf{A} whose square root is not positive definite. It would follow that there necessarily exists a nonzero $\mathbf{x} \in \mathbf{R}^n$ such that $\mathbf{x}' \mathbf{A}^{1/2} \mathbf{x} = \|\mathbf{D}^{1/4} \mathbf{P} \mathbf{x}\|^2 = 0$. Since $\mathbf{D}^{1/4} \mathbf{P} \mathbf{x} = \mathbf{0}$,

$$\mathbf{D}^{1/2} \mathbf{P} \mathbf{x} = \mathbf{D}^{1/4} \mathbf{D}^{1/4} \mathbf{P} \mathbf{x} = \mathbf{0} \quad \Rightarrow \quad \mathbf{x}' \mathbf{A} \mathbf{x} = \|\mathbf{D}^{1/2} \mathbf{P} \mathbf{x}\|^2 = 0.$$

And this contradicts the assumption that \mathbf{A} is positive definite. \square

3. The rank of a matrix

Recall that vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ are *linearly independent* if

$$c_1 \mathbf{v}_1 + \dots + c_k \mathbf{v}_k = \mathbf{0} \quad \Rightarrow \quad c_1 = \dots = c_k = 0.$$

For instance, $\mathbf{v}_1 := (1, 0)'$ and $\mathbf{v}_2 := (0, 1)'$ are linearly independent 2-vectors.

The *column rank* of a matrix \mathbf{A} is the maximum number of linearly independent column vectors of \mathbf{A} . The *row rank* of a matrix \mathbf{A} is the maximum number of linearly independent row vectors of \mathbf{A} . We can interpret these definitions geometrically as follows: First, suppose \mathbf{A} is $m \times n$ and define $\mathcal{C}(\mathbf{A})$ denote the linear space of all vectors of the form $c_1 \mathbf{v}_1 + \dots + c_n \mathbf{v}_n$, where $\mathbf{v}_1, \dots, \mathbf{v}_n$ are the column vectors of \mathbf{A} and c_1, \dots, c_n are real numbers. We call $\mathcal{C}(\mathbf{A})$ the *column space* of \mathbf{A} .

We can define the *row space* $\mathcal{R}(\mathbf{A})$, of \mathbf{A} similarly, or simply define $\mathcal{R}(\mathbf{A}) := \mathcal{C}(\mathbf{A}')$.

Lemma 9. *For every $m \times n$ matrix \mathbf{A} ,*

$$\mathcal{C}(\mathbf{A}) = \{\mathbf{A} \mathbf{x} : \mathbf{x} \in \mathbf{R}^n\}, \quad \mathcal{R}(\mathbf{A}) := \{\mathbf{x}' \mathbf{A} : \mathbf{x} \in \mathbf{R}^m\}.$$

We can think of an $m \times n$ matrix \mathbf{A} as a mapping from \mathbf{R}^n into \mathbf{R}^m ; namely, we can think of matrix \mathbf{A} also as the function $\mathbf{f}_{\mathbf{A}}(\mathbf{x}) := \mathbf{x} \mapsto \mathbf{Ax}$. In this way we see that $\mathcal{C}(\mathbf{A})$ is also the “range” of the function $\mathbf{f}_{\mathbf{A}}$.

Proof. Let us write the columns of \mathbf{A} as $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$. Note that $\mathbf{y} \in \mathcal{C}(\mathbf{A})$ if and only if there exist c_1, \dots, c_n such that $\mathbf{y} = c_1\mathbf{a}_1 + \dots + c_n\mathbf{a}_n = \mathbf{Ac}$, where $\mathbf{c} := (c_1, \dots, c_n)'$. This shows that $\mathcal{C}(\mathbf{A})$ is the collection of all vectors of the form \mathbf{Ax} , for $\mathbf{x} \in \mathbf{R}^n$. The second assertion [about $\mathcal{R}(\mathbf{A})$] follows from the definition of $\mathcal{R}(\mathbf{A})$ equalling $\mathcal{C}(\mathbf{A}')$ and the already-proven first assertion. \square

It then follows, from the definition of dimension, that

$$\text{column rank of } \mathbf{A} = \dim \mathcal{C}(\mathbf{A}), \quad \text{row rank of } \mathbf{A} = \dim \mathcal{R}(\mathbf{A}).$$

Proposition 10. *Given any matrix \mathbf{A} , its row rank and column rank are the same. We write their common value as $\text{rank}(\mathbf{A})$.*

Proof. Suppose \mathbf{A} is $m \times n$ and its column rank is r . Let $\mathbf{b}_1, \dots, \mathbf{b}_r$ denote a basis for $\mathcal{C}(\mathbf{A})$ and consider the matrix $m \times r$ matrix $\mathbf{B} := (\mathbf{b}_1, \dots, \mathbf{b}_r)$. Write \mathbf{A} , columnwise, as $\mathbf{A} := (\mathbf{a}_1, \dots, \mathbf{a}_n)$. For every $1 \leq j \leq n$, there exists $c_{1,j}, \dots, c_{r,j}$ such that $\mathbf{a}_j = c_{1,j}\mathbf{b}_1 + \dots + c_{r,j}\mathbf{b}_r$. Let $\mathbf{C} := (c_{i,j})$ be the resulting $r \times n$ matrix, and note that $\mathbf{A} = \mathbf{BC}$. Because $A_{i,j} = \sum_{k=1}^r B_{i,k}C_{k,j}$, every row of \mathbf{A} is a linear combination of the rows of \mathbf{C} . In other words, $\mathcal{R}(\mathbf{A}) \subseteq \mathcal{R}(\mathbf{C})$ and hence the row rank of \mathbf{A} is $\leq \dim \mathcal{R}(\mathbf{C}) = r =$ the column rank of \mathbf{A} . Apply this fact to \mathbf{A}' to see that also the row rank of \mathbf{A}' is \leq the column rank of \mathbf{A}' ; equivalently that the column rank of \mathbf{A} is \leq the row rank of \mathbf{A} . \square

Proposition 11. *If \mathbf{A} is $n \times m$ and \mathbf{B} is $m \times k$, then*

$$\text{rank}(\mathbf{AB}) \leq \min(\text{rank}(\mathbf{A}), \text{rank}(\mathbf{B})).$$

Proof. The proof uses an idea that we exploited already in the proof of Proposition 10: Since $(\mathbf{AB})_{j,l} = \sum_{\nu=1}^m A_{j,\nu}B_{\nu,l}$, the rows of \mathbf{AB} are linear combinations of the rows of \mathbf{B} ; that is $\mathcal{R}(\mathbf{AB}) \subseteq \mathcal{R}(\mathbf{B})$, whence $\text{rank}(\mathbf{AB}) \leq \text{rank}(\mathbf{B})$. Also, $\mathcal{C}(\mathbf{AB}) \subseteq \mathcal{C}(\mathbf{A})$, whence $\text{rank}(\mathbf{AB}) \leq \text{rank}(\mathbf{A})$. These observations complete the proof. \square

Proposition 12. *If \mathbf{A} and \mathbf{C} are nonsingular, then*

$$\text{rank}(\mathbf{ABC}) = \text{rank}(\mathbf{B}),$$

provided that the dimensions match up so that \mathbf{ABC} makes sense.

Proof. Let $\mathbf{D} := \mathbf{ABC}$; our goal is to show that $\text{rank}(\mathbf{D}) = \text{rank}(\mathbf{B})$.

Two applications of the previous proposition together yield $\text{rank}(\mathbf{D}) \leq \text{rank}(\mathbf{AB}) \leq \text{rank}(\mathbf{B})$. And since $\mathbf{B} = \mathbf{A}^{-1}\mathbf{DC}^{-1}$, we have also $\text{rank}(\mathbf{B}) \leq \text{rank}(\mathbf{A}^{-1}\mathbf{D}) \leq \text{rank}(\mathbf{D})$. \square

Corollary 13. *If \mathbf{A} is an $n \times n$ real and symmetric matrix, then $\text{rank}(\mathbf{A}) =$ the total number of nonzero eigenvalues of \mathbf{A} . In particular, \mathbf{A} has full rank if and only if \mathbf{A} is nonsingular. Finally, $\mathcal{C}(\mathbf{A})$ is the linear space spanned by the eigenvectors of \mathbf{A} that correspond to nonzero eigenvalues.*

Proof. We write \mathbf{A} , in spectral form, as $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1}$, and apply the preceding proposition to see that $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{D})$, which is clearly the total number of nonzero eigenvalue of \mathbf{A} . Since \mathbf{A} is nonsingular if and only if all of its eigenvalues are nonzero, \mathbf{A} has full rank if and only if \mathbf{A} is nonsingular.

Finally, suppose \mathbf{A} has rank $k \leq n$; this is the number of its nonzero eigenvalues $\lambda_1, \dots, \lambda_k$. Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ denote orthonormal eigenvectors such that $\mathbf{v}_1, \dots, \mathbf{v}_k$ are eigenvectors that correspond to $\lambda_1, \dots, \lambda_k$ and $\mathbf{v}_{k+1}, \dots, \mathbf{v}_n$ are eigenvectors that correspond to eigenvalues 0 [Gram–Schmidt]. And define \mathcal{E} to be the span of $\mathbf{v}_1, \dots, \mathbf{v}_k$; i.e.,

$$\mathcal{E} := \{c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k : c_1, \dots, c_k \in \mathbf{R}\}.$$

Our final goal is to prove that $\mathcal{E} = \mathcal{C}(\mathbf{A})$, which we know is equal to the linear space of all vectors of the form $\mathbf{A}\mathbf{x}$.

Clearly, $c_1\mathbf{v}_1 + \dots + c_k\mathbf{v}_k = \mathbf{A}\mathbf{x}$, where $\mathbf{x} = \sum_{j=1}^k (c_j/\lambda_j)\mathbf{v}_j$. Therefore, $\mathcal{E} \subseteq \mathcal{C}(\mathbf{A})$. If $k = n$, then this suffices because in that case $\mathbf{v}_1, \dots, \mathbf{v}_n$ is a basis for \mathbf{R}^n , hence $\mathcal{E} = \mathcal{C}(\mathbf{A}) = \mathbf{R}^n$. If $k < n$, then we can write every $\mathbf{x} \in \mathbf{R}^n$ as $a_1\mathbf{v}_1 + \dots + a_n\mathbf{v}_n$, so that $\mathbf{A}\mathbf{x} = \sum_{j=1}^k a_j\lambda_j\mathbf{v}_j \in \mathcal{E}$. Thus, $\mathcal{C}(\mathbf{A}) \subseteq \mathcal{E}$ and we are done. \square

Let \mathbf{A} be $m \times n$ and define the *null space* [or “kernel”] of \mathbf{A} as

$$\mathcal{N}(\mathbf{A}) := \{\mathbf{x} \in \mathbf{R}^n : \mathbf{A}\mathbf{x} = \mathbf{0}\}.$$

Note that $\mathcal{N}(\mathbf{A})$ is the linear span of the eigenvectors of \mathbf{A} that correspond to eigenvalue 0. The other eigenvectors can be chosen to be orthogonal to these, and hence the preceding proof contains the facts that: (i) Nonzero elements of $\mathcal{N}(\mathbf{A})$ are orthogonal to nonzero elements of $\mathcal{C}(\mathbf{A})$; and (ii)

$$\dim \mathcal{N}(\mathbf{A}) + \text{rank}(\mathbf{A}) = n \quad (= \text{the number of columns of } \mathbf{A}). \quad (5)$$

Proposition 14. $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}'\mathbf{A}) = \text{rank}(\mathbf{A}\mathbf{A}')$ for every $m \times n$ matrix \mathbf{A} .

Proof. If $\mathbf{A}\mathbf{x} = \mathbf{0}$ then $\mathbf{A}'\mathbf{A}\mathbf{x} = \mathbf{0}$, and if $\mathbf{A}'\mathbf{A}\mathbf{x} = \mathbf{0}$, then $\|\mathbf{A}\mathbf{x}\|^2 = \mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} = 0$. In other words, $\mathcal{N}(\mathbf{A}) = \mathcal{N}(\mathbf{A}'\mathbf{A})$. Because $\mathbf{A}'\mathbf{A}$ and \mathbf{A} both have n columns, it follows from (5) that $\text{rank}(\mathbf{A}'\mathbf{A}) = \text{rank}(\mathbf{A})$. Apply this observation to \mathbf{A}' to see that $\text{rank}(\mathbf{A}') = \text{rank}(\mathbf{A}\mathbf{A}')$ as well. The result follows from this and the fact that \mathbf{A} and \mathbf{A}' have the same rank (Proposition 10). \square

4. Projection matrices

A matrix \mathbf{A} is said to be a *projection matrix* if: (i) \mathbf{A} is symmetric; and (ii) \mathbf{A} is “idempotent”; that is, $\mathbf{A}^2 = \mathbf{A}$.

Note that projection matrices are always positive semidefinite. Indeed, $\mathbf{x}'\mathbf{A}\mathbf{x} = \mathbf{x}'\mathbf{A}^2\mathbf{x} = \mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} = \|\mathbf{A}\mathbf{x}\|^2 \geq 0$

Proposition 15. *If \mathbf{A} is an $n \times n$ projection matrix, then so is $\mathbf{I} - \mathbf{A}$. Moreover, all eigenvalues of \mathbf{A} are zeros and ones, and $\text{rank}(\mathbf{A}) = \text{the number of eigenvalues that are equal to one}$.*

Proof. $(\mathbf{I} - \mathbf{A})^2 = \mathbf{I} - 2\mathbf{A} + \mathbf{A}^2 = \mathbf{I} - \mathbf{A}$. Since $\mathbf{I} - \mathbf{A}$ is symmetric also, it is a projection. If λ is an eigenvalue of \mathbf{A} and \mathbf{x} is a corresponding eigenvector, then $\lambda\mathbf{x} = \mathbf{A}\mathbf{x} = \mathbf{A}^2\mathbf{x} = \lambda\mathbf{A}\mathbf{x} = \lambda^2\mathbf{x}$. Multiply both sides by \mathbf{x}' to see that $\lambda\|\mathbf{x}\|^2 = \lambda^2\|\mathbf{x}\|^2$. Since $\|\mathbf{x}\| > 0$, it follows that $\lambda \in \{0, 1\}$. The total number of nonzero eigenvalues is then the total number of eigenvalues that are ones. Therefore, the rank of \mathbf{A} is the total number of eigenvalues that are one. \square

Corollary 16. *If \mathbf{A} is a projection matrix, then $\text{rank}(\mathbf{A}) = \text{tr}(\mathbf{A})$.*

Proof. Simply recall that $\text{tr}(\mathbf{A})$ is the sum of the eigenvalues, which for a projection matrix, is the total number of eigenvalues that are one. \square

Why are they called “projection” matrices? Or, perhaps even more importantly, what is a “projection”?

Lemma 17. *Let Ω denote a linear subspace of \mathbf{R}^n , and $\mathbf{x} \in \mathbf{R}^n$ be fixed. Then there exists a unique element $\mathbf{y} \in \Omega$ that is closest to \mathbf{x} ; that is,*

$$\|\mathbf{y} - \mathbf{x}\| = \min_{\mathbf{z} \in \Omega} \|\mathbf{z} - \mathbf{x}\|.$$

The point \mathbf{y} is called the projection of \mathbf{x} onto Ω .

Proof. Let $k := \dim \Omega$, so that there exists an orthonormal basis $\mathbf{b}_1, \dots, \mathbf{b}_k$ for Ω . Extend this to a basis $\mathbf{b}_1, \dots, \mathbf{b}_n$ for all of \mathbf{R}^n by the Gram–Schmidt method.

Given a fixed vector $\mathbf{x} \in \mathbf{R}^n$, we can write it as $\mathbf{x} := c_1\mathbf{b}_1 + \dots + c_n\mathbf{b}_n$ for some $c_1, \dots, c_n \in \mathbf{R}$. Define $\mathbf{y} := c_1\mathbf{b}_1 + \dots + c_k\mathbf{b}_k$. Clearly, $\mathbf{y} \in \Omega$ and $\|\mathbf{y} - \mathbf{x}\|^2 = \sum_{i=k+1}^n c_i^2$. Any other $\mathbf{z} \in \Omega$ can be written as $\mathbf{z} = \sum_{i=1}^k d_i\mathbf{b}_i$, and hence $\|\mathbf{z} - \mathbf{x}\|^2 = \sum_{i=1}^k (d_i - c_i)^2 + \sum_{i=k+1}^n c_i^2$, which is strictly greater than $\|\mathbf{y} - \mathbf{x}\|^2 = \sum_{i=k+1}^n c_i^2$ unless $d_i = c_i$ for all $i = 1, \dots, k$; i.e., unless $\mathbf{z} = \mathbf{y}$. \square

Usually, we have a k -dimensional linear subspace Ω of \mathbf{R}^n that is the range of some $n \times k$ matrix \mathbf{A} . That is, $\Omega = \{\mathbf{A}\mathbf{y} : \mathbf{y} \in \mathbf{R}^k\}$. Equivalently, $\Omega = \mathcal{C}(\mathbf{A})$. In that case,

$$\min_{\mathbf{z} \in \Omega} \|\mathbf{z} - \mathbf{x}\|^2 = \min_{\mathbf{y} \in \mathbf{R}^k} \|\mathbf{A}\mathbf{y} - \mathbf{x}\|^2 = \min_{\mathbf{y} \in \mathbf{R}^k} [\mathbf{y}'\mathbf{A}'\mathbf{A}\mathbf{y} - \mathbf{y}'\mathbf{A}'\mathbf{x} - \mathbf{x}'\mathbf{A}\mathbf{y} + \mathbf{x}'\mathbf{x}].$$

Because $\mathbf{y}'\mathbf{A}'\mathbf{x}$ is a scalar, the preceding is simplified to

$$\min_{\mathbf{z} \in \Omega} \|\mathbf{z} - \mathbf{x}\|^2 = \min_{\mathbf{y} \in \mathbb{R}^k} [\mathbf{y}'\mathbf{A}'\mathbf{A}\mathbf{y} - 2\mathbf{y}'\mathbf{A}'\mathbf{x} + \mathbf{x}'\mathbf{x}].$$

Suppose that the $k \times k$ positive semidefinite matrix $\mathbf{A}'\mathbf{A}$ is nonsingular [so that $\mathbf{A}'\mathbf{A}$ and hence also $(\mathbf{A}'\mathbf{A})^{-1}$ are both positive definite]. Then, we can relabel variables $[\boldsymbol{\alpha} := \mathbf{A}'\mathbf{A}\mathbf{y}]$ to see that

$$\min_{\mathbf{z} \in \Omega} \|\mathbf{z} - \mathbf{x}\|^2 = \min_{\boldsymbol{\alpha} \in \mathbb{R}^k} [\boldsymbol{\alpha}'(\mathbf{A}'\mathbf{A})^{-1}\boldsymbol{\alpha} - 2\boldsymbol{\alpha}'(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{x} + \mathbf{x}'\mathbf{x}].$$

A little arithmetic shows that

$$\begin{aligned} & (\boldsymbol{\alpha} - \mathbf{A}'\mathbf{x})'(\mathbf{A}'\mathbf{A})^{-1}(\boldsymbol{\alpha} - \mathbf{A}'\mathbf{x}) \\ &= \boldsymbol{\alpha}'(\mathbf{A}'\mathbf{A})^{-1}\boldsymbol{\alpha} - 2\boldsymbol{\alpha}'(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{x} + \mathbf{x}'\mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{x}. \end{aligned}$$

Consequently,

$$\begin{aligned} & \min_{\mathbf{z} \in \Omega} \|\mathbf{z} - \mathbf{x}\|^2 \\ &= \min_{\boldsymbol{\alpha} \in \mathbb{R}^k} [(\boldsymbol{\alpha} - \mathbf{A}'\mathbf{x})'(\mathbf{A}'\mathbf{A})^{-1}(\boldsymbol{\alpha} - \mathbf{A}'\mathbf{x}) - \mathbf{x}'\mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{x} + \mathbf{x}'\mathbf{x}]. \end{aligned}$$

The first term in the parentheses is ≥ 0 ; in fact it is > 0 unless we select $\boldsymbol{\alpha} = \mathbf{A}'\mathbf{x}$. This proves that the projection of \mathbf{x} onto Ω is obtained by setting $\boldsymbol{\alpha} := \mathbf{A}'\mathbf{x}$, in which case the projection itself is $\mathbf{A}\mathbf{y} = \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{x}$ and the distance between \mathbf{y} and \mathbf{x} is the square root of $\|\mathbf{x}\|^2 - \mathbf{x}'\mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'\mathbf{x}$.

Let $\mathbf{P}_\Omega := \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'$. It is easy to see that \mathbf{P}_Ω is a projection matrix. The preceding shows that $\mathbf{P}_\Omega\mathbf{x}$ is the projection of \mathbf{x} onto Ω for every $\mathbf{x} \in \mathbb{R}^n$. That is, we can think of \mathbf{P}_Ω as the matrix that projects onto Ω . Moreover, the distance between \mathbf{x} and the linear subspace Ω [i.e., $\min_{\mathbf{z} \in \mathbb{R}^k} \|\mathbf{z} - \mathbf{x}\|$] is exactly the square root of $\mathbf{x}'\mathbf{x} - \mathbf{x}'\mathbf{P}_\Omega\mathbf{x} = \mathbf{x}'(\mathbf{I} - \mathbf{P}_\Omega)\mathbf{x} = \|(\mathbf{I} - \mathbf{P}_\Omega)\mathbf{x}\|^2$, because $\mathbf{I} - \mathbf{P}_\Omega$ is a projection matrix. What space does it project into?

Let Ω^\perp denote the collection of all n -vectors that are perpendicular to every element of Ω . If $\mathbf{z} \in \Omega^\perp$, then we can write, for all $\mathbf{x} \in \mathbb{R}^n$,

$$\begin{aligned} \|\mathbf{z} - \mathbf{x}\|^2 &= \|\mathbf{z} - (\mathbf{I} - \mathbf{P}_\Omega)\mathbf{x} + \mathbf{P}_\Omega\mathbf{x}\|^2 \\ &= \|\mathbf{z} - (\mathbf{I} - \mathbf{P}_\Omega)\mathbf{x}\|^2 + \|\mathbf{P}_\Omega\mathbf{x}\|^2 - 2\{\mathbf{z} - (\mathbf{I} - \mathbf{P}_\Omega)\mathbf{x}\}'\mathbf{P}_\Omega\mathbf{x} \\ &= \|\mathbf{z} - (\mathbf{I} - \mathbf{P}_\Omega)\mathbf{x}\|^2 + \|\mathbf{P}_\Omega\mathbf{x}\|^2, \end{aligned}$$

since \mathbf{z} is orthogonal to every element of Ω including $\mathbf{P}_\Omega\mathbf{x}$, and $\mathbf{P}_\Omega = \mathbf{P}_\Omega^2$. Take the minimum over all $\mathbf{z} \in \Omega^\perp$ to find that $\mathbf{I} - \mathbf{P}_\Omega$ is the projection onto Ω^\perp . Let us summarize our findings.

Proposition 18. *If $\mathbf{A}'\mathbf{A}$ is nonsingular [equivalently, has full rank], then $\mathbf{P}_{\mathcal{C}(\mathbf{A})} := \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'$ is the projection onto $\mathcal{C}(\mathbf{A})$, $\mathbf{I} - \mathbf{P}_{\mathcal{C}(\mathbf{A})} = \mathbf{P}_{\mathcal{C}(\mathbf{A})^\perp}$ is the projection onto Ω^\perp , and we have*

$$\mathbf{x} = \mathbf{P}_{\mathcal{C}(\mathbf{A})}\mathbf{x} + \mathbf{P}_{\mathcal{C}(\mathbf{A})^\perp}\mathbf{x}, \quad \text{and} \quad \|\mathbf{x}\|^2 = \|\mathbf{P}_{\mathcal{C}(\mathbf{A})}\mathbf{x}\|^2 + \|\mathbf{P}_{\mathcal{C}(\mathbf{A})^\perp}\mathbf{x}\|^2.$$

The last result is called the “Pythagorean property.”

Quadratic forms

Let \mathbf{A} be a real and symmetric $n \times n$ matrix. Then the *quadratic form* associated to \mathbf{A} is the function $Q_{\mathbf{A}}$ defined by

$$Q_{\mathbf{A}}(\mathbf{x}) := \mathbf{x}'\mathbf{A}\mathbf{x} \quad (\mathbf{x} \in \mathbf{R}^n).$$

We have seen quadratic forms already, particularly in the context of positive-semidefinite matrices.

1. Random quadratic forms

Let $\mathbf{X} := (X_1, \dots, X_n)'$ be an n -dimensional random vector. We are interested in the random quadratic form $Q_{\mathbf{A}}(\mathbf{X}) := \mathbf{X}'\mathbf{A}\mathbf{X}$.

Proposition 1. If $E\mathbf{X} := \boldsymbol{\mu}$ and $\text{Var}(\mathbf{X}) := \boldsymbol{\Sigma}$, then

$$E(\mathbf{X}'\mathbf{A}\mathbf{X}) = \text{tr}(\mathbf{A}\boldsymbol{\Sigma}) + \boldsymbol{\mu}'\mathbf{A}\boldsymbol{\mu}.$$

In symbols, $E(Q_{\mathbf{A}}(\mathbf{X})) = \text{tr}(\mathbf{A}\boldsymbol{\Sigma}) + Q_{\mathbf{A}}(\boldsymbol{\mu})$.

Proof. We can write

$$\begin{aligned} \mathbf{X}'\mathbf{A}\mathbf{X} &= (\mathbf{X} - \boldsymbol{\mu})'\mathbf{A}\mathbf{X} + \boldsymbol{\mu}'\mathbf{A}\mathbf{X} \\ &= (\mathbf{X} - \boldsymbol{\mu})'\mathbf{A}(\mathbf{X} - \boldsymbol{\mu}) + \boldsymbol{\mu}'\mathbf{A}\mathbf{X} + (\mathbf{X} - \boldsymbol{\mu})'\mathbf{A}\boldsymbol{\mu}. \end{aligned}$$

If we take expectations, then the last term vanishes and we obtain

$$E(\mathbf{X}'\mathbf{A}\mathbf{X}) = E[(\mathbf{X} - \boldsymbol{\mu})'\mathbf{A}(\mathbf{X} - \boldsymbol{\mu})] + \boldsymbol{\mu}'\mathbf{A}\boldsymbol{\mu}.$$

It suffices to verify that the expectation on the right-hand side is the trace of $\mathbf{A}\boldsymbol{\Sigma}$. But this is a direct calculation: Let $Y_j := X_j - \mu_j$, so that $\mathbf{Y} = \mathbf{X} - \boldsymbol{\mu}$

and hence

$$\begin{aligned}
E[(\mathbf{X} - \boldsymbol{\mu})' \mathbf{A} (\mathbf{X} - \boldsymbol{\mu})] &= E(\mathbf{Y}' \mathbf{A} \mathbf{Y}) \\
&= \sum_{i=1}^n \sum_{j=1}^n E(Y_i A_{i,j} Y_j) = \sum_{i=1}^n \sum_{j=1}^n A_{i,j} [\text{Var}(\mathbf{Y})]_{i,j} \\
&= \sum_{i=1}^n \sum_{j=1}^n A_{i,j} [\text{Var}(\mathbf{X} - \boldsymbol{\mu})]_{i,j} = \sum_{i=1}^n \sum_{j=1}^n A_{i,j} [\text{Var}(\mathbf{X})]_{i,j} \\
&= \sum_{i=1}^n \sum_{j=1}^n A_{i,j} \Sigma_{i,j} = \sum_{i=1}^n \sum_{i=1}^n A_{i,j} \Sigma_{j,i} \\
&= \sum_{i=1}^n [\mathbf{A} \boldsymbol{\Sigma}]_{i,i} = \text{tr}(\mathbf{A} \boldsymbol{\Sigma}),
\end{aligned}$$

as desired. \square

We easily get the following by a relabeling ($\mathbf{X} \Leftrightarrow \mathbf{X} - \mathbf{b}$):

Corollary 2. For every nonrandom $\mathbf{b} \in \mathbf{R}^n$,

$$E[(\mathbf{X} - \mathbf{b})' \mathbf{A} (\mathbf{X} - \mathbf{b})] = \text{tr}(\mathbf{A} \boldsymbol{\Sigma}) + (\boldsymbol{\mu} - \mathbf{b})' \mathbf{A} (\boldsymbol{\mu} - \mathbf{b}).$$

In particular, $E[(\mathbf{X} - \boldsymbol{\mu})' \mathbf{A} (\mathbf{X} - \boldsymbol{\mu})] = \text{tr}(\mathbf{A} \boldsymbol{\Sigma})$.

2. Examples of quadratic forms

What do quadratic forms look like? It is best to proceed by example.

Example 3. If $\mathbf{A} := \mathbf{I}_{n \times n}$, then $Q_{\mathbf{A}}(\mathbf{x}) = \sum_{i=1}^n x_i^2$. Because $\text{tr}(\mathbf{A} \boldsymbol{\Sigma}) = \text{tr}(\boldsymbol{\Sigma}) = \sum_{i=1}^n \text{Var}(X_i)$, it follows that

$$E\left(\sum_{i=1}^n X_i^2\right) = \sum_{i=1}^n \text{Var}(X_i) + \left(\sum_{i=1}^n \mu_i^2\right).$$

This ought to be a familiar formula. \square

Example 4. If

$$\mathbf{A} := \mathbf{1}_{m \times m} := \begin{pmatrix} 1 & \cdots & 1 \\ \vdots & & \vdots \\ 1 & \cdots & 1 \end{pmatrix}_{m \times m},$$

then $Q_{\mathbf{A}}(\mathbf{x}) = (\sum_{i=1}^n x_i)^2$. Note that

$$\text{tr}(\mathbf{A} \boldsymbol{\Sigma}) = \sum_{i=1}^n \sum_{j=1}^n A_{i,j} \Sigma_{i,j} = \sum_{i=1}^n \sum_{j=1}^n \text{Cov}(X_i, X_j).$$

Therefore,

$$\mathbb{E} \left[\left(\sum_{i=1}^n X_i \right)^2 \right] = \sum_{i=1}^n \sum_{j=1}^n \text{Cov}(X_i, X_j) + \left(\sum_{i=1}^n \mu_i \right)^2 ;$$

this is another familiar formula. \square

Example 5. One can combine matrices in a natural way to obtain new quadratic forms from old ones. Namely, if $a, b \in \mathbf{R}$ and \mathbf{A} and \mathbf{B} are real and symmetric $n \times n$ matrices, then $Q_{a\mathbf{A}+b\mathbf{B}}(\mathbf{x}) = aQ_{\mathbf{A}}(\mathbf{x}) + bQ_{\mathbf{B}}(\mathbf{x})$. For instance, suppose $\mathbf{A} := \mathbf{I}_{n \times n}$ and $\mathbf{B} := \mathbf{1}_{n \times n}$. Then,

$$a\mathbf{A} + b\mathbf{B} = \begin{pmatrix} a+b & b & b & \cdots & b \\ b & a+b & b & \cdots & b \\ b & b & a+b & \cdots & b \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b & b & b & \cdots & a+b \end{pmatrix},$$

and, thanks to the preceding two examples,

$$Q_{a\mathbf{A}+b\mathbf{B}}(\mathbf{x}) = a \sum_{i=1}^n x_i^2 + b \left(\sum_{i=1}^n x_i \right)^2.$$

An important special case is when $a := 1$ and $b := -1/n$. In that case,

$$\mathbf{A} - \frac{1}{n}\mathbf{B} = \begin{pmatrix} 1-1/n & -1/n & -1/n & \cdots & -1/n \\ -1/n & 1-1/n & -1/n & \cdots & -1/n \\ -1/n & -1/n & 1-1/n & \cdots & -1/n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1/n & -1/n & -1/n & \cdots & 1-1/n \end{pmatrix},$$

and

$$Q_{\mathbf{A}-(1/n)\mathbf{B}}(\mathbf{x}) = \sum_{i=1}^n x_i^2 - \frac{1}{n} \left(\sum_{i=1}^n x_i \right)^2 = \sum_{i=1}^n (x_i - \bar{x})^2.$$

Note that

$$\text{tr}(\mathbf{A}\Sigma) = \sum_{i=1}^n \text{Var}(X_i) + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \text{Cov}(X_i, X_j).$$

Consider the special case that the X_i 's are uncorrelated. In that case, $\text{tr}(\mathbf{A}\Sigma) = (1 - 1/n) \sum_{i=1}^n \text{Var}(X_i)$, and hence

$$\mathbb{E} \left[\sum_{i=1}^n (X_i - \bar{X})^2 \right] = (1 - 1/n) \sum_{i=1}^n \text{Var}(X_i) + \sum_{i=1}^n (\mu_i - \bar{\mu})^2.$$

When the X_i 's are i.i.d. this yields $\mathbb{E} \sum_{i=1}^n (X_i - \bar{X})^2 = (n-1)\text{Var}(X_1)$, which is a formula that you have seen in the context of the unbiasedness of the sample variance estimator $S^2 := (n-1)^{-1} \sum_{i=1}^n (X_i - \bar{X})^2$. \square

Example 6. Consider a symmetric matrix of the form

$$\mathbf{A} := \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 1 & 0 & 1 & \cdots & 0 & 0 \\ 0 & 1 & 0 & & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & 0 & \cdots & 1 & 0 \end{pmatrix}.$$

That is, the super- and sub-diagonal entries are all ones and all other entries are zeros. Then,

$$Q_{\mathbf{A}}(\mathbf{x}) = 2 \sum_{i=1}^{n-2} x_i x_{i+2}.$$

Other examples can be constructed in this way as well, and by also combining such examples. \square

3. The variance of a random quadratic form

In the previous section we computed the expectation of $\mathbf{X}'\mathbf{A}\mathbf{X}$ where \mathbf{X} is a random vector. Here let us say a few things about the variance of the same random vector, under some conditions on \mathbf{X} .

Proposition 7. Suppose $\mathbf{X} := (X_1, \dots, X_n)'$ where the X_j 's are i.i.d. with mean zero and four finite moments. Then,

$$\text{Var}(\mathbf{X}'\mathbf{A}\mathbf{X}) = (\mu_4 - 3\mu_2^2) \sum_{i=1}^n A_{i,i}^2 + (\mu_2^2 - 1) (\text{tr}(\mathbf{A}))^2 + 2\mu_2^2 \text{tr}(\mathbf{A}^2),$$

where $\mu_2 := E(X_1^2)$ and $\mu_4 := E(X_1^4)$.

One can generalize this a little more as well, with more or less the same set of techniques, in order to compute the variance of $\mathbf{X}'\mathbf{A}\mathbf{X}$ in the case that the X_i 's are independent, with common first four moments, and not necessarily mean zero.

Proof. Suppose $\mathbf{X} := (X_1, \dots, X_n)'$, where X_1, \dots, X_n are independent and mean zero. Suppose $\mu_2 := E(X_i^2)$ and $\mu_4 := E(X_i^4)$ do not depend on i [e.g., because the X_j 's are independent]. Then we can write

$$(\mathbf{X}'\mathbf{A}\mathbf{X})^2 = \sum_{1 \leq i,j,k,\ell \leq n} \sum_{1 \leq i,j,k,\ell \leq n} A_{i,j} A_{k,\ell} X_i X_j X_k X_\ell.$$

Note that

$$\mathbb{E}(X_i X_j X_k X_\ell) = \begin{cases} \mu_4 & \text{if } i = j = k = \ell, \\ \mu_2^2 & \text{if } i = j \neq k = \ell \text{ or} \\ & \text{if } i = k \neq j = \ell \text{ or} \\ & \text{if } i = \ell \neq k = j, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore,

$$\begin{aligned} \mathbb{E}[(\mathbf{X}'\mathbf{A}\mathbf{X})^2] &= \sum_{i=1}^n A_{i,i}^2 \mu_4 + \sum_{1 \leq i \neq k \leq n} A_{i,i} A_{k,k} \mu_2^2 + \sum_{1 \leq i \neq j \leq n} A_{i,j} A_{j,i} \mu_2^2 + \sum_{1 \leq i \neq k \leq n} A_{i,k} A_{k,i} \mu_2^2 \\ &= \mu_4 \sum_{i=1}^n A_{i,i}^2 + \mu_2^2 \left[\sum_{1 \leq i \neq k \leq n} A_{i,i} A_{k,k} + 2 \sum_{1 \leq i \neq j \leq n} A_{i,j}^2 \right]. \end{aligned}$$

Next, we identify the double sums in turn:

$$\begin{aligned} \sum_{1 \leq i \neq k \leq n} A_{i,i} A_{k,k} &= \sum_{i=1}^n A_{i,i} \sum_{k=1}^n A_{k,k} - \sum_{i=1}^n A_{i,i}^2 = (\text{tr}(\mathbf{A}))^2 - \sum_{i=1}^n A_{i,i}^2, \\ \sum_{1 \leq i \neq j \leq n} A_{i,j}^2 &= \sum_{i=1}^n \sum_{j=1}^n A_{i,j}^2 - \sum_{i=1}^n A_{i,i}^2 = \sum_{i=1}^n \sum_{j=1}^n A_{i,j} A_{j,i} - \sum_{i=1}^n A_{i,i}^2 \\ &= \sum_{i=1}^n (A^2)_{i,i} - \sum_{i=1}^n A_{i,i}^2 = \text{tr}(\mathbf{A}^2) - \sum_{i=1}^n A_{i,i}^2. \end{aligned}$$

Consequently,

$$\begin{aligned} \mathbb{E}[(\mathbf{X}'\mathbf{A}\mathbf{X})^2] &= \mu_4 \sum_{i=1}^n A_{i,i}^2 + \mu_2^2 \left[(\text{tr}(\mathbf{A}))^2 - \sum_{i=1}^n A_{i,i}^2 + 2\text{tr}(\mathbf{A}^2) - 2 \sum_{i=1}^n A_{i,i}^2 \right] \\ &= (\mu_4 - 3\mu_2^2) \sum_{i=1}^n A_{i,i}^2 + \mu_2^2 \left[(\text{tr}(\mathbf{A}))^2 + 2\text{tr}(\mathbf{A}^2) \right]. \end{aligned}$$

Therefore, in this case,

$$\text{Var}(\mathbf{X}'\mathbf{A}\mathbf{X}) = (\mu_4 - 3\mu_2^2) \sum_{i=1}^n A_{i,i}^2 + \mu_2^2 \left[(\text{tr}(\mathbf{A}))^2 + 2\text{tr}(\mathbf{A}^2) \right] - [\mathbb{E}(\mathbf{X}'\mathbf{A}\mathbf{X})]^2.$$

This proves the result because $\mathbb{E}(\mathbf{X}'\mathbf{A}\mathbf{X}) = \text{tr}(\mathbf{A})$. \square

Moment-generating functions and independence

Let $\mathbf{X} := (X_1, \dots, X_n)'$ be a random vector. Its *moment generating function* [written MGF for short] $M_{\mathbf{X}}$ is defined as

$$M_{\mathbf{X}}(\mathbf{t}) := \mathbb{E}e^{\mathbf{t}'\mathbf{X}} \quad (\mathbf{t} \in \mathbf{R}^n).$$

It is the case that $M_{\mathbf{X}}(\mathbf{t})$ is a well-defined quantity, but it might be infinite for some, and even all, values of $\mathbf{t} \in \mathbf{R}^n$. The following is a hard fact from classical analysis:

Theorem 1 (Uniqueness theorem of MGFs). *Suppose there exists $t > 0$ such that $M_{\mathbf{X}}(\mathbf{t}) < \infty$ for all $\mathbf{t} \in \mathbf{R}^n$ with $\|\mathbf{t}\| \leq r$. Then, the distribution of \mathbf{X} is determined uniquely by the function $M_{\mathbf{X}}$. That is, if \mathbf{Y} is any random vector whose MGF is the same as $M_{\mathbf{X}}$, then \mathbf{Y} has the same distribution as \mathbf{X} .*

We are interested in examples, and primarily those that involve normal distributions in one form or another.

Example 2. If $X \sim N(\mu, \sigma^2)$, then

$$\begin{aligned} M_X(t) &= Ee^{tX} = \int_{-\infty}^{\infty} e^{tx} \frac{e^{-(x-\mu)^2/(2\sigma^2)}}{\sqrt{2\pi\sigma^2}} dx \\ &= \int_{-\infty}^{\infty} e^{t(\sigma y + \mu)} \frac{e^{-y^2/2}}{\sqrt{2\pi}} dy \quad (y := (x - \mu)/\sigma) \\ &= e^{t\mu} \int_{-\infty}^{\infty} \frac{e^{t\sigma y - y^2/2}}{\sqrt{2\pi}} dy = e^{t\mu} \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2}[y^2 - 2yt\sigma]}}{\sqrt{2\pi}} dy. \end{aligned}$$

We complete the square $[y^2 - 2yt\sigma = (y - t\sigma)^2 - (t\sigma)^2]$ in order to see that

$$M_X(t) = \exp\left(t\mu + \frac{t^2\sigma^2}{2}\right).$$

Therefore, the uniqueness theorem (Theorem 1) tells us that any random variable Y whose MGF is $M_Y(t) = \exp(t\mu + \frac{1}{2}t^2\sigma^2)$ is distributed according to $N(\mu, \sigma^2)$. \square

Example 3 (MGF of a simple multivariable normal). Suppose $X_i \sim N(\mu_i, \sigma_i^2)$ ($1 \leq i \leq n$) are independent. Then, the MGF of $\mathbf{X} := (X_1, \dots, X_n)'$ is

$$M_{\mathbf{X}}(\mathbf{t}) = Ee^{\mathbf{t}'\mathbf{X}} = \prod_{j=1}^n Ee^{t_j X_j} = \prod_{j=1}^n e^{t_j \mu_j + \frac{1}{2} t_j^2 \sigma_j^2} = \exp\left(\sum_{j=1}^n t_j \mu_j + \frac{1}{2} \sum_{j=1}^n t_j^2 \sigma_j^2\right)$$

for all $\mathbf{t} \in \mathbf{R}^n$. \square

Example 4 (MGF of χ_1^2). If X is standard normal, then $Y := X^2 \sim \chi_1^2$. Now

$$\begin{aligned} M_Y(t) &= Ee^{tX^2} = \int_{-\infty}^{\infty} \frac{e^{tx^2 - x^2/2}}{\sqrt{2\pi}} dx \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2}(1 - 2t)x^2\right) dx. \end{aligned}$$

If $t \geq 1/2$, then the preceding is infinite. Otherwise, a change of variables $[y = \sqrt{1 - 2t}x]$ tells us that it is equal to

$$\int_{-\infty}^{\infty} \frac{e^{-y^2/2}}{\sqrt{2\pi}} \frac{dy}{\sqrt{1 - 2t}} = \frac{1}{\sqrt{1 - 2t}}.$$

In other words, $M_{X^2}(t) = \infty$ if $t \geq 1/2$ and $M_{X^2}(t) = (1 - 2t)^{-1/2}$ if $t < 1/2$. \square

Example 5 (MGF of χ_n^2). Let $X_1, \dots, X_n \sim N(0, 1)$ be independent, and consider the χ_n^2 random variable $Y := \sum_{i=1}^n X_i^2$. Its MGF is

$$M_Y(t) = \prod_{j=1}^n M_{X_j^2}(t) = \begin{cases} (1 - 2t)^{-n/2} & \text{if } t < 1/2, \\ \infty & \text{if } t \geq 1/2. \end{cases}$$

According to Theorem 1, this is a formula for the MGF of the χ_n^2 distribution, and identifies that distribution uniquely. \square

Theorem 6 (Independence theorem of MGFs). *Let \mathbf{X} be a random n -vector with a MGF that is finite in an open neighborhood of the origin $\mathbf{0} \in \mathbf{R}^n$. Suppose there exists $r = 1, \dots, n$ such that*

$$M_{\mathbf{X}}(\mathbf{t}) = M_{\mathbf{X}}(t_1, \dots, t_r, 0, \dots, 0) \cdot M_{\mathbf{X}}(0, \dots, 0, t_{r+1}, \dots, t_n)$$

for all $\mathbf{t} \in \mathbf{R}^n$. Then, (X_1, \dots, X_r) and (X_{r+1}, \dots, X_n) are independent.

Proof. Let $\tilde{\mathbf{X}}$ denote an independent copy of \mathbf{X} . Define a new random vector \mathbf{Y} as follows:

$$\mathbf{Y} := \begin{pmatrix} X_1 \\ \vdots \\ X_r \\ \tilde{X}_{r+1} \\ \vdots \\ \tilde{X}_n \end{pmatrix}.$$

Then,

$$M_{\mathbf{Y}}(\mathbf{t}) = M_{\mathbf{X}}(t_1, \dots, t_r, 0, \dots, 0) \cdot M_{\mathbf{X}}(0, \dots, 0, t_{r+1}, \dots, t_n).$$

According to the condition of this theorem, \mathbf{X} and \mathbf{Y} have the same MGF's, and therefore they have the same distribution (Theorem 1). That is, for all sets A_1, \dots, A_n ,

$$P\{X_1 \in A_1, \dots, X_n \in A_n\} = P\{Y_1 \in A_1, \dots, Y_n \in A_n\},$$

which is, by construction equal to

$$P\{X_1 \in A_1, \dots, X_r \in A_r\} \cdot P\{\tilde{X}_{r+1} \in A_{r+1}, \dots, \tilde{X}_n \in A_n\}.$$

Since $\tilde{\mathbf{X}}$ has the same distribution as \mathbf{X} , this proves that

$$P\{X_1 \in A_1, \dots, X_n \in A_n\} = P\{X_1 \in A_1, \dots, X_r \in A_r\} \cdot P\{X_{r+1} \in A_{r+1}, \dots, X_n \in A_n\},$$

which has the desired result. \square

Gaussian Random Vectors

1. The multivariate normal distribution

Let $\mathbf{X} := (X_1, \dots, X_n)'$ be a random vector. We say that \mathbf{X} is a *Gaussian random vector* if we can write

$$\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z},$$

where $\boldsymbol{\mu} \in \mathbf{R}^n$, \mathbf{A} is an $n \times k$ matrix and $\mathbf{Z} := (Z_1, \dots, Z_k)'$ is a k -vector of i.i.d. standard normal random variables.

Proposition 1. *Let \mathbf{X} be a Gaussian random vector, as above. Then,*

$$\mathbf{E}\mathbf{X} = \boldsymbol{\mu}, \text{Var}(\mathbf{X}) := \boldsymbol{\Sigma} = \mathbf{A}\mathbf{A}', \text{ and } M_{\mathbf{X}}(\mathbf{t}) = e^{\mathbf{t}'\boldsymbol{\mu} + \frac{1}{2}\mathbf{t}'\boldsymbol{\Sigma}\mathbf{t}} = e^{\mathbf{t}'\boldsymbol{\mu} + \frac{1}{2}\mathbf{t}'\boldsymbol{\Sigma}\mathbf{t}},$$

for all $\mathbf{t} \in \mathbf{R}^n$.

Thanks to the uniqueness theorem of MGF's it follows that the distribution of \mathbf{X} is determined by $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$, and the fact that it is multivariate normal. From now on, we sometimes write $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, when we mean that $M_{\mathbf{X}}(\mathbf{t}) = \exp(\mathbf{t}'\boldsymbol{\mu} + \frac{1}{2}\mathbf{t}'\boldsymbol{\Sigma}\mathbf{t})$. Interestingly enough, the choice of \mathbf{A} and \mathbf{Z} are typically not unique; only $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ influences the distribution of \mathbf{X} .

Proof. The expectation of \mathbf{X} is $\boldsymbol{\mu}$, since $\mathbf{E}(\mathbf{A}\mathbf{Z}) = \mathbf{A}\mathbf{E}(\mathbf{Z}) = \mathbf{0}$. Also,

$$\mathbf{E}(\mathbf{X}\mathbf{X}') = \mathbf{E}([\boldsymbol{\mu} + \mathbf{A}\mathbf{Z}][\boldsymbol{\mu} + \mathbf{A}\mathbf{Z}]') = \boldsymbol{\mu}\boldsymbol{\mu}' + \mathbf{A}\mathbf{E}(\mathbf{Z}\mathbf{Z}')\mathbf{A}'.$$

Since $\mathbf{E}(\mathbf{Z}\mathbf{Z}') = \mathbf{I}$, the variance-covariance of \mathbf{X} is $\mathbf{E}(\mathbf{X}\mathbf{X}') - (\mathbf{E}\mathbf{X})(\mathbf{E}\mathbf{X})' = \mathbf{E}(\mathbf{X}\mathbf{X}') - \boldsymbol{\mu}\boldsymbol{\mu}' = \mathbf{A}\mathbf{A}'$, as desired. Finally, note that $M_{\mathbf{X}}(\mathbf{t}) = \exp(\mathbf{t}'\boldsymbol{\mu}) \cdot M_{\mathbf{Z}}(\mathbf{A}'\mathbf{t})$. This establishes the result on the MGF of \mathbf{X} , since $M_{\mathbf{Z}}(\mathbf{s}) = \prod_{j=1}^n \exp(s_j^2/2) = \exp(\frac{1}{2}\|\mathbf{s}\|^2)$ for all $\mathbf{s} \in \mathbf{R}^n$. \square

We say that \mathbf{X} has the *multivariate normal distribution* with parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma} := \mathbf{A}\mathbf{A}'$, and write this as $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \mathbf{A}\mathbf{A}')$.

Theorem 2. $\mathbf{X} := (X_1, \dots, X_n)'$ has a multivariate normal distribution if and only if $\mathbf{t}'\mathbf{X} = \sum_{i=1}^n t_i X_i$ has a normal distribution on the line for every $\mathbf{t} \in \mathbf{R}^n$. That is, X_1, \dots, X_n are jointly normally distributed if and only if all of their linear combinations are normally distributed.

Note that the distribution of \mathbf{X} depends on \mathbf{A} only through the positive semidefinite $n \times n$ matrix $\boldsymbol{\Sigma} := \mathbf{A}\mathbf{A}'$. Sometimes we say also that X_1, \dots, X_n are *jointly normal* [or *Gaussian*] when $\mathbf{X} := (X_1, \dots, X_n)'$ has a multivariate normal distribution.

Proof. If $\mathbf{X} \in N_n(\boldsymbol{\mu}, \mathbf{A}\mathbf{A}')$ then we can write it as $\mathbf{X} = \boldsymbol{\mu} + \mathbf{A}\mathbf{Z}$, we as before. In that case, $\mathbf{t}'\mathbf{X} = \mathbf{t}'\boldsymbol{\mu} + \mathbf{t}'\mathbf{A}\mathbf{Z}$ is a linear combination of Z_1, \dots, Z_k , whence has a normal distribution with mean $t_1\mu_1 + \dots + t_n\mu_n$ and variance $\mathbf{t}'\mathbf{A}\mathbf{A}'\mathbf{t} = \|\mathbf{A}'\mathbf{t}\|^2$.

For the converse, suppose that $\mathbf{t}'\mathbf{X}$ has a normal distribution for every $\mathbf{t} \in \mathbf{R}^n$. Let $\boldsymbol{\mu} := E\mathbf{X}$ and $\boldsymbol{\Sigma} := \text{Var}(\mathbf{X})$, and observe that $\mathbf{t}'\mathbf{X}$ has mean vector $\mathbf{t}'\boldsymbol{\mu}$ and variance-covariance matrix $\mathbf{t}'\boldsymbol{\Sigma}\mathbf{t}$. Therefore, the MGF of the univariate normal $\mathbf{t}'\mathbf{X}$ is $M_{\mathbf{t}'\mathbf{X}}(s) = \exp(s\mathbf{t}'\boldsymbol{\mu} + \frac{1}{2}s^2\mathbf{t}'\boldsymbol{\Sigma}\mathbf{t})$ for all $s \in \mathbf{R}$. Note that $M_{\mathbf{t}'\mathbf{X}}(s) = E \exp(s\mathbf{t}'\mathbf{X})$. Therefore, apply this with $s := 1$ to see that $M_{\mathbf{t}'\mathbf{X}}(1) = M_{\mathbf{X}}(\mathbf{t})$ is the MGF of a multivariate normal. The uniqueness theorem for MGF's (Theorem 1, p. 27) implies the result. \square

2. The nondegenerate case

Suppose $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and recall that $\boldsymbol{\Sigma}$ is always positive semidefinite. We say that \mathbf{X} is *nondegenerate* when $\boldsymbol{\Sigma}$ is positive definite (equivalently, invertible).

Take, in particular, $X \sim N_1(\mu, \Sigma)$; μ can be any real number and Σ is a positive semidefinite 1×1 matrix; i.e., $\Sigma \geq 0$. The distribution of X is defined via its MGF as

$$M_X(t) = e^{t\mu + \frac{1}{2}t^2\Sigma}.$$

When X is nondegenerate ($\Sigma > 0$), $X \sim N(\mu, \Sigma)$. If $\Sigma = 0$, then $M_X(t) = \exp(t\mu)$; therefore by the uniqueness theorem of MGFs, $P\{X = \mu\} = 1$. Therefore, $N_1(\mu, \sigma^2)$ is the generalization of $N(\mu, \sigma^2)$ in order to include the case that $\sigma = 0$. We will not write $N_1(\mu, \sigma^2)$; instead we always write $N(\mu, \sigma^2)$ as no confusion should arise.

Theorem 3. $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ has a probability density function if and only if it is nondegenerate. In that case, the pdf of \mathbf{X} is

$$f_{\mathbf{X}}(\mathbf{a}) = \frac{1}{(2\pi)^{n/2} (\det \boldsymbol{\Sigma})^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{a} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\mathbf{a} - \boldsymbol{\mu})\right)$$

for all $\mathbf{a} \in \mathbf{R}^n$.

Proof. First of all let us consider the case that \mathbf{X} is degenerate. In that case Σ has some number $k < n$ of strictly-positive eigenvalues. The proof of Theorem 2 tells us that we can write $\mathbf{X} = \mathbf{AZ} + \boldsymbol{\mu}$, where \mathbf{Z} is a k -dimensional vector of i.i.d. standard normals and \mathbf{A} is an $n \times k$ matrix. Consider the k -dimensional space

$$\mathbf{E} := \{\mathbf{x} \in \mathbf{R}^n : \mathbf{x} = \mathbf{Az} + \boldsymbol{\mu} \text{ for some } \mathbf{z} \in \mathbf{R}^k\}.$$

Because $P\{\mathbf{Z} \in \mathbf{R}^k\} = 1$, it follows that $P\{\mathbf{X} \in \mathbf{E}\} = 1$. If \mathbf{X} had a pdf $f_{\mathbf{X}}$, then

$$1 = P\{\mathbf{X} \in \mathbf{E}\} = \int_{\mathbf{E}} f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x}.$$

But the n -dimensional volume of \mathbf{E} is zero since the dimension of \mathbf{E} is $k < n$. This creates a contradiction [unless \mathbf{X} did not have a pdf, that is].

If \mathbf{X} is nondegenerate, then we can write $\mathbf{X} = \mathbf{AZ} + \boldsymbol{\mu}$, where \mathbf{Z} is an n -vector of i.i.d. standard normals and $\Sigma = \mathbf{AA}'$ is invertible; see the proof of Theorem 2. Recall that the choice of \mathbf{A} is not unique; in this case, we can always choose $\mathbf{A} := \Sigma^{1/2}$ because $\Sigma^{1/2}\mathbf{Z} + \boldsymbol{\mu} \sim N_n(\boldsymbol{\mu}, \Sigma)$. In other words,

$$X_i = \sum_{j=1}^n A_{ij}Z_j + \mu_i = \sum_{j=1}^n \Sigma_{ij}^{1/2}Z_j + \mu_i := g_i(Z_1, \dots, Z_n) \quad (1 \leq i \leq n).$$

If $\mathbf{a} = \Sigma^{1/2}\mathbf{z} + \boldsymbol{\mu}$, then $\mathbf{z} = \Sigma^{-1/2}(\mathbf{a} - \boldsymbol{\mu})$. Therefore, the change of variables formula of elementary probability implies that

$$f_{\mathbf{X}}(\mathbf{a}) = \frac{f_{\mathbf{Z}}(\Sigma^{-1/2}(\mathbf{a} - \boldsymbol{\mu}))}{|\det J|},$$

as long as $\det J \neq 0$, where

$$J := \begin{pmatrix} \frac{\partial g_1}{\partial z_1} & \cdots & \frac{\partial g_1}{\partial z_n} \\ \vdots & & \vdots \\ \frac{\partial g_n}{\partial z_1} & \cdots & \frac{\partial g_n}{\partial z_n} \end{pmatrix} = \begin{pmatrix} A_{1,1} & \cdots & A_{1,n} \\ \vdots & & \vdots \\ A_{n,1} & \cdots & A_{n,n} \end{pmatrix} = \mathbf{A}.$$

Because $\det(\Sigma) = \det(\mathbf{AA}') = (\det \mathbf{A})^2$, it follows that $\det \mathbf{A} = (\det \Sigma)^{1/2}$, and hence

$$f_{\mathbf{X}}(\mathbf{a}) = \frac{1}{(\det \Sigma)^{1/2}} f_{\mathbf{Z}}(\Sigma^{-1/2}(\mathbf{a} - \boldsymbol{\mu})).$$

Because of the independence of the Z_j 's,

$$f_{\mathbf{Z}}(\mathbf{z}) = \prod_{j=1}^n \frac{e^{-z_j^2/2}}{\sqrt{2\pi}} = \frac{1}{(2\pi)^{n/2}} e^{-\mathbf{z}'\mathbf{z}/2}$$

for all $\mathbf{z} \in \mathbf{R}^n$. Therefore,

$$f_{\mathbf{Z}}(\Sigma^{-1/2}(\mathbf{a} - \boldsymbol{\mu})) = \frac{1}{(2\pi)^{n/2}} \exp\left(-\frac{1}{2}(\mathbf{a} - \boldsymbol{\mu})'\Sigma^{-1}(\mathbf{a} - \boldsymbol{\mu})\right),$$

and the result follows. \square

3. The bivariate normal distribution

A *bivariate normal distribution* has the form $N_2(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\mu_1 = EX_1$, $\mu_2 = EX_2$, $\Sigma_{1,1} = \text{Var}(X_1) := \sigma_1^2 > 0$, $\Sigma_{2,2} = \text{Var}(X_2) := \sigma_2^2 > 0$, and $\Sigma_{1,2} = \Sigma_{2,1} = \text{Cov}(X_1, X_2)$. Let

$$\rho := \text{Corr}(X_1, X_2) := \frac{\text{Cov}(X_1, X_2)}{\sqrt{\text{Var}(X_1) \cdot \text{Var}(X_2)}}$$

denote the correlation between X_1 and X_2 , and recall that $-1 \leq \rho \leq 1$. Then, $\Sigma_{1,2} = \Sigma_{2,1} = \rho\sigma_1\sigma_2$, whence

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}.$$

Since $\det \boldsymbol{\Sigma} = \sigma_1^2\sigma_2^2(1 - \rho^2)$, it follows immediately that our bivariate normal distribution is non-degenerate if and only if $-1 < \rho < 1$, in which case

$$\boldsymbol{\Sigma}^{-1} = \begin{pmatrix} \frac{1}{\sigma_1^2(1 - \rho^2)} & -\frac{\rho}{1 - \rho^2} \cdot \frac{1}{\sigma_1\sigma_2} \\ -\frac{\rho}{1 - \rho^2} \cdot \frac{1}{\sigma_1\sigma_2} & \frac{1}{\sigma_2^2(1 - \rho^2)} \end{pmatrix}.$$

Because

$$\mathbf{z}'\boldsymbol{\Sigma}^{-1}\mathbf{z} = \left(\frac{z_1}{\sigma_1}\right)^2 - 2\rho\left(\frac{z_1}{\sigma_1}\right)\left(\frac{z_2}{\sigma_2}\right) + \left(\frac{z_2}{\sigma_2}\right)^2$$

for all $\mathbf{z} \in \mathbf{R}^n$, the pdf of $\mathbf{X} = (X_1, X_2)'$ —in the non-degenerate case where there is a pdf—is

$$\begin{aligned} f_{\mathbf{X}}(x_1, x_2) \\ = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1 - \rho^2}} \exp\left(-\frac{1}{2(1 - \rho^2)} \left[\left(\frac{x_1 - \mu_1}{\sigma_1}\right)^2 - 2\rho\left(\frac{x_1 - \mu_1}{\sigma_1}\right)\left(\frac{x_2 - \mu_2}{\sigma_2}\right) + \left(\frac{x_2 - \mu_2}{\sigma_2}\right)^2 \right]\right). \end{aligned}$$

But of course non-degenerate cases are also possible. For instance, suppose $Z \sim N(0, 1)$ and define $\mathbf{X} := (Z, -Z)$. Then $\mathbf{X} = \mathbf{A}Z$ where $\mathbf{A} := (1, -1)'$, whence

$$\boldsymbol{\Sigma} = \mathbf{A}\mathbf{A}' = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$$

is singular. In general, if $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and the rank of $\boldsymbol{\Sigma}$ is $k < n$, then \mathbf{X} depends only on k [and not n] i.i.d. $N(0, 1)$'s. This can be gleaned from the proof of Theorem 2.

4. A few important properties of multivariate normal distributions

Proposition 4. Let $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. If \mathbf{C} is an $m \times n$ matrix and \mathbf{d} is an m -vector, then $\mathbf{CX} + \mathbf{d} \sim N_m(\mathbf{C}\boldsymbol{\mu} + \mathbf{d}, \mathbf{C}\boldsymbol{\Sigma}\mathbf{C}')$. In general, $\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}'$ is positive semidefinite; it is positive definite if and only if it has full rank m .

In particular, if \mathbf{a} is a nonrandom n -vector, then $\mathbf{a}'\mathbf{X} \sim N(\mathbf{a}'\boldsymbol{\mu}, \mathbf{a}'\boldsymbol{\Sigma}\mathbf{a})$.

Proof. We compute the MGF of $\mathbf{CX} + \mathbf{d}$ as follows:

$$M_{\mathbf{CX}+\mathbf{d}}(\mathbf{t}) = E \exp(\mathbf{t}'[\mathbf{CX} + \mathbf{d}]) = e^{\mathbf{t}'\mathbf{d}} M_{\mathbf{X}}(\mathbf{s}),$$

where $\mathbf{s} := \mathbf{C}'\mathbf{t}$. Therefore,

$$M_{\mathbf{CX}+\mathbf{d}}(\mathbf{t}) = \exp\left(\mathbf{t}'\mathbf{d} + \mathbf{s}'\boldsymbol{\mu} + \frac{1}{2}\mathbf{s}'\boldsymbol{\Sigma}\mathbf{s}\right) = \exp\left(\mathbf{t}'\boldsymbol{\nu} + \frac{1}{2}\mathbf{t}'\mathbf{Q}\mathbf{t}\right),$$

where $\boldsymbol{\nu} := \mathbf{C}\boldsymbol{\mu} + \mathbf{d}$ and $\mathbf{Q} := \mathbf{C}\boldsymbol{\Sigma}\mathbf{C}'$. Finally, a general fact about symmetric matrices (Corollary 13, p. 16) implies that the symmetric $m \times m$ matrix $\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}'$ is nonsingular if and only if it has full rank m . \square

Proposition 5. If $\mathbf{X} \in N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, for a nonsingular variance-covariance matrix $\boldsymbol{\Sigma}$, and $\mathbf{C}_{m \times n}$ and $\mathbf{d}_{n \times 1}$ are nonrandom, then $\mathbf{CX} + \mathbf{d}$ is nonsingular if and only if $\text{rank}(\mathbf{C}) = m$.

Proof. Recall that the nonsingularity of $\boldsymbol{\Sigma}$ is equivalent to it being positive definite. Now $\mathbf{CX} + \mathbf{d}$ is multivariate normal by the preceding result. It is nondegenerate if and only if $\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}'$ is positive definite. But $\mathbf{x}'\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}'\mathbf{x} = (\mathbf{C}'\mathbf{x})'\boldsymbol{\Sigma}(\mathbf{C}'\mathbf{x}) > 0$ if and only if $(\mathbf{C}'\mathbf{x}) \neq \mathbf{0}$, since $\boldsymbol{\Sigma}$ is positive definite. Therefore, $\mathbf{CX} + \mathbf{d}$ is nondegenerate if and only if $\mathbf{C}'\mathbf{x} \neq \mathbf{0}$ whenever $\mathbf{x} \neq \mathbf{0}$. This is equivalent to $\mathbf{x}'\mathbf{C} \neq \mathbf{0}$ for all nonzero vectors \mathbf{x} ; that is, \mathbf{C} has row rank—hence rank— m . \square

The following is an easy corollary of the previous proposition, and identifies the “standard multivariate normal” distribution as the distribution of i.i.d. standard univariate normal distributions. It also states that we do not change the distribution of a standard multivariate normal if we apply to it an orthogonal matrix.

Corollary 6. $\mathbf{Z} \sim N_n(\mathbf{0}, \mathbf{I})$ if and only if Z_1, \dots, Z_n are i.i.d. $N(0, 1)$ ’s. Moreover, if $\mathbf{Z} \sim N_n(\mathbf{0}, \mathbf{I})$ and $\mathbf{A}_{n \times n}$ is orthogonal then $\mathbf{AZ} \sim N_n(\mathbf{0}, \mathbf{I})$ also.

Next we state another elementary fact, derived by looking only at the MGF’s. It states that a subset of a multivariate normal vector itself is multivariate normal.

Proposition 7. Suppose $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $1 \leq i_1 < i_2 < \dots < i_k \leq n$ is a subsequence of $1, \dots, n$. Then, $(X_{i_1}, \dots, X_{i_k})' \sim N_k(\boldsymbol{\nu}, \mathbf{Q})$, where

$$\boldsymbol{\nu} := E \begin{pmatrix} X_{i_1} \\ \vdots \\ X_{i_k} \end{pmatrix} = \begin{pmatrix} \mu_{i_1} \\ \vdots \\ \mu_{i_k} \end{pmatrix}, \quad \mathbf{Q} := \text{Var} \begin{pmatrix} X_{i_1} \\ \vdots \\ X_{i_k} \end{pmatrix} = \begin{pmatrix} \Sigma_{i_1, i_1} & \dots & \Sigma_{i_1, i_k} \\ \vdots & & \vdots \\ \Sigma_{i_k, i_1} & \dots & \Sigma_{i_k, i_k} \end{pmatrix}.$$

Proposition 8. Suppose $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and assume that we can divide the X_i 's into two groups: $(X_i)_{i \in G}$ and $(X_j)_{j \notin G}$, where G is a subset of the index set $\{1, \dots, n\}$. Suppose in addition that $\text{Cov}(X_i, X_j) = 0$ for all $i \in G$ and $j \notin G$. Then, $(X_i)_{i \in G}$ is independent from $(X_j)_{j \notin G}$.

Thus, for example, if $(X_1, X_2, X_3)'$ has a trivariate normal distribution and X_1 is uncorrelated from X_2 and X_3 , then X_1 is independent of (X_2, X_3) . For a second example suppose that (X_1, X_2, X_3, X_4) has a multivariate normal distribution and: $E(X_1 X_2) = E(X_1)E(X_2)$, $E(X_1 X_3) = E(X_1)E(X_3)$, $E(X_4 X_2) = E(X_4)E(X_2)$, and $E(X_4 X_3) = E(X_4)E(X_3)$, then (X_1, X_4) and (X_2, X_3) are two independent bivariate normal random vectors.

Proof. I will prove the following special case of the proposition; the general case follows from a similar reasoning, but the notation is messier.

Suppose (X_1, X_2) has a bivariate normal distribution and $E(X_1 X_2) = E(X_1)E(X_2)$. Then, X_1 and X_2 are independent. In order to prove this we write the MGF of $\mathbf{X} := (X_1, X_2)'$:

$$\begin{aligned} M_{\mathbf{X}}(\mathbf{t}) &= e^{\mathbf{t}'\boldsymbol{\mu} + \frac{1}{2}\mathbf{t}'\boldsymbol{\Sigma}\mathbf{t}} \\ &= e^{t_1\mu_1 + t_2\mu_2} \cdot \exp\left(\frac{1}{2}(t_1, t_2) \begin{pmatrix} \text{Var}(X_1) & 0 \\ 0 & \text{Var}(X_2) \end{pmatrix} \begin{pmatrix} t_1 \\ t_2 \end{pmatrix}\right) \\ &= e^{t_1\mu_1 + \frac{1}{2}t_1^2\text{Var}(X_1)} \cdot e^{t_2\mu_2 + \frac{1}{2}t_2^2\text{Var}(X_2)} \\ &= M_{X_1}(t_1) \cdot M_{X_2}(t_2). \end{aligned}$$

The result follows from the independence theorem for MGF's (Theorem 6, p. 29). \square

Remark 9. The previous proposition has generalizations. For instance, suppose we could decompose $\{1, \dots, n\}$ into k disjoint groups G_1, \dots, G_k [so $G_i \cap G_j = \emptyset$ if $i \neq j$, and $G_1 \cup \dots \cup G_k = \{1, \dots, n\}$] such that X_{i_1}, \dots, X_{i_k} are [pair-wise] uncorrelated for all $i_1 \in G_1, \dots, i_k \in G_k$. Then, $(X_i)_{i \in G_1}, \dots, (X_i)_{i \in G_k}$ are independent multivariate normal random vectors. The proof is the same as in the case $k = 2$. \square

Remark 10. It is important that \mathbf{X} has a multivariate normal distribution. For instance, we can construct two standard-normal random variables X and Y , on the same probability space, such that X and Y are uncorrelated but dependent. Here is one way to do this: Let $Y \sim N(0, 1)$ and $S = \pm 1$ with probability $1/2$

each. Assume that S and Y are independent, and define $X := S|Y|$. Note that

$$\begin{aligned} P\{X \leq a\} &= P\{X \leq a, S = 1\} + P\{X \leq a, S = -1\} \\ &= \frac{1}{2}P\{|Y| \leq a\} + \frac{1}{2}P\{-|Y| \leq a\}. \end{aligned}$$

If $a \geq 0$, then $P\{X \leq a\} = \frac{1}{2}P\{|Y| \leq a\} + \frac{1}{2} = P\{Y \leq a\}$. Similarly, $P\{X \leq a\} = P\{Y \leq a\}$ if $a \leq 0$. Therefore, $X, Y \sim N(0, 1)$. Furthermore, X and Y are uncorrelated because S has mean zero; here is why: $E(XY) = E(SY|Y|) = E(S)E(Y|Y|) = 0 = E(X)E(Y)$. But X and Y are not independent because $|X| = |Y|$: For instance, $P\{|X| < 1\} > 0$, but $P\{|X| < 1 \mid |Y| \geq 2\} = 0$. The problem is [and can only be] that $(X, Y)'$ is *not* bivariate normal. \square

5. Quadratic forms

Given a multivariate-normal random variable $\mathbf{X} \sim N_n(\mathbf{0}, \mathbf{I})$ and an $n \times n$ positive semidefinite matrix $\mathbf{A} := (A_{ij})$, we can consider the random quadratic form

$$Q_{\mathbf{A}}(\mathbf{X}) := \mathbf{X}'\mathbf{A}\mathbf{X}.$$

We can write \mathbf{A} , in spectral form, as $\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}'$, so that

$$Q_{\mathbf{A}}(\mathbf{X}) = \mathbf{X}'\mathbf{P}\mathbf{D}\mathbf{P}'\mathbf{X}.$$

Since \mathbf{P} is orthogonal and $\mathbf{X} \sim N_n(\mathbf{0}, \mathbf{I})$, $\mathbf{Z} := \mathbf{P}'\mathbf{X} \sim N_n(\mathbf{0}, \mathbf{I})$ as well. Therefore,

$$Q_{\mathbf{A}}(\mathbf{X}) = \mathbf{Z}'\mathbf{D}\mathbf{Z} = \sum_{i=1}^n D_{i,i}Z_i^2.$$

If \mathbf{A} is a projection matrix, then all of the $D_{i,i}$'s are ones and zeros. In that case, $Q_{\mathbf{A}}(\mathbf{X}) \sim \chi_r^2$, where $r :=$ the number of eigenvalues of \mathbf{A} that are ones; i.e., $r = \text{rank}(\mathbf{A})$. Finally, recall that the rank of a projection matrix is equal to its trace (Corollary 16, p. 16). Let us summarize our findings.

Proposition 11. *If $\mathbf{X} \sim N_n(\mathbf{0}, \mathbf{I})$ and \mathbf{A} is a projection matrix, then $\mathbf{X}'\mathbf{A}\mathbf{X} \sim \chi_{\text{rank}(\mathbf{A})}^2 = \chi_{\text{tr}(\mathbf{A})}^2 = \chi_r^2$, where $r :=$ the total number of nonzero [i.e., one] eigenvalues of \mathbf{A} .*

Example 12. Let

$$\mathbf{A} := \begin{pmatrix} 1 - 1/n & 1/n & -1/n & \cdots & 1/n \\ -1/n & 1 - 1/n & -1/n & \cdots & -1/n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1/n & -1/n & -1/n & \cdots & 1 - 1/n \end{pmatrix}.$$

Then we have seen (Example 5, p. 23) that

$$\mathbf{x}'\mathbf{A}\mathbf{x} = \sum_{i=1}^n (x_i - \bar{x})^2 \quad \text{for all } \mathbf{x} \in \mathbf{R}^n.$$

Now let us observe that \mathbf{A} has the form

$$\mathbf{A} = \mathbf{I} - \mathbf{B},$$

where $\mathbf{B} := (1/n)\mathbf{1}_{n \times n}$. Note that \mathbf{B} is symmetric and $\mathbf{B}^2 = \mathbf{B}$. Therefore, \mathbf{B} is a projection, and hence so is $\mathbf{A} = \mathbf{I} - \mathbf{B}$. Clearly, $\text{tr}(\mathbf{A}) = n - 1$. Therefore, Proposition 11 implies the familiar fact that if X_1, \dots, X_n are i.i.d. standard normals, then $\sum_{i=1}^n (X_i - \bar{X})^2 \sim \chi_{n-1}^2$. \square

Example 13. If \mathbf{A} is an $n \times n$ projection matrix of rank [or trace] r , then $\mathbf{I} - \mathbf{A}$ is an $n \times n$ projection matrix of rank [or trace] $n - r$. Therefore, $\mathbf{X}'(\mathbf{I} - \mathbf{A})\mathbf{X} \sim \chi_{n-r}^2$, whenever $\mathbf{X} \sim N_n(\mathbf{0}, \mathbf{I})$. \square

Example 14. What is the distribution of \mathbf{X} is a nonstandard multivariate normal? Suppose $\mathbf{X} \sim N_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and \mathbf{A} is a projection matrix. If \mathbf{X} is nondegenerate, then $\boldsymbol{\Sigma}^{-1/2}(\mathbf{X} - \boldsymbol{\mu}) \sim N_n(\mathbf{0}, \mathbf{I})$. Therefore,

$$(\mathbf{X} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1/2}\mathbf{A}\boldsymbol{\Sigma}^{-1/2}(\mathbf{X} - \boldsymbol{\mu}) \sim \chi_{\text{rank}(\mathbf{A})}^2 = \chi_{\text{tr}(\mathbf{A})}^2,$$

for every $n \times n$ projection matrix \mathbf{A} . In particular,

$$(\mathbf{X} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\mathbf{X} - \boldsymbol{\mu}) \sim \chi_n^2,$$

which can be seen by specializing the preceding to the projection matrix $\mathbf{A} := \mathbf{I}$. Specializing further still, we see that if X_1, \dots, X_n are independent normal random variables, then we obtain the familiar fact that

$$\sum_{i=1}^n \left(\frac{X_i - \mu_i}{\sigma_i} \right)^2 \sim \chi_n^2,$$

where $\mu_i := EX_i$ and $\sigma_i^2 := \text{Var}(X_i)$. \square

Linear Models

1. The basic model

We now study a *linear statistical model*. That is, we study the models where the observations $\mathbf{Y} := (Y_1, \dots, Y_n)'$ has the following assumed property:

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

where $\boldsymbol{\beta} := (\beta_0, \beta_1, \dots, \beta_{p-1})$ is a vector of p unknown parameters, and

$$\mathbf{X} := \begin{pmatrix} x_{1,0} & \cdots & x_{1,p-1} \\ \vdots & & \vdots \\ x_{n,0} & \cdots & x_{n,p-1} \end{pmatrix}$$

is the so-called “regression matrix,” or “design matrix.” The elements of the $n \times p$ matrix \mathbf{X} are assumed to be known; these are the “descriptive” or “explanatory” variables, and the randomness of the observed values is inherited from the “noise vector,” $\boldsymbol{\varepsilon} := (\varepsilon_1, \dots, \varepsilon_n)'$, which we may think of as being “typically small.” Note that we are changing our notation slightly; \mathbf{X} is no longer assumed to be a random vector [this is done in order to conform with the historical development of the subject].

Throughout, we assume always that the ε_i ’s are independent with mean zero and common variance σ^2 , where $\sigma > 0$ is possibly [in fact, typically] unknown.

In particular, it follows from this assumption that

$$\mathbf{E}\boldsymbol{\varepsilon} = \mathbf{0} \text{ and } \text{Var}(\boldsymbol{\varepsilon}) = \sigma^2 \mathbf{I}. \quad (1)$$

Let us emphasize that our linear model, once written out coordinatewise, is

$$Y_i = \beta_0 x_{i,0} + \cdots + \beta_{p-1} x_{i,p-1} + \varepsilon_i \quad (1 \leq i \leq n).$$

It is intuitively clear that unless $n \geq p$, we cannot hope to effectively use our n observed values in order to estimate the $p + 1$ unknowns $\sigma^2, \beta_0, \dots, \beta_{p-1}$. Therefore, we assume always that

$$n \geq p + 1. \quad (2)$$

This condition guarantees that our linear model is not overspecified.

The best-studied linear models are the *normal models*. Those are linear models for which we assume the more stringent condition that

$$\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}, \sigma^2 \mathbf{I}). \quad (3)$$

Example 1 (A measurement-error model). Here we study a so-called measurement-error model: Suppose the observations Y_1, \dots, Y_n satisfy

$$Y_i = \mu + \varepsilon_i \quad (1 \leq i \leq n)$$

for an unknown parameter μ . This is a simplest example of a linear model, where $\boldsymbol{\beta} = \mu$ is 1×1 , and $\mathbf{X} := \mathbf{1}_{n \times 1}$ is a vector of n ones. \square

Example 2 (Simple linear regression). In simple linear regression we assume that the observed values have the form

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i \quad (1 \leq i \leq n),$$

where x_i is the predictive variable the corresponds to observation i , and β_0, β_1 are unknown. Simple linear regression fits into our theory of linear models, once we set the design matrix as

$$\mathbf{X} := \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix}.$$

Example 3 (Polynomial regression). Consider a nonlinear regression model

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \dots + \beta_{p-1} x_i^{p-1} + \varepsilon_i \quad (1 \leq i \leq n),$$

where p is a known integer ≥ 1 [$p - 1$ denotes the degree of the polynomial approximation to the observed y 's]. Then polynomial regression models are linear models with design matrices of the form

$$\mathbf{X} := \begin{pmatrix} 1 & x_1 & x_1^2 & \dots & x_1^{p-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & x_n & x_n^2 & \dots & x_n^{p-1} \end{pmatrix}.$$

Example 4 (One-way layout with equal observations). In the simplest case of one-way layout, in analysis of variance, our observations are indexed by vectors themselves as follows:

$$Y_{i,j} = \mu_i + \varepsilon_{i,j} \quad (1 \leq i \leq I, 1 \leq j \leq J).$$

For instance, suppose we are interested in the effect of I different fertilizers. We apply these fertilizers to J different blocks, independently, and "measure

the effect.” Then, $Y_{i,j}$ is the effect of fertilizer i in block j . The preceding model is assuming that, up to sampling error, the effect of fertilizer i is μ_i . This is a linear model. Indeed, we can create a new random vector \mathbf{Y} of IJ observations by simply “vectorizing” the $Y_{i,j}$ ’s:

$$\mathbf{Y} := (Y_{1,1}, \dots, Y_{1,J}, Y_{2,1}, \dots, Y_{2,J}, \dots, Y_{I,1}, \dots, Y_{I,J})'.$$

The vector β of unknowns is $\beta := (\mu_1, \dots, \mu_I)'$, and the design matrix is the following $IJ \times I$ matrix:

$$\mathbf{X} := \begin{pmatrix} \mathbf{1}_{J \times 1} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{1}_{J \times 1} & \mathbf{0} & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{1}_{J \times 1} \end{pmatrix},$$

where $\mathbf{1}_{J \times 1} := (1, \dots, 1)'$ is a J -vector of all ones. It is possible to show that *one-way layout with unequal number of observations* is also a linear model. That is the case where $Y_{i,j} = \mu_i + \varepsilon_{i,j}$, where $1 \leq i \leq I$ and $1 \leq j \leq J_i$ [the number of observed values might differ from block to block].

2. The least-squares estimator of $\theta := \mathbf{X}\beta$

Let us return to our general linear model

$$\mathbf{Y} = \mathbf{X}\beta + \varepsilon.$$

Ultimately, our goal is to first find and then analyze the least-squares estimator $\hat{\beta}$ of β . But first, let us find the least-squares estimate for $\theta := \mathbf{X}\beta$. In other words, we wish to perform the following optimization problem:

$$\min_{\beta \in \mathbb{R}^p} \|\mathbf{Y} - \mathbf{X}\beta\| = \min_{\theta \in \mathcal{C}(\mathbf{X})} \|\mathbf{Y} - \theta\|.$$

Abstractly speaking, the minimizer solves

$$\hat{\theta} = P_{\mathcal{C}(\mathbf{X})}\mathbf{Y}.$$

But is there an optimal β ? As we shall see next, there certainly is a unique $\hat{\beta}$ when \mathbf{X} has full rank.

3. The least-squares estimator of β

Our *least-squares estimator* $\hat{\beta}$ of the vector parameter β is defined via

$$\min_{\beta \in \mathbb{R}^p} \|\mathbf{Y} - \mathbf{X}\beta\| = \|\mathbf{Y} - \mathbf{X}\hat{\beta}\|.$$

We aim to solve this minimization problem under natural conditions on the design matrix \mathbf{X} . But first, let us introduce some notation. The vector

$$\hat{\mathbf{Y}} := \hat{\theta} := \mathbf{X}\hat{\beta}$$

is called the vector of *fitted values*, and the coordinates of

$$\mathbf{e} := \mathbf{Y} - \hat{\mathbf{Y}}$$

are the so-called *residuals*.

Now we write our minimization problem in the following form: First find the minimizing $\hat{\boldsymbol{\beta}} \in \mathbf{R}^p$ that solves

$$\min_{\mathbf{z} \in \mathcal{C}(\mathbf{X})} \|\mathbf{Y} - \mathbf{z}\| = \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|.$$

Now we know from Proposition 18 (p. 19) that the vector \mathbf{z} that achieves this minimum does so uniquely, and is given by $\mathbf{P}_{\mathcal{C}(\mathbf{X})}\mathbf{Y}$, where we recall $\mathbf{P}_{\mathcal{C}(\mathbf{X})} := \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ is projection onto the column space of \mathbf{X} ; this of course is valid provided that $(\mathbf{X}'\mathbf{X})^{-1}$ exists. Now the matrix $\mathbf{P}_{\mathcal{C}(\mathbf{X})}$ plays such an important role that it has its own name: It is called the *hat matrix*, and is denoted as

$$\mathbf{H} := \mathbf{P}_{\mathcal{C}(\mathbf{X})} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'.$$

\mathbf{H} is called the *hat matrix* because it maps the observations \mathbf{Y} to the fitted values $\hat{\mathbf{Y}}$ [informally, it puts a “hat” over \mathbf{Y}]. More precisely, the defining feature of \mathbf{H} is that

$$\hat{\mathbf{Y}} = \mathbf{H}\mathbf{Y},$$

once again provided that $\mathbf{X}'\mathbf{X}$ is nonsingular. The following gives us a natural method for checking this nonsingularity condition in terms of \mathbf{X} directly.

Lemma 5. $\mathbf{X}'\mathbf{X}$ is nonsingular if and only if $\text{rank}(\mathbf{X}) = p$.

Proof. Basic linear algebra tells us that the positive semidefinite $\mathbf{X}'\mathbf{X}$ is nonsingular if and only if it is positive definite; i.e., if and only if it has full rank. Since the rank of $\mathbf{X}'\mathbf{X}$ is the same as the rank of \mathbf{X} , and since $n > p$, $\mathbf{X}'\mathbf{X}$ has full rank if and only if its rank, which is the same as $\text{rank}(\mathbf{X})$, is p . \square

From now on we always assume [unless we state explicitly otherwise] that $\text{rank}(\mathbf{X}) = p$. We have shown that, under this condition, if there is a $\hat{\boldsymbol{\beta}}$, then certainly $\hat{\mathbf{Y}} = \mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{H}\mathbf{Y} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$. In order to find $\hat{\boldsymbol{\beta}}$ from this, multiply both sides by $(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ to see that $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$.

The quantity $\text{RSS} := \mathbf{e}'\mathbf{e} = \|\mathbf{e}\|^2 := \|\mathbf{Y} - \hat{\mathbf{Y}}\|^2$ is called the *sum of squared residuals*, also known as the *residual sum of squared errors*, and is given by $\|(I - \mathbf{H})\mathbf{Y}\|^2 = \|\mathbf{P}_{\mathcal{C}(\mathbf{X})^\perp}\mathbf{Y}\|^2$. In particular, we have the following:

Proposition 6. If $\text{rank}(\mathbf{X}) = p$, then the least-squares estimator of $\boldsymbol{\beta}$ is

$$\hat{\boldsymbol{\beta}} := (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}, \tag{4}$$

and $\text{RSS} = \|(I - \mathbf{H})\mathbf{Y}\|^2$.

Let us make some elementary computations with the least squares estimator of $\boldsymbol{\beta}$.

Lemma 7. $\hat{\beta}$ is an unbiased estimator of β , and $\text{Var}(\hat{\beta}) = \sigma^2(\mathbf{X}'\mathbf{X})^{-1}$.

Proof. Because $E\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'E\mathbf{Y} = \beta$, it follows that $\hat{\beta}$ is unbiased. Also, $\text{Var}(\hat{\beta}) = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\text{Var}(\mathbf{Y})\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} = \sigma^2(\mathbf{X}'\mathbf{X})^{-1}$, because of (1). \square

4. Optimality

Theorem 8. Let $\theta := \mathbf{X}\beta$ be estimated, via least squares, by $\hat{\theta} := \mathbf{H}\mathbf{Y}$. Then, for all linear unbiased estimates of $\mathbf{c}'\theta$, the estimator $\mathbf{c}'\hat{\theta}$ uniquely minimizes the variance.

By a “linear unbiased estimator of $\mathbf{c}'\theta$ ” we mean an estimator of the form $\sum_{j=1}^n a_j Y_j$ whose expectation is $\mathbf{c}'\theta$. In this sense, $\mathbf{c}'\hat{\theta}$ is the *best linear unbiased estimator* [or “BLUE”] of $\mathbf{c}'\theta$. The preceding can be improved upon as follows, though we will not prove it:

Theorem 9 (Rao). Under the normal model (3), $\mathbf{c}'\hat{\theta}$ is the unique UMVUE of $\mathbf{c}'\theta$, for every nonrandom vector $\mathbf{c} \in \mathbf{R}^n$.

Let us consider Theorem 8 next.

Proof of Theorem 8. We saw on page 41 that $\hat{\theta} := \mathbf{H}\mathbf{Y}$ irrespective of whether or not $(\mathbf{X}'\mathbf{X})^{-1}$ exists. Therefore,

$$\begin{aligned} E(\mathbf{c}'\hat{\theta}) &= \mathbf{c}'E\hat{\theta} = \mathbf{c}'\theta, \quad \text{Var}(\mathbf{c}'\hat{\theta}) = \mathbf{c}'\text{Var}(\mathbf{H}\mathbf{Y})\mathbf{c} = \mathbf{c}'\mathbf{H}'\text{Var}(\mathbf{Y})\mathbf{H}\mathbf{c} \\ &= \sigma^2\|\mathbf{H}\mathbf{c}\|^2. \end{aligned}$$

Any other linear estimator has the form $\mathbf{a}'\mathbf{Y}$, and satisfies

$$E(\mathbf{a}'\mathbf{Y}) = \mathbf{a}'E\mathbf{Y} = \mathbf{a}'\theta, \quad \text{Var}(\mathbf{a}'\mathbf{Y}) = \mathbf{a}'\text{Var}(\mathbf{Y})\mathbf{a} = \sigma^2\|\mathbf{a}\|^2.$$

If, in addition, $\mathbf{a}'\mathbf{Y}$ is unbiased, then it follows that $\mathbf{a}'\theta = \mathbf{c}'\theta$; i.e., $\mathbf{a} - \mathbf{c}$ is orthogonal to θ . This should hold no matter what value β [and hence θ] takes. Since $\mathcal{C}(\mathbf{X})$ is the collection of all possible values of θ , it follows that $\mathbf{a} - \mathbf{c}$ is orthogonal to $\mathcal{C}(\mathbf{X})$. Because \mathbf{H} is projection onto $\mathcal{C}(\mathbf{X})$, it follows that $\mathbf{H}(\mathbf{a} - \mathbf{c}) = \mathbf{0}$; equivalently, $\mathbf{H}\mathbf{c} = \mathbf{H}\mathbf{a}$. Therefore, $\text{Var}(\mathbf{c}'\hat{\theta}) = \sigma^2\|\mathbf{H}\mathbf{a}\|^2$ and

$$\text{Var}(\mathbf{a}'\mathbf{Y}) - \text{Var}(\mathbf{c}'\hat{\theta}) = \sigma^2\{\|\mathbf{a}\|^2 - \|\mathbf{H}\mathbf{a}\|^2\} = \sigma^2\|(\mathbf{I} - \mathbf{H})\mathbf{a}\|^2 \geq 0,$$

thanks to the Pythagorean property. \square

5. Regression and prediction

Now that we have the least-squares estimate for β , let us use it in order to make prediction.

Recall that our model is $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. In applications, Y_i is the i th observation for the y variable, and the linear model is really saying that given an explanatory variable $\mathbf{x} = (x_0, \dots, x_{p-1})'$,

$$y = \beta_0 x_0 + \dots + \beta_{p-1} x_{p-1} + \text{"noise."}$$

Therefore, our prediction, for a given \mathbf{x} , is

$$[\text{predicted value}] y = \hat{\beta}_0 x_0 + \dots + \hat{\beta}_{p-1} x_{p-1}, \quad (5)$$

where $\hat{\boldsymbol{\beta}} = (\hat{\beta}_0, \dots, \hat{\beta}_{p-1})'$ is the least-squares estimate of $\boldsymbol{\beta}$. We may view the right-hand side of (5) as a function of \mathbf{x} , and call (5) the equation for the "regression line."

6. Estimation of σ^2

We wish to also estimate σ^2 . The estimator of interest to us turns out to be the following:

$$S^2 := \frac{1}{n-p} \text{RSS}, \quad (6)$$

as the next lemma suggests.

Lemma 10. S^2 is an unbiased estimator of σ^2 .

Proof. Recall that $\text{RSS} = \mathbf{e}'\mathbf{e} = \|\mathbf{Y} - \mathbf{H}\mathbf{Y}\|^2$. We can write the RSS as $\|(I - \mathbf{H})\mathbf{Y}\|^2 = \mathbf{Y}'(I - \mathbf{H})'(I - \mathbf{H})\mathbf{Y} = \mathbf{Y}'(I - \mathbf{H})\mathbf{Y}$. In other words, the RSS is a random quadratic form for the matrix $\mathbf{A} := I - \mathbf{H}$, and hence

$$\begin{aligned} E(\text{RSS}) &= \text{tr}((I - \mathbf{H})\text{Var}(\mathbf{Y})) + (\mathbf{E}\mathbf{Y})'(I - \mathbf{H})(\mathbf{E}\mathbf{Y}) \\ &= \sigma^2 \text{tr}(I - \mathbf{H}) + (\mathbf{X}\boldsymbol{\beta})'(I - \mathbf{H})\mathbf{X}\boldsymbol{\beta}. \end{aligned}$$

Because $\mathbf{X}\boldsymbol{\beta} \in \mathcal{C}(\mathbf{X})$, $I - \mathbf{H}$ projects onto the orthogonal subspace of where it is, therefore $(I - \mathbf{H})\mathbf{X}\boldsymbol{\beta} = \mathbf{0}$. And the trace of the projection matrix $(I - \mathbf{H})$ is its rank, which is $n - \text{tr}(\mathbf{H}) = n - p$, since \mathbf{X} has full rank p . It follows that $E(\text{RSS}) = \sigma^2(n - p)$, and therefore $E(S^2) = \sigma^2$. \square

7. The normal model

In the important special case of the normal model,

$$\mathbf{Y} \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2 I).$$

Therefore, $\hat{\boldsymbol{\beta}} \sim N_p(\boldsymbol{\beta}, \sigma^2(\mathbf{X}'\mathbf{X})^{-1})$. And the vector $(I - \mathbf{H})\mathbf{Y} = \mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}$ of residual errors is also a multivariate normal:

$$(I - \mathbf{H})\mathbf{Y} \sim N_{n-p}((I - \mathbf{H})\mathbf{X}\boldsymbol{\beta}, \sigma^2(I - \mathbf{H})(I - \mathbf{H})') = N_{n-p}(\mathbf{0}, \sigma^2(I - \mathbf{H})).$$

Therefore, in the normal model,

$$S^2 = \frac{1}{n-p} \|(I - \mathbf{H})\mathbf{Y}\|^2 \sim \sigma^2 \frac{\chi_{n-p}^2}{n-p}.$$

Finally, we note that $\mathbf{t}'\hat{\boldsymbol{\beta}} + \mathbf{s}'(\mathbf{I} - \mathbf{H})\mathbf{Y}$ is a matrix times \mathbf{Y} for all $\mathbf{t} \in \mathbf{R}^p$ and $\mathbf{s} \in \mathbf{R}^{n-p}$. Therefore, $(\hat{\boldsymbol{\beta}}, (\mathbf{I} - \mathbf{H})\mathbf{Y})$ is also a multivariate normal. But

$$\text{Cov}(\hat{\boldsymbol{\beta}}, (\mathbf{I} - \mathbf{H})\mathbf{Y}) = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\text{Var}(\mathbf{Y})(\mathbf{I} - \mathbf{H})' = \sigma^2(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'(\mathbf{I} - \mathbf{H}) = \mathbf{0},$$

since the columns of \mathbf{X} are obviously orthogonal to every element in $\mathcal{C}(\mathbf{X})^\perp$ and $\mathbf{I} - \mathbf{H} = \mathbf{P}_{\mathcal{C}(\mathbf{X})^\perp}$. This shows that $\hat{\boldsymbol{\beta}}$ and $(\mathbf{I} - \mathbf{H})\mathbf{Y}$ are independent, and hence $\hat{\boldsymbol{\beta}}$ and S^2 are also independent. Thus, we summarize our findings.

Theorem 11. *The least-squares estimator $\hat{\boldsymbol{\beta}}$ of $\boldsymbol{\beta}$ is given by (4); it is always unbiased. Moreover, S^2 is an unbiased estimator of σ^2 . Under the normal model, S and $\hat{\boldsymbol{\beta}}$ are independent, and*

$$\hat{\boldsymbol{\beta}} \sim N_p(\boldsymbol{\beta}, \sigma^2(\mathbf{X}'\mathbf{X})^{-1}), \quad S^2 \sim \sigma^2 \frac{\chi_{n-p}^2}{n-p}.$$

Recall that \mathbf{Y} has the nondegenerate multivariate normal distribution $N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. Therefore, its pdf is

$$f_{\mathbf{Y}}(\mathbf{y}) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2\right).$$

This shows readily the following.

Lemma 12. *In the normal model, $\hat{\boldsymbol{\beta}}$ is the MLE of $\boldsymbol{\beta}$ and $(\frac{n-p}{n}) S^2$ is the MLE for σ^2 .*

Proof. Clearly, maximizing the likelihood function, over all $\boldsymbol{\beta}$, is the same as minimizing $\|\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}\|^2$. Therefore, MLE = least squares for $\boldsymbol{\beta}$. As for σ^2 , we write the log likelihood function:

$$L(\sigma) = -\frac{n}{2} \ln(2\pi) - n \ln \sigma - \frac{1}{2\sigma^2} \|\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}\|^2.$$

Then,

$$L'(\sigma) = -\frac{n}{\sigma} + \frac{1}{\sigma^3} \|\mathbf{Y} - \mathbf{X}\boldsymbol{\beta}\|^2.$$

Set $L'(\sigma) = 0$ and solve to see that the MLE of σ^2 is $\frac{1}{n} \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2 = (\frac{n-p}{n}) S^2$, thanks to the MLE principle and the already-proven fact that the MLE of $\boldsymbol{\beta}$ is $\hat{\boldsymbol{\beta}}$. \square

8. Some examples

1. A measurement-error model. Recall the measurement-error model

$$Y_i = \mu + \varepsilon_i \quad (1 \leq i \leq n).$$

We have seen that this is a linear model with $p = 1$, $\boldsymbol{\beta} = \mu$, and $\mathbf{X} := \mathbf{1}_{n \times 1}$. Since $\mathbf{X}'\mathbf{X} = n$ and $\mathbf{X}'\mathbf{Y} = \sum_{i=1}^n Y_i$, we have

$$\hat{\boldsymbol{\beta}} = \hat{\mu} := \bar{Y},$$

and

$$\frac{1}{n-1}S^2 = \frac{1}{n-1} \|\mathbf{Y} - \bar{Y}\mathbf{1}_{n \times 1}\|^2 = \frac{1}{n-1} \sum_{i=1}^n (Y_i - \bar{Y})^2 := S^2.$$

These are unbiased estimators for μ and σ^2 , respectively. Under the normal model, \bar{Y} and S^2 are independent, $\bar{Y} \sim N(\mu, \sigma^2)$ and $(n-1)S^2 \sim \sigma^2 \chi_{n-1}^2$. These are some of the highlights of Math. 5080.

2. Simple linear regression. Recall that simple linear regression is our linear model in the special case that $p = 2$, $\boldsymbol{\beta} = (\beta_0, \beta_1)'$, and

$$\mathbf{X} := \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix}.$$

We have

$$\mathbf{X}'\mathbf{X} = \begin{pmatrix} n & \sum_{i=1}^n x_i \\ \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 \end{pmatrix}, \quad \mathbf{X}'\mathbf{Y} = \begin{pmatrix} \sum_{i=1}^n Y_i \\ \sum_{i=1}^n x_i Y_i \end{pmatrix},$$

and

$$\begin{aligned} (\mathbf{X}'\mathbf{X})^{-1} &= \frac{1}{n \sum_{i=1}^n x_i^2 - (\sum_{i=1}^n x_i)^2} \begin{pmatrix} \sum_{i=1}^n x_i^2 & -\sum_{i=1}^n x_i \\ -\sum_{i=1}^n x_i & n \end{pmatrix} \\ &= \frac{1}{\sum_{i=1}^n (x_i - \bar{x})^2} \begin{pmatrix} \frac{1}{n} \sum_{i=1}^n x_i^2 & -\bar{x} \\ -\bar{x} & 1 \end{pmatrix}. \end{aligned}$$

Therefore,

$$\begin{pmatrix} \hat{\beta}_0 \\ \hat{\beta}_1 \end{pmatrix} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{Y},$$

which leads to

$$\hat{\beta}_1 = \frac{\sum_{i=1}^n (x_i - \bar{x})Y_i}{\sum_{i=1}^n (x_i - \bar{x})^2} = \frac{\sum_{i=1}^n (x_i - \bar{x})(Y_i - \bar{Y})}{\sum_{i=1}^n (x_i - \bar{x})^2} \text{ and } \hat{\beta}_0 = \bar{Y} - \hat{\beta}_1 \bar{x}.$$

We have derived these formulas by direct computation already. In this way we find that the fitted values are

$$\hat{\mathbf{Y}} = \mathbf{X}\hat{\boldsymbol{\beta}} = \begin{pmatrix} \hat{\beta}_0 + \hat{\beta}_1 x_1 \\ \vdots \\ \hat{\beta}_0 + \hat{\beta}_1 x_n \end{pmatrix}.$$

Also,

$$S^2 = \frac{1}{n-2} \|\mathbf{Y} - \hat{\mathbf{Y}}\|^2 = \frac{1}{n-2} \sum_{i=1}^n (Y_i - \hat{\beta}_0 - \hat{\beta}_1 x_i)^2,$$

and this is independent of $(\hat{\beta}_0, \hat{\beta}_1)$ under the normal model.

Recall that our linear model is, at the moment, the simple regression model,

$$Y_i = \beta_0 + \beta_1 x_i + \varepsilon_i.$$

Perhaps the most important first question that we can ask in this context is $\beta_1 = 0$; that is, we wish to know whether the x variables are [linearly] independent of the Y 's. Let us try to find a confidence interval for β_1 in order to answer this question. From now on, we work under the normal model. Recall that under the normal model,

$$\hat{\beta} \sim N_2(\beta, \sigma^2(\mathbf{X}'\mathbf{X})^{-1}) \Rightarrow \hat{\beta}_1 \sim N(\beta_1, \sigma^2 [(\mathbf{X}'\mathbf{X})^{-1}]_{2,2}) = N\left(\beta_1, \frac{\sigma^2}{\sum_{i=1}^n (x_i - \bar{x})^2}\right).$$

Equivalently,

$$Z := \frac{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2}}{\sigma} (\hat{\beta}_1 - \beta_1) \sim N(0, 1).$$

Now, S^2/σ^2 is independent of Z and is distributed as $\chi_{n-2}^2/(n-2)$. Therefore,

$$\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \left(\frac{\hat{\beta}_1 - \beta_1}{S} \right) = \frac{Z}{\sqrt{S^2/\sigma^2}} \sim t_{n-2}.$$

Therefore, a $(1 - \alpha) \times 100\%$ confidence interval for β_1 is

$$\hat{\beta}_1 \pm \frac{S}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2}} t_{n-2}^{\alpha/2},$$

where t_{ν}^r is the point whose right area, under the t_{ν} pdf, is r . If zero is not in this confidence interval, then our statistical prediction is that β_1 is not zero [at the confidence level α].

3. A remark on polynomial regression. Recall that, in polynomial regression, we have $\mathbf{Y} = \mathbf{X}\beta + \epsilon$, where \mathbf{X} is the following design matrix:

$$\mathbf{X} := \begin{pmatrix} 1 & x_1 & x_1^2 & \cdots & x_1^{p-1} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & x_n & x_n^2 & \cdots & x_n^{p-1} \end{pmatrix}.$$

If $p = 2$, then this is simple linear regression. Next consider the “quadratic regression” model where $p = 3$. That is,

$$\mathbf{X} := \begin{pmatrix} 1 & x_1 & x_1^2 \\ \vdots & \vdots & \vdots \\ 1 & x_n & x_n^2 \end{pmatrix} \Rightarrow \mathbf{X}'\mathbf{X} = \begin{pmatrix} n & \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 \\ \sum_{i=1}^n x_i & \sum_{i=1}^n x_i^2 & \sum_{i=1}^n x_i^3 \\ \sum_{i=1}^n x_i^2 & \sum_{i=1}^n x_i^3 & \sum_{i=1}^n x_i^4 \end{pmatrix}.$$

Because $n \geq 4$, \mathbf{X} is nonsingular if and only if \mathbf{x} is not a vector of constants [a natural condition]. But you see that already it is painful to invert $\mathbf{X}'\mathbf{X}$. This example shows the importance of using a computer in linear models: Even fairly simple models are hard to work with, using only direct calculations.

Assessing Normality

1. Visual data exploration

A big part of our theory of linear models has been under the normal model; that is the most successful part of the theory applies when $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$ where we assumed that $\boldsymbol{\epsilon} \sim N_n(\mathbf{0}, \sigma^2 \mathbf{I})$; equivalently, that $\epsilon_1, \dots, \epsilon_n$ are independent $N(0, \sigma^2)$'s.

A natural question, for a given data set, is to ask, “is the noise coming as i.i.d. $N(0, \sigma^2)$'s”?

Since the noise is not observable in our model, it seems natural that we “estimate” it using the residuals $\mathbf{e} := \mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}$. It stands to reason that if $\boldsymbol{\epsilon}$ is a vector of n i.i.d. $N(0, \sigma^2)$'s, then the histogram of e_1, \dots, e_n should look like a normal density. One should not underestimate the utility of this simple idea; for instance, we should see, roughly, that:

- approximately 68.3% of the e_i 's should fall in $[-S^2, S^2]$,
- approximately 95.4% of the e_i 's should fall in $[-2S^2, 2S^2]$, etc.

This is very useful as a first attempt to assess the normality of the noise in our problem. But it is not conclusive since we cannot assign confidence levels [the method is a little heuristic]. It turns out that this method can be improved upon in several directions, and with only a little more effort.

2. General remarks

We can use the Pearson's χ^2 -test in order to test whether a certain data comes from the $N(\mu_0, \sigma_0^2)$ distribution, where μ_0 and σ_0 are known. Now we wish to address the same problem, but in the more interesting case that μ_0 and σ_0 are *unknown*. [For instance, you may wish to know whether or not you are allowed

to use the usual homoskedasticity assumption in the usual measurement-error model of linear models.]

Here we discuss briefly some solutions to this important problem. Although we write our approach specifically in the context of linear models, these ideas can be developed more generally to test for normality of data in other settings.

3. Histograms

Consider the linear model $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. The pressing question is, “is it true that $\boldsymbol{\epsilon} \sim N_n(\mathbf{0}, \sigma^2 \mathbf{I}_n)$ ”?

To answer this, consider the “residuals,”

$$\mathbf{e} = \mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}.$$

If $\boldsymbol{\epsilon} \sim N_n(\mathbf{0}, \sigma^2 \mathbf{I}_n)$ then one would like to think that the histogram of the e_i 's should look like a normal pdf with mean 0 and variance σ^2 (why?). How close is close? It helps to think more generally.

Consider a sample U_1, \dots, U_n (e.g., $U_i = e_i$). We wish to know where the U_i 's are coming from a normal distribution. The first thing to do is to plot the histogram. In R you type,

```
hist(u,nclass=n)
```

where u denotes the vector of the samples U_1, \dots, U_n and n denotes the number of bins in the histogram.

For instance, consider the following exam data:

```
16.8 9.2 0.0 17.6 15.2 0.0 0.0 10.4 10.4 14.0 11.2 13.6
12.4 14.8 13.2 17.6 9.2 7.6 9.2 14.4 14.8 15.6 14.4 4.4
14.0 14.4 0.0 0.0 10.8 16.8 0.0 15.2 12.8 14.4 14.0 17.2
0.0 14.4 17.2 0.0 0.0 0.0 14.0 5.6 0.0 0.0 13.2 17.6 16.0
16.0 0.0 12.0 0.0 13.6 16.0 8.4 11.6 0.0 10.4 0.0 14.4 0.0
18.4 17.2 14.8 16.0 16.0 0.0 10.0 13.6 12.0 15.2
```

The command `hist(f1.dat,nclass=15)` produces Figure 8.1(a).¹

Try this for different values of `nclass` to see what types of histograms you can obtain. You should always ask, “which one represents the truth the best”? Is there a unique answer?

Now the data U_1, \dots, U_n is probably not coming from a normal distribution if the histogram does not have the “right” shape. Ideally, it would be symmetric, and the tails of the distribution taper off rapidly.

In Figure 8.1(a), there were many students who did not take the exam in question. They received a ‘0’ but this grade should probably not contribute to our knowledge of the distribution of all such grades. Figure 8.1(b) shows the

¹You can obtain this data freely from the website below:
<http://www.math.utah.edu/~davar/math6010/2011/Notes/f1.dat>.

histogram of the same data set when the zeros are removed. [This histogram appears to be closer to a normal density.]

4. QQ-Plots

QQ-plots are a better way to assess how closely a sample follows a certain distribution.

To understand the basic idea note that if U_1, \dots, U_n is a sample from a normal distribution with mean μ and variance σ^2 , then about 68.3% of the sample points should fall in $[\mu - \sigma, \mu + \sigma]$, 95.4% should fall in $[\mu - 2\sigma, \mu + 2\sigma]$, etc.

Now let us be more careful still. Let $U_{1:n} \leq \dots \leq U_{n:n}$ denote the order statistics of U_1, \dots, U_n . Then no matter how you make things precise, the fraction of data “below” $U_{j:n}$ is $(j \pm 1)/n$. So we make a continuity correction and *define* the fraction of the data below $U_{j:n}$ to be $(j - \frac{1}{2})/n$.

Consider the normal “quantiles,” q_1, q_2, \dots, q_n :

$$\Phi(q_j) := \int_{-\infty}^{q_j} \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx = \frac{j - \frac{1}{2}}{n}; \quad \text{i.e., } q_j := \Phi^{-1} \left(\frac{j - \frac{1}{2}}{n} \right).$$

Now suppose once again that $U_1, \dots, U_n \sim N(\mu, \sigma^2)$ is a random [i.i.d.] sample. Let $Z_j := (U_j - \mu)/\sigma$, so that $Z_1, \dots, Z_n \sim N(0, 1)$. The Z ’s are standardized data, and we expect the fraction of the standardized data that fall below q_j to be about $(j - \frac{1}{2})/n$. In other words, we can put together our observations to deduce that $Z_{j:n} \approx q_j$. Because $U_{j:n} = \sigma Z_{j:n} + \mu$, it follows that $U_{j:n} \approx \sigma q_j + \mu$. In other words, we expect the sample order statistics $U_{1:n}, \dots, U_{n:n}$ to be very close to some linear function of the normal quantiles q_1, \dots, q_n . In other words, if U_1, \dots, U_n is a random sample from some normal distribution, then we expect the scatterplot of the pairs $(q_1, U_{1:n}), \dots, (q_n, U_{n:n})$ to follow closely a line. [The slope and intercept are σ and μ , respectively.]

QQ-plots are simply the plots of the $N(0, 1)$ -quantiles q_1, \dots, q_n versus the order statistics $U_{1:n}, \dots, U_{n:n}$. To draw the qqplot of a vector \mathbf{u} in R, you simply type

`qqnorm(u).`

Figure 8.2(a) contains the qq-plot of the exam data we have been studying here.

5. The Correlation Coefficient of the QQ-Plot

In its complete form, the R-command `qqnorm` has the following syntax:

`qqnorm(u, datax = FALSE, plot = TRUE).`

The parameter `u` denotes the data; `datax` is “FALSE” if the data values are drawn on the y -axis (default). It is “TRUE” if you wish to plot $(U_{j:n}, q_j)$ instead of the more traditional $(q_j, U_{j:n})$. The option `plot=TRUE` (default) tells R to plot the qq-plot, whereas `plot=FALSE` produces a vector. So for instance, try

```
V = qqnorm(u, plot = FALSE).
```

This creates two vectors: `V$x` and `V$y`. The first contains the values of all q_j 's, and the second all of the $U_{j:n}$'s. So now you can compute the correlation coefficient of the qq-plot by typing:

```
V = qqnorm(u, plot = FALSE)
cor(V$x, V$y).
```

If you do this for the qq-plot of the grade data, then you will find a correlation of ≈ 0.910 . After censoring out the no-show exams, we obtain a correlation of ≈ 0.971 . This produces a noticeable difference, and shows that the grades are indeed normal.

In fact, one can analyse this procedure statistically [“is the sample correlation coefficient corresponding to the line sufficiently close to ± 1 ”?].

6. Some benchmarks

Figures 8.3 and 8.4 contain four distinct examples. I have used qq-plot in the program environment “R.” The image on the left-hand side of Figure 8.3 shows a simulation of 10000 standard normal random variables (in R, you type `x=rnorm(10000,0,1)`), and its qq-plot is drawn on typing `qqnorm(x)`. In a very strong sense, **this figure is the benchmark**.

The image on the right-hand side of Figure 8.3 shows a simulation of 10000 standard Cauchy random variables. That is, the density function is $f(x) = (1/\pi)(1 + x^2)^{-1}$. This is done by typing `y=rcauchy(10000,0,1)`, and the resulting qq-plot is produced upon typing `qqnorm(y)`. We know that the Cauchy has much fatter tails than normals. For instance,

$$P\{\text{Cauchy} > a\} = \frac{1}{\pi} \int_a^\infty \frac{dx}{1+x^2} \sim \frac{1}{\pi a} \quad (\text{as } a \rightarrow \infty),$$

whereas $P\{N(0,1) > a\}$ decays faster than exponentially.² Therefore, for a large,

$$P\{N(0,1) > a\} \ll P\{\text{Cauchy} > a\}.$$

This heavy-tailedness can be read off in Figure 8.3(b): The Cauchy qq-plot grows faster than linearly on the right-hand side. *And this means that the*

²In fact it can be shown that $\bar{\Phi}(a) := \int_a^\infty \varphi(x) dx \approx a^{-1}\varphi(a)$ as $a \rightarrow \infty$, where φ denotes the $N(0,1)$ pdf. Here is why: Let $G(a) := a^{-1}\varphi(a)$. We know from the fundamental theorem of calculus that $\bar{\Phi}'(a) = -\varphi(a) = -aG(a)$. Also, $G'(a) = -a^{-1}G(a) - aG(a) \approx -aG(a)$ as $a \rightarrow \infty$. In summary: $\bar{\Phi}(a), G(a) \approx$ and $\bar{\Phi}'(a) \approx G'(a)$. Therefore, $\bar{\Phi}(a) \approx G(a)$, thanks to the L'Hôpital's rule of calculus.

standard Cauchy distribution has fatter right-tails. Similar remarks apply to the left tails.

Figure 8.4(a) shows the result of the qq-plot of a simulation of 10000 iid uniform-(0,1) random variables. [To generate these uniform random variables you type, `runif(10000,0,1)`.]


Now uniform-(0,1) random variables have much smaller tails than normals because uniforms are in fact *bounded*. This fact manifests itself in Figure 8.4(a). For instance, we can see that the right-tail of the qq-plot for uniform-(0,1) grows less rapidly than linearly. And this shows that the right-tail of a uniform is much smaller than that of a normal. Similar remarks apply to the left tails.

A comparison of the figures mentioned so far should give you a feeling for how sensitive qq-plots are to the effects of tails. [All are from distributions that are symmetric about their median.] Finally, let us consider Figure 8.4(a), which shows an example of 10000 Gamma random variables with $\alpha = \beta = 1$. You generate them in R by typing `x=rgamma(10000,1,1)`. Gamma distributions are inherently *asymmetric*. You can see this immediately in the qq-plot for Gammas; see Figure 8.4(b). Because Gamma random variables are non-negative, the left tail is much smaller than that of a normal. Hence, the left tail of the qq-plot grows more slowly than linearly. The right tail however is fatter. [This is always the case. However, for the sake of simplicity consider the special case where Gamma=Exponential.] This translates to the faster-than-linear growth of the right-tail of the corresponding qq-plot (Figure 8.4(b)).

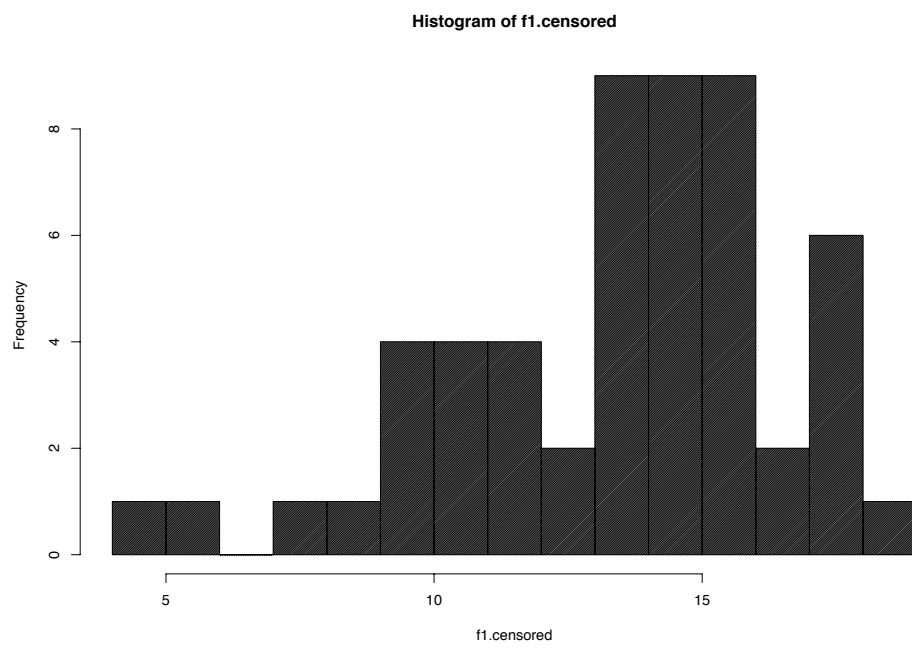
I have shown you Figures 8.3 and 8.4 in order to high-light the basic features of qq-plots in ideal settings. By “ideal” I mean “simulated data,” of course.

Real data does not generally lead to such sleek plots. Nevertheless one learns a lot from simulated data, mainly because simulated data helps identify key issues without forcing us to have to deal with imperfections and other flaws.

But it is important to keep in mind that it is real data that we are ultimately after. And so the histogram and qq-plot of a certain real data set are depicted in Figure 8.5. Have a careful look and ask yourself a number of questions: Is the data normally distributed? Can you see how the shape of the histogram manifests itself in the shape and gaps of the qq-plot? Do the tails look like those of a normal distribution? To what extent is the “gap” in the histogram “real”? By this I mean to ask what do you think might happen if we change the bin-size of the histogram in Figure 8.5?




(a) Grades




(b) Censored Grades

Figure 8.1. Histogram of grades and censored grades



(a) QQ-plot of grades



(b) QQ-plot of censored grades

Figure 8.2. QQ-plots of grades and censored grades

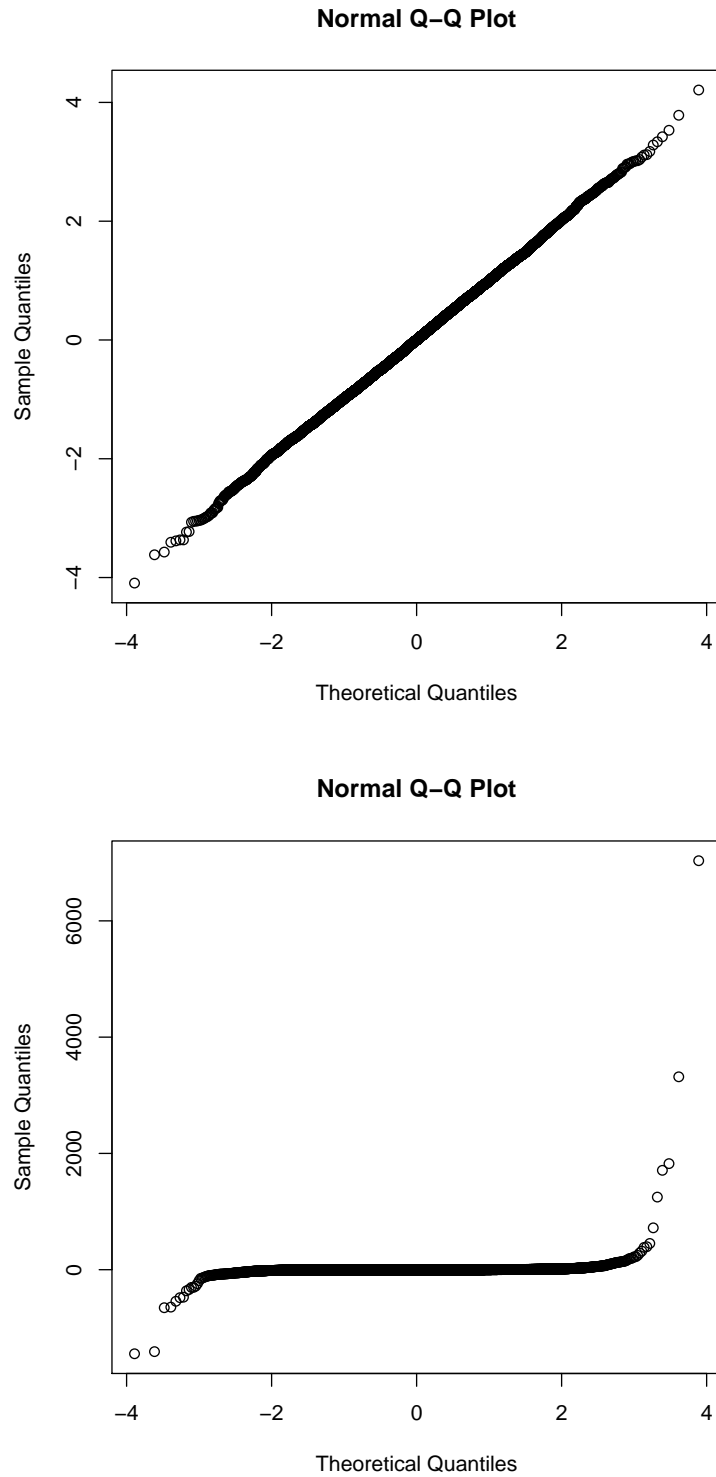


Figure 8.3. (a) is $N(0, 1)$ data; (b) is Cauchy data

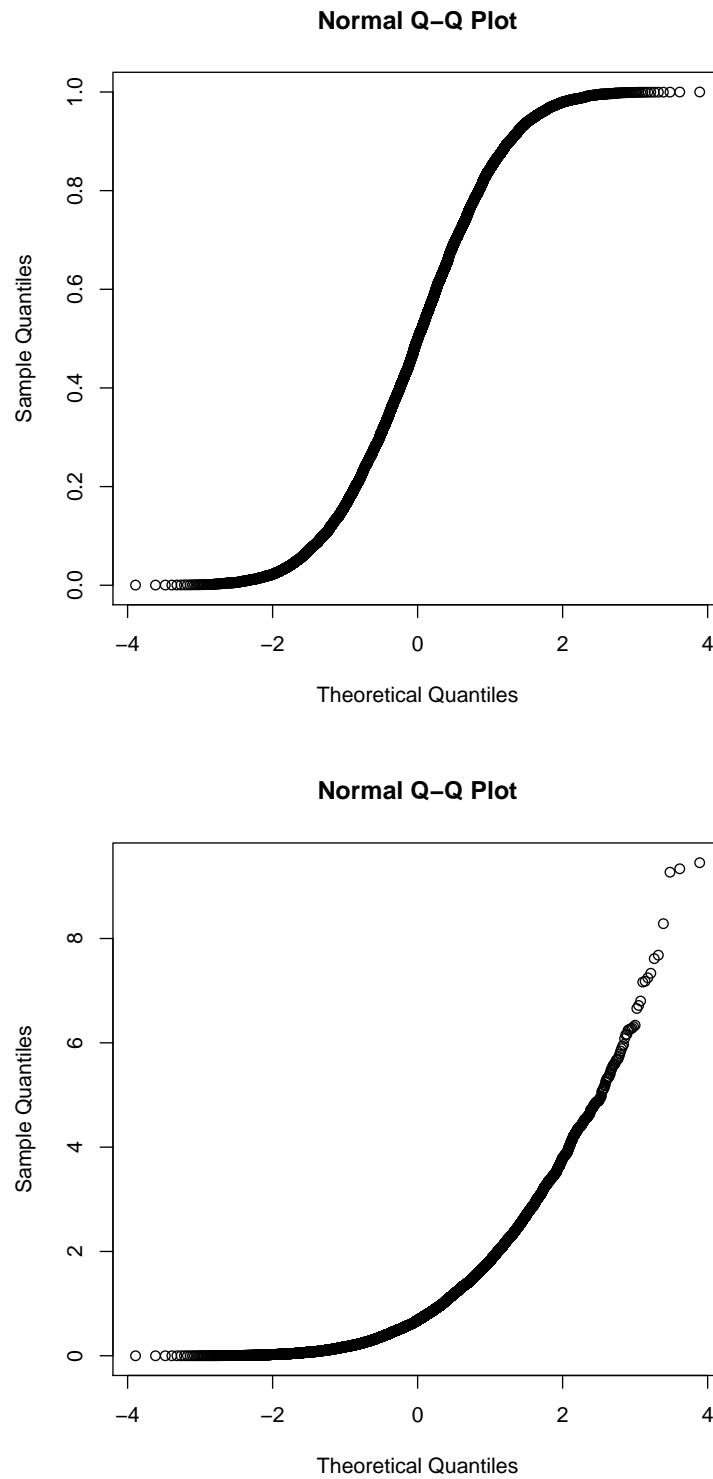


Figure 8.4. (a) is the qq-plot of $\text{unif}(0, 1)$; (b) is the qq-plot of a $\text{Gamma}(1, 1)$.

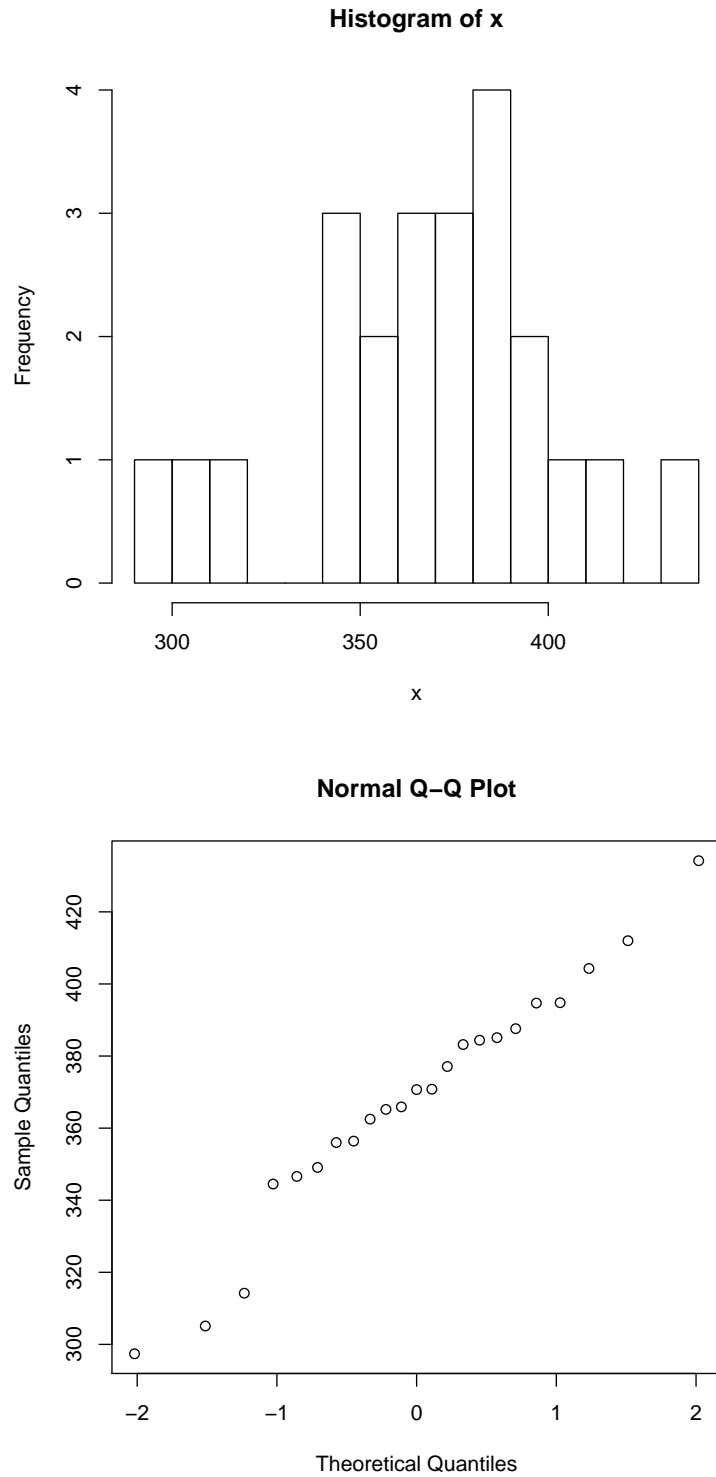


Figure 8.5. Histogram and qq-plot of the data

Hypothesis Testing

Throughout, we assume the normal-error linear model that is based on the model

$$y = \beta_1 x_1 + \cdots + \beta_p x_p + \text{noise}.$$

[Note the slight change in the notation.]

1. A test for one parameter

Suppose we want to test to see whether or not the $(\ell + 1)$ st x -variable has a [linear] effect on the y variable. Of course, $1 \leq \ell \leq p$, so we are really testing the statistical hypothesis

$$H_0 : \beta_\ell = 0.$$

Since $\hat{\boldsymbol{\beta}} \sim N_p(\boldsymbol{\beta}, \sigma^2(\mathbf{X}'\mathbf{X})^{-1})$, it follows that

$$\hat{\beta}_\ell \sim N(\beta_\ell, \sigma^2 [(\mathbf{X}'\mathbf{X})^{-1}]_{\ell,\ell}).$$

Because S is independent of $\hat{\boldsymbol{\beta}}$ and hence $\hat{\beta}_\ell$, and since $S^2/\sigma^2 \sim \chi_{n-p}^2/(n-p)$,

$$\frac{\hat{\beta}_\ell - \beta_\ell}{S \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{\ell,\ell}}} = \frac{\sigma}{S} \cdot \frac{\hat{\beta}_\ell - \beta_\ell}{\sigma \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{\ell,\ell}}} \sim t_{n-p}.$$

Therefore, it is now a routine matter to set up a t -test for $H_0 : \beta_\ell = 0$. As usual, testing has implications that are unattractive; it is much better to present a confidence interval [which you can then use for a test if you want, any way]: A $(1 - \alpha) \times 100\%$ confidence interval for β_ℓ is

$$\hat{\beta}_\ell \pm S t_{n-p}^{\alpha/2} \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{\ell,\ell}}.$$

If you really insist on performing a level- α test for β_ℓ , it suffices to check to see if this confidence interval contains 0. If 0 is not in the confidence interval then you reject. Otherwise, you do nothing.

2. Least-squares estimates for contrasts

We wish to study a more general problem. Recall that our model has the form

$$y = \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_p x_p + \text{noise}.$$

If we sample then the preceding becomes $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$, as before. As part of model verification, we might ask to see if $(x_i)_{i \in J}$ should be excised from the model, where $J := \{\ell, \ell + 1, \dots, r\}$ is a subset of the index $\{1, \dots, n\}$. In other words, we ask

$$H_0 : \beta_\ell = \cdots = \beta_r = 0.$$

Note that we can translate the preceding, using the language of matrix analysis, as $H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{0}$, where

$$\mathbf{A} := \begin{pmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix},$$

where the identity matrix in the middle is $(r - \ell + 1) \times (r - \ell + 1)$; it starts on position (ℓ, ℓ) and runs $r - \ell$ units in rows and in columns.

Now we ask a slightly more general question [it pays to do this, as it turns out]: Suppose \mathbf{A} is a $q \times p$ matrix of full rank $q \leq p$, and we are interested in testing the hypothesis,

$$H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{0}. \quad (1)$$

The first question to ask is, “how can we estimate $\boldsymbol{\beta}$ ”? The answer is given to us by the principle of least squares: We write [as before]

$$\mathbf{Y} = \boldsymbol{\theta} + \boldsymbol{\epsilon}, \quad \text{where } \boldsymbol{\theta} := \mathbf{X}\boldsymbol{\beta},$$

and $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_n)'$ are mean-zero random variables, and first find the least-squares estimate $\hat{\boldsymbol{\theta}}_{H_0}$ of $\boldsymbol{\theta}$, under the assumption that H_0 is valid. That is, we seek to minimize $\|\mathbf{Y} - \mathbf{X}\mathbf{b}\|^2$ over all p -vectors \mathbf{b} such that $\mathbf{A}\mathbf{b} = \mathbf{0}$. The optimal value yields $\hat{\boldsymbol{\theta}}_{H_0} = \mathbf{X}\hat{\boldsymbol{\beta}}_{H_0}$. Then we obtain $\hat{\boldsymbol{\beta}}_{H_0}$ by noticing that if \mathbf{X} has full rank, then $\hat{\boldsymbol{\beta}}_{H_0} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\hat{\boldsymbol{\theta}}_{H_0}$.

Now it follows by differentiation [or just geometrically] that $\hat{\boldsymbol{\theta}}_{H_0}$ is the projection of \mathbf{Y} onto the subspace \mathcal{G} of all vectors of the form $\boldsymbol{\vartheta} = \mathbf{X}\mathbf{b}$ that satisfy $\mathbf{A}\mathbf{b} = \mathbf{0}$, where \mathbf{b} is a p -vector. We can simplify this description a little when \mathbf{X} has full rank. Note that whenever $\boldsymbol{\vartheta} = \mathbf{X}\mathbf{b}$, we can solve to get $\mathbf{b} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\vartheta}$. Therefore, it follows that—when \mathbf{X} has full rank— $\hat{\boldsymbol{\theta}}_{H_0}$ is the projection of the observations vector \mathbf{Y} onto the subspace \mathcal{G} of all vectors of the form $\boldsymbol{\vartheta}$ that satisfy

$$\mathbf{A}_1\boldsymbol{\vartheta} = \mathbf{0}, \quad \text{where } \mathbf{A}_1 := \mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'.$$

In other words, \mathcal{G} is the subspace of $\mathcal{C}(\mathbf{X})$, whose every element $\boldsymbol{\vartheta}$ is orthogonal to every row of \mathbf{A}_1 . In symbols,

$$\mathcal{G} = \mathcal{C}(\mathbf{X}) \cap [\mathcal{C}(\mathbf{A}'_1)]^\perp.$$

Because $\mathbf{A}'_1 \mathbf{b} = \mathbf{X} \mathbf{c}$ for $\mathbf{c} := (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A} \mathbf{b}$, it follows that $\mathcal{C}(\mathbf{A}'_1)$ is a subspace of $\mathcal{C}(\mathbf{X})$. Therefore, we can apply the Pythagorean property to see that

$$\begin{aligned} \hat{\boldsymbol{\theta}}_{H_0} &= \mathbf{P}_{\mathcal{G}} \mathbf{Y} = \mathbf{P}_{\mathcal{C}(\mathbf{X}) \cap [\mathcal{C}(\mathbf{A}'_1)]^\perp} \mathbf{Y} \\ &= \mathbf{P}_{\mathcal{C}(\mathbf{X})} \mathbf{Y} - \mathbf{P}_{\mathcal{C}(\mathbf{A}'_1)} \mathbf{Y} \\ &= \hat{\boldsymbol{\theta}} - \mathbf{A}'_1 (\mathbf{A}_1 \mathbf{A}'_1)^{-1} \mathbf{A}_1 \mathbf{Y}. \end{aligned}$$

Now

$$\mathbf{A}_1 \mathbf{A}'_1 = \mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{X} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}' = \mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}'.$$

Therefore,

$$\hat{\boldsymbol{\theta}}_{H_0} = \hat{\boldsymbol{\theta}} - \mathbf{X} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}' [\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}']^{-1} \mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{Y}.$$

Aside: How do we know that $\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}'$ is nonsingular? Note that $\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}'$ is positive semidefinite. Now $\mathbf{X}' \mathbf{X}$ is positive definite; therefore, so is its inverse. Therefore, we can write $(\mathbf{X}' \mathbf{X})^{-1} = \mathbf{B}^2 = \mathbf{B} \mathbf{B}'$, where $\mathbf{B} := (\mathbf{X}' \mathbf{X})^{-1/2}$ is the square root of $(\mathbf{X}' \mathbf{X})^{-1}$. In this way we find that the rank of $\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}'$ is the same as the rank of $\mathbf{A} \mathbf{A}'$. Since \mathbf{A} has full rank, $\mathbf{A} \mathbf{A}'$ is invertible. Equivalently, full rank. Equivalently, $\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}'$ is a full-rank positive definite matrix; hence nonsingular.

The vector $\hat{\boldsymbol{\theta}}_{H_0} := \hat{\mathbf{Y}}_{H_0} = \mathbf{X} \hat{\boldsymbol{\beta}}_{H_0}$ is the vector of fitted values, assuming that H_0 is correct. Therefore, the least-squares estimate for $\boldsymbol{\beta}$ —under H_0 —is

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{H_0} &:= (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \hat{\boldsymbol{\theta}}_{H_0} = \hat{\boldsymbol{\beta}} - (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}' [\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}']^{-1} \mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{Y} \\ &= \hat{\boldsymbol{\beta}} - (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}' [\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}']^{-1} \mathbf{A} \hat{\boldsymbol{\beta}} \\ &= \left(\mathbf{I} - (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}' [\mathbf{A} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{A}']^{-1} \mathbf{A} \right) \hat{\boldsymbol{\beta}}. \end{aligned}$$

This can be generalized further as follows: Suppose we wish to test

$$H_0 : \mathbf{A} \boldsymbol{\beta} = \mathbf{c},$$

where \mathbf{c} is a known q -vector [we just studied this in the case that $\mathbf{c} = \mathbf{0}$]. Then we reduce the problem to the previous one as follows: First find a known p -vector $\boldsymbol{\beta}_0$ such that $\mathbf{A} \boldsymbol{\beta}_0 = \mathbf{c}$. Then, create a new parametrization of our problem by setting

$$\boldsymbol{\gamma} := \boldsymbol{\beta} - \boldsymbol{\beta}_0,$$

and

$$\tilde{\mathbf{Y}} := \mathbf{X} \boldsymbol{\gamma} + \boldsymbol{\varepsilon}, \quad \text{equivalently} \quad \tilde{\mathbf{Y}} := \mathbf{Y} - \mathbf{X} \boldsymbol{\beta}_0.$$

Since $\mathbf{A}\boldsymbol{\gamma} = \mathbf{0}$, we know the least-squares estimate $\hat{\boldsymbol{\gamma}}_{H_0}$ is given by

$$\hat{\boldsymbol{\gamma}}_{H_0} = \left(\mathbf{I} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{A} \right) \hat{\boldsymbol{\gamma}},$$

where

$$\hat{\boldsymbol{\gamma}} := (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\tilde{\mathbf{Y}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} - \boldsymbol{\beta}_0 = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0.$$

In other words,

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{H_0} - \boldsymbol{\beta}_0 &= \left(\mathbf{I} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{A} \right) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \\ &= \left(\mathbf{I} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{A} \right) \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{A} \boldsymbol{\beta}_0 \\ &= \left(\mathbf{I} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{A} \right) \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A} [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{c}. \end{aligned}$$

In this way, we have discovered the following:

Theorem 1. Consider once again the general linear model $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. If $\mathbf{A}_{q \times p}$ and $\mathbf{c}_{q \times 1}$ are known, and \mathbf{A} has full rank $q \leq p$, then the least-squares estimate for $\boldsymbol{\beta}$ —under the null hypothesis $H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{c}$ —is

$$\hat{\boldsymbol{\beta}}_{H_0} = \boldsymbol{\Theta}\hat{\boldsymbol{\beta}} + \boldsymbol{\mu},$$

where

$$\boldsymbol{\Theta} := \left(\mathbf{I} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{A} \right), \quad (2)$$

and

$$\boldsymbol{\mu} = \boldsymbol{\mu}(\mathbf{c}) := (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} \mathbf{c}, \quad (3)$$

provided that \mathbf{X} has full rank.

3. The normal model

Now consider the same problem under the normal model. That is, we consider $H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{c}$ under the assumption that $\boldsymbol{\epsilon} \sim N_p(\mathbf{0}, \sigma^2 \mathbf{I})$.

Theorem 2. Consider the normal-error linear model $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. If $\mathbf{A}_{q \times p}$ and $\mathbf{c}_{q \times 1}$ are known, and \mathbf{A} has full rank $q \leq p$, then the least-squares estimate for $\boldsymbol{\beta}$ —under the null hypothesis $H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{c}$ —satisfies

$$\hat{\boldsymbol{\beta}}_{H_0} \sim N_p(\boldsymbol{\beta}, \sigma^2 \boldsymbol{\Theta}(\mathbf{X}'\mathbf{X})^{-1} \boldsymbol{\Theta}'),$$

provided that \mathbf{X} has full rank.

Indeed, since

$$\hat{\boldsymbol{\beta}} \sim N_p(\boldsymbol{\beta}, \sigma^2(\mathbf{X}'\mathbf{X})^{-1}) \quad \text{and} \quad \hat{\boldsymbol{\beta}}_{H_0} = \boldsymbol{\Theta}\hat{\boldsymbol{\beta}} + \boldsymbol{\mu},$$

it follows that

$$\hat{\boldsymbol{\beta}}_{H_0} \sim N_p(\boldsymbol{\Theta}\boldsymbol{\beta} + \boldsymbol{\mu}, \sigma^2 \boldsymbol{\Theta}(\mathbf{X}'\mathbf{X})^{-1} \boldsymbol{\Theta}').$$

Therefore, it remains to check that $\boldsymbol{\Theta}\boldsymbol{\beta} + \boldsymbol{\mu} = \boldsymbol{\beta}$ when $\mathbf{A}\boldsymbol{\beta} = \mathbf{c}$. But this is easy to see directly.

Next we look into inference for σ^2 . Recall that our estimation of σ^2 was based on $\text{RSS} := \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2$. Under H_0 , we do the natural thing and estimate σ^2 instead by

$$\begin{aligned} \text{RSS}_{H_0} &:= \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}_{H_0}\|^2 \\ &= \left\| \underbrace{\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}}_{\mathcal{T}_1} - \underbrace{\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1}[\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]}_{\mathcal{T}_2} \right\|^2. \end{aligned}$$

I claim that \mathcal{T}_2 is orthogonal to \mathcal{T}_1 ; indeed,

$$\begin{aligned} \mathcal{T}_2'\mathcal{T}_1 &= [\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]'\left[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'\right]^{-1}\mathbf{A}\overbrace{(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}}^{\hat{\boldsymbol{\beta}}} \\ &\quad - [\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]'\left[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'\right]^{-1}\mathbf{A}\underbrace{(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}}_I\hat{\boldsymbol{\beta}} \\ &= 0. \end{aligned}$$

Therefore, the Pythagorean property tells us that

$$\begin{aligned} \text{RSS}_{H_0} &= \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2 + \|\mathcal{T}_2\|^2 \\ &= \text{RSS} + \|\mathcal{T}_2\|^2. \end{aligned}$$

Next we compute

$$\begin{aligned} \|\mathcal{T}_2\|^2 &= \mathcal{T}_2'\mathcal{T}_2 \\ &= [\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]'\left[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'\right]^{-1}\underbrace{\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\underbrace{\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}}_I\mathbf{A}'}_{\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'}\left[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'\right]^{-1}[\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}] \\ &= [\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]'\left[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'\right]^{-1}[\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]. \end{aligned}$$

In other words,

$$\text{RSS}_{H_0} = \text{RSS} + [\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]'\left[\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'\right]^{-1}[\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}]. \quad (4)$$

Moreover, the two terms on the right-hand side are independent because $\hat{\boldsymbol{\beta}}$ and $\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}$ —hence $\hat{\boldsymbol{\beta}}$ and $\text{RSS} = \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2$ —are independent. Now we know the distribution of $\text{RSS} := (n - p)S^2 \sim \sigma^2\chi_{n-p}^2$. Therefore, it remains

to find the distribution of the second term on the right-hand side of (4). But

$$\mathbf{A}\hat{\boldsymbol{\beta}} \sim N_q(\mathbf{A}\boldsymbol{\beta}, \sigma^2 \mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}') \stackrel{H_0}{=} N_q\left(\mathbf{c}, \sigma^2 \underbrace{\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}'}_{:=\boldsymbol{\Sigma}}\right).$$

Therefore, $\mathbf{Z} := \sigma^{-1}\boldsymbol{\Sigma}^{-1/2}(\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c}) \sim N_q(\mathbf{0}, \mathbf{I}_{q \times q})$. Also, we can write the second term on the right-hand side of (4) as

$$[\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c}]' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} [\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c}] = \sigma^2 \mathbf{Z}'\mathbf{Z} = \sigma^2 \|\mathbf{Z}\|^2 \sim \sigma^2 \chi_q^2.$$

Let us summarize our efforts.

Theorem 3. Consider normal-error linear model $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. Suppose $\mathbf{A}_{q \times p}$ and $\mathbf{c}_{q \times 1}$ are known, and \mathbf{A} has full rank $q \leq p$. Then under the null hypothesis $H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{c}$, we can write

$$\text{RSS}_{H_0} = \text{RSS} + W,$$

provided that \mathbf{X} has full rank, where RSS and W are independent, we recall that $\text{RSS} \sim \sigma^2(n-p)\chi_{n-p}^2$, and $W \sim \sigma^2\chi_q^2$. In particular,

$$\begin{aligned} \frac{(\text{RSS}_{H_0} - \text{RSS})/q}{\text{RSS}/(n-p)} &\stackrel{H_0}{\sim} \frac{\chi_q^2/q}{\chi_{n-p}^2/(n-p)} \quad [\text{the two } \chi^2\text{'s are independent}] \\ &= F_{q, n-p}. \end{aligned}$$

See your textbook for the distribution of this test statistic under the alternative [this is useful for power computations]. The end result is a “noncentral F distribution.”

4. Examples

1. A measurement-error model. For our first example, consider a random sample $Y_1, \dots, Y_n \sim N(\mu, \sigma^2)$; equivalently,

$$Y_i = \mu + \epsilon_i \quad (1 \leq i \leq n),$$

where $\boldsymbol{\epsilon} \sim N_n(\mathbf{0}, \sigma^2 \mathbf{I})$. This is a linear model with $p = 1$, $\mathbf{X} := \mathbf{1}_{n \times 1}$, and $\boldsymbol{\beta} := \mu$. Recall that $(\mathbf{X}'\mathbf{X})^{-1} = 1/n$ and hence $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} = \bar{Y}$.

If we test $H_0 : \mu = \mu_0$ for a μ_0 that is known, then $\mathbf{A} = 1$ is a 1×1 matrix ($q = 1$) and $\mathbf{A}\boldsymbol{\beta} = \mathbf{c}$ with $\mathbf{c} = \mu_0$.

Given that H_0 is true, the least-squares estimator of μ $[\hat{\boldsymbol{\beta}}_{H_0}]$ is

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{H_0} &:= \hat{\mu}_{H_0} = \hat{\boldsymbol{\beta}} + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} (\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}) \\ &= \bar{Y} + \frac{1}{n} \cdot n \cdot (\mu_0 - \bar{Y}) = \mu_0. \end{aligned}$$

[Is this sensible?] And

$$\text{RSS} = \|\mathbf{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2 = \sum_{i=1}^n (Y_i - \bar{Y})^2 = ns_y^2.$$

Therefore,

$$\begin{aligned} \text{RSS}_{H_0} - \text{RSS} &= (\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c})' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} (\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c}) \\ &= n(\bar{Y} - \mu_0)^2. \end{aligned}$$

And

$$\frac{(\text{RSS}_{H_0} - \text{RSS})/q}{\text{RSS}/(n-p)} = \frac{(\bar{Y} - \mu_0)^2}{s_y^2/(n-1)} \stackrel{H_0}{\sim} F_{1,n-1}.$$

But

$$\frac{\bar{Y} - \mu_0}{s_y/\sqrt{n-1}} \stackrel{H_0}{\sim} t_{n-1}.$$

Therefore, in particular, $t_k^2 = F_{1,k}$.

2. Simple linear regression. Here,

$$Y_i = \alpha + \beta x_i + \varepsilon_i \quad (1 \leq i \leq n).$$

Therefore, $p = 2$,

$$\boldsymbol{\beta} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \quad \text{and} \quad \mathbf{X} = \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix}.$$

Recall that the least-squares estimates of α and β are

$$\hat{\alpha} = \bar{Y} - \hat{\beta}\bar{x}, \quad \hat{\beta} = \frac{rs_y}{s_x}.$$

Now consider testing the hypothesis,

$$H_0 : \beta = 0, \alpha = \mu_0,$$

where μ_0 is known.

Let $\mathbf{c} = (\mu_0, 0)'$ and $\mathbf{A} = \mathbf{I}_2$, so that $q = 2$. Then, H_0 is the same as $H_0 : \mathbf{A}\boldsymbol{\beta} = \mathbf{c}$. We have

$$\hat{\boldsymbol{\beta}}_{H_0} = \hat{\boldsymbol{\beta}} + (\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\beta}}) = \mathbf{c} = \begin{pmatrix} \mu_0 \\ 0 \end{pmatrix}.$$

[Is this sensible?]

Now,

$$\text{RSS}_{H_0} - \text{RSS} = (\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c})' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} (\mathbf{A}\hat{\boldsymbol{\beta}} - \mathbf{c}) = (\hat{\boldsymbol{\beta}} - \mathbf{c})' (\mathbf{X}'\mathbf{X}) (\hat{\boldsymbol{\beta}} - \mathbf{c}).$$

Now,

$$\mathbf{X}'\mathbf{X} = n \begin{pmatrix} 1 & \bar{x} \\ \bar{x} & \bar{x}^2 \end{pmatrix} \Rightarrow \hat{\boldsymbol{\beta}} - \mathbf{c} = \begin{pmatrix} \bar{Y} - \hat{\beta}\bar{x} - \mu_0 \\ \hat{\beta} \end{pmatrix}.$$

Therefore,

$$(\hat{\beta} - c)'(\mathbf{X}'\mathbf{X}) = n(\bar{Y} - \mu_0)(1, \bar{x}),$$

whence

$$\text{RSS}_{H_0} - \text{RSS} = n(\bar{Y} - \mu_0)^2.$$

Next we compute

$$\text{RSS} = \|\mathbf{Y} - \mathbf{X}\hat{\beta}\|^2 = \sum_{i=1}^n [Y_i - (\mathbf{X}\hat{\beta})_i]^2.$$

Since

$$\mathbf{X}\hat{\beta} = \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix} \begin{pmatrix} \bar{Y} - \hat{\beta}\bar{x} \\ \hat{\beta} \end{pmatrix} = (\bar{Y} + \hat{\beta}(x_i - \bar{x}))_{i=1}^n = \left(\bar{Y} + \frac{r s_y}{s_x}(x_i - \bar{x}) \right)_{i=1}^n,$$

it follows that

$$\begin{aligned} \text{RSS} &= \sum_{i=1}^n \left[(Y_i - \bar{Y})^2 + \frac{r^2 s_y^2}{s_x^2} (x_i - \bar{x})^2 - 2 \frac{r s_y}{s_x} (Y_i - \bar{Y})(x_i - \bar{x}) \right] \\ &= n s_y^2 + n r^2 s_y^2 - \frac{2 n r s_y}{s_x} \sum_{i=1}^n (Y_i - \bar{Y})(x_i - \bar{x}) \\ &= n s_y^2 + n r^2 s_y^2 - 2 n r^2 s_y^2 \\ &= n s_y^2 (1 - r^2). \end{aligned}$$

Therefore,

$$\frac{(\bar{Y} - \mu_0)^2}{s_y^2 (1 - r^2)} \stackrel{H_0}{\sim} F_{2, n-2}.$$

3. Two-sample mean. Consider two populations: $N(\mu_1, \sigma^2)$ and $N(\mu_2, \sigma^2)$ with equal variances. We wish to know if $\mu_1 = \mu_2$. Take two independent random samples,

$$y_{1,1}, \dots, y_{1,n_1} \sim N(\mu_1, \sigma^2),$$

$$y_{2,1}, \dots, y_{2,n_2} \sim N(\mu_2, \sigma^2).$$

We have a linear model: $p = 2$, $n = n_1 + n_2$,

$$\mathbf{Y} = \begin{pmatrix} y_{1,1} \\ \vdots \\ y_{1,n_1} \\ y_{2,1} \\ \vdots \\ y_{2,n_2} \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} 1 & 0 \\ \vdots & \vdots \\ 1 & 0 \\ 0 & 1 \\ \vdots & \vdots \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} \mathbf{1}_{n_1 \times 1} & \mathbf{0}_{n_1 \times 1} \\ \mathbf{0}_{n_2 \times 1} & \mathbf{1}_{n_2 \times 1} \end{pmatrix}.$$

In particular,

$$\mathbf{X}'\mathbf{X} = \begin{pmatrix} n_1 & 0 \\ 0 & n_2 \end{pmatrix} \Rightarrow (\mathbf{X}'\mathbf{X})^{-1} = \begin{pmatrix} n_1^{-1} & 0 \\ 0 & n_2^{-1} \end{pmatrix}.$$

So now consider

$$H_0 : \mu_1 = \mu_2 \iff H_0 : \mu_1 = \mu_2 \iff H_0 : \underbrace{(1, -1)}_{\mathbf{A}} \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} = \mathbf{0}.$$

That is, $q = 1$, $\mathbf{A} := (1, -1)$, and $\mathbf{c} = 0$. In this way we find that

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} = \begin{pmatrix} \bar{y}_{1,\bullet} \\ \bar{y}_{2,\bullet} \end{pmatrix}.$$

[Does this make intuitive sense?]

In order to find $\hat{\boldsymbol{\beta}}_{H_0}$, we first compute

$$\mathbf{A}\hat{\boldsymbol{\beta}} = (1, -1) \begin{pmatrix} \bar{y}_{1,\bullet} \\ \bar{y}_{2,\bullet} \end{pmatrix} = \bar{y}_{1,\bullet} - \bar{y}_{2,\bullet}.$$

Also,

$$(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' = \begin{pmatrix} n_1^{-1} & 0 \\ 0 & n_2^{-1} \end{pmatrix} \begin{pmatrix} 1 \\ -1 \end{pmatrix} = \begin{pmatrix} n_1^{-1} \\ -n_2^{-1} \end{pmatrix},$$

so that

$$\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' = \frac{1}{n_2} + \frac{1}{n_2} = \frac{n}{n_1 n_2}.$$

Now we put things together:

$$\begin{aligned} \hat{\boldsymbol{\beta}}_{H_0} &= \hat{\boldsymbol{\beta}} + \begin{pmatrix} n_1^{-1} \\ -n_2^{-1} \end{pmatrix} \frac{n_1 n_2}{n} (\bar{y}_{1,\bullet} - \bar{y}_{2,\bullet}) \\ &= \hat{\boldsymbol{\beta}} + \begin{pmatrix} \frac{n_2}{2} (\bar{y}_{1,\bullet} - \bar{y}_{2,\bullet}) \\ -\frac{n_1}{n} (\bar{y}_{1,\bullet} - \bar{y}_{2,\bullet}) \end{pmatrix} \\ &= \begin{pmatrix} \frac{n_1}{n} \bar{y}_{1,\bullet} + \frac{n_2}{n} \bar{y}_{2,\bullet} \\ \frac{n_1}{n} \bar{y}_{1,\bullet} + \frac{n_2}{n} \bar{y}_{2,\bullet} \end{pmatrix}. \end{aligned}$$

Since $n_2 \bar{y}_{2,\bullet} = \sum_{j=1}^{n_2} y_{j2,j}$ and $n_1 \bar{y}_{1,\bullet} = \sum_{i=1}^{n_1} y_{i1,i}$, it follows that

$$\hat{\boldsymbol{\beta}}_{H_0} = \begin{pmatrix} \bar{y}_{\bullet,\bullet} \\ \bar{y}_{\bullet,\bullet} \end{pmatrix}.$$

[Does this make sense?] Since

$$\mathbf{X}\hat{\boldsymbol{\beta}} = \begin{pmatrix} 1 & 0 \\ \vdots & \vdots \\ 1 & 0 \\ 0 & 1 \\ \vdots & \vdots \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \bar{y}_{1,\bullet} \\ \bar{y}_{2,\bullet} \end{pmatrix} = \begin{pmatrix} \bar{y}_{1,\bullet} \\ \vdots \\ \bar{y}_{1,\bullet} \\ \bar{y}_{2,\bullet} \\ \vdots \\ \bar{y}_{2,\bullet} \end{pmatrix} = \begin{pmatrix} \bar{y}_{1,\bullet} \mathbf{1}_{n_1 \times 1} \\ \bar{y}_{2,\bullet} \mathbf{1}_{n_2 \times 1} \end{pmatrix},$$

we have

$$\begin{aligned} \text{RSS} &= \sum_{j=1}^{n_1} (y_{1j} - \bar{y}_{1,\bullet})^2 + \sum_{j=1}^{n_2} (y_{2j} - \bar{y}_{2,\bullet})^2 \\ &= n_1 s_1^2 + n_2 s_2^2. \end{aligned}$$

In particular,

$$\frac{\text{RSS}}{n-p} = \frac{n_1}{n-2} s_1^2 + \frac{n_2}{n-2} s_2^2 := s_p^2$$

is the so-called “pooled variance.”

Similarly,

$$\text{RSS}_{H_0} - \text{RSS} = \frac{n_1 n_2}{n} (\bar{y}_{1,\bullet} - \bar{y}_{2,\bullet})^2.$$

Therefore,

$$\frac{\frac{n_1 n_2}{n} (\bar{y}_{1,\bullet} - \bar{y}_{2,\bullet})^2}{s_p^2} \stackrel{H_0}{\sim} F_{1, n-2} \Rightarrow \frac{\bar{y}_{1,\bullet} - \bar{y}_{2,\bullet}}{s_p \sqrt{\frac{n}{n_1 n_2}}} \stackrel{H_0}{\sim} t_{n-2}.$$

4. ANOVA: One-way layout. Consider p populations that are respectively distributed as $N(\mu_1, \sigma^2), \dots, N(\mu_p, \sigma^2)$. We wish to test

$$H_0 : \mu_1 = \dots = \mu_p.$$

We have seen that we are in the setting of linear models, so we can compute $\hat{\boldsymbol{\beta}}_{H_0}$ etc. that way. I will leave this up to you and compute directly instead: Sample $y_{j,1}, \dots, y_{j,n_j}$ i.i.d. $N(\mu_j, \sigma^2)$ [independent also as j varies]. Then we vectorize:

$$\mathbf{Y} := \begin{pmatrix} y_{1,1} \\ \vdots \\ y_{1,n_1} \\ \vdots \\ y_{p,1} \\ \vdots \\ y_{p,n_p} \end{pmatrix}; \quad \text{etc.}$$

Instead we now find $\hat{\beta}$ directly by solving

$$\min_{\mu} \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \mu_i)^2.$$

That is, compute

$$\frac{\partial}{\partial \mu_i} \sum_{j=1}^{n_i} (y_{i,j} - \mu_i)^2 = - \sum_{j=1}^{n_i} 2 (y_{i,j} - \mu_i) \equiv 0 \implies \hat{\mu}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} y_{i,j} = \bar{y}_{i,\bullet}.$$

This yields

$$\hat{\beta} = \begin{pmatrix} \bar{y}_{1,\bullet} \\ \vdots \\ \bar{y}_{n,\bullet} \end{pmatrix}.$$

What about $\hat{\beta}_{H_0}$? Under H_0 , $\mu_1 = \dots = \mu_p \equiv \mu$ and so $q = p - 1$. So we have

$$\min_{\mu} \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \mu)^2 \implies \hat{\beta}_{H_0} = \begin{pmatrix} \bar{y}_{\bullet,\bullet} \\ \vdots \\ \bar{y}_{\bullet,\bullet} \end{pmatrix}.$$

Also,

$$\text{RSS} = \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet})^2,$$

and

$$\begin{aligned} \text{RSS}_{H_0} &= \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{\bullet,\bullet})^2 = \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet} + \bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet})^2 \\ &= \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet})^2 + 2 \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet}) (\bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet}) \\ &\quad + \sum_{i=1}^p \sum_{j=1}^{n_i} (\bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet})^2 \\ &= \sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet})^2 + \sum_{i=1}^p \sum_{j=1}^{n_i} (\bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet})^2 \\ &= \text{RSS} + \sum_{i=1}^p n_i (\bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet})^2. \end{aligned}$$

It follows from the general theory that

$$\frac{\sum_{i=1}^p n_i (\bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet})^2 / (p-1)}{\sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet})^2 / (n-p)} \stackrel{H_0}{\sim} F_{p-1, n-p}.$$

“Statistical interpretation”:

$$\frac{\sum_{i=1}^p n_i (\bar{y}_{i,\bullet} - \bar{y}_{\bullet,\bullet})^2}{p - 1} = \text{The variation between the samples;}$$

whereas

$$\frac{\sum_{i=1}^p \sum_{j=1}^{n_i} (y_{i,j} - \bar{y}_{i,\bullet})^2}{n - p} = \text{The variation within the samples.}$$

Therefore,

$$\text{RSS}_{H_0} = \text{Variation between} + \text{Variation within} = \text{Total variation.}$$

Confidence Intervals and Sets

Throughout we adopt the normal-error model, and wish to say some things about the construction of confidence intervals [and sets] for the parameters $\beta_0, \dots, \beta_{p-1}$.

1. Confidence intervals for one parameter

Suppose we want a confidence intervals for the i th parameter β_i , where $1 \leq i \leq p$ is fixed. Recall that

$$\frac{\hat{\beta}_i - \beta_i}{\sqrt{\sigma^2 [(\mathbf{X}'\mathbf{X})^{-1}]_{i,i}}} \sim N(0, 1),$$

and

$$S^2 := \frac{\text{RSS}}{n-p} = \frac{1}{n-p} \sum_{i=1}^n (Y_i - (\mathbf{x}\hat{\boldsymbol{\beta}})_i)^2 \sim \frac{\sigma^2 \chi_{n-p}^2}{n-p},$$

and is independent of $\hat{\boldsymbol{\beta}}$. Therefore,

$$\frac{\hat{\beta}_i - \beta_i}{\sqrt{S^2 [(\mathbf{X}'\mathbf{X})^{-1}]_{i,i}}} \sim \frac{N(0, 1)}{\sqrt{\chi_{n-p}^2/(n-p)}} = t_{n-p},$$

where the normal and χ^2 random variables are independent. Therefore,

$$\hat{\beta}_i \pm t_{n-p}^{(\alpha/2)} S \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{i,i}}$$

is a $(1 - \alpha) \times 100\%$ confidence interval for β_i . This yields a complete analysis of confidence intervals for a univariate parameter; these confidence intervals can also be used for testing, of course.

2. Confidence ellipsoids

The situation is more interesting if we wish to say something about more than one parameter at the same time. For example, suppose we want to know about (β_0, β_1) in the *overly-simplified case that $\sigma = 1$* . The general philosophy of confidence intervals [for univariate parameters] suggests that we look for a random set Ω such that

$$P \left\{ \begin{pmatrix} \hat{\beta}_1 \\ \hat{\beta}_2 \end{pmatrix} - \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} \in \Omega \right\} = 1 - \alpha.$$

And we know that

$$\begin{pmatrix} \hat{\beta}_1 \\ \hat{\beta}_2 \end{pmatrix} \sim N \left(\begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix}, \Sigma \right),$$

where Σ is a 2×2 matrix with

$$\Sigma_{i,j} = [(\mathbf{X}'\mathbf{X})^{-1}]_{i,j}.$$

Here is a possible method: Let

$$\hat{\boldsymbol{\gamma}} := \begin{pmatrix} \hat{\beta}_1 - \beta_1 \\ \hat{\beta}_2 - \beta_2 \end{pmatrix}, \quad \text{so that} \quad \hat{\boldsymbol{\gamma}} \sim N_2(\mathbf{0}, \Sigma).$$

Also, recall that

$$\hat{\boldsymbol{\gamma}}' \Sigma^{-1} \hat{\boldsymbol{\gamma}} \sim \chi_2^2.$$

Therefore, one natural choice for Ω is

$$\Omega := \{ \mathbf{x} \in \mathbf{R}^2 : \mathbf{x}' \Sigma^{-1} \mathbf{x} \leq \chi_2^2(\alpha/2) \}.$$

What does Ω look like? In order to answer this, let us apply the spectral theorem:

$$\Sigma = \mathbf{P} \mathbf{D} \mathbf{P}' \quad \implies \quad \Sigma^{-1} = \mathbf{P} \mathbf{D}^{-1} \mathbf{P}'.$$

Then we can represent Ω as follows:

$$\begin{aligned} \Omega &= \{ \mathbf{x} \in \mathbf{R}^2 : \mathbf{x}' \mathbf{P} \mathbf{D}^{-1} \mathbf{P}' \mathbf{x} \leq \chi_2^2(\alpha/2) \} \\ &= \{ \mathbf{x} \in \mathbf{R}^2 : (\mathbf{P}' \mathbf{x})' \mathbf{D}^{-1} (\mathbf{P}' \mathbf{x}) \leq \chi_2^2(\alpha/2) \} \\ &= \{ \mathbf{P} \mathbf{y} \in \mathbf{R}^2 : \mathbf{y}' \mathbf{D}^{-1} \mathbf{y} \leq \chi_2^2(\alpha/2) \} \\ &= \left\{ \mathbf{P} \mathbf{y} \in \mathbf{R}^2 : \frac{y_1^2}{\lambda_1} + \frac{y_2^2}{\lambda_2} \leq \chi_2^2(\alpha/2) \right\}. \end{aligned}$$

Consider

$$\mathcal{E} := \left\{ \mathbf{y} \in \mathbf{R}^2 : \frac{y_1^2}{\lambda_1} + \frac{y_2^2}{\lambda_2} \leq \chi_2^2(\alpha/2) \right\}. \quad (1)$$

This is the interior of an ellipsoid, and $\Omega = \mathbf{P}\mathcal{E}$ is the image of the ellipsoid under the “linear orthogonal map” \mathbf{P} . Such sets are called “generalized ellipsoids,” and we have found a $(1 - \alpha) \times 100\%$ confidence [generalized] ellipsoid for $(\beta_1, \beta_2)'$.

The preceding can be generalized to any number of the parameters $\beta_{i_1}, \dots, \beta_{i_k}$, but is hard to work with, as the geometry of Ω can be complicated [particularly if $k \gg 2$]. Therefore, instead we might wish to look for approximate confidence sets that are easier to work with. Before we move on though, let me mention that if you want to know whether or not $H_0 : \beta_1 = \beta_{1,0}, \beta_2 = \beta_{2,0}$, then we can use these confidence bounds fairly easily, since it is not hard to check whether or not $(\beta_{1,0}, \beta_{2,0})'$ is in Ω : You simply compute the scalar quantity

$$(\beta_{1,0}, \beta_{2,0}) \Sigma^{-1} \begin{pmatrix} \beta_{1,0} \\ \beta_{2,0} \end{pmatrix},$$

and check to see if it is $\leq \chi_2^2(\alpha/2)$! But if you really need to imagine or see the confidence set[s], then this exact method can be unwieldy [particularly in higher dimensions than 2].

3. Bonferonni bounds

Our approximate confidence intervals are based on a fact from general probability theory.

Proposition 1 (Bonferonni’s inequality). *Let E_1, \dots, E_k be k events. Then,*

$$P(E_1 \cap \dots \cap E_k) \geq 1 - \sum_{j=1}^k P(E_j^c) = 1 - \sum_{j=1}^k (1 - P(E_j)).$$

Proof. The event $E_1^c \cup \dots \cup E_k^c$ is the complement of $E_1 \cap \dots \cap E_k$. Therefore,

$$P(E_1 \cap \dots \cap E_k) = 1 - P(E_1^c \cup \dots \cup E_k^c),$$

and this is $\geq 1 - \sum_{j=1}^k P(E_j^c)$ because the probability of a union is at most the sum of the individual probabilities. \square

Here is how we can use Bonferonni’s inequality. Define

$$C_j := \hat{\beta}_j \pm t_{n-p}^{(\alpha/4)} S \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{j,j}} \quad (j = 1, 2).$$

We have seen already that

$$P\{\beta_j \in C_j\} = 1 - \frac{\alpha}{2}.$$

[This is why we used $\alpha/4$ in the definition of C_j .] Therefore, Bonferonni’s inequality implies that

$$P\{\beta_1 \in C_1, \beta_2 \in C_2\} \geq 1 - \left[\frac{\alpha}{2} + \frac{\alpha}{2} \right] = 1 - \alpha.$$

In other words, (C_1, C_2) forms a “conservative” $(1 - \alpha) \times 100\%$ simultaneous confidence interval for (β_1, β_2) . This method becomes very inaccurate quickly as the number of parameters of interest grows. For instance, if you want Bonferonni confidence sets for $(\beta_1, \beta_2, \beta_3)$, then you need to use individual confidence intervals with confidence level $1 - (\alpha/3)$ each. And for k parameters you need individual confidence level $1 - (\alpha/k)$, which can yield bad performance when k is large.

This method is easy to implement, but usually **very** conservative.¹

4. Scheffé’s simultaneous conservative confidence bounds

There is a lovely method, due to Scheffé, that works in a similar fashion to the Bonferonni method; but has also the advantage of being often [far] less conservative! [We are now talking strictly about our linear model.] The starting point of this discussion is a general fact from matrix analysis.

Proposition 2 (The Rayleigh—Ritz inequality). *If $L_{p \times p}$ is positive definite, then for every p -vector b ,*

$$b' L^{-1} b = \max_{h \neq 0} \left[\frac{(h' b)^2}{h' L h} \right].$$

The preceding is called an inequality because it says that

$$b' L^{-1} b \leq \frac{(h' b)^2}{h' L h} \quad \text{for all } h_{p \times 1},$$

and it also tells us that the inequality is achieved for some $h \in \mathbf{R}^p$.

Proof of Rayleigh—Ritz inequality. Recall the Cauchy—Schwarz inequality from your linear algebra course: $(u' v)^2 \leq \|u\|^2 \cdot \|v\|^2$ with equality if and only if $v = a u$ for some $a \in \mathbf{R}$. It follows then that

$$\max_{v \neq 0} \left[\frac{(u' v)^2}{v' v} \right] = \|u\|^2.$$

Now write $L := P D P'$, in spectral form, and change variables:

$$v := (L^{1/2})' h,$$

in order to see that

$$u' v = u' (L^{1/2})' h, \quad v' v = h' L h,$$

¹On the other hand, the Bonferonni method can be applied to a wide range of statistical problems that involve simultaneous confidence intervals [and is not limited to the theory of linear models. So it is well worth your while to understand this important method.

and so

$$\|\mathbf{u}\|^2 = \max_{\mathbf{h} \neq \mathbf{0}} \left[\frac{(\mathbf{u}' (\mathbf{L}^{1/2})' \mathbf{h})^2}{\mathbf{h}' \mathbf{L} \mathbf{h}} \right].$$

Change variables again: $\mathbf{b} := \mathbf{L}^{1/2} \mathbf{u}$ to see that

$$\|\mathbf{u}\|^2 = \|\mathbf{L}^{-1/2} \mathbf{b}\|^2 = \mathbf{b}' \mathbf{L}^{-1} \mathbf{b}.$$

This does the job. □

Now let us return to Scheffé's method. Recall that

$$\hat{\boldsymbol{\beta}} - \boldsymbol{\beta} \sim N_p(\mathbf{0}, \sigma^2(\mathbf{X}'\mathbf{X})^{-1}),$$

so that

$$\frac{1}{\sigma^2} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' (\mathbf{X}'\mathbf{X}) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \sim \chi_p^2,$$

and hence

$$\frac{(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' (\mathbf{X}'\mathbf{X}) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})}{pS^2} = \frac{(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' (\mathbf{X}'\mathbf{X}) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) / p}{S^2} \sim F_{p, n-p}.$$

We apply the Rayleigh–Ritz inequality with $\mathbf{b} := \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}$ and $\mathbf{L} := (\mathbf{X}'\mathbf{X})^{-1}$ in order to find that

$$(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' (\mathbf{X}'\mathbf{X}) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) = \max_{\mathbf{h} \neq \mathbf{0}} \left[\frac{\{\mathbf{h}' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})\}^2}{\mathbf{h}' (\mathbf{X}'\mathbf{X})^{-1} \mathbf{h}} \right].$$

Therefore,

$$P \left\{ \frac{1}{pS^2} \max_{\mathbf{h} \neq \mathbf{0}} \left[\frac{\{\mathbf{h}' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})\}^2}{\mathbf{h}' (\mathbf{X}'\mathbf{X})^{-1} \mathbf{h}} \right] \leq F_{p, n-p}(\alpha) \right\} = 1 - \alpha.$$

Equivalently,

$$P \left\{ \frac{|\mathbf{h}' \hat{\boldsymbol{\beta}} - \mathbf{h}' \boldsymbol{\beta}|}{\sqrt{\mathbf{h}' (\mathbf{X}'\mathbf{X})^{-1} \mathbf{h}}} \leq \sqrt{pS^2 F_{p, n-p}(\alpha)} \text{ for all } \mathbf{h} \in \mathbf{R}^p \right\} = 1 - \alpha.$$

If we restrict attention to a subcollection of \mathbf{h} 's then the probability is even more. In particular, consider only \mathbf{h} 's that are the standard basis vectors of \mathbf{R}^p , in order to deduce from the preceding that

$$P \left\{ \frac{|\hat{\beta}_j - \beta_j|}{\sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{jj}}} \leq \sqrt{pS^2 F_{p, n-p}(\alpha)} \text{ for all } 1 \leq j \leq p \right\} \geq 1 - \alpha.$$

In other words, we have demonstrated the following.

Theorem 3 (Scheffé). *The following are conservative $(1 - \alpha) \times 100\%$ simultaneous confidence bounds for $(\beta_1, \dots, \beta_p)'$:*

$$\hat{\beta}_j \pm \sqrt{pS^2 [(\mathbf{X}'\mathbf{X})^{-1}]_{jj} F_{p, n-p}(\alpha)} \quad (1 \leq j \leq p).$$

When the sample size n is very large, the preceding yields an asymptotic simplification. Recall that $\chi_{n-p}^2/(n-p) \rightarrow 1$ as $n \rightarrow \infty$. Therefore, $F_{p, n-p}(\alpha) \approx \chi_p^2(\alpha)/p$ for $n \gg 1$. Therefore, the following are asymptotically-conservative simultaneous confidence bounds:

$$\hat{\beta}_j \pm S \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{jj} \chi_p^2(\alpha)} \quad (1 \leq j \leq p) \text{ for } n \gg 1.$$

5. Confidence bounds for the regression surface

Given a vector $\mathbf{x} \in \mathbf{R}^p$ of predictor variables, our linear model yields $E\mathbf{y} = \mathbf{x}'\boldsymbol{\beta}$. In other words, we can view our efforts as one about trying to understand the unknown function

$$f(\mathbf{x}) := \mathbf{x}'\boldsymbol{\beta}.$$

And among other things, we have found the following estimator for f :

$$\hat{f}(\mathbf{x}) := \mathbf{x}'\hat{\boldsymbol{\beta}}.$$

Note that if $\mathbf{x} \in \mathbf{R}^p$ is held fixed, then

$$\hat{f}(\mathbf{x}) \sim N(f(\mathbf{x}), \sigma^2 \mathbf{x}'(\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}).$$

Therefore, a $(1 - \alpha) \times 100\%$ confidence interval for $f(\mathbf{x})$ [for a fixed predictor variable \mathbf{x}] is

$$\hat{f}(\mathbf{x}) \pm S \sqrt{\mathbf{x}'(\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}} t_{n-p}^{(\alpha/2)}.$$

That is, if we are interested in a confidence interval for $f(\mathbf{x})$ for a fixed \mathbf{x} , then we have

$$P \left\{ f(\mathbf{x}) = \hat{f}(\mathbf{x}) \pm S \sqrt{\mathbf{x}'(\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}} t_{n-p}^{(\alpha/2)} \right\} = 1 - \alpha.$$

On the other hand, we can also apply Scheffé's method and obtain the following simultaneous $(1 - \alpha) \times 100\%$ confidence set:

$$P \left\{ f(\mathbf{x}) = \hat{f}(\mathbf{x}) \pm S \sqrt{p \mathbf{x}'(\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}} F_{p, n-p}(\alpha) \text{ for all } \mathbf{x} \in \mathbf{R}^p \right\} \geq 1 - \alpha.$$

Example 4 (Simple linear regression). Consider the basic regression model

$$Y_i = \alpha + \beta x_i + \varepsilon_i \quad (1 \leq i \leq n).$$

If $\mathbf{w} = (1, w)'$, then a quick computation yields

$$\mathbf{w}'(\mathbf{X}'\mathbf{X})^{-1} \mathbf{w} = \frac{[\bar{x}^2 - 2w\bar{x} + w^2]}{ns_x^2},$$

where we recall $s_x^2 := \sum_{i=1}^p (x_i - \bar{x})^2$. Therefore, a simultaneous confidence interval for all of $\alpha + \beta w$'s, as w ranges over \mathbf{R} , is

$$\hat{\alpha} + \hat{\beta}w \pm S \sqrt{\frac{2}{ns_x^2} [\bar{x}^2 - 2w\bar{x} + w^2]} F_{2,n-2}(\alpha).$$

This expression can be simplified further because:

$$\begin{aligned} \bar{x}^2 - 2w\bar{x} + w^2 &= \bar{x}^2 - (\bar{x})^2 + (\bar{x})^2 - 2w\bar{x} + w^2 \\ &= s_x^2 + (\bar{x} - w)^2. \end{aligned}$$

Therefore, with probability $1 - \alpha$, we have

$$\alpha + \beta w = \hat{\alpha} + \hat{\beta}w \pm S \sqrt{2 \left(\frac{1}{n} + \frac{(\bar{x} - w)^2}{ns_x^2} \right)} F_{2,n-2}(\alpha) \quad \text{for all } w.$$

On the other hand, if we want a confidence interval for $\alpha + \beta w$ for a fixed w , then we can do better using our t -test:

$$\alpha + \beta w = \hat{\alpha} + \hat{\beta}w \pm S \sqrt{\left(\frac{1}{n} + \frac{(\bar{x} - w)^2}{ns_x^2} \right)} t_{n-2}^{(\alpha/2)}.$$

You should check the details of this computation. \square

6. Prediction intervals

The difference between confidence and prediction intervals is this: For a confidence interval we try to find an interval around the **parameter** $\mathbf{x}'\boldsymbol{\beta}$. For a prediction interval we do so for the **random variable** $y_0 := \mathbf{x}_0'\boldsymbol{\beta} + \varepsilon_0$, where \mathbf{x}_0 is known and fixed and ε_0 is the “noise,” which is hitherto unobserved (i.e., independent of the vector \mathbf{Y} of observations).

It is not hard to construct a good prediction interval of this type: Note that

$$\hat{y}_0 := \mathbf{x}_0'\hat{\boldsymbol{\beta}}$$

satisfies

$$\hat{y}_0 - y_0 = \mathbf{x}_0'(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) - \varepsilon_0 \sim N(0, \sigma^2 [\mathbf{x}_0'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{x}_0 + 1]).$$

Therefore, a prediction interval is

$$\hat{y}_0 \pm S t_{n-p}^{(\alpha/2)} \sqrt{\mathbf{x}_0'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{x}_0 + 1}.$$

This means that

$$P \left\{ y_0 \in \hat{y}_0 \pm S t_{n-p}^{(\alpha/2)} \sqrt{\mathbf{x}_0'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{x}_0 + 1} \right\} = 1 - \alpha,$$

but note that both y_0 and the prediction interval are now random.

Polynomial Regression

1. Introduction

Here is a careful statement of the Weierstrass approximation theorem, which states in loose terms that every continuous function on a bounded set can be approximated arbitrarily well by a polynomial.

Theorem 1 (Weierstrass Approximation Theorem). *If $f : [a, b] \rightarrow \mathbf{R}$ is continuous, then for all $\epsilon > 0$ there exists a polynomial P such that*

$$\max_{a \leq x \leq b} |f(x) - P(x)| < \epsilon.$$

Therefore, if we can try to estimate a variable y using only a variable x , then it makes good sense to try to fit polynomials to our regression fit. The general linear model becomes, in this case,

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \cdots + \beta_k x_i^k + \varepsilon_i \quad (1 \leq i \leq n).$$

This is indeed a linear model, as we have seen, with design matrix

$$\mathbf{X} = \begin{pmatrix} 1 & x_1 & x_1^2 & \cdots & x_1^k \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & x_n & x_n^2 & \cdots & x_n^k \end{pmatrix}$$

A practical problem that arises here is this: If k is large then $\mathbf{X}'\mathbf{X}$ is “ill-conditioned” and therefore very hard to invert accurately. For instance, suppose x_1, \dots, x_n are spaced uniformly in $(0, 1)$, so that $x_1 = 1/n, x_2 = 2/n, \dots, x_n = n/n = 1$. In that case,

$$(\mathbf{X}'\mathbf{X})_{i,j} = \sum_{k=1}^n X_{k,i} X_{k,j} = \sum_{k=1}^n x_k^i x_k^j = \sum_{k=1}^n \left(\frac{k}{n}\right)^i \left(\frac{k}{n}\right)^j = n \times \frac{1}{n} \sum_{k=1}^n \left(\frac{k}{n}\right)^{i+j}.$$

Riemann's theory of integral tells us that

$$\frac{1}{n} \sum_{k=1}^n \left(\frac{k}{n}\right)^{i+j} \simeq \int_0^1 x^{i+j} dx = \frac{1}{i+j+1},$$

when n is large. Therefore, the $k \times k$ matrix $\mathbf{X}'\mathbf{X}$ satisfies

$$\mathbf{X}'\mathbf{X} \simeq n \begin{pmatrix} \frac{1}{3} & \frac{1}{4} & \cdots & \frac{1}{k+2} \\ \frac{1}{4} & \frac{1}{5} & \cdots & \frac{1}{k+3} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{k+2} & \frac{1}{k+3} & \cdots & \frac{1}{2k+1} \end{pmatrix} := n\mathbf{M},$$

for large values of n . For example,

$$\mathbf{M} = \begin{pmatrix} 1/3 & 1/4 & 1/5 & 1/6 & 1/7 & 1/8 & 1/9 \\ 1/4 & 1/5 & 1/6 & 1/7 & 1/8 & 1/9 & 1/10 \\ 1/5 & 1/6 & 1/7 & 1/8 & 1/9 & 1/10 & 1/11 \\ 1/6 & 1/7 & 1/8 & 1/9 & 1/10 & 1/11 & 1/12 \\ 1/7 & 1/8 & 1/9 & 1/10 & 1/11 & 1/12 & 1/13 \\ 1/8 & 1/9 & 1/10 & 1/11 & 1/12 & 1/13 & 1/14 \\ 1/9 & 1/10 & 1/11 & 1/12 & 1/13 & 1/14 & 1/15 \end{pmatrix}.$$

in the $k = 7$ case. Already k is only modestly sized, but you see that almost half of the entries in \mathbf{M} are below 0.01. Therefore, a small error in the computation \mathbf{M}^{-1} leads to a massive error in the computation of $(\mathbf{X}'\mathbf{X})^{-1} \simeq (1/n)\mathbf{M}^{-1}$ for modest-to-large values of n . Your textbook recommends that you keep polynomial regression to below polynomials of degree 6.

2. Orthogonal polynomials

A more robust approach is to use *orthogonal polynomials*. That is, reparametrize our polynomial regression model such as

$$Y_i = \gamma_0 \phi_0(x_i) + \gamma_1 \phi_1(x_i) + \cdots + \gamma_k \phi_k(x_i) + \varepsilon_i,$$

where the γ_i 's are unknown parameters, ϕ_j is a j th-order polynomial, and

$$\sum_{i=1}^n \phi_m(x_i) \phi_l(x_i) = \begin{cases} 1 & \text{whenever } m = l, \\ 0 & \text{whenever } m \neq l \end{cases} = I_{m,l}.$$

[In other words, ϕ_0, \dots, ϕ_k are *orthonormal* over $\{x_i\}_{i=1}^n$.]

Note that this is a linear model with $p := k + 1$. Indeed,

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

where

$$\mathbf{X} = \begin{pmatrix} \phi_0(x_1) & \cdots & \phi_k(x_1) \\ \vdots & & \vdots \\ \phi_0(x_n) & \cdots & \phi_k(x_n) \end{pmatrix}.$$

Before we discuss how we can find these ϕ_i 's, let us see what we have gained by using them. Notice that $X_{i,j} = \phi_j(x_i)$ ($1 \leq i \leq n$, $0 \leq j \leq k$) [note that the matrix is now indexed in a slightly funny way, but this is of course ok]. Therefore,

$$(\mathbf{X}'\mathbf{X})_{i,j} = \sum_k X_{k,i}X_{k,j} = \sum_k \phi_i(x_k)\phi_j(x_k) = I_{i,j} \implies \mathbf{X}'\mathbf{X} = \mathbf{I}.$$

In particular, it is now a trivial task to invert $\mathbf{X}'\mathbf{X} = \mathbf{I}$! Our parameter estimates are also simplified: $\hat{\boldsymbol{\gamma}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} = \mathbf{X}'\mathbf{Y}$. Therefore,

$$\hat{\gamma}_i = \sum_j X_{j,i}Y_j = \sum_{j=1}^n \phi_i(x_j)Y_j,$$

and

$$\hat{\boldsymbol{\gamma}} \sim N_{k+1}(\boldsymbol{\gamma}, \sigma^2 \mathbf{I}) \quad \text{under the normal-error model.}$$

In other words, under the normal model, the $\hat{\gamma}_i$'s are independent!

Let us compute the fitted values next: $\hat{\mathbf{Y}} = \mathbf{X}\hat{\boldsymbol{\gamma}} = \mathbf{X}\mathbf{X}'\mathbf{Y}$; therefore,

$$\begin{aligned} \text{RSS} &= \|\mathbf{Y} - \hat{\mathbf{Y}}\|^2 = \|\mathbf{Y}\|^2 - 2\mathbf{Y}'\hat{\mathbf{Y}} + \|\hat{\mathbf{Y}}\|^2 \\ &= \|\mathbf{Y}\|^2 - 2(\mathbf{X}'\mathbf{Y})^2 + (\mathbf{X}'\mathbf{Y})^2 = \|\mathbf{Y}\|^2 - (\mathbf{X}'\mathbf{Y})^2 \\ &= \|\mathbf{Y}\|^2 - \|\hat{\boldsymbol{\gamma}}\|^2. \end{aligned}$$

Now suppose we are testing for contrasts:

$$H_0 : \mathbf{A}\boldsymbol{\gamma} = \mathbf{c},$$

where $\mathbf{A}_{q \times (k+1)}$ is full rank with $q \leq k$, and $\mathbf{c}_{q \times 1}$ is a known vector. Then,

$$\begin{aligned} \hat{\boldsymbol{\gamma}}_{H_0} &= \hat{\boldsymbol{\gamma}} + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} (\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\gamma}}) \\ &= \hat{\boldsymbol{\gamma}} + \mathbf{A}'(\mathbf{A}\mathbf{A}')^{-1} (\mathbf{c} - \mathbf{A}\hat{\boldsymbol{\gamma}}), \end{aligned}$$

and

$$\begin{aligned} \text{RSS}_{H_0} - \text{RSS} &= (\mathbf{A}\hat{\boldsymbol{\gamma}} - \mathbf{c})' [\mathbf{A}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{A}']^{-1} (\mathbf{A}\hat{\boldsymbol{\gamma}} - \mathbf{c}) \\ &= (\mathbf{A}\hat{\boldsymbol{\gamma}} - \mathbf{c})' (\mathbf{A}\mathbf{A}')^{-1} (\mathbf{A}\hat{\boldsymbol{\gamma}} - \mathbf{c}). \end{aligned}$$

In particular,

$$\frac{(\mathbf{A}\hat{\boldsymbol{\gamma}} - \mathbf{c})' (\mathbf{A}\mathbf{A}')^{-1} (\mathbf{A}\hat{\boldsymbol{\gamma}} - \mathbf{c}) / q}{(\|\mathbf{Y}\|^2 - \|\hat{\boldsymbol{\gamma}}\|^2) / (n - k - 1)} \stackrel{H_0}{\sim} F_{q, n-p}.$$

One can work out, using similar means, Scheffé-type conservative confidence intervals, etc.

3. The Forsythe–Hayes 2-step method

It remains to introduce a way of finding our orthonormal polynomials ϕ_0, \dots, ϕ_k . Remember that our goal is to find these so that, given x_1, \dots, x_n ,

$$\sum_{i=1}^n \phi_m(x_i) \phi_\ell(x_i) = I_{m,\ell}.$$

Forsythe (1957) and Hayes (1969) adapted a classical method from analysis (Christoffel, 1858; Darboux, 1878) to the present setting in order to find an algorithm for finding such polynomials. We follow their method next.

Suppose we could find polynomials ψ_0, \dots, ψ_k such that

$$\sum_{i=1}^n \psi_m(x_i) \psi_\ell(x_i) = 0 \quad \text{whenever } m \neq \ell.$$

That is, ψ 's are merely orthogonal polynomials. Then,

$$\phi_m(x) := \frac{\psi_m(x)}{\sqrt{\sum_{i=1}^n [\psi_m(x_i)]^2}} \quad (0 \leq m \leq k, x \in \mathbf{R})$$

defines our orthonormal polynomials. It remains to compute the ψ_j 's.

We start with

$$\psi_0(x) := 1,$$

and ψ_1 has the form $\psi_1(x) = x - a_1$. [The slightly more general form $\psi_1(x) = c(a - x)$ does not gain us any further insight; why?] The constant a_1 has to be chosen so that $\sum_{i=1}^n \psi_1(x_i) \psi_0(x_i) = 0 \implies \sum_{i=1}^n x_i - na_1 = 0 \implies a_1 = \bar{x}$. That is,

$$\psi_1(x) = x - \bar{x}.$$

From here on we can describe things algorithmically. Suppose we have defined ψ_0, \dots, ψ_r so that they are orthogonal to each other [*the induction hypothesis*]. Then for all $r \geq 1$,

$$\psi_{r+1}(x) := (x - a_{r+1})\psi_r(x) - b_r\psi_{r-1}(x),$$

where a_{r+1} and b_r are chosen so that ψ_{r+1} is orthogonal to ψ_0, \dots, ψ_r . In particular, we must have

$$\begin{aligned} 0 &= \sum_{i=1}^n \psi_{r+1}(x_i) \psi_r(x_i) \\ &= \sum_{i=1}^n (x_i - a_{r+1}) [\psi_r(x_i)]^2 - b_r \sum_{i=1}^n \psi_{r-1}(x_i) \psi_r(x_i) \\ &= \sum_{i=1}^n x_i [\psi_r(x_i)]^2 - a_{r+1} \sum_{i=1}^n [\psi_r(x_i)]^2 \quad (\text{by induction}). \end{aligned}$$

Therefore, if we choose

$$a_{r+1} = \frac{\sum_{i=1}^n x_i [\psi_r(x_i)]^2}{\sum_{i=1}^n [\psi_r(x_i)]^2},$$

then ψ_{r+1} is orthogonal to ψ_r .

In order to find b_r we note that we must set things up—at the very least—so that ψ_{r+1} is orthogonal to ψ_{r-1} . That is,

$$\begin{aligned} 0 &= \sum_{i=1}^n \psi_{r+1}(x_i) \psi_{r-1}(x_i) \\ &= \sum_{i=1}^n (x_i - a_{r+1}) \psi_r(x_i) \psi_{r-1}(x_i) - b_r \sum_{i=1}^n [\psi_{r-1}(x_i)]^2 \\ &= \sum_{i=1}^n x_i \psi_r(x_i) \psi_{r-1}(x_i) - a_{r+1} \sum_{i=1}^n \psi_r(x_i) \psi_{r-1}(x_i) - b_r \sum_{i=1}^n [\psi_{r-1}(x_i)]^2 \\ &= \sum_{i=1}^n x_i \psi_r(x_i) \psi_{r-1}(x_i) - b_r \sum_{i=1}^n [\psi_{r-1}(x_i)]^2; \end{aligned}$$

thanks to induction. We can simplify this a little more still, mainly because $\psi_r(x) = (x - a_r) \psi_{r-1}(x) - b_r \psi_{r-2}(x)$. In this way we find that

$$\begin{aligned} \sum_{i=1}^n x_i \psi_r(x_i) \psi_{r-1}(x_i) &= \sum_{i=1}^n (x_i - a_r) \psi_r(x_i) \psi_{r-1}(x_i) \quad (\text{induction}) \\ &= \sum_{i=1}^n [\psi_r(x_i)]^2 + b_r \sum_{i=1}^n \psi_{r-2}(x_i) \psi_r(x_i) = \sum_{i=1}^n [\psi_r(x_i)]^2. \end{aligned}$$

Therefore, our candidate for b_r is

$$b_r = \frac{\sum_{i=1}^n [\psi_r(x_i)]^2}{\sum_{i=1}^n [\psi_{r-1}(x_i)]^2}.$$

With the present choice of a_{r+1} and b_r , ψ_{r+1} is orthogonal to both ψ_r and ψ_{r-1} . Now it remains to prove that this choice of a_{r+1} and b_r actually works. That is, we proceed to show that, for our construction, ψ_{r+1} is orthogonal to ψ_ℓ for every $\ell \leq r-2$. Note that

$$\begin{aligned} \sum_{i=1}^n \psi_{r+1}(x_i) \psi_\ell(x_i) &= \sum_{i=1}^n (x_i - a_{r+1}) \psi_r(x_i) \psi_\ell(x_i) - b_r \sum_{i=1}^n \psi_{r-1}(x_i) \psi_\ell(x_i) \\ &= \sum_{i=1}^n x_i \psi_r(x_i) \psi_\ell(x_i) \\ &= \sum_{i=1}^n (x_i - a_{\ell+1}) \psi_r(x_i) \psi_\ell(x_i) - b_\ell \sum_{i=1}^n \psi_r(x_i) \psi_{\ell+1}(x_i). \end{aligned}$$

=0 by induction

The defining property of the ψ 's tells us that for $\ell \leq r - 2$,

$$\sum_{i=1}^n \psi_{r+1}(x_i) \psi_{\ell}(x_i) = \sum_{i=1}^n \psi_r(x_i) \psi_{\ell+1}(x_i) = 0,$$

thanks to induction.

Splines

1. Introduction

One way to think of linear regression is this: We have data of pairs of points $(x_1, y_1), \dots, (x_n, y_n)$, and try to fit a line through this. We have added to this procedure by studying what happens if the error in this approximation $[y = \beta_0 + \beta_1 x + \text{noise}]$ has a normal distribution etc. Here we study a different angle on these problems: Are there better curves than lines that can be fitted through the points $(x_1, y_1), \dots, (x_n, y_n)$? One answer is polynomial regression. Another is splines, which I will say a few things about.

Suppose we have a fixed interval $[a, b]$ in which we are seeing numbers, and also have a subdivision of this interval,

$$a := \xi_0 \leq \xi_1 < \dots < \xi_k \leq \xi_{k+1} := b.$$

We say that a function $S : [a, b] \rightarrow \mathbf{R}$ is a *spline* with *knots* $\xi_1 < \dots < \xi_k$ if there exists an integer $M \geq 1$ —called the *order of S*—such that: (i) S is a polynomial of degree $M - 1$ in $[\xi_i, \xi_{i+1}]$ for every $1 \leq i \leq k - 1$; and (ii) S and its first $M - 2$ derivatives all exist and are continuous everywhere in $[a, b]$. When $a = -\infty$ and $b = \infty$, we think of ξ_0 and ξ_{k+1} also as knots of S .

In other words, S is a spline if it is smooth function that is piecewise polynomial.

2. Linear splines

Suppose our data is $(\xi_1, y_1), \dots, (\xi_k, y_k)$ and we wish to do a piecewise linear fit. The following is an order $M = 2$ spline that does the job:

$$S(x) = y_i + (y_{i+1} - y_i) \left(\frac{x - \xi_i}{\xi_{i+1} - \xi_i} \right) \quad \text{for every } \xi_i \leq x < \xi_{i+1}.$$

This does the job. (Draw a picture!) But we need to understand how it was derived.

First you reparametrize the interval $[\xi_i, \xi_{i+1})$ by $[0, 1)$. That is,

$$t := \frac{x - \xi_i}{\xi_{i+1} - \xi_i} \quad \text{for } \xi_i \leq x < \xi_{i+1}. \quad (1)$$

As the variable x increases from ξ_i to ξ_{i+1} , the variable t increases from 0 to 1. That is, with this parametrization,

$$S(x) = Y_i(t) \quad \text{when } \xi_i \leq x < \xi_{i+1}. \quad (2)$$

In parameter t , the i th piece of the spline should look like a line:

$$Y_i(t) = a_i + b_i t \quad (0 \leq t < 1),$$

subject to $Y_i(0) = y_i$ and $Y_i(1) = y_{i+1}$. Therefore, $a_i = y_i$ and $a_i + b_i = y_{i+1}$. This leads us to

$$Y_i(t) = y_i + (y_{i+1} - y_i)t \quad (0 \leq t < 1).$$

Use this together with (1) and (2) in order to compute the spline at any point x .

3. Quadratic splines

Here, $M = 3$. We apply the same reparametrization (1) and (2) as before, but now use the following quadratic interpolation formula in place of the previous linear one:

$$Y_i(t) = a_i + b_i t + c_i t^2 \quad (0 \leq t < 1);$$

subject to $Y_i(0) = y_i$, $Y_i(1) := \lim_{t \uparrow 1} Y_i(t) = y_{i+1}$. This yields,

$$a_i = y_i, \quad a_i + b_i + c_i = y_{i+1} \implies \begin{cases} y_i = a_i, \\ y_{i+1} - y_i = b_i + c_i. \end{cases}$$

This identifies the a_i 's and shows also that if we can compute all of the b_i 's then we will have explicit formulas for the c_i 's as well.

We also have a continuity condition on the derivative of order $M - 2 = 1$. To see how that manifests itself note that

$$Y_i'(t) = b_i + 2c_i t \quad (0 \leq t < 1).$$

In particular, $Y_i'(0) = b_i$ and $Y_i'(1) = b_i + 2c_i$. If $\xi_i \leq x < \xi_{i+1}$, then

$$S'(x) = \frac{dS(x)}{dx} = \frac{dY_i(t)}{dt} \cdot \frac{dt}{dx} = \frac{Y_i'(t)}{\xi_{i+1} - \xi_i} = \frac{b_i + 2c_i t}{\xi_{i+1} - \xi_i},$$

thanks to the chain rule. Since we want S' to be continuous everywhere, it is enough to check that $S'(\xi_{i+1}-) = S'(\xi_{i+1}+)$. But

$$S'(\xi_{i+1}-) = \frac{Y_i'(1)}{\xi_{i+1} - \xi_i} = \frac{b_i + 2c_i}{\xi_{i+1} - \xi_i},$$

and

$$S'(\xi_{i+1}+) = \frac{Y'_{i+1}(0)}{\xi_{i+2} - \xi_{i+1}} = \frac{b_{i+1}}{\xi_{i+2} - \xi_{i+1}}.$$

Therefore, we also want

$$\frac{b_i + 2c_i}{\xi_{i+1} - \xi_i} = \frac{b_{i+1}}{\xi_{i+2} - \xi_{i+1}}.$$

Equivalently,

$$c_i = \frac{1}{2} \left[\left(\frac{\xi_{i+1} - \xi_i}{\xi_{i+2} - \xi_{i+1}} \right) b_{i+1} - b_i \right].$$

Since $y_{i+1} - y_i = b_i + c_i$, the preceding shows that

$$\begin{aligned} y_{i+1} - y_i &= b_i + \frac{1}{2} \left[\left(\frac{\xi_{i+1} - \xi_i}{\xi_{i+2} - \xi_{i+1}} \right) b_{i+1} - b_i \right] \\ &= \frac{b_{i+1}}{2} \left(\frac{\xi_{i+1} - \xi_i}{\xi_{i+2} - \xi_{i+1}} \right) + \frac{b_i}{2}. \end{aligned}$$

Equivalently, that

$$b_{i+1} = [2(y_{i+1} - y_i) - b_i] \left(\frac{\xi_{i+2} - \xi_{i+1}}{\xi_{i+1} - \xi_i} \right).$$

That is, if we knew b_0 then we could iteratively compute all of the b_i 's, and be finished.

In order to find b_0 we need one more condition. Usually, people assume a zero derivative condition at the left-most knot:

$$S'(a) = 0 \implies Y'_0(0) = 0 \implies b_0 = 0.$$

This completes the algorithm for computing our spline. Any other choice of $S'(a)$, than zero, would yield a different [perfectly ok] spline as well.

In practice, it is not enough to use quadratic splines, since splines start to look smooth at the cubic level. This requires a great deal more computation, but is still a manageable problem.

4. Cubic splines

Now we consider the still-more interesting case that $M = 4$ [cubic splines]. Recall that we want a piece-wise cubic fit S of our data such that S , S' , and S'' are continuous [not piecewise continuous] functions. We will adopt the boundary conditions

$$S''(a) = S''(b) = 0,$$

as it is customary. As before, we work on each piece $[\xi_i, \xi_{i+1}]$ separately by first changing variables for the i th piece:

$$t := \frac{x - \xi_i}{\xi_{i+1} - \xi_i} \quad (\xi_i \leq x < \xi_{i+1}).$$

Then in t coordinates, our spline model is

$$Y_i(t) = a_i + b_i t + c_i t^2 + d_i t^3 \quad (0 \leq t < 1),$$

in the i th piece. First of all, our spline function has to go through the data points. That is, $Y_i(0) = y_i$ for all $0 \leq i \leq k$. In other words,

$$a_i = y_i \quad (0 \leq i \leq k). \quad (3)$$

This shows that the a_i 's are determined. We need to compute the b 's, c 's, and d 's.

The statement that S is continuous is equivalent to $Y_{i+1}(0) = Y_i(1)$ for all i . That is,

$$y_{i+1} = y_i + b_i + c_i + d_i \quad (0 \leq i \leq k-1). \quad (4)$$

We want also that S' to be continuous. Now, if $\xi_i \leq x < \xi_{i+1}$, then

$$S'(x) = \frac{dS(x)}{dx} = \frac{dY_i(t)}{dt} \cdot \frac{dt}{dx} = \frac{Y'_i(t)}{\xi_{i+1} - \xi_i} = \frac{b_i + 2c_i t + 3d_i t^2}{\xi_{i+1} - \xi_i}.$$

The continuity of S' is equivalent to the statement that $S'(\xi_{i+1}-) = S'(\xi_{i+1}+)$ for all i . Put in yet another way, we want

$$S'(\xi_{i+1}-) = \frac{Y'_i(1)}{\xi_{i+1} - \xi_i} = S'(\xi_{i+1}+) = \frac{Y'_{i+1}(0)}{\xi_{i+2} - \xi_{i+1}}.$$

Plug in the value of $Y'_i(1)$ and $Y'_{i+1}(0)$ to see that

$$\frac{b_i + 2c_i + 3d_i}{\xi_{i+1} - \xi_i} = \frac{b_{i+1}}{\xi_{i+2} - \xi_{i+1}}. \quad (5)$$

This shows that if we could solve the c_i 's and the d_i 's, as well as b_0 , then we could also find all of the b_i 's. In fact, the way we do so is as follows: Set

$$B_i := \frac{b_i}{\xi_{i+1} - \xi_i}, \quad C_i := \frac{c_i}{\xi_{i+1} - \xi_i}, \quad D_i := \frac{d_i}{\xi_{i+1} - \xi_i}.$$

Then (5) says that

$$B_i + 2C_i + 3D_i = B_{i+1} \iff B_{i+1} - B_i = 2C_i + 3D_i.$$

Add from $i = 0$ to $i = j$ in order to see that for all j ,

$$B_j = B_0 + 2 \sum_{i=0}^j C_i + 3 \sum_{i=0}^j D_i. \quad (6)$$

Our second-derivative requirement is that S'' is continuous. If $\xi_i \leq x < \xi_{i+1}$, then the chain rule tells us that

$$S''(x) = Y''_i(t) \cdot \frac{dt}{dx} = \frac{2c_i + 6d_i t}{\xi_{i+1} - \xi_i},$$

since $d^2t/dx^2 = 0$. Thus, we set $S''(\xi_{i+1}-) = S''(\xi_{i+1}+)$ in order to see that

$$\frac{Y_i''(1)}{\xi_{i+1} - \xi_i} = \frac{Y_{i+1}''(0)}{\xi_{i+2} - \xi_{i+1}},$$

which is to say that

$$\frac{2c_i + 6d_i}{\xi_{i+1} - \xi_i} = \frac{2c_{i+1}}{\xi_{i+2} - \xi_{i+1}}.$$

I.e.,

$$D_i = \frac{d_i}{\xi_{i+1} - \xi_i} = \frac{1}{3} \left[\frac{c_{i+1}}{\xi_{i+2} - \xi_{i+1}} - \frac{c_i}{\xi_{i+1} - \xi_i} \right] = \frac{C_{i+1} - C_i}{3}. \quad (7)$$

Thus, if we could compute all of the c_i 's, then we will know all of the d_i 's as well, and thereby all b_i 's, which ends our task. In fact, by (7),

$$\sum_{\ell=0}^i D_\ell = \frac{1}{3} \sum_{\ell=0}^i (C_{\ell+1} - C_\ell) = \frac{C_{i+1} - C_0}{3}.$$

Plug this into (6) and change variables [$j \leftrightarrow i$] to obtain

$$\begin{aligned} B_i &= B_0 + 2 \sum_{\ell=0}^i C_\ell + 3[C_{i+1} - C_0] \\ &= B_0 - C_0 + 2 \sum_{\ell=1}^i C_\ell + 3C_{i+1}. \end{aligned}$$

Now, by (4),

$$\begin{aligned} E_i &:= \frac{y_{i+1} - y_i}{\xi_{i+1} - \xi_i} = B_i + C_i + D_i \\ &= B_0 - C_0 + 2 \sum_{\ell=1}^i C_\ell + 3C_{i+1} + C_i + \frac{C_{i+1} - C_i}{3} \\ &= B_0 - C_0 + 2 \sum_{\ell=1}^{i-1} C_\ell + \frac{8}{3}C_i + \frac{10}{3}C_{i+1}, \end{aligned}$$

where $\sum_{\ell=1}^0 C_\ell := 0$. In other words, given B_0 and C_0 [oftentimes, people choose them to be zero], the vector $\mathbf{C} := (C_1, \dots, C_{k+1})'$ solves

$$\mathbf{E} = (B_0 - C_0)\mathbf{1}_{(k+1) \times 1} + \mathbf{MC},$$

where \mathbf{M} is the following $(k+1) \times (k+1)$ matrix:

$$\mathbf{M} := \begin{pmatrix} 8/3 & 10/3 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 2 & 8/3 & 10/3 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 2 & 2 & 8/3 & 10/3 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ 2 & 2 & 2 & 2 & \cdots & 2 & 8/3 & 10/3 & 0 \\ 2 & 2 & 2 & 2 & \cdots & 2 & 2 & 8/3 & 10/3 \\ 2 & 2 & 2 & 2 & \cdots & 2 & 2 & 2 & 8/3 \end{pmatrix}.$$

The matrix \mathbf{M} turns out to be invertible. Therefore, we obtain \mathbf{C} by

$$\mathbf{C} = \mathbf{M}^{-1} [\mathbf{E} - (\mathbf{B}_0 - \mathbf{C}_0)\mathbf{1}_{(k+1) \times 1}].$$

Now, (7) yields $\mathbf{D} := (D_1, \dots, D_{k+1})'$, and (6) allows us to finally compute $\mathbf{B} := (B_1, \dots, B_{k+1})'$. Now backsolve for the b_i 's, c_i 's, and the d_i 's to finish [$b_i = (\xi_{i+1} - \xi_i)B_i$, etc.]