

# Statistics 849 notes – Fall 2010

Douglas Bates

Based on notes by Sündüz Keleş from Fall 2009  
that were transcribed by Brittany Schwefel

September 17, 2010



# Contents

<b>1</b>	<b>The Gaussian Linear Model</b>	<b>1</b>
1.1	Gaussian Linear Model . . . . .	1
1.2	Linear algebra of least squares . . . . .	2
1.3	Matrix decompositions . . . . .	3
1.3.1	Orthogonal matrices . . . . .	3
1.3.2	The QR decomposition . . . . .	3
1.3.3	Comparison to the usual text-book formulas . . . . .	12
1.3.4	R functions related to the QR decomposition . . . . .	15
1.4	Related matrix decompositions . . . . .	16
1.4.1	The Cholesky decomposition . . . . .	16
1.4.2	Evaluation of the Cholesky decomposition . . . . .	16
1.4.3	The singular value decomposition . . . . .	17
1.5	Theoretical results on the eigendecomposition . . . . .	21
1.5.1	Eigenvalues and Eigenvectors . . . . .	21
1.5.2	Diagonalization of a Symmetric Matrix . . . . .	22
1.5.3	Spectral Decomposition . . . . .	23
1.5.4	Trace and Determinant of A . . . . .	23
<b>2</b>	<b>Quadratic Forms</b>	<b>25</b>
2.1	Quadratic Forms . . . . .	25
2.2	Idempotent Matrices . . . . .	26
2.3	Expected Values and Covariance Matrices of Random Vectors . . . . .	27
2.4	Mean and Variance of Quadratic Forms . . . . .	28
2.5	Distribution of Quadratic Forms in Normal Random Variables . . . . .	28
<b>3</b>	<b>Properties of coefficient estimates</b>	<b>35</b>
3.1	Geometric Properties . . . . .	35
3.2	Calculus Approach . . . . .	36
3.3	Algebraic Properties of $\hat{\beta}$ . . . . .	36
3.4	Rank deficient cases and the Moore-Penrose inverse . . . . .	37
3.4.1	Properties of Generalized Inverses . . . . .	37
3.5	Properties of $\hat{\beta}$ . . . . .	38

**4 The Gauss-Markov Theorem**

**47**

# Chapter 1

## The Gaussian Linear Model

Statistics 849 is part of a two-semester sequence, 849 & 850, on *Theory and Applications of Regression and Analysis of Variance*. In these courses we study statistical models that relate a *response* to values of one or more *covariates*, which are variables that are observed in conjunction with the response.

The statistical inferences are based on a probability model that characterizes the distribution of the vector-valued *random variable*,  $\mathcal{Y}$ , as it depends on values of the covariates. We build the model based on observed values of the responses, represented by the vector  $\mathbf{y}$ , and corresponding values of the covariates.

All the models we will study are based on a *linear predictor* expression,  $\mathbf{X}\boldsymbol{\beta}$ , where the  $n \times p$  matrix  $\mathbf{X}$  is the *model matrix* created from a model specification and the values of the covariates. Here  $n$  is the number of observations and  $p$  is the dimension of the *coefficient vector*,  $\boldsymbol{\beta}$ . The coefficients are *parameters* in the model. We form *estimates*,  $\hat{\boldsymbol{\beta}}$ , of these parameters from the observed data.

We assume that  $n \geq p$ . That is, we have at least as many observations as we have coefficients in the model.

### 1.1 Gaussian Linear Model

A basic model for a response,  $\mathcal{Y}$ , that is measured on a continuous scale, is the *Gaussian Linear Model*

$$\mathcal{Y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}_T, \sigma^2 \mathbf{I}_n) \quad (1.1)$$

where  $\mathbf{I}_n$  is the  $n$ -dimensional identity matrix,  $\boldsymbol{\beta}_T$  is the “true”, but unknown, value of the coefficient vector and  $\mathcal{N}$  denotes the multivariate Gaussian (also called *normal*) distribution.

The probability density of  $\mathcal{Y}$ ,

$$f_{\mathcal{Y}}(\mathbf{y}) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(\frac{-\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}_T\|^2}{2\sigma^2}\right), \quad (1.2)$$

is called a *spherical normal* density, because contours of constant density are concentric spheres centered at  $\mathbf{X}\boldsymbol{\beta}_T$ .

The *likelihood*,  $L(\boldsymbol{\beta}, \sigma | \mathbf{y})$ , of the parameters,  $\boldsymbol{\beta}$  and  $\sigma$ , given the observed responses,  $\mathbf{y}$ , and the model matrix,  $\mathbf{X}$ , is the same expression as the probability density,  $f_{\mathcal{Y}}(\mathbf{y})$ , but regarded as a function of the parameters given the data, as opposed to the density, which is a function of  $\mathbf{y}$  for known values of the parameters.

$$L(\boldsymbol{\beta}, \sigma | \mathbf{y}) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left(-\frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2}{2\sigma^2}\right). \quad (1.3)$$

The *maximum likelihood estimates (mles)* of the parameters are, as the name suggests, the values of the parameters that maximize the likelihood

$$\left(\widehat{\boldsymbol{\beta}}', \widehat{\sigma}_L\right)' = \arg \max_{\boldsymbol{\beta}, \sigma} L(\boldsymbol{\beta}, \sigma | \mathbf{y}) \quad (1.4)$$

As often happens it is much easier to maximize the expression for the *log-likelihood*

$$\ell(\boldsymbol{\beta}, \sigma | \mathbf{y}) = \log(L(\boldsymbol{\beta}, \sigma | \mathbf{y})) = -\frac{n}{2} \log(2\pi\sigma^2) - \frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2}{2\sigma^2} \quad (1.5)$$

than to maximize the likelihood. Because the logarithm function is monotonic increasing, the values of the parameters that maximize the log-likelihood, written  $\arg \max_{\boldsymbol{\beta}, \sigma} \ell(\boldsymbol{\beta}, \sigma | \mathbf{y})$ , are exactly the same as the values that maximize the likelihood,  $\arg \max_{\boldsymbol{\beta}, \sigma} L(\boldsymbol{\beta}, \sigma | \mathbf{y})$ .

The expression can be simplified further by converting to the *deviance*, which is negative twice the log-likelihood,

$$d(\boldsymbol{\beta}, \sigma | \mathbf{y}) = -2\ell(\boldsymbol{\beta}, \sigma | \mathbf{y}) = n \log(2\pi\sigma^2) + \frac{\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2}{\sigma^2}. \quad (1.6)$$

Because of the negative sign, the mle's are the values that minimize the deviance. For any fixed value of  $\sigma^2$ , the deviance is minimized with respect to  $\boldsymbol{\beta}$  when the *residual sum of squares*,

$$S(\boldsymbol{\beta} | \mathbf{y}) = \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2,$$

is minimized. Thus the mle of the coefficient vector,  $\widehat{\boldsymbol{\beta}}$ , in the Gaussian linear model is the *least squares estimate*

$$\widehat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2. \quad (1.7)$$

## 1.2 Linear algebra of least squares

Because the Gaussian Linear Model,  $\mathcal{Y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}_T, \sigma^2 \mathbf{I}_n)$ , is intimately tied to the Euclidean distance,  $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$ , and because the set of all possible fitted values,  $\{\mathbf{X}\boldsymbol{\beta} : \boldsymbol{\beta} \in \mathbb{R}^p\}$ , which is called the *column span* of  $\mathbf{X}$  and written  $\text{col}(\mathbf{X})$ , is a linear subspace of  $\mathbb{R}^n$ , linear algebra, especially as related to the model matrix,  $\mathbf{X}$ , and the response vector,  $\mathbf{y}$ , is fundamental to the theory and practice of linear regression analysis.

We will concentrate on the theoretical and computational aspects of linear algebra as related to the linear model and the implementation of such models in R.

## 1.3 Matrix decompositions

### 1.3.1 Orthogonal matrices

An *orthogonal*  $n \times n$  matrix,  $\mathbf{Q}$  has the property that its transpose is its inverse,

$$\mathbf{Q}'\mathbf{Q} = \mathbf{Q}\mathbf{Q}' = \mathbf{I}_n.$$

These properties imply that the columns of  $\mathbf{Q}$  must be orthogonal to each other and must all have unit length. The same is true for the rows.

An orthogonal matrix has a special property that it preserves lengths.

#### Preserving lengths

For any  $\mathbf{x} \in \mathbb{R}^n$

$$\|\mathbf{Q}\mathbf{x}\|^2 = (\mathbf{Q}\mathbf{x})'\mathbf{Q}\mathbf{x} = \mathbf{x}'\mathbf{Q}'\mathbf{Q}\mathbf{x} = \mathbf{x}'\mathbf{x} = \|\mathbf{x}\|^2$$

Thus the linear transformation determined by  $\mathbf{Q}$  or by  $\mathbf{Q}'$  must be a *rigid* transformation, composed of reflections or rotations.

Orthogonal transformations of the response space,  $\mathbb{R}^n$ , will be important to us because they preserve lengths and because the likelihood of the parameters,  $\boldsymbol{\beta}$ , is related to the squared length of the residual vector,  $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$ .

### 1.3.2 The QR decomposition

Any  $n \times p$  matrix  $\mathbf{X}$  has a QR decomposition consisting of an orthogonal  $n \times n$  matrix  $\mathbf{Q}$  and a  $p \times p$  matrix  $\mathbf{R}$  that is zero below the main diagonal (in other words, it is *upper triangular*). The QR decomposition of the model matrix  $\mathbf{X}$  is written

$$\mathbf{X} = \mathbf{Q} \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix} = [\mathbf{Q}_1 \quad \mathbf{Q}_2] \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix} = \mathbf{Q}_1\mathbf{R} \tag{1.8}$$

where  $\mathbf{Q}_1$  is the first  $p$  columns of  $\mathbf{Q}$  and  $\mathbf{Q}_2$  is the last  $n - p$  columns of  $\mathbf{Q}$ .

That fact that matrices  $\mathbf{Q}$  and  $\mathbf{R}$  must exist is proved by construction. The matrix  $\mathbf{Q}$  is the product of  $p$  *Householder reflections* (see the Wikipedia page for the QR decomposition). The process of generating the upper triangular matrix  $\mathbf{R}$  is similar to the Gram-Schmidt orthogonalization process, but more flexible and more numerically stable. If the diagonal elements of  $\mathbf{R}$  are all non-zero (in practice this means that none of them are very small in absolute value) then  $\mathbf{X}$  has *full column rank* and the columns of  $\mathbf{Q}_1$  form an *orthonormal basis* for  $\text{col}(\mathbf{X})$ .

The implementation of the QR decomposition in R guarantees that any elements on the diagonal of  $\mathbf{R}$  that are considered effectively zero are rearranged by column permutation to occur in the trailing columns. That is, if the rank of  $\mathbf{X}$  is  $k < p$  then the first  $k$  columns of  $\mathbf{Q}$  form an orthonormal basis for  $\text{col}(\mathbf{X}\mathbf{P})$  where  $\mathbf{P}$  is a  $p \times p$  permutation matrix, which means that it is a rearrangement of the columns of  $\mathbf{I}_p$ .

Our text often mentions rank-deficient cases where  $\text{rank}(\mathbf{X}) = k < p$ . In practice the rank deficient case rarely occurs because the process of building the model matrix in R involves a considerable amount of analysis of the model formula to remove the most common cases of rank deficiency. Nevertheless, rank deficiency can occur and is detected and handled in the `lm` function in R.

Because multiplication by an orthogonal matrix like  $\mathbf{Q}'$  preserves lengths we can write

$$\begin{aligned}\hat{\boldsymbol{\beta}} &= \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 \\ &= \arg \min_{\boldsymbol{\beta}} \|\mathbf{Q}'(\mathbf{y} - \mathbf{X}\boldsymbol{\beta})\|^2 \\ &= \arg \min_{\boldsymbol{\beta}} \|\mathbf{Q}'\mathbf{y} - \mathbf{Q}'\mathbf{X}\boldsymbol{\beta}\|^2 \\ &= \arg \min_{\boldsymbol{\beta}} \|\mathbf{c}_1 - \mathbf{R}\boldsymbol{\beta}\|^2 + \|\mathbf{c}_2\|^2\end{aligned}\tag{1.9}$$

where  $\mathbf{c}_1 = \mathbf{Q}'_1\mathbf{y}$  is the first  $p$  elements of  $\mathbf{Q}'\mathbf{y}$  and  $\mathbf{c}_2 = \mathbf{Q}'_2\mathbf{y}$  is the last  $n - p$  elements. If  $\text{rank}(\mathbf{X}) = p$  then  $\text{rank}(\mathbf{R}) = p$  and  $\mathbf{R}^{-1}$  exists so we can write  $\hat{\boldsymbol{\beta}} = \mathbf{R}^{-1}\mathbf{c}_1$  (although you don't actually calculate  $\mathbf{R}^{-1}$  to solve the triangular linear system  $\mathbf{R}\hat{\boldsymbol{\beta}} = \mathbf{c}_1$  for  $\hat{\boldsymbol{\beta}}$ ).

In a model fit by the `lm()` or `aov()` functions in R there is a component `$effects` which is  $\mathbf{Q}'\mathbf{y}$ . The component `$qr` is a condensed form of the QR decomposition of the model matrix  $\mathbf{X}$ . The matrix  $\mathbf{R}$  is embedded in there but the matrix  $\mathbf{Q}$  is a virtual matrix represented as a product of Householder reflections and not usually evaluated explicitly.

**R Exercise:** To see this theory in action, we will start with a very simple linear model. The Formaldehyde data are six observations from a calibration experiment. The response, `optden`, is the optical density. Only one covariate, `carb`, which is the carbohydrate concentration, is included in the data frame.

```
> str(Formaldehyde)

'data.frame':      6 obs. of  2 variables:
 $ carb   : num  0.1 0.3 0.5 0.6 0.7 0.9
 $ optden: num  0.086 0.269 0.446 0.538 0.626 0.782
```

The `model.matrix()` function extracts the model matrix,  $\mathbf{X}$ , from a fitted linear model object

```
> (X <- model.matrix(lm1 <- lm(optden ~ 1 + carb, Formaldehyde)))

  (Intercept) carb
1           1 0.1
2           1 0.3
3           1 0.5
4           1 0.6
5           1 0.7
6           1 0.9
attr(,"assign")
[1] 0 1
```



The `$qr` component is an object of class "qr"

```
> class(qr1m1 <- lm1$qr)
```

```
[1] "qr"
```

for which there are many extractor functions and methods (see `?qr`).

```
> (R <- qr.R(qr1m1))
```

```
      (Intercept)      carb
1 -2.449490 -1.2655697
2  0.000000  0.6390097
```

produces the  $\mathbf{R}$  matrix while

```
> (Q1 <- qr.Q(qr1m1))
```

```
      [,1]      [,2]
[1,] -0.4082483 -0.65205066
[2,] -0.4082483 -0.33906635
[3,] -0.4082483 -0.02608203
[4,] -0.4082483  0.13041013
[5,] -0.4082483  0.28690229
[6,] -0.4082483  0.59988661
```

by default produces  $\mathbf{Q}_1$ . First we should check that their product is indeed  $\mathbf{X}$

```
> (Q1R <- Q1 %*% R)
```

```
      (Intercept) carb
[1,]           1  0.1
[2,]           1  0.3
[3,]           1  0.5
[4,]           1  0.6
[5,]           1  0.7
[6,]           1  0.9
```

It seems to be the same, although as in all floating point calculations on a computer, there may be some small imprecision caused by round-off error in the calculations. This is why we don't use exact comparisons on the results of floating point calculations

```
> all(X == Q1R)
```

```
[1] FALSE
```

but instead compare results using

```
> all.equal(X, Q1R, check.attr = FALSE)
```

```
[1] TRUE
```

(As a model matrix,  $\mathbf{x}$  has some additional attributes that are not present in the product  $\mathbf{Q}\mathbf{1R}$ , which is why we turn off checking of the attributes of the two objects.)

Notice that  $\mathbf{X}$  and  $\mathbf{y}$  are not explicitly part of the fitted model object, `lm1`.

```
> names(lm1)
```

```
[1] "coefficients" "residuals"      "effects"        "rank"          "fitted.values"
[6] "assign"       "qr"             "df.residual"   "xlevels"      "call"
[11] "terms"       "model"
```

Both are generated from the *model.frame*, which is stored as the component `$model`. Although this is getting into more detail than is needed at present, the reason for introducing the model frame is to say that the safe way of extracting the response vector,  $\mathbf{y}$ , is

```
> (y <- model.response(model.frame(lm1)))
```

```
      1      2      3      4      5      6
0.086 0.269 0.446 0.538 0.626 0.782
```

We have already seen that `model.matrix` returns the matrix  $\mathbf{X}$  from the fitted model object.

We can produce the full  $n \times n$  orthogonal matrix  $\mathbf{Q}$  from `qr.Q()` by setting the optional argument `complete=TRUE`. We do this for illustration only. In practice the matrix  $\mathbf{Q}$  is never explicitly created — it is a “virtual” matrix in the sense that it is a product of Householder reflections that are stored much more compactly than  $\mathbf{Q}$  would be stored.

```
> (Q <- qr.Q(qr1m1, complete=TRUE))
```

```
      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,] -0.4082483 -0.65205066 -0.37370452 -0.3405290 -0.3073534 -0.2410023
[2,] -0.4082483 -0.33906635  0.05460995  0.2207196  0.3868293  0.7190487
[3,] -0.4082483 -0.02608203  0.86857638 -0.1439791 -0.1565346 -0.1816455
[4,] -0.4082483  0.13041013 -0.15359661  0.8125532 -0.2212971 -0.2889976
[5,] -0.4082483  0.28690229 -0.17576960 -0.2309146  0.7139404 -0.3963496
[6,] -0.4082483  0.59988661 -0.22011559 -0.3178501 -0.4155847  0.3889463
```

The `$effects` vector should be the product  $\mathbf{Q}'\mathbf{y}$  but it happens that they are stored differently. The `$effects` vector is a vector of length  $n$  and the product  $\mathbf{Q}'\mathbf{y}$  is an  $n \times 1$  matrix. To compare them, we need to make `$effects` an  $n \times 1$  matrix or make  $\mathbf{Q}'\mathbf{y}$  into a vector. A convenient way of making an  $n \times 1$  matrix from an  $n$ -vector is the function `cbind()`, which creates matrices or data frames by binding columns together. If we give it a single vector argument it creates an  $n \times 1$  matrix

```
> str(cbind(lm1$effects))
```

```

num [1:6, 1] -1.12146 0.55996 0.00514 0.00992 0.01069 ...
- attr(*, "dimnames")=List of 2
..$ : chr [1:6] "(Intercept)" "carb" "" "" ...
..$ : NULL

> str(crossprod(Q, y))

num [1:6, 1] -1.12146 0.55996 0.00514 0.00992 0.01069 ...

> all.equal(cbind(lm1$effects), crossprod(Q, y), check.attr=FALSE)

[1] TRUE

```

(The function `crossprod(A,B)` creates  $A'B$  directly, without creating  $A'$  from  $A$ . It is most commonly used to create matrices like  $X'X$  as

```

> crossprod(X)

              (Intercept) carb
(Intercept)      6.0 3.10
carb              3.1 2.01

```

The companion function, `tcrossprod`, creates  $XX'$ .)

If we wish to do the comparison by converting  $Q'y$  to a vector, we can use

```

> all.equal(lm1$effects, as.vector(crossprod(Q, y)), check.attr=FALSE)

[1] TRUE

```

I find `cbind` easier to type than `as.vector`.

Another way of generating  $Q'y$  is with the function `qr.qty()`

```

> all.equal(lm1$effects, qr.qty(qr1m1, y), check.attr=FALSE)

[1] TRUE

```

We should check that  $Q$  is indeed orthogonal and that  $Q_1'Q_1 = I_p$ . The matrix  $I_k$  is generated by `diag(nrow=k)`.

```

> all.equal(crossprod(Q1), diag(nrow=ncol(Q1)))

[1] TRUE

> all.equal(crossprod(Q), diag(nrow=nrow(Q)))

[1] TRUE

> all.equal(tcrossprod(Q), diag(nrow=nrow(Q)))

```

```
[1] TRUE
```

When we print a matrix that may have negligibly small non-zero values in it

```
> crossprod(Q1)
```

```
      [,1]      [,2]
[1,] 1.000000e+00 1.197637e-16
[2,] 1.197637e-16 1.000000e+00
```

we can clean up the output with `zapsmall()` which, as the name suggests, zeros the very small values.

```
> zapsmall(crossprod(Q1))
```

```
      [,1] [,2]
[1,]    1    0
[2,]    0    1
```

```
> zapsmall(crossprod(Q))
```

```
      [,1] [,2] [,3] [,4] [,5] [,6]
[1,]    1    0    0    0    0    0
[2,]    0    1    0    0    0    0
[3,]    0    0    1    0    0    0
[4,]    0    0    0    1    0    0
[5,]    0    0    0    0    1    0
[6,]    0    0    0    0    0    1
```

Because the diagonal elements of  $\mathbf{R}$  are all safely non-zero, we can solve the system  $\mathbf{R}\hat{\boldsymbol{\beta}} = \mathbf{Q}'_1\mathbf{y}$  for the coefficient estimates,  $\hat{\boldsymbol{\beta}}$ . We could use the function `solve()` to do this but it is better to use `backsolve()` for the solution to upper triangular systems,

```
> coef(lm1)
```

```
(Intercept)      carb
0.005085714 0.876285714
```

```
> backsolve(R, crossprod(Q1, y))
```

```
      [,1]
[1,] 0.005085714
[2,] 0.876285714
```

```
> all.equal(coef(lm1), as.vector(backsolve(R, crossprod(Q1, y))), check.attr=FALSE)
```

```
[1] TRUE
```

The function `qr.coef` combines the multiplication of  $\mathbf{y}$  by  $\mathbf{Q}'_1$  and the `backsolve` step

```
> qr.coef(qr1m1, y)
```

```
(Intercept)      carb
0.005085714 0.876285714
```

**R Exercise:** As seen above, a linear model is specified as a model formula and the data frame in which to evaluate the formula. Because the formula is analyzed for conditions that may introduce rank deficiency and consequently removes those conditions, rank deficient cases occur infrequently. Of course, it is possible to artificially generate data with a built-in rank dependency

```
> set.seed(1234) # allow for reproducible "random" numbers
> badDat <- within(data.frame(x1=1:20, x2=rnorm(20,mean=6,sd=0.2),
+ x4=rexp(20,rate=0.02),
+ y=runif(20,min=18,max=24)),
+ x3 <- x1 + 2*x2) # create linear combination
> (summary(lm2 <- lm(y ~ x1 + x2 + x3 + x4, badDat)))
```

Call:

```
lm(formula = y ~ x1 + x2 + x3 + x4, data = badDat)
```

Residuals:

```
   Min       1Q   Median       3Q      Max
-2.3444 -1.7670 -0.3585  1.6159  3.0292
```

Coefficients: (1 not defined because of singularities)

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	1.793e+01	1.390e+01	1.290	0.215
x1	4.140e-02	8.553e-02	0.484	0.635
x2	3.822e-01	2.358e+00	0.162	0.873
x3	NA	NA	NA	NA
x4	3.901e-05	8.680e-03	0.004	0.996

Residual standard error: 2.02 on 16 degrees of freedom

Multiple R-squared: 0.021, Adjusted R-squared: -0.1626

F-statistic: 0.1144 on 3 and 16 DF, p-value: 0.9504

```
> (lm2qr <- lm2$qr)$rank
```

```
[1] 4
```

```
> qr.R(lm2qr) # the columns are rearranged
```

	(Intercept)	x1	x2	x4	x3
1	-4.472136	-46.95743	-26.6086150	-182.87421	-1.001747e+02
2	0.000000	25.78759	0.1721295	83.85993	2.613185e+01
3	0.000000	0.00000	-0.8668932	35.62409	-1.733786e+00
4	0.000000	0.00000	0.0000000	232.75139	2.005660e-15
5	0.000000	0.00000	0.0000000	0.00000	5.179752e-15

```
> lm2qr$pivot # the permutation vector
```

```
[1] 1 2 3 5 4
```

but we don't do this in practice.

The common case of an analysis of variance model which, when written as

$$y_{i,j} = \mu + \alpha_i + \epsilon_{i,j} \quad i = 1, \dots, I, \quad j = 1, \dots, n_i$$

would generate linearly dependent columns for  $\mu$  and the  $\alpha_i$ ,  $i = 1, \dots, I$  is analyzed and represented by the intercept column and  $I - 1$  columns for the factor.

```
> str(InsectSprays)

'data.frame':      72 obs. of  2 variables:
 $ count: num  10 7 20 14 14 12 10 23 17 20 ...
 $ spray: Factor w/ 6 levels "A","B","C","D",...: 1 1 1 1 1 1 1 1 1 1 ...

> unique(mm <- model.matrix(lm3 <- lm(count ~ spray, InsectSprays)))

  (Intercept) sprayB sprayC sprayD sprayE sprayF
1             1      0      0      0      0      0
13            1      1      0      0      0      0
25            1      0      1      0      0      0
37            1      0      0      1      0      0
49            1      0      0      0      1      0
61            1      0      0      0      0      1

> attr(mm, "assign")

[1] 0 1 1 1 1 1

> qr.R(lm3[["qr"]])

  (Intercept)  sprayB  sprayC  sprayD  sprayE  sprayF
1  -8.485281 -1.414214 -1.4142136 -1.4142136 -1.4142136 -1.4142136
2   0.000000  3.162278 -0.6324555 -0.6324555 -0.6324555 -0.6324555
3   0.000000  0.000000  3.0983867 -0.7745967 -0.7745967 -0.7745967
4   0.000000  0.000000  0.0000000  3.0000000 -1.0000000 -1.0000000
5   0.000000  0.000000  0.0000000  0.0000000  2.8284271 -1.4142136
6   0.000000  0.000000  0.0000000  0.0000000  0.0000000  2.4494897
```

This shows only the unique rows in the model matrix. The six levels of the `spray` factor are represented by 5 indicator columns.

Because we are discussing an analysis of variance model we also show the `"assign"` attribute of the model matrix. This indicates that the first column is associated with the 0th term, which is the intercept, and the second through sixth columns are associated with the first term, which is `spray`.

In general, a factor with  $I$  levels is converted to a set of  $I - 1$  columns. These are called *contrasts*, but be warned that these do not fulfill the definition of contrasts as used in some texts. You should think of them as being a set of columns representing changes between levels of the factor.

The type of contrasts generated is controlled by the option called `"contrasts"`.

```
> getOption("contrasts")

      unordered      ordered
"contr.treatment"  "contr.poly"
```

(By the way, plotting these data first would show that this is not a good model. The count variable is, not surprisingly, a count and does not have constant variance. A better model would use the square root of the count as a response.)

### The determinant of an orthogonal matrix

As described on its Wikipedia page, the *determinant*,  $|\mathbf{A}|$ , of the  $k \times k$  square matrix,  $\mathbf{A}$ , is the volume of the parallelepiped spanned by its columns (or, equivalently, the volume spanned by its rows). Because we can consider either the rows or the columns when evaluating the determinant, we must have

$$|\mathbf{A}| = |\mathbf{A}'|.$$

We can regard  $|\mathbf{A}|$  as the magnification factor in the transformation  $\mathbf{x} \rightarrow \mathbf{A}\mathbf{x}$  from  $\mathbb{R}^k$  to  $\mathbb{R}^k$ . This transformation takes the unit cube to a parallelepiped with volume  $|\mathbf{A}|$ . Composing transformations will just multiply the magnification factors so we must have

$$|\mathbf{AB}| = |\mathbf{A}||\mathbf{B}|$$

We know that the columns of an orthogonal matrix  $\mathbf{Q}$  are *orthonormal* hence they span a unit volume. That is, for an  $n \times n$  matrix  $\mathbf{Q}$

$$\mathbf{Q}'\mathbf{Q} = \mathbf{I}_n \quad \Rightarrow \quad |\mathbf{Q}| = \pm 1.$$

The sign indicates whether the transformation preserves orientation. In two dimensions a rotation preserves orientation and a reflection reverses orientation.

Furthermore, the determinant of a *diagonal* matrix or a *triangular* matrix is simply the product of its diagonal elements. (For a triangular matrix, first consider the  $2 \times 2$  case and the shape of the parallelogram spanned by the columns. The width of the parallelogram is the (1,1) element and the height is the (2,2) element so the area is the product of the diagonal elements (up to sign). Then convince yourself that this property scales to a parallelepiped in  $k$  dimensions.)

From these properties we can formally derive

$$1 = |\mathbf{I}_n| = |\mathbf{Q}\mathbf{Q}'| = |\mathbf{Q}||\mathbf{Q}'| = |\mathbf{Q}|^2 \quad \Rightarrow \quad |\mathbf{Q}| = \pm 1.$$

Interestingly, one way that the determinant,  $|\mathbf{A}|$  is evaluated in practice is by forming the QR decomposition of  $\mathbf{A}$ , taking the product of the diagonal elements of  $\mathbf{R}$ , and determining whether  $|\mathbf{Q}|$  has a plus or a minus sign.

**R Exercise:** The `det()` function evaluates the determinant of a square matrix (although if you check its definition you will find that it just calls another function `determinant`, which is the preferred approach).

```
> all.equal(det(R), prod(diag(R)))
```

```
[1] TRUE
```

```
> det(crossprod(Q1))
```

```
[1] 1
```

```
> det(crossprod(Q))
```

```
[1] 1
```

### 1.3.3 Comparison to the usual text-book formulas

Most text books state that the least squares estimates are

$$\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y} \quad (1.10)$$

giving the impression that  $\hat{\beta}$  is calculated this way. It isn't.

If you substitute  $\mathbf{X} = \mathbf{Q}_1\mathbf{R}$  in eqn. 1.10 you get

$$(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y} = (\mathbf{R}'\mathbf{R})^{-1}\mathbf{R}'\mathbf{Q}'_1\mathbf{y} = \mathbf{R}^{-1}(\mathbf{R}')^{-1}\mathbf{R}'\mathbf{Q}'_1\mathbf{y} = \mathbf{R}^{-1}\mathbf{Q}'_1\mathbf{y},$$

our previous result.

Whenever you see  $\mathbf{X}'\mathbf{X}$  in a formula you should mentally replace it by  $\mathbf{R}'\mathbf{R}$  and similarly replace  $(\mathbf{X}'\mathbf{X})^{-1}$  by  $\mathbf{R}^{-1}(\mathbf{R}')^{-1}$  then see if you can simplify the result.

For example, the variance of the least squares estimator  $\hat{\beta}$  is

$$\text{Var}(\hat{\beta}) = \sigma^2(\mathbf{X}'\mathbf{X})^{-1} = \sigma^2\mathbf{R}^{-1}(\mathbf{R}')^{-1}$$

The R function `chol2inv` calculates  $\mathbf{R}^{-1}(\mathbf{R}')^{-1}$  directly from  $\mathbf{R}$  without evaluating  $\mathbf{R}^{-1}$  explicitly (not a big deal in most cases but when  $p$  is very large it should be faster and more accurate than evaluating  $\mathbf{R}^{-1}$  explicitly).

Also, the determinant of  $\mathbf{X}'\mathbf{X}$  is

$$|\mathbf{X}'\mathbf{X}| = |\mathbf{R}'\mathbf{R}| = |\mathbf{R}|^2 = \left( \prod_{i=1}^p r_{i,i} \right)^2$$

The fitted values  $\hat{\mathbf{y}}$  are  $\mathbf{Q}_1\mathbf{Q}'_1\mathbf{y}$  and thus the *hat matrix* (which puts a “hat” on  $\mathbf{y}$  by transforming it to  $\hat{\mathbf{y}}$ ) is the  $n \times n$  matrix  $\mathbf{Q}_1\mathbf{Q}'_1$ . Often we are interested in the diagonal elements of the hat matrix, which are the sums of the squares of rows of  $\mathbf{Q}_1$ . (In practice you don't want to calculate the entire  $n \times n$  hat matrix just to get the diagonal elements when  $n$  could be very large.)



The residuals,  $\hat{\mathbf{e}} = \mathbf{y} - \hat{\mathbf{y}}$ , are calculated as  $\hat{\mathbf{e}} = \mathbf{Q}_2\mathbf{Q}'_2\mathbf{y}$ .

The matrices  $\mathbf{Q}_1\mathbf{Q}'_1$  and  $\mathbf{Q}_2\mathbf{Q}'_2$  are *projection matrices*, which means that they are symmetric and *idempotent*. (A square matrix  $\mathbf{A}$  is idempotent if  $\mathbf{A}\mathbf{A} = \mathbf{A}$ .) When  $\text{rank}(\mathbf{X}) = p$ , the hat matrix  $\mathbf{Q}_1\mathbf{Q}'_1$  projects any vector in  $\mathbb{R}^n$  onto the column span of  $\mathbf{X}$ . The other projection,  $\mathbf{Q}_2\mathbf{Q}'_2$ , is onto the subspace orthogonal to the column span of  $\mathbf{X}$  (see the figure on the front cover of the text).

**R Exercise:** We have already seen that  $\hat{\boldsymbol{\beta}}$  can be calculated as

```
> backsolve(R, crossprod(Q1, y))
```

```
      [,1]
[1,] 0.005085714
[2,] 0.876285714
```

or as

```
> qr.coef(qr1m1, y)
```

```
(Intercept)      carb
0.005085714 0.876285714
```

The functions `qr.fitted()` and `qr.resid()` perform projection onto  $\text{col}(\mathbf{X})$  and onto its orthogonal complement, respectively.

The fitted values are, unsurprisingly, calculated as

```
> qr.fitted(qr1m1, y)
```

```
      1      2      3      4      5      6
0.09271429 0.26797143 0.44322857 0.53085714 0.61848571 0.79374286
```

```
> all.equal(qr.fitted(qr1m1, y), fitted(lm1))
```

```
[1] TRUE
```

and the residuals as

```
> qr.resid(qr1m1, y)
```

```
      1      2      3      4      5      6
-0.006714286 0.001028571 0.002771429 0.007142857 0.007514286 -0.011742857
```

```
> all.equal(qr.resid(qr1m1, y), residuals(lm1))
```

```
[1] TRUE
```

We use the explicit calculations for illustration only. In practice, use of the “extractor” methods, `fitted()` and `residuals()`, is preferred.

If we wanted the projection matrices  $\mathbf{P}_1$  for projection onto  $\text{col } \mathbf{X}$  and onto the residual space we could form them as

```
> (P1 <- tcrossprod(Q1))
```

```
      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,] 0.59183673 0.38775510 0.1836735 0.08163265 -0.02040816 -0.22448980
[2,] 0.38775510 0.28163265 0.1755102 0.12244898 0.06938776 -0.03673469
[3,] 0.18367347 0.17551020 0.1673469 0.16326531 0.15918367 0.15102041
[4,] 0.08163265 0.12244898 0.1632653 0.18367347 0.20408163 0.24489796
[5,] -0.02040816 0.06938776 0.1591837 0.20408163 0.24897959 0.33877551
[6,] -0.22448980 -0.03673469 0.1510204 0.24489796 0.33877551 0.52653061
```

and

```
> (P2 <- tcrossprod(Q[, -(1:2)]))
```

```
      [,1]      [,2]      [,3]      [,4]      [,5]      [,6]
[1,] 0.40816327 -0.38775510 -0.1836735 -0.08163265 0.02040816 0.22448980
[2,] -0.38775510 0.71836735 -0.1755102 -0.12244898 -0.06938776 0.03673469
[3,] -0.18367347 -0.17551020 0.8326531 -0.16326531 -0.15918367 -0.15102041
[4,] -0.08163265 -0.12244898 -0.1632653 0.81632653 -0.20408163 -0.24489796
[5,] 0.02040816 -0.06938776 -0.1591837 -0.20408163 0.75102041 -0.33877551
[6,] 0.22448980 0.03673469 -0.1510204 -0.24489796 -0.33877551 0.47346939
```

respectively (the expression `Q[, -(1:2)]` drops the first two columns of  $Q$  producing  $Q_2$ ). (And, of course, we don't do this in practice, especially if  $n$  is large. Instead we use `qr.fitted` or `qr.resid` if we want to project vectors other than  $\mathbf{y}$ .)

We can check that  $P_1$  and  $P_2$  are projection matrices. They are obviously symmetric by construction so we check the idempotent property

```
> all.equal(P1 %*% P1, P1)
```

```
[1] TRUE
```

```
> all.equal(P2 %*% P2, P2)
```

```
[1] TRUE
```

Because  $P_1$  is projection onto  $\text{col}(\mathbf{X})$  it should take  $\mathbf{X}$  to itself

```
> all.equal(P1 %*% X, X, check.attr=FALSE)
```

```
[1] TRUE
```

and  $P_2$  should take  $\mathbf{X}$  to zeros, although in practice we expect very small but possibly non-zero values.

```
> all.equal(P2 %*% X, 0 * X, check.attr=FALSE)
```

```
[1] TRUE
```

(The weird construction,  $0 * X$ , create a matrix of zeros that is the same size as  $X$ .)

Because  $P1$  is the hat matrix, we can get its diagonal elements as

```
> diag(P1)
```

```
[1] 0.5918367 0.2816327 0.1673469 0.1836735 0.2489796 0.5265306
```

As mentioned, the alternative calculation is

```
> rowSums(Q1^2)
```

```
[1] 0.5918367 0.2816327 0.1673469 0.1836735 0.2489796 0.5265306
```

and a third way, preferred in practice, is

```
> hatvalues(lm1)
```

```
      1      2      3      4      5      6
0.5918367 0.2816327 0.1673469 0.1836735 0.2489796 0.5265306
```

$\text{rank}(X)$ , which is the number of linearly independent columns in  $X$ , is calculated during the decomposition and also stored as the `$rank` component of the fitted model

```
> lm1$rank
```

```
[1] 2
```

```
> qr1m1$rank
```

```
[1] 2
```

### 1.3.4 R functions related to the QR decomposition

To review, every time you fit a linear model with `lm` or `aov` or `lm.fit`, the returned object contains a `$qr` component. This is a condensed form of the QR decomposition of  $X$ , only slightly larger than  $X$  itself. Its class is "qr".

There are several extractor functions for a "qr" object: `qr.R()`, `qr.Q()` and `qr.X()`, which regenerates the original matrix. By default `qr.Q()` returns the matrix called  $Q_1$  above with  $p$  columns but you can specify the number of columns desired. Typical alternative choices are  $n$  or  $\text{rank}(X)$ .

The `$rank` component of a "qr" object is the computed rank of  $X$  (and, hence, of  $R$ ). The `$pivot` component is the permutation applied to the columns. It will be `1:p` when  $\text{rank}(X) = p$  but when  $\text{rank}(X) < p$  it may be other than the identity permutation.

Several functions are applied to a "qr" object and a vector or matrix. These include `qr.coef()`, `qr.qy()`, `qr.qty()`, `qr.resid()` and `qr.fitted()`. The `qr.qy()` and `qr.qty()` functions multiply an  $n$ -vector or an  $n \times m$  matrix by  $Q$  or  $Q'$  without ever forming  $Q$ . Similarly, `qr.fitted()` creates  $Q_1 Q_1' x$  and `qr.resid()` creates  $Q_2 Q_2' x$  without forming  $Q$ .

The `is.qr()` function tests an object to determine if it is of class "qr".

## 1.4 Related matrix decompositions

### 1.4.1 The Cholesky decomposition

The Cholesky decomposition of a positive definite symmetric matrix, which means a  $p \times p$  symmetric matrix  $\mathbf{A}$  such that  $\mathbf{x}'\mathbf{A}\mathbf{x} > 0$  for all non-zero  $\mathbf{x} \in \mathbb{R}^p$  is of the form

$$\mathbf{A} = \mathbf{R}'\mathbf{R} = \mathbf{L}\mathbf{L}'$$

where  $\mathbf{R}$  is an upper triangular  $p \times p$  matrix and  $\mathbf{L} = \mathbf{R}'$  is lower triangular. The two forms are the same decomposition: it is just a matter of whether you want  $\mathbf{R}$ , the factor on the right, or  $\mathbf{L}$ , the factor on the left. Generally statisticians write the decomposition as  $\mathbf{R}'\mathbf{R}$ .

The decomposition is only determined up to changes in sign of the rows of  $\mathbf{R}$  (or, equivalently, the columns of  $\mathbf{L}$ ). For definiteness we require positive diagonal elements in  $\mathbf{R}$ .

When  $\text{rank}(\mathbf{X}) = p$  the Cholesky decomposition  $\mathbf{R}$  of  $\mathbf{X}'\mathbf{X}$  is equal to the matrix  $\mathbf{R}$  from the QR decomposition up to changes in sign of rows. The matrix  $\mathbf{X}'\mathbf{X}$  matrix is obviously symmetric and it is positive definite because

$$\mathbf{x}'(\mathbf{X}'\mathbf{X})\mathbf{x} = \mathbf{x}'(\mathbf{R}'\mathbf{R})\mathbf{x} = \|\mathbf{R}\mathbf{x}\|^2 \geq 0$$

with equality only when  $\mathbf{R}\mathbf{x} = \mathbf{0}$ , which, when  $\text{rank}(\mathbf{R}) = p$ , implies that  $\mathbf{x} = \mathbf{0}$ .

### 1.4.2 Evaluation of the Cholesky decomposition

The R function `chol()` evaluates the Cholesky decomposition. As mentioned above `chol2inv()` creates  $(\mathbf{X}'\mathbf{X})^{-1}$  directly from the Cholesky decomposition of  $\mathbf{X}'\mathbf{X}$ .

Generally the QR decomposition is preferred to the Cholesky decomposition for least squares problems because there is a certain loss of precision when forming  $\mathbf{X}'\mathbf{X}$ . However, when  $n$  is very large you may want to build up  $\mathbf{X}'\mathbf{X}$  using blocks of rows. Also, if  $\mathbf{X}$  is *sparse* it is an advantage to use sparse matrix techniques to evaluate and store the Cholesky decomposition.

The `Matrix` package for R provides even more capabilities related to the Cholesky decomposition, especially for sparse matrices.

For everything we will do in Statistics 849 the QR decomposition should be the method of choice.

#### R Exercises:

```
> chol(crossprod(X))
```

```

              (Intercept)      carb
(Intercept)  2.449490  1.2655697
carb         0.000000  0.6390097
```

### 1.4.3 The singular value decomposition

Another decomposition related to orthogonal matrices is the singular value decomposition (or SVD) in which the matrix  $\mathbf{X}$  is reduced to a diagonal form

$$\mathbf{X} = \mathbf{U}_1 \mathbf{D} \mathbf{V}' = \mathbf{U} \begin{bmatrix} \mathbf{D} \\ \mathbf{0} \end{bmatrix} \mathbf{V}' \quad (1.11)$$

where  $\mathbf{U}$  is an  $n \times n$  orthogonal matrix,  $\mathbf{D}$  is a  $p \times p$  diagonal matrix with non-negative diagonal elements (which are called the *singular values* of  $\mathbf{X}$ ) and  $\mathbf{V}$  is a  $p \times p$  orthogonal matrix. As for  $\mathbf{Q}$  and  $\mathbf{Q}_1$ ,  $\mathbf{U}_1$  consists of the first  $p$  columns of  $\mathbf{U}$ . For definiteness we order the diagonal elements of  $\mathbf{D}$ , which must be non-negative, in decreasing order.

Just like  $\mathbf{Q}_1$ , the columns of  $\mathbf{U}_1$  form an orthonormal basis for  $\text{col}(\mathbf{X})$  when  $\mathbf{X}$  has full column rank (which means that the singular values are all safely positive). If  $\text{rank}(\mathbf{X}) = r < p$  then the first  $r$  columns of  $\mathbf{U}$  form the orthonormal basis.

One way to visualize the singular value decomposition of  $\mathbf{X}$  is to remember that a  $p$ -sphere in  $\mathbb{R}^p$  will get mapped to an ellipsoid in  $\text{col}(\mathbf{X})$  by  $\mathbf{X}$ . The singular values are the lengths of the principal axes of this ellipsoid. The right singular vectors (columns of  $\mathbf{V}$ ) are the directions in the parameter space that map onto the principal axes of the ellipsoid. The first  $\text{rank}(\mathbf{X})$  left singular vectors (columns of  $\mathbf{U}$ ) are the principal axes of the ellipsoid.

The singular value decomposition of  $\mathbf{X}$  is related to the eigendecomposition or spectral decomposition of  $\mathbf{X}'\mathbf{X}$  because

$$\mathbf{X}'\mathbf{X} = \mathbf{V} \mathbf{D} \mathbf{U}'_1 \mathbf{U}_1 \mathbf{D} \mathbf{V}' = \mathbf{V} \mathbf{D}^2 \mathbf{V}'$$

implying that the eigenvalues of  $\mathbf{X}'\mathbf{X}$  are the squares of the singular values of  $\mathbf{X}$  and the right singular vectors, which are the columns of  $\mathbf{V}$ , are also the eigenvectors of  $\mathbf{X}'\mathbf{X}$

Calculation of the SVD is an iterative (as opposed to a direct) computation and potentially more computing intensive than the QR decomposition, although modern methods for evaluating the SVD are very good indeed.

Symbolically we can write the least squares solution in the full-rank case as

$$\hat{\boldsymbol{\beta}} = \mathbf{V} \mathbf{D}^{-1} \mathbf{U}'_1 \mathbf{y}$$

where  $\mathbf{D}^{-1}$  is a diagonal matrix whose diagonal elements are the inverses of the diagonal elements of  $\mathbf{D}$ .

The pseudo-inverse or *generalized inverse* of  $\mathbf{X}$ , written  $\mathbf{X}^-$ , is calculated from the pseudo-inverse of the diagonal matrix,  $\mathbf{D}$ . In theory the diagonal elements of  $\mathbf{D}^-$  are  $1/d_{i,i}$  when  $d_{i,i} \neq 0$  and 0 when  $d_{i,i} = 0$ . However, we can't count on  $d_{i,i}$  being 0 even when, in theory, it should be. We need to decide when the singular values are close to zero, which is actually a very difficult problem. At best we can use some heuristics, based on the ratio of  $d_{i,i}/d_{1,1}$ , to decide when a diagonal element is "effectively zero".

The use of the pseudo-inverse seems to be a convenient way to handle rank-deficient  $\mathbf{X}$  matrices but, as mentioned above, the best way to handle rank-deficient  $\mathbf{X}$  matrices is not to produce them in the first place. Even when a rank-deficient  $\mathbf{X}$  is produced we use a pivoted QR decomposition rather than a pseudo-inverse.

**R Exercises:** The SVD of our model matrix  $\mathbf{X}$  is

```
> str(Xsv <- svd(X))
```

```
List of 3
```

```
$ d: num [1:2] 2.773 0.564
$ u: num [1:6, 1:2] 0.334 0.368 0.403 0.42 0.437 ...
$ v: num [1:2, 1:2] 0.878 0.479 -0.479 0.878
```

We see that, by default, the `svd()` function produces the diagonal of  $\mathbf{D}$ , the matrix  $\mathbf{U}_1$  and the matrix  $\mathbf{V}$ . We should check that  $\mathbf{X} = \mathbf{U}_1 \mathbf{D} \mathbf{V}'$ , as advertised. We could form a diagonal matrix  $\mathbf{D}$  from the `$d` component of `Xsv` but multiplication of a matrix on the left by a diagonal matrix corresponds to scaling its rows, so we write the reconstruction as

```
> Xsv$u %*% (Xsv$d * t(Xsv$v))
```

```
      [,1] [,2]
[1,]    1  0.1
[2,]    1  0.3
[3,]    1  0.5
[4,]    1  0.6
[5,]    1  0.7
[6,]    1  0.9
```

```
> all.equal(Xsv$u %*% (Xsv$d * t(Xsv$v)), X, check.attr=FALSE)
```

```
[1] TRUE
```

We check that the matrix  $\mathbf{U}_1$  has orthonormal columns and that  $\mathbf{V}$  is orthogonal

```
> zapsmall(crossprod(Xsv$u))
```

```
      [,1] [,2]
[1,]    1    0
[2,]    0    1
```

```
> zapsmall(crossprod(Xsv$v))
```

```
      [,1] [,2]
[1,]    1    0
[2,]    0    1
```

The squares of the singular values should be the eigenvalues of  $\mathbf{X}'\mathbf{X}$  and the eigenvectors of  $\mathbf{X}'\mathbf{X}$  should be the columns of  $\mathbf{V}$ , up to changes in sign along columns. (The eigenvectors, which are really just directions, are only determined up to changes in sign, and in the case of repeated eigenvalues, only up to orthogonal transformation within the repeated eigenvalue's eigenspace.)

```
> str(ev <- eigen(crossprod(X)))
```

```

List of 2
 $ values : num [1:2] 7.691 0.319
 $ vectors: num [1:2, 1:2] -0.878 -0.479 0.479 -0.878

> Xsv$d^2

[1] 7.6914651 0.3185349

> all.equal(ev$values, Xsv$d^2)

[1] TRUE

> ev$vectors

      [,1]      [,2]
[1,] -0.8778294  0.4789735
[2,] -0.4789735 -0.8778294

> Xsv$v

      [,1]      [,2]
[1,] 0.8778294 -0.4789735
[2,] 0.4789735  0.8778294

> all.equal(-Xsv$v, ev$vectors)

[1] TRUE

```

In practice, you never need to calculate the eigenvalues and eigenvectors of  $\mathbf{X}'\mathbf{X}$ . It is more effective and more stable to calculate the singular value decomposition of  $\mathbf{X}$  and use the squares of the singular values and the  $\$v$  component (assuming that you really do need the eigenvalues and eigenvectors which, most of the time, you don't).

The reason that it is preferable to work with decompositions of  $\mathbf{X}$  rather than forming  $\mathbf{X}'\mathbf{X}$  is related to the condition number of these matrices. As described on the Wikipedia page, the condition number of a matrix, written  $\kappa(\mathbf{X})$ , is the ratio of its largest and smallest singular values. Obviously we must have  $\kappa(\mathbf{X}) \geq 1$ . A matrix with  $\kappa$  close to 1 is well-conditioned. A matrix with a very large condition number is close to being singular, in that spheres are mapped to highly elongated ellipsoids.

An orthogonal matrix or a rectangular matrix with orthonormal columns must have a condition number of 1 because it maps a sphere to a sphere. (Recall that, for us, rectangular matrices like  $\mathbf{X}$  have more rows than columns. In the opposite case, more columns than rows, it would be the rows that are orthonormal.) In fact, all the singular values of an orthogonal matrix must be unity because it preserves lengths so the unit sphere gets mapped to the unit sphere.

We can check that matrices like  $\mathbf{Q}$ ,  $\mathbf{Q}_1$  and  $\mathbf{U}_1$  have a condition number of 1.

```
> svd(Q, nu=0, nv=0)$d
```

```
[1] 1 1 1 1 1 1
```

```
> kappa(Q)
```

```
[1] 1
```

```
> svd(Q1, nu=0, nv=0)$d
```

```
[1] 1 1
```

```
> kappa(Q1)
```

```
[1] 1
```

```
> kappa(Xsv$u)
```

```
[1] 1
```

The condition number of  $\mathbf{X}$  can be explicitly calculated as

```
> Xsv$d
```

```
[1] 2.773349 0.564389
```

```
> (kappaX <- Xsv$d[1]/Xsv$d[length(Xsv$d)])
```

```
[1] 4.913897
```

(The complicated expression in the last line is to generalize the calculation. It will give the correct answer when there are more than two singular values.) The `kappa()` function, by default, produces an upper bound on the condition number, because this upper bound can be calculated directly. To get the exact value set the optional argument `exact=TRUE`.

```
> kappa(X)
```

```
[1] 5.1073
```

```
> kappa(X, exact=TRUE)
```

```
[1] 4.913897
```

In practice we usually calculate the reciprocal of the condition number because its value is in  $[0, 1]$  and it is easier to decide when it is close to zero instead of trying to decide when  $\kappa(\mathbf{X})$  is “close to”  $\infty$ . We compare the reciprocal condition number to the relative machine precision,

```
> .Machine$double.eps
```

```
[1] 2.220446e-16
```






Figure 1.1: The image of the unit circle in  $\mathbb{R}^2$  after mapping by  $U_1'X$  (left panel) and by  $Q_1'X$  (right panel)

A matrix is considered computationally singular when its reciprocal condition number is within some multiple, typical values are 10 or 100, of this number.

Getting back to the question of why we prefer to work with  $X$  directly, instead of forming  $X'X$ , it is because  $\kappa(X'X) = \kappa(X)^2$ . If  $\kappa(X) = 10^6$ , which is large but not catastrophically so, then  $\kappa(X'X)$  will be  $10^{12}$ , which means it is very close to being singular.

Finally, let's revisit the idea of the singular values being the lengths of the principal axes of the image of the unit sphere in the map  $\beta \rightarrow X\beta$ . When  $p = 2$  the unit sphere is the circle of radius 1 centered at the origin and the ellipsoid mentioned above will be an ellipse.

A convenient way of creating a  $2 \times N$  matrix whose columns are the points on the unit circle is to start with a sequence of values from 0 to  $2\pi$  and use its sines and cosines

```
> str(rad <- seq(0, 2*pi, len=201))
num [1:201] 0 0.0314 0.0628 0.0942 0.1257 ...
> str(circ <- rbind(cos(rad), sin(rad)))
num [1:2, 1:201] 1 0 0.9995 0.0314 0.998 ...
```

The  $n$ -dimensional response vectors corresponding to these points on the circle are

```
> fits <- X %% circ
```

To plot this image in two dimensions (Fig. 1.1) we need to represent these points with respect to an orthogonal basis for  $\text{col}(X)$ . Fortunately we have two such bases: the columns of  $Q_1$  and of  $U_1$ . In the  $U_1$  basis the principal axes of the ellipse correspond to the coordinate axes. In the  $Q_1$  basis the principal axes are skewed.

## 1.5 Theoretical results on the eigendecomposition

### 1.5.1 Eigenvalues and Eigenvectors

For any  $k \times k$  matrix  $A$ , the roots of the  $k^{\text{th}}$  degree polynomial equation in  $\lambda$ ,  $|\lambda I_k - A| = 0$ , which we will write as  $\lambda_1, \dots, \lambda_k$  are called the *eigenvalues* of  $A$ . The polynomial is called the

characteristic polynomial of  $\mathbf{A}$ .

Any nonzero  $n \times 1$  vector  $\mathbf{v}_i \neq \mathbf{0}$  such that  $\mathbf{A}\mathbf{v}_i = \lambda_i\mathbf{v}_i$  is an *eigenvector* of  $\mathbf{A}$  corresponding to the eigenvalue  $\lambda_i$ .

For any diagonal matrix  $\mathbf{D} = \text{diag}(d_1, \dots, d_k)$ ,  $|\lambda\mathbf{I}_k - \mathbf{D}| = \prod_{i=1}^k (\lambda - d_i) = 0$  has roots  $d_i$ , therefore the diagonal elements  $d_i$ ,  $i = 1, \dots, n$  are the eigenvalues of  $\mathbf{D}$ .

If  $\mathbf{Q}$  is an orthogonal matrix, then  $\mathbf{Q}\mathbf{A}\mathbf{Q}'$  and  $\mathbf{A}$  have the same eigenvalues.

*Proof.*

$$\begin{aligned}
 |\lambda\mathbf{I} - \mathbf{Q}\mathbf{A}\mathbf{Q}'| &= |\lambda\mathbf{Q}\mathbf{Q}' - \mathbf{Q}\mathbf{A}\mathbf{Q}'| \\
 &= |\mathbf{Q}||\lambda\mathbf{Q}' - \mathbf{A}\mathbf{Q}'| \\
 &= |\mathbf{Q}||\lambda\mathbf{I} - \mathbf{A}||\mathbf{Q}'| \\
 &= |\mathbf{Q}|^2|\lambda\mathbf{I} - \mathbf{A}| = 1|\lambda\mathbf{I} - \mathbf{A}| \\
 &= |\lambda\mathbf{I} - \mathbf{A}|
 \end{aligned} \tag{1.12}$$

□

**Note:** Although the eigenvalues are defined as the roots of the characteristic polynomial, in practice they are not calculated this way. In fact, if you check the documentation for function `solve.polynomial()` in the `polynom` package for R you will find that it uses the numerical methods for evaluating the eigenvalues of the *companion matrix* of a polynomial to solve for the polynomial's roots.

### 1.5.2 Diagonalization of a Symmetric Matrix

For any  $k \times k$  symmetric matrix  $\mathbf{A}$  (i.e.  $\mathbf{A}' = \mathbf{A}$ ), there exists an orthogonal matrix  $\mathbf{Q}$  such that  $\mathbf{Q}\mathbf{A}\mathbf{Q}'$  is a diagonal matrix  $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$  where  $\lambda_i$  are the eigenvalues of  $\mathbf{A}$ . The corresponding eigenvectors of  $\mathbf{A}$  are the column vectors of  $\mathbf{Q}$ .

*Proof.* Let  $\mathbf{e}_i$ ,  $i = 1, \dots, n$  be  $n \times 1$  unit vectors that form the canonical basis of  $\mathbb{R}^n$ , (i.e.  $\mathbf{e}_i = (0, \dots, 1, \dots, 0)'$  where the 1 is in the  $i^{\text{th}}$  position) and  $\mathbf{q}_i$  be the  $i^{\text{th}}$  column of  $\mathbf{Q}$ . That is,  $\mathbf{q}_i = \mathbf{Q}\mathbf{e}_i$ . Then

$$\mathbf{Q}'\mathbf{A}\mathbf{Q} = \mathbf{\Lambda} \Rightarrow \mathbf{Q}'\mathbf{A}\mathbf{Q}\mathbf{e}_i = \mathbf{\Lambda}\mathbf{e}_i = \lambda_i\mathbf{e}_i$$

Multiplying on the left by  $\mathbf{Q}$  produces

$$\mathbf{A}\mathbf{q}_i = \underbrace{\mathbf{Q}\mathbf{Q}'}_{\mathbf{I}}\mathbf{A}\mathbf{Q}\mathbf{e}_i = \lambda_i\mathbf{Q}\mathbf{e}_i = \lambda_i\mathbf{q}_i.$$

□

### 1.5.3 Spectral Decomposition

From the relationship  $\mathbf{Q}'\mathbf{A}\mathbf{Q} = \mathbf{\Lambda}$  just established for a  $k \times k$  symmetric matrix  $\mathbf{A}$  we can compute its spectral decomposition,

$$\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}' = \sum_{i=1}^k \lambda_i \mathbf{q}_i \mathbf{q}_i'$$

where  $\mathbf{q}_i$  is the  $i^{\text{th}}$  column of  $\mathbf{Q}$ .

$$\mathbf{Q}\mathbf{Q}' = \sum_{i=1}^k \mathbf{q}_i \mathbf{q}_i' = \mathbf{I}$$

### 1.5.4 Trace and Determinant of $\mathbf{A}$

The relationship  $\mathbf{Q}'\mathbf{A}\mathbf{Q} = \mathbf{\Lambda}$  for symmetric  $\mathbf{A}$  implies that the trace,  $\text{tr}(\mathbf{A})$ , and the determinant,  $|\mathbf{A}|$ , are the same as those of  $\mathbf{\Lambda}$ .

$$\text{tr}(\mathbf{A}) = \text{tr}(\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}') = \text{tr}(\mathbf{\Lambda}\mathbf{Q}\mathbf{Q}') = \text{tr}(\mathbf{\Lambda}) = \sum_{i=1}^k \lambda_i$$

where we have used the property that  $\text{tr}(\mathbf{C}\mathbf{D}) = \text{tr}(\mathbf{D}\mathbf{C})$  for any conformable matrices  $\mathbf{C}$  and  $\mathbf{D}$  (meaning that if  $\mathbf{C}$  is  $m \times n$  then  $\mathbf{D}$  must be  $n \times m$ ).

$$|\mathbf{A}| = |\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}'| = |\mathbf{Q}||\mathbf{\Lambda}||\mathbf{Q}'| = |\mathbf{Q}|^2|\mathbf{\Lambda}| = \prod_{i=1}^k \lambda_i$$



## Chapter 2

# Quadratic Forms of Random Variables

### 2.1 Quadratic Forms

For a  $k \times k$  symmetric matrix  $\mathbf{A} = \{a_{ij}\}$  the quadratic function of  $k$  variables  $\mathbf{x} = (x_1, \dots, x_n)'$  defined by

$$Q(\mathbf{x}) = \mathbf{x}'\mathbf{A}\mathbf{x} = \sum_{i=1}^k \sum_{j=1}^k a_{i,j}x_i x_j$$

is called the *quadratic form* with matrix  $\mathbf{A}$ .

If  $\mathbf{A}$  is not symmetric, we can have an equivalent expression/quadratic form replacing  $\mathbf{A}$  by  $(\mathbf{A} + \mathbf{A}')/2$ .

**Definition 1.**  $Q(\mathbf{x})$  and the matrix  $\mathbf{A}$  are called positive definite if

$$Q(\mathbf{x}) = \mathbf{x}'\mathbf{A}\mathbf{x} > 0, \quad \forall \mathbf{x} \in \mathbb{R}^k, \mathbf{x} \neq \mathbf{0}$$

and positive semi-definite if

$$Q(\mathbf{x}) \geq \forall \mathbf{x} \in \mathbb{R}^k$$

For negative definite and negative semi-definite, replace the  $>$  and  $\geq$  in the above definitions by  $<$  and  $\leq$ , respectively.

**Theorem 1.** A symmetric matrix  $\mathbf{A}$  is positive definite if and only if it has a Cholesky decomposition  $\mathbf{A} = \mathbf{R}'\mathbf{R}$  with strictly positive diagonal elements in  $\mathbf{R}$ , so that  $\mathbf{R}^{-1}$  exists. (In practice this means that none of the diagonal elements of  $\mathbf{R}$  are very close to zero.)

*Proof.* The “if” part is proven by construction. The Cholesky decomposition,  $\mathbf{R}$ , is constructed a row at a time and the diagonal elements are evaluated as the square roots of expressions calculated from the current row of  $\mathbf{A}$  and previous rows of  $\mathbf{R}$ . If the expression whose square root is to be calculated is not positive then you can determine a non-zero  $\mathbf{x} \in \mathbb{R}^k$  for which  $\mathbf{x}'\mathbf{A}\mathbf{x} \leq 0$ .

Suppose that  $\mathbf{A} = \mathbf{R}'\mathbf{R}$  with  $\mathbf{R}$  invertible. Then

$$\mathbf{x}'\mathbf{A}\mathbf{x} = \mathbf{x}'\mathbf{R}'\mathbf{R}\mathbf{x} = \|\mathbf{R}\mathbf{x}\|^2 \geq 0$$

with equality only if  $\mathbf{R}\mathbf{x} = \mathbf{0}$ . But if  $\mathbf{R}^{-1}$  exists then  $\mathbf{x} = \mathbf{R}^{-1}\mathbf{0}$  must also be zero.  $\square$

**Transformation of Quadratic Forms:**

**Theorem 2.** Suppose that  $\mathbf{B}$  is a  $k \times k$  nonsingular matrix. Then the quadratic form  $Q^*(\mathbf{y}) = \mathbf{y}'\mathbf{B}'\mathbf{A}\mathbf{B}\mathbf{y}$  is positive definite if and only if  $Q(\mathbf{x}) = \mathbf{x}'\mathbf{A}\mathbf{x}$  is positive definite. Similar results hold for positive semi-definite, negative definite and negative semi-definite.

*Proof.*

$$Q^*(\mathbf{y}) = \mathbf{y}'\mathbf{B}'\mathbf{A}\mathbf{B}\mathbf{y} = \mathbf{x}'\mathbf{A}\mathbf{x} > 0$$

where  $\mathbf{x} = \mathbf{B}\mathbf{y} \neq \mathbf{0}$  because  $\mathbf{y} \neq \mathbf{0}$  and  $\mathbf{B}$  is nonsingular.  $\square$

**Theorem 3.** For any  $k \times k$  symmetric matrix  $\mathbf{A}$  the quadratic form defined by  $\mathbf{A}$  can be written using its spectral decomposition as

$$Q(\mathbf{x}) = \mathbf{x}'\mathbf{A}\mathbf{x} = \sum_{i=1}^k \lambda_i \|\mathbf{q}'_i \mathbf{x}\|^2$$

where the eigendecomposition of  $\mathbf{A}$  is  $\mathbf{Q}'\mathbf{\Lambda}\mathbf{Q}$  with  $\mathbf{\Lambda}$  diagonal with diagonal elements  $\lambda_i$ ,  $i = 1, \dots, k$ ,  $\mathbf{Q}$  is the orthogonal matrix with the eigenvectors,  $\mathbf{q}_i$ ,  $i = 1, \dots, k$  as its columns. (Be careful to distinguish the bold face  $\mathbf{Q}$ , which is a matrix, from the unbolded  $Q(\mathbf{x})$ , which is the quadratic form.)

*Proof.* For any  $\mathbf{x} \in \mathbb{R}^k$  let  $\mathbf{y} = \mathbf{Q}'\mathbf{x} = \mathbf{Q}^{-1}\mathbf{x}$ . Then

$$Q(\mathbf{x}) = \text{tr}(\mathbf{x}'\mathbf{A}\mathbf{x}) = \text{tr}(\mathbf{x}'\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}'\mathbf{x}) = \text{tr}(\mathbf{y}'\mathbf{\Lambda}\mathbf{y}) = \text{tr}(\mathbf{\Lambda}\mathbf{y}\mathbf{y}') = \sum_{i=1}^k \lambda_i y_i^2 = \sum_{i=1}^k \lambda_i \|\mathbf{q}'_i \mathbf{x}\|^2$$

This proof uses a common “trick” of expressing the scalar  $Q(\mathbf{x})$  as the trace of a  $1 \times 1$  matrix so we can reverse the order of some matrix multiplications.  $\square$

**Corollary 1.** A symmetric matrix  $\mathbf{A}$  is positive definite if and only if its eigenvalues are all positive, negative definite if and only if its eigenvalues are all negative, and positive semi-definite if all its eigenvalues are non-negative.

**Corollary 2.**  $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{\Lambda})$  hence  $\text{rank}(\mathbf{A})$  equals the number of non-zero eigenvalues of  $\mathbf{A}$

## 2.2 Idempotent Matrices

**Definition 2** (Idempotent). The  $k \times k$  matrix  $\mathbf{A}$ , is idempotent if  $\mathbf{A}^2 = \mathbf{A}\mathbf{A} = \mathbf{A}$ .

**Definition 3** (Projection matrices). A symmetric, idempotent matrix  $\mathbf{A}$  is a projection matrix. The effect of the mapping  $\mathbf{x} \rightarrow \mathbf{A}\mathbf{x}$  is orthogonal projection of  $\mathbf{x}$  onto  $\text{col}(\mathbf{A})$ .

**Theorem 4.** All the eigenvalues of an idempotent matrix are either zero or one.

*Proof.* Suppose that  $\lambda$  is an eigenvalue of the idempotent matrix  $\mathbf{A}$ . Then there exists a non-zero  $\mathbf{x}$  such that  $\mathbf{Ax} = \lambda\mathbf{x}$ . But  $\mathbf{Ax} = \mathbf{AAx}$  because  $\mathbf{A}$  is idempotent. Thus

$$\lambda\mathbf{x} = \mathbf{Ax} = \mathbf{AAx} = \mathbf{A}(\lambda\mathbf{x}) = \lambda(\mathbf{Ax}) = \lambda^2\mathbf{x}$$

and

$$\mathbf{0} = \lambda^2\mathbf{x} - \lambda\mathbf{x} = \lambda(\lambda - 1)\mathbf{x}$$

for some non-zero  $\mathbf{x}$ , which implies that  $\lambda = 0$  or  $\lambda = 1$ .  $\square$

**Corollary 3.** *The  $k \times k$  symmetric matrix  $\mathbf{A}$  is idempotent of rank( $\mathbf{A}$ ) =  $r$  iff  $\mathbf{A}$  has  $r$  eigenvalues equal to 1 and  $k - r$  eigenvalues equal to 0*

*Proof.* A matrix  $\mathbf{A}$  with  $r$  eigenvalues of 1 and  $k - r$  eigenvalues of zero has  $r$  non-zero eigenvalues and hence rank( $\mathbf{A}$ ) =  $r$ . Because  $\mathbf{A}$  is symmetric its eigendecomposition is  $\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}'$  for an orthogonal  $\mathbf{Q}$  and a diagonal  $\mathbf{\Lambda}$ . Because the eigenvalues of  $\mathbf{\Lambda}$  are the same as those of  $\mathbf{A}$ , they must be all zeros or ones. That is all the diagonal elements of  $\mathbf{\Lambda}$  are zero or one. Hence  $\mathbf{\Lambda}$  is idempotent,  $\mathbf{\Lambda}\mathbf{\Lambda} = \mathbf{\Lambda}$ , and

$$\mathbf{AA} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}'\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}' = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}' = \mathbf{A}$$

is also idempotent.  $\square$

**Corollary 4.** *For a symmetric idempotent matrix  $\mathbf{A}$ , we have  $\text{tr}(\mathbf{A}) = \text{rank}(\mathbf{A})$ , which is the dimension of  $\text{col}(\mathbf{A})$ , the space into which  $\mathbf{A}$  projects.*

## 2.3 Expected Values and Covariance Matrices of Random Vectors

An  $k$ -dimensional *vector-valued random variable* (or, more simply, a *random vector*),  $\mathcal{X}$ , is a  $k$ -vector composed of  $k$  scalar random variables

$$\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_k)'$$

If the expected values of the component random variables are  $\mu_i = E(\mathcal{X}_i)$ ,  $i = 1, \dots, k$  then

$$E(\mathcal{X}) = \boldsymbol{\mu}_{\mathcal{X}} = (\mu_1, \dots, \mu_k)'$$

Suppose that  $\mathcal{Y} = (\mathcal{Y}_1, \dots, \mathcal{Y}_m)'$  is an  $m$ -dimensional random vector, then the *covariance* of  $\mathcal{X}$  and  $\mathcal{Y}$ , written  $\text{Cov}(\mathcal{X}, \mathcal{Y})$  is

$$\boldsymbol{\Sigma}_{\mathcal{XY}} = \text{Cov}(\mathcal{X}, \mathcal{Y}) = E[(\mathcal{X} - \boldsymbol{\mu}_{\mathcal{X}})(\mathcal{Y} - \boldsymbol{\mu}_{\mathcal{Y}})']$$

The *variance-covariance* matrix of  $\mathcal{X}$  is

$$\text{Var}(\mathcal{X}) = \boldsymbol{\Sigma}_{\mathcal{XX}} = E[(\mathcal{X} - \boldsymbol{\mu}_{\mathcal{X}})(\mathcal{X} - \boldsymbol{\mu}_{\mathcal{X}})']$$

Suppose that  $\mathbf{c}$  is a constant  $m$ -vector,  $\mathbf{A}$  is a constant  $m \times k$  matrix and  $\mathcal{Z} = \mathbf{AZ} + \mathbf{c}$  is a linear transformation of  $\mathcal{X}$ . Then

$$E(\mathcal{Z}) = \mathbf{AE}(\mathcal{X}) + \mathbf{c}$$

and

$$\text{Var}(\mathcal{Z}) = \mathbf{A} \text{Var}(\mathcal{X}) \mathbf{A}'$$

If we let  $\mathcal{W} = \mathbf{B}\mathcal{Y} + \mathbf{d}$  be a linear transformation of  $\mathcal{Y}$  for suitably sized  $\mathbf{B}$  and  $\mathbf{d}$  then

$$\text{Cov}(\mathcal{Z}, \mathcal{W}) = \mathbf{A} \text{Cov}(\mathcal{X}, \mathcal{Y}) \mathbf{B}'$$

**Theorem 5.** *The variance-covariance matrix  $\Sigma_{\mathcal{X}, \mathcal{X}}$  of  $\mathcal{X}$  is a symmetric and positive semi-definite matrix*

*Proof.* The result follows from the property that the variance of a scalar random variable is non-negative. Suppose that  $\mathbf{b}$  is any nonzero, constant  $k$ -vector. Then

$$0 \leq \text{Var}(\mathbf{b}'\mathcal{X}) = \mathbf{b}'\Sigma_{\mathcal{X}, \mathcal{X}}\mathbf{b}$$

which is the positive, semi-definite condition. □

## 2.4 Mean and Variance of Quadratic Forms

**Theorem 6.** *Let  $\mathcal{X}$  be a  $k$ -dimensional random vector and  $\mathbf{A}$  be a constant  $k \times k$  symmetric matrix. If  $E(\mathcal{X}) = \boldsymbol{\mu}$  and  $\text{Var}(\mathcal{X}) = \Sigma$ , then*

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

*Proof.*

$$\begin{aligned} E(\mathcal{X}'\mathbf{A}\mathcal{X}) &= \text{tr}(E(\mathcal{X}'\mathbf{A}\mathcal{X})) \\ &= E[\text{tr}(\mathcal{X}'\mathbf{A}\mathcal{X})] \\ &= E[\text{tr}(\mathbf{A}\mathcal{X}\mathcal{X}')] \\ &= \text{tr}(\mathbf{A}E[\mathcal{X}\mathcal{X}']) \\ &= \text{tr}(\mathbf{A}(\text{Cov}(\mathcal{X}) + \boldsymbol{\mu}\boldsymbol{\mu}')) \\ &= \text{tr}(\mathbf{A}\Sigma_{\mathcal{X}, \mathcal{X}}) + \text{tr}(\mathbf{A}\boldsymbol{\mu}\boldsymbol{\mu}') \\ &= \text{tr}(\mathbf{A}\Sigma_{\mathcal{X}, \mathcal{X}}) + \boldsymbol{\mu}'\mathbf{A}\boldsymbol{\mu} \end{aligned}$$

□

## 2.5 Distribution of Quadratic Forms in Normal Random Variables

**Definition 4** (Non-Central  $\chi^2$ ). *If  $\mathcal{X}$  is a (scalar) normal random variable with  $E(\mathcal{X}) = \mu$  and  $\text{Var}(\mathcal{X}) = 1$ , then the random variable  $\mathcal{V} = \mathcal{X}^2$  is distributed as  $\chi_1^2(\lambda^2)$ , which is called the noncentral  $\chi^2$  distribution with 1 degree of freedom and non-centrality parameter  $\lambda^2 = \mu^2$ . The mean and variance of  $\mathcal{V}$  are*

$$E[\mathcal{V}] = 1 + \lambda^2 \text{ and } \text{Var}[\mathcal{V}] = 2 + 4\lambda^2$$



As described in the previous chapter, we are particularly interested in random  $n$ -vectors,  $\mathbf{Y}$ , that have a *spherical normal distribution*.

**Theorem 7.** Let  $\mathcal{Y} \sim \mathcal{N}(\boldsymbol{\mu}, \sigma^2 \mathbf{I}_n)$  be an  $n$ -vector with a spherical normal distribution and  $\mathbf{A}$  be an  $n \times n$  symmetric matrix. Then the ratio  $\mathcal{Y}'\mathbf{A}\mathcal{Y}/\sigma^2$  will have a  $\chi_r^2(\lambda^2)$  distribution with  $\lambda^2 = \boldsymbol{\mu}'\mathbf{A}\boldsymbol{\mu}/\sigma^2$  if and only if  $\mathbf{A}$  is idempotent with  $\text{rank}(\mathbf{A}) = r$

*Proof.* Suppose that  $\mathbf{A}$  is idempotent (which, in combination with being symmetric, means that it is a projection matrix) and has  $\text{rank}(\mathbf{A}) = r$ . Its eigendecomposition,  $\mathbf{A} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}'$ , is such that  $\mathbf{V}$  is orthogonal and  $\boldsymbol{\Lambda}$  is  $n \times n$  diagonal with exactly  $r = \text{rank}(\mathbf{A})$  ones and  $n - r$  zeros on the diagonal. Without loss of generality we can (and do) arrange the eigenvalues in decreasing order so that  $\lambda_j = 1, j = 1, \dots, r$  and  $\lambda_j = 0, j = r + 1, \dots, n$ . Let  $\mathcal{X} = \mathbf{V}'\mathcal{Y}$

$$\begin{aligned} \frac{\mathcal{Y}'\mathbf{A}\mathcal{Y}}{\sigma^2} &= \frac{\mathcal{Y}'\mathbf{V}\boldsymbol{\Lambda}\mathbf{V}'\mathcal{Y}}{\sigma^2} \\ &= \frac{\mathcal{X}'\boldsymbol{\Lambda}\mathcal{X}}{\sigma^2} \\ &= \sum_{j=1}^n \lambda_j \frac{\mathcal{X}_j^2}{\sigma^2} \\ &= \sum_{j=1}^r \frac{\mathcal{X}_j^2}{\sigma^2} \end{aligned}$$

(Notice that the last sum is to  $j = r$ , not  $j = n$ .) However,  $\frac{\mathcal{X}_j}{\sigma} \sim \mathcal{N}(\mathbf{v}_j'\boldsymbol{\mu}/\sigma, 1)$  so  $\frac{\mathcal{X}_j^2}{\sigma^2} \sim \chi_1^2((\mathbf{v}_j'\boldsymbol{\mu}/\sigma)^2)$ . Therefore

$$\sum_{j=1}^r \frac{\mathcal{X}_j^2}{\sigma^2} \sim \chi_{(r)}^2(\lambda^2) \text{ where } \lambda^2 = \frac{\boldsymbol{\mu}'\mathbf{V}\boldsymbol{\Lambda}\mathbf{V}'\boldsymbol{\mu}}{\sigma^2} = \frac{\boldsymbol{\mu}'\mathbf{A}\boldsymbol{\mu}}{\sigma^2}$$

□

**Corollary 5.** For  $\mathbf{A}$  a projection of rank  $r$ ,  $(\mathcal{Y}'\mathbf{A}\mathcal{Y})/\sigma^2$  has a central  $\chi_r^2$  distribution if and only if  $\mathbf{A}\boldsymbol{\mu} = \mathbf{0}$

*Proof.* The  $\chi_r^2$  distribution will be central if and only if

$$0 = \boldsymbol{\mu}'\mathbf{A}\boldsymbol{\mu} = \boldsymbol{\mu}'\mathbf{A}\mathbf{A}\boldsymbol{\mu} = \boldsymbol{\mu}'\mathbf{A}'\mathbf{A}\boldsymbol{\mu} = \|\mathbf{A}\boldsymbol{\mu}\|^2$$

□

**Corollary 6.** In the full-rank Gaussian linear model,  $\mathbf{Y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I}_n)$ , the residual sum of squares,  $\|\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2$  has a central  $\sigma^2 \chi_{n-r}^2$  distribution.

*Proof.* In the full rank model with the QR decomposition of  $\mathbf{X}$  given by

$$\mathbf{X} = [\mathbf{Q}_1 \quad \mathbf{Q}_2] \begin{bmatrix} \mathbf{R} \\ \mathbf{0} \end{bmatrix}$$

and  $\mathbf{R}$  invertible, the fitted values are  $\mathbf{Q}_1\mathbf{Q}'_1\mathcal{Y}$  and the residuals are  $\mathbf{Q}_2\mathbf{Q}'_2\mathbf{y}$  so the residual sum of squares is the quadratic form  $\mathcal{Y}'\mathbf{Q}_2\mathbf{Q}'_2\mathcal{Y}$ . The matrix defining the quadratic form,  $\mathbf{Q}_2\mathbf{Q}'_2$ , is a projection matrix. It is obviously symmetric and it is idempotent because  $\mathbf{Q}_2\mathbf{Q}'_2\mathbf{Q}_2\mathbf{Q}'_2 = \mathbf{Q}_2\mathbf{Q}'_2$ . As

$$\mathbf{Q}'_2\boldsymbol{\mu} = \mathbf{Q}'_2\mathbf{X}\boldsymbol{\beta}_0 = \mathbf{Q}'_2\mathbf{Q}_1\mathbf{R}\boldsymbol{\beta}_0 = \underbrace{\mathbf{0}}_{(n-p)\times n} \mathbf{R}\boldsymbol{\beta}_0 = \underbrace{\mathbf{0}}_{(n-p)\times p} \boldsymbol{\beta}_0 = \mathbf{0}_{n-p}$$

the ratio

$$\frac{\mathcal{Y}'\mathbf{Q}_2\mathbf{Q}'_2\mathcal{Y}}{\sigma^2} \sim \chi^2_{n-p}$$

and the RSS has a central  $\sigma^2\chi^2_{n-p}$  distribution. □

**R Exercises:** Let's check some of these results by simulation. First we claim that if  $\mathcal{X} \sim \mathcal{N}(\mu, 1)$  then  $\mathcal{X}^2 \sim \chi^2(\lambda^2)$  where  $\lambda^2 = \mu^2$ . First simulate from a standard normal distribution

```
> set.seed(1234) # reproducible "random" values
> X <- rnorm(100000) # standard normal values
> zapsmall(summary(V <- X^2)) # a very skew distribution
```

```
   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.0000  0.1026  0.4521  0.9989  1.3190 20.3300
```

```
> var(V)
```

```
[1] 1.992403
```

The mean and variance of the simulated values agree quite well with the theoretical values of 1 and 2, respectively.

To check the form of the distribution we could plot an empirical density function but this distribution has its maximum density at 0 and is zero to the left of 0 so an empirical density is a poor indication of the actual shape of the density. Instead, in Fig. 2.1, we present the quantile-quantile plot for this sample versus the (theoretical) quantiles of the  $\chi^2_1$  distribution.

Now simulate a non-central  $\chi^2$  with non-centrality parameter  $\lambda^2 = 4$

```
> V1 <- rnorm(100000, mean=2)^2
> zapsmall(summary(V1))
```

```
   Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
0.000  1.773  3.994  5.003  7.144 39.050
```

```
> var(V1)
```




Figure 2.1: A quantile-quantile plot of the squares of simulated  $\mathcal{N}(0, 1)$  random variables versus the quantiles of the  $\chi_1^2$  distribution. The dashed line is a reference line through the origin with a slope of 1.

```
[1] 17.95924
```

The sample mean is close to the theoretical value of  $5 = 1 + \lambda^2$  and the sample variance is close to the theoretical value of  $2 + 4\lambda^2$  although perhaps not as close as one would hope in a sample of size 100,000.

A quantile-quantile plot versus the non-central distribution,  $\chi_1^2(4)$ , (Fig. 2.2) and versus the central distribution,  $\chi_1^2$ , shows that the sample does follow the claimed distribution  $\chi_1^2(4)$  and is stochastically larger than the  $\chi_1^2$  distribution.

More interesting, perhaps is the distribution of the residual sum of squares from a regression model. We simulate from our previously fitted model `lm1`

```
> lm1 <- lm(optden ~ carb, Formaldehyde)
> str(Ymat <- data.matrix(unname(simulate(lm1, 10000))))

num [1:6, 1:10000] 0.088 0.258 0.444 0.521 0.619 ...
- attr(*, "dimnames")=List of 2
..$ : chr [1:6] "1" "2" "3" "4" ...
..$ : NULL
```




Figure 2.2: Quantile-quantile plots of a sample of squares of  $\mathcal{N}(2, 1)$  random variables versus the quantiles of a  $\chi_1^2(4)$  non-central distribution (left panel) and a  $\chi_1^2$  central distribution (right panel)

```
> str(RSS <- deviance(fits <- lm(Ymat ~ carb, Formaldehyde)))
num [1:10000] 0.000104 0.000547 0.00055 0.000429 0.000228 ...
> fits[["df.residual"]]
[1] 4
```

Here the `Ymat` matrix is 10,000 simulated response vectors from model `lm1` using the estimated parameters as the true values of  $\beta$  and  $\sigma^2$ . Notice that we can fit the model to **all** 10,000 response vectors in a single call to the `lm()` function.

The `deviance()` function applied to a model fit by `lm()` returns the residual sum of square, which is not technically the deviance but is often the quantity of interest.

These simulated residual sums of squares should have a  $\sigma^2\chi_4^2$  distribution where  $\sigma^2$  is the residual sum of squares in model `lm1` divided by 4.

```
> (sigsq <- deviance(lm1)/4)
[1] 7.48e-05
> summary(RSS)
      Min.   1st Qu.   Median     Mean   3rd Qu.    Max.
6.997e-07 1.461e-04 2.537e-04 3.026e-04 4.114e-04 1.675e-03
```




Figure 2.3: Quantile-quantile plot of the scaled residual sum of squares,  $RSS_{sq}$ , from simulated responses versus the quantiles of a  $\chi_4^2$  distribution (left panel) and the corresponding probability-probability plot on the right panel.

We expect a mean of  $4\sigma^2$  and a variance of  $2 \cdot 4(\sigma^2)^2$ . It is easier to see this if we divide these values by  $\sigma^2$

```
> summary(RSSsc <- RSS/sigsq)
```

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.009354	1.953000	3.392000	4.045000	5.500000	22.390000

```
> var(RSSsc)
```

```
[1] 8.000057
```

A quantile-quantile plot with respect to the  $\chi_4^2$  distribution (Fig. 2.3) shows very good agreement between the empirical and theoretical quantiles. Also shown in Fig. 2.3 is the probability-probability plot. Instead of plotting the sample quantiles versus the theoretical quantiles we take equally spaced values on the probability scale (function `ppoints()`), evaluate the sample quantiles and then apply the theoretical cdf to the empirical quantiles. This should also produce a straight line. It has the advantage that the points are equally spaced on the x-axis.

We could also plot the empirical density of these simulated values and overlay it with the theoretical density (Fig. 2.4).

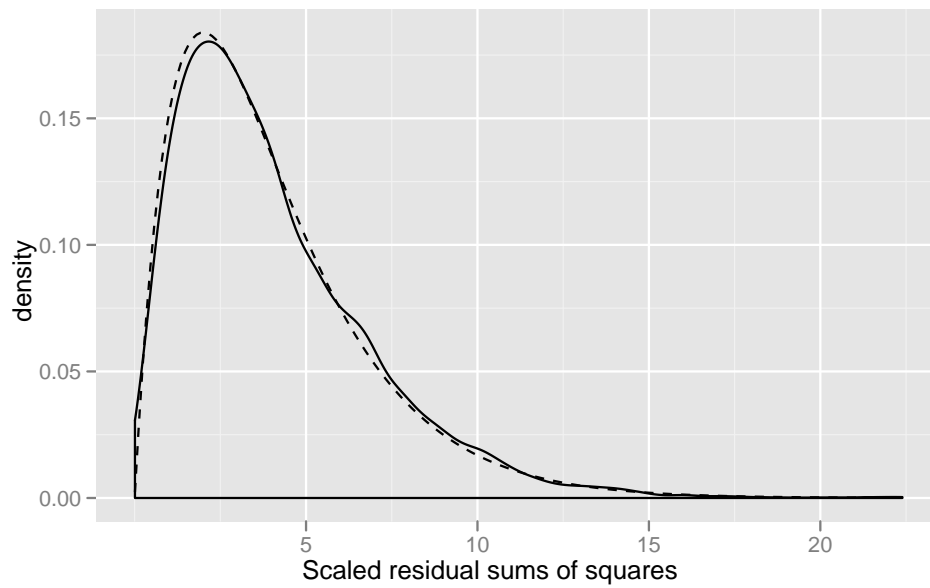


Figure 2.4: Empirical density plot of the scaled residual sums of squares,  $\text{RSSsq}$ , from simulated responses. The overlaid dashed line is the density of a  $\chi_4^2$  random variable. The peak of the empirical density gets shifted a bit to the right because of the way the empirical density is calculated. It uses a symmetric kernel which is not a good choice for a skewed density like this.

## Chapter 3

# Properties of coefficient estimates

In chapter 1 we described properties of the coefficient estimates,  $\hat{\boldsymbol{\beta}}$ , in the Gaussian linear model

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

The estimates are called the *least squares estimates* because they minimize the sum of squared residuals from the observed responses,  $\mathbf{y}$ . That is,

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} S(\boldsymbol{\beta}) = \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$$

### 3.1 Geometric Properties

Recall that  $\text{col}(\mathbf{X})$ , the *column span* of the  $n \times p$  model matrix  $\mathbf{X}$  is a linear subspace of the response space,  $\mathbb{R}^n$ ,

$$\text{col}(\mathbf{X}) = \{\mathbf{X}\boldsymbol{\beta} : \boldsymbol{\beta} \in \mathbb{R}^p\}$$

The dimension of  $\text{col}(\mathbf{X})$  is  $k = \text{rank}(\mathbf{X})$  and the QR decomposition used in R uses column pivoting to ensure that the first  $k$  columns of  $\mathbf{Q}$  are an orthonormal basis for  $\text{col}(\mathbf{X})$ .

At the risk of some confusion, we will refer to these  $k$  columns as  $\mathbf{Q}_1$  which is equivalent to our previous definition in the most common case of full column rank for  $\mathbf{X}$ .

The fitted values,  $\hat{\mathbf{y}}$ , are the (orthogonal) projection of  $\mathbf{y}$  onto  $\text{col}(\mathbf{X})$

$$\hat{\mathbf{y}} = \mathbf{H}\mathbf{y} = \mathbf{Q}_1\mathbf{Q}'_1\mathbf{y}$$

where the “hat matrix”,  $\mathbf{H} = \mathbf{Q}_1\mathbf{Q}'_1$ , is a projection matrix of rank

$$\text{tr}(\mathbf{Q}_1\mathbf{Q}'_1) = \text{tr}(\mathbf{Q}'_1\mathbf{Q}_1) = \text{tr}(\mathbf{I}_k) = k$$

The diagonal matrix,  $\mathbf{D}$ , in the singular value decomposition (sect. 1.4.3, p. 17),  $\mathbf{X} = \mathbf{U}_1\mathbf{D}\mathbf{V}'$ , has exactly  $p - k$  values that are (effectively) zero and these will be in the last  $p - k$  positions. (Recall that the singular values, which must be non-negative, are in decreasing order.) Thus the first  $k$  columns of  $\mathbf{U}_1$  also form an orthonormal basis for  $\text{col}(\mathbf{X})$ .

The residual at the parameter estimates,  $\hat{\mathbf{e}} = \mathbf{y} - \hat{\mathbf{y}}$  is orthogonal to  $\text{col}(\mathbf{X})$ . We can prove this by showing that  $\hat{\mathbf{e}}$  is orthogonal to the  $k$  columns of  $\mathbf{Q}_1$  which form a basis for  $\text{col}(\mathbf{X})$ .

$$\mathbf{Q}'_1 \hat{\mathbf{e}} = \mathbf{Q}'_1 (\mathbf{y} - \hat{\mathbf{y}}) = \mathbf{Q}'_1 (\mathbf{I}_n - \mathbf{Q}_1 \mathbf{Q}'_1) \mathbf{y} = \left( \mathbf{Q}'_1 - \underbrace{\mathbf{Q}'_1 \mathbf{Q}_1}_{\mathbf{I}_p} \mathbf{Q}'_1 \right) \mathbf{y} = \mathbf{0}$$

This is also an obvious geometric property that to minimize the distance between a point on a hyperplane and a general point in the response space,  $\arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$ , you use orthogonal projection of  $\mathbf{y}$  onto  $\text{col}(\mathbf{X})$  which implies that the residual is orthogonal to  $\text{col}(\mathbf{X})$ .

Often this relationship is characterized as the *normal equations*. The residual will be orthogonal to  $\text{col}(\mathbf{X})$  if it is orthogonal to all the columns of  $\mathbf{X}$ , which is to say

$$\mathbf{X}' (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}) = \mathbf{0} \Rightarrow (\mathbf{X}'\mathbf{X})\hat{\boldsymbol{\beta}} = \mathbf{X}'\mathbf{y}$$

## 3.2 Calculus Approach

The function

$$\begin{aligned} S(\boldsymbol{\beta}) &= \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 \\ &= (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) \\ &= \mathbf{y}'\mathbf{y} - \mathbf{y}'\mathbf{X}\boldsymbol{\beta} - \boldsymbol{\beta}'\mathbf{X}'\mathbf{y} + \boldsymbol{\beta}'\mathbf{X}'\mathbf{X}\boldsymbol{\beta} \\ &= \mathbf{y}'\mathbf{y} - 2\boldsymbol{\beta}'\mathbf{X}'\mathbf{y} + \boldsymbol{\beta}'\mathbf{X}'\mathbf{X}\boldsymbol{\beta} \end{aligned}$$

is a real-valued function of the  $p$ -vector,  $\boldsymbol{\beta}$ , ( $S : \mathbb{R}^p \rightarrow \mathbb{R}$ ), with gradient vector

$$\frac{dS}{d\boldsymbol{\beta}} = -2\mathbf{X}'\mathbf{y} + 2\mathbf{X}'\mathbf{X}\boldsymbol{\beta}.$$

Thus a critical point,  $\boldsymbol{\beta}_c$ , at which the gradient is zero, satisfies

$$\mathbf{X}'\mathbf{X}\boldsymbol{\beta}_c = \mathbf{X}'\mathbf{y}.$$

The Hessian matrix of  $S(\boldsymbol{\beta})$ ,

$$\frac{d^2 S}{d\boldsymbol{\beta} d\boldsymbol{\beta}'} = 2\mathbf{X}'\mathbf{X}$$

is positive semi-definite. If  $\mathbf{X}$  is full rank then  $\mathbf{X}'\mathbf{X}$  is positive definite and the critical point will be the minimizer of  $S(\boldsymbol{\beta})$ .

## 3.3 Algebraic Properties of $\hat{\boldsymbol{\beta}}$

1.  $\hat{\boldsymbol{\beta}}$  satisfies  $\mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}'\mathbf{y}$  and minimizes  $S(\boldsymbol{\beta}) = \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$
2. If  $\mathbf{X}$  is of rank  $p$ , then  $\hat{\boldsymbol{\beta}}$  is unique, satisfying  $\mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}'\mathbf{y}$  or, equivalently,  $\mathbf{R}\hat{\boldsymbol{\beta}} = \mathbf{Q}'_1\mathbf{y}$  for an invertible, upper-triangular  $p \times p$  matrix  $\mathbf{R}$ .



3. If  $\text{rank}(\mathbf{X}) < p$ ,  $\mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}'\mathbf{y}$  has multiple solutions for  $\hat{\boldsymbol{\beta}}$  but  $\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}}$  is the same for all such  $\hat{\boldsymbol{\beta}}$

*Proof.* To prove item 3: Suppose that  $\hat{\boldsymbol{\beta}}_1$  and  $\hat{\boldsymbol{\beta}}_2$  are such that  $\mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}}_1 = \mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}}_2 = \mathbf{X}'\mathbf{y}$ . Then

$$\mathbf{X}'(\mathbf{X}\hat{\boldsymbol{\beta}}_1 - \mathbf{X}\hat{\boldsymbol{\beta}}_2) = \mathbf{0}$$

which implies that

$$0 = (\hat{\boldsymbol{\beta}}_1 - \hat{\boldsymbol{\beta}}_2)\mathbf{X}'\mathbf{X}(\hat{\boldsymbol{\beta}}_1 - \hat{\boldsymbol{\beta}}_2) = \|\mathbf{X}\hat{\boldsymbol{\beta}}_1 - \mathbf{X}\hat{\boldsymbol{\beta}}_2\|^2 \Rightarrow \mathbf{X}\hat{\boldsymbol{\beta}}_1 = \mathbf{X}\hat{\boldsymbol{\beta}}_2$$

□

### 3.4 Rank deficient cases and the Moore-Penrose inverse

In practice a rank-deficient model matrix,  $\mathbf{X}$ , is handled by two methods

1. Don't create it in the first place, use  $I - 1$  *contrasts* for a factor with  $I$  levels.
2. Use the pivoted QR decomposition that retains the original order of the columns except that columns whose diagonal elements in  $\mathbf{R}$  would be effectively zero are moved to trailing positions.

After that the calculation proceeds as in the full-rank case except that only the first  $k = \text{rank}(\mathbf{X})$  columns are used in  $\mathbf{Q}_1$  and  $\mathbf{R}$  is taken as the  $k \times k$  upper-left submatrix of the calculated  $p \times p$   $\mathbf{R}$ .

When discussing the singular value decomposition in Chap. 1, we mentioned the pseudo-inverse or *generalized inverse* of  $\mathbf{X}$ , written  $\mathbf{X}^-$ , which formally is called the *Moore-Penrose generalized inverse*. If  $\text{rank}(\mathbf{X}) = k < p$  then there are  $p - k$  singular values of zero (in practice, very close to zero). The SVD is

$$\mathbf{X} = \mathbf{U}_1\mathbf{D}\mathbf{V}' = \tilde{\mathbf{U}}\tilde{\mathbf{D}}\mathbf{V}'$$

where  $\tilde{\mathbf{U}}$  is the first  $k$  columns of  $\mathbf{U}_1$  and  $\tilde{\mathbf{D}}$  is the first  $k$  rows of  $\mathbf{D}$ .  $\mathbf{D}^-$ , the Moore-Penrose inverse of  $\mathbf{D}$ , is also a diagonal matrix with diagonal elements  $1/d_{i,i}$ ,  $i = 1, \dots, k$  and zero thereafter. The Moore-Penrose inverse of  $\tilde{\mathbf{D}}$  is the first  $k$  columns of  $\mathbf{D}^-$ . Finally, the Moore-Penrose generalized inverse of  $\mathbf{X}$  is

$$\mathbf{X}^- = \mathbf{V}\mathbf{D}^-\mathbf{U}_1' = \mathbf{V}\tilde{\mathbf{D}}^-\tilde{\mathbf{U}}'$$

This is an interesting theoretical tool but in practice it is not necessary to form the SVD in order to solve rank-deficient least squares problems.

#### 3.4.1 Properties of Generalized Inverses

Let  $\mathbf{A}$  be an  $n \times p$  matrix and  $\mathbf{A}^-$  be its  $p \times n$  pseudo-inverse. The conditions that  $\mathbf{A}$  and  $\mathbf{A}^-$  must satisfy are

1.  $\mathbf{A}\mathbf{A}^-\mathbf{A} = \mathbf{A}$  (i.e.  $\mathbf{A}\mathbf{A}^-$  maps the columns of  $\mathbf{A}$  to themselves.)

2.  $\mathbf{A}^{-}\mathbf{A}\mathbf{A}^{-} = \mathbf{A}^{-}$  (i.e.  $\mathbf{A}^{-}\mathbf{A}$  maps the columns of  $\mathbf{A}^{-}$  to themselves.)
3. Both  $\mathbf{A}\mathbf{A}^{-}$  and  $\mathbf{A}^{-}\mathbf{A}$  are symmetric

(It is easy to verify these conditions for our case of  $\mathbf{X}^{-} = \mathbf{V}\mathbf{D}^{-}\mathbf{U}$  where  $\mathbf{X}$  is  $n \times p$  with  $\text{rank}(\mathbf{X}) \leq p \leq n$ . In fact, you will do so on a homework assignment.)

Let  $\mathbf{H} = \mathbf{A}^{-}\mathbf{A}$  be the associated projection in  $\mathbb{R}^p$ . Then the condition  $\mathbf{A}\mathbf{A}^{-}\mathbf{A} = \mathbf{A}$  implies

1.  $\mathbf{H}$  is idempotent because  $\mathbf{H}\mathbf{H} = \mathbf{A}^{-}\mathbf{A}\mathbf{A}^{-}\mathbf{A} = \mathbf{A}^{-}\mathbf{A} = \mathbf{H}$ .
2.  $\mathbf{A}\mathbf{H} = \mathbf{A}$  (just plug in the definition of  $\mathbf{H}$ ) so  $\text{rank}(\mathbf{A}) \leq \text{rank}(\mathbf{H})$ . However, we also have  $\text{rank}(\mathbf{H}) \leq \text{rank}(\mathbf{A})$  because  $\mathbf{H} = \mathbf{A}^{-}\mathbf{A}$ . Thus  $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{H}) = \text{tr}(\mathbf{H})$
3. A general solution of  $\mathbf{A}\mathbf{x} = \mathbf{0}$  is

$$\mathbf{x} = (\mathbf{H} - \mathbf{I}_p)\mathbf{z}$$

where  $\mathbf{z}$  is any vector in  $\mathbb{R}^p$

$$\begin{aligned} \mathbf{A}\mathbf{x} &= \mathbf{A}(\mathbf{H} - \mathbf{I}_p)\mathbf{z} \\ &= (\mathbf{A}\mathbf{H} - \mathbf{A})\mathbf{z} \\ &= (\mathbf{A} - \mathbf{A})\mathbf{z} = \mathbf{0} \end{aligned}$$

4. A general solution to  $\mathbf{A}\mathbf{x} = \mathbf{y}$  is

$$\mathbf{x} = \mathbf{A}^{-}\mathbf{y} + (\mathbf{H} - \mathbf{I}_p)\mathbf{z}$$

$$\mathbf{A}\mathbf{A}^{-}\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x} \Rightarrow \mathbf{A}\mathbf{A}^{-}\mathbf{y} = \mathbf{y}$$

$$\mathbf{x} = \mathbf{A}\mathbf{A}^{-}\mathbf{y} + \underbrace{\mathbf{A}(\mathbf{H} - \mathbf{I}_p)\mathbf{z}}_{\mathbf{0}} \text{ and } \mathbf{A}\mathbf{A}^{-}\mathbf{y} = \mathbf{y} \Rightarrow \mathbf{y} = \mathbf{A}\mathbf{x}$$

$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-}\mathbf{X}'\mathbf{y}$  is a particular least squares solution in the rank deficient case. The general solution is

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-}\mathbf{X}'\mathbf{y} + (\mathbf{H} - \mathbf{I}_p)\mathbf{z}, \quad \mathbf{z} \in \mathbb{R}^p$$

where  $\mathbf{H} = (\mathbf{X}'\mathbf{X})^{-}(\mathbf{X}'\mathbf{X})$ .

### 3.5 Properties of $\hat{\boldsymbol{\beta}}$

In the full-rank Gaussian linear model

$$\begin{aligned} E[\hat{\boldsymbol{\beta}}] &= E[(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathcal{Y}] \\ &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'E[\mathcal{Y}] \\ &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}\boldsymbol{\beta} \\ &= \boldsymbol{\beta} \end{aligned}$$

which is to say that the least squares estimator is an *unbiased estimator* of  $\beta$ . Furthermore

$$\begin{aligned}\text{Var}(\hat{\beta}) &= \text{Var}((\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathcal{Y}) \\ &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\text{Var}(\mathcal{Y})\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} \\ &\quad \sigma^2(\mathbf{X}'\mathbf{X})^{-1}\end{aligned}$$

**R Exercises:** Consider the models fit in Chap. 1

```
> lm1 <- lm(optden ~ 1 + carb, Formaldehyde)
> set.seed(1234) # allow for reproducible "random" numbers
> badDat <- within(data.frame(x1=1:20, x2=rnorm(20,mean=6,sd=0.2),
+ x4=rexp(20,rate=0.02),
+ y=runif(20,min=18,max=24)),
+ x3 <- x1 + 2*x2) # create linear combination
> lm2 <- lm(y ~ x1 + x2 + x3 + x4, badDat)
> lm3 <- lm(count ~ spray, InsectSprays)
> lmlst <- list(lm1=lm1, lm2=lm2, lm3=lm3)
> mmlst <- lapply(lmlst, model.matrix)
```

We know that models `lm1` and `lm3` are full-rank but model `lm2` is rank-deficient.

```
> sapply(lmlst, function(fm) c(rank=fm[["rank"]], p=length(coef(fm))))
```

```
      lm1 lm2 lm3
rank   2   4   6
p       2   5   6
```

which is reflected in the diagonal elements of the  $\mathbf{R}$  matrices and in the singular values of the model matrices and in the condition number

```
> lapply(lmlst, function(fm) diag(fm[["qr"]][["qr"]]))
```

```
$lm1
[1] -2.4494897  0.6390097
$lm2
[1] -4.472136e+00  2.578759e+01 -8.668932e-01  2.327514e+02  5.179752e-15
```

```
$lm3
[1] -8.485281  3.162278  3.098387  3.000000  2.828427  2.449490
```

```
> lapply(mmlst, function(mm) svd(mm, nu=0, nv=0)[["d"]])
```

```
$lm1
[1] 2.773349 0.564389
$lm2
[1] 3.197628e+02 8.828492e+01 1.396147e+01 1.446151e-01 3.340320e-15
```

```
$lm3
[1] 9.069136 3.464102 3.464102 3.464102 3.464102 1.323169
```

```
> sapply(lmlst, kappa, exact=TRUE)

      lm1      lm2      lm3
4.913897e+00 1.512149e+17 6.854102e+00
```

For the full-rank models, `lm1` and `lm3`, the pseudo-inverse,  $\mathbf{X}^-$  is simply the matrix that creates the estimated coefficients,  $\hat{\boldsymbol{\beta}}$  from the observed response vector  $\mathbf{y}$ . We can write it in various forms as

$$\mathbf{X}^- = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' = \mathbf{R}^{-1}\mathbf{Q}'_1 = \mathbf{V}\mathbf{D}^{-1}\mathbf{U}'_1$$

For full-rank models like these the pseudo-inverse,  $\mathbf{X}^-$  is unique and

$$\mathbf{X}^-\mathbf{X} = \mathbf{R}^{-1}\underbrace{\mathbf{Q}'_1\mathbf{Q}_1}_{\mathbf{I}_p}\mathbf{R} = \mathbf{I}_p$$

We can verify the conditions for the Moore-Penrose inverse symbolically. For example,

$$\mathbf{X}\mathbf{X}^-\mathbf{X} = \mathbf{Q}_1\underbrace{\mathbf{R}\mathbf{R}^{-1}}_{\mathbf{I}_p}\underbrace{\mathbf{Q}'_1\mathbf{Q}_1}_{\mathbf{I}_p}\mathbf{R} = \mathbf{Q}_1\mathbf{R} = \mathbf{X}$$

or numerically

```
> X <- mmlst[[3]]
> lm3qr <- lmlst[[3]]$qr
> Q1 <- qr.Q(lm3qr)
> R <- qr.R(lm3qr)
> Xpinv <- backsolve(R, t(Q1))
> zapsmall(Xpinv %*% X)

      (Intercept) sprayB sprayC sprayD sprayE sprayF
[1,]           1         0         0         0         0         0
[2,]           0         1         0         0         0         0
[3,]           0         0         1         0         0         0
[4,]           0         0         0         1         0         0
[5,]           0         0         0         0         1         0
[6,]           0         0         0         0         0         1

> all.equal(X %*% Xpinv %*% X, X, check.attr=FALSE)

[1] TRUE

> all.equal(Xpinv %*% X %*% Xpinv, Xpinv, check.attr=FALSE)

[1] TRUE
```

For the rank-deficient model, `lm2`, there are many pseudo-inverses.

```

> X <- mmlst[[2]]
> lm2qr <- lmlst[[2]]$qr
> SVD <- svd(X)
> (rr <- lm2qr$rank)                # rank

[1] 4

> (rrind <- seq_len(rr))           # safer than 1:rr

[1] 1 2 3 4

> (dpinv <- c(1/SVD$d[rrind], rep(0, ncol(X) - rr)))

[1] 0.003127319 0.011326963 0.071625693 6.914905158 0.000000000

> str(Xpinv1 <- SVD$v %*% (dpinv * t(SVD$u)))

num [1:5, 1:20] 1.245927 0.048805 -0.055057 -0.061309 -0.000183 ...

> zapsmall(Xpinv1 %*% X)

      (Intercept)      x1      x2      x3 x4
[1,]          1 0.0000000 0.0000000 0.0000000 0
[2,]          0 0.8333333 -0.3333333 0.1666667 0
[3,]          0 -0.3333333 0.3333333 0.3333333 0
[4,]          0 0.1666667 0.3333333 0.8333333 0
[5,]          0 0.0000000 0.0000000 0.0000000 1

> all.equal(X %*% Xpinv1 %*% X, X, check.attr=FALSE)

[1] TRUE

> all.equal(Xpinv1 %*% X %*% Xpinv1, Xpinv1, check.attr=FALSE)

[1] TRUE

> ## An alternative construction is to reduce the SVD components to the first 4 columns
> str(SVDred <- list(d=SVD$d[rrind], u=SVD$u[,rrind], v=SVD$v[,rrind]))

List of 3
 $ d: num [1:4] 319.763 88.285 13.961 0.145
 $ u: num [1:20, 1:4] 0.0262 0.0559 0.155 0.0498 0.1678 ...
 $ v: num [1:5, 1:4] 0.00901 0.11732 0.05349 0.2243 0.96591 ...

> str(Xpinv2 <- with(SVDred, v %*% (1/d * t(u))))

num [1:5, 1:20] 1.245927 0.048805 -0.055057 -0.061309 -0.000183 ...

> all.equal(Xpinv2, Xpinv1)

```

```
[1] TRUE

> ## Finally, we can use a similar construction on the QR decomposition
> ## taking into account the rearrangement of the columns of X
> Xpiv <- X[, lm2qr$pivot]
> str(Xpinv3 <- rbind(backsolve(qr.R(lm2qr)[rrind, rrind], t(qr.Q(lm2qr)[, rrind])), 0))

  num [1:5, 1:20] 1.245927 -0.012504 -0.177675 -0.000183 0 ...

> all.equal(Xpiv %*% Xpinv3 %*% Xpiv, Xpiv, check.attr=FALSE)

[1] TRUE

> all.equal(Xpinv3 %*% Xpiv %*% Xpinv3, Xpinv3, check.attr=FALSE)

[1] TRUE
```

The last two constructions show that the Moore-Penrose pseudo-inverse is a matter of collecting the independent columns at the left hand side of the matrix and the linearly-dependent columns on the right hand side, then truncating the decomposition. In other words, if  $\mathbf{X}$  is less than full rank then you just find a set of full-rank columns and proceed as before.

**R Exercise: (Simulating linear model fits)** The `simulate` functions allow us to simulate a matrix of responses based on a fitted model, then fit all the simulated responses in a single call to `lm`. This is much, much faster than any loop-based approach would be.

The result of `simulate` is a named list of response vectors so we drop the names and convert the list to a matrix.

```
> str(Ymat <- data.matrix(unname(simulate(lm1, 10000))))

  num [1:6, 1:10000] 0.0843 0.2584 0.4324 0.5263 0.6142 ...
- attr(*, "dimnames")=List of 2
..$ : chr [1:6] "1" "2" "3" "4" ...
..$ : NULL

> fits <- lm(Ymat ~ carb, Formaldehyde)
> str(coefs <- coef(fits))

  num [1:2, 1:10000] -0.00201 0.87284 -0.00922 0.90215 0.00369 ...
- attr(*, "dimnames")=List of 2
..$ : chr [1:2] "(Intercept)" "carb"
..$ : NULL
```

Most of the time we want the coefficients to be a data frame instead with columns corresponding to the coefficient names.

```
> str(coefs <- data.frame(t(coef(fits)), check.names=FALSE))
```

```
'data.frame':      10000 obs. of  2 variables:
 $ (Intercept): num  -0.002005 -0.009224 0.003691 -0.001113 0.000131 ...
 $ carb       : num  0.873 0.902 0.887 0.896 0.892 ...
```

Recall that the “true” coefficients for this model are

```
> printCoefmat(coef(summary(lm1)))

              Estimate Std. Error t value Pr(>|t|)
(Intercept) 0.0050857  0.0078337  0.6492   0.5516
carb        0.8762857  0.0135345 64.7444 3.409e-07
```

For an unbiased estimator the mean of the distribution of the estimator should be the parameter value.

```
> sapply(coefs, mean)
```

```
(Intercept)      carb
0.005078325 0.876313251
```

and the standard deviations should be close to the standard errors

```
> sapply(coefs, sd)
```

```
(Intercept)      carb
0.007943042 0.013647517
```

The correlation of sample of coefficient estimates should be close to the value for the fitted model

```
> summary(lm1, corr=TRUE)$correlation
```

```
(Intercept)      carb
(Intercept)  1.000000 -0.892664
carb        -0.892664  1.000000
```

```
> cor(coefs)
```

```
(Intercept)      carb
(Intercept)  1.000000 -0.896498
carb        -0.896498  1.000000
```

If, instead, we wish to consider the variance-covariance matrices, we use

```
> vcov(lm1)
```

```
(Intercept)      carb
(Intercept) 6.136653e-05 -0.0000946449
carb      -9.464490e-05  0.0001831837
```

```
> var(coefs)
```




Figure 3.1: Empirical density plots of coefficient estimates from data simulated according to the estimated parameters in model `lm1`

```

              (Intercept)          carb
(Intercept) 6.309192e-05 -0.0000971829
carb       -9.718290e-05  0.0001862547

```

In Fig. 3.1 we show the empirical density plots for the coefficients separately. Alternatively, we could examine the normal Q-Q plots (Fig. 3.2).

We could also plot contours of the estimated 2-dimensional density (Fig. 3.3). The background of the empirical density contours is like a two-dimensional histogram but using hexagonal shaped bins instead of rectangles.






Figure 3.2: Normal quantile-quantile plots of coefficient estimates from responses simulated according to the estimated parameters in model `lm1`.




Figure 3.3: Normal quantile-quantile plots of coefficient estimates from responses simulated according to the estimated parameters in model `lm1`.



## Chapter 4

# The Gauss-Markov Theorem

In Chap. 3 we showed that the least squares estimator,  $\hat{\boldsymbol{\beta}}_{LSE}$ , in a Gaussian linear model has is *unbiased*, meaning that  $E[\hat{\boldsymbol{\beta}}_{LSE}] = \boldsymbol{\beta}$ , and that its variance-covariance matrix is

$$\text{Var} \hat{\boldsymbol{\beta}}_{LSE} = \sigma^2 (\mathbf{X}'\mathbf{X})^{-1} = \sigma^2 \mathbf{R}^{-1}(\mathbf{R}^{-1})'$$

The Gauss-Markov theorem says that this variance-covariance (or *dispersion*) is the best that we can do when we restrict ourselves to *linear unbiased estimators*, which means estimators that are linear functions of  $\mathcal{Y}$  and are unbiased.

To make these definitions more formal:

**Definition 5** (Minimum Dispersion). *Let  $\mathcal{T} = (\mathcal{T}_1, \dots, \mathcal{T}_p)'$  be an estimator of  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)'$ . The dispersion of  $\mathcal{T}$  is  $\mathbf{D}(\mathcal{T}) = E[(\mathcal{T} - \boldsymbol{\theta})(\mathcal{T} - \boldsymbol{\theta})']$ . If  $\mathcal{T}$  is unbiased then its dispersion is simply its variance-covariance matrix,  $\mathbf{D}(\mathcal{T}) = \text{Var}(\mathcal{T})$ .  $\mathcal{T}$  is minimum dispersion unbiased estimator of  $\boldsymbol{\theta}$  if  $\mathbf{D}(\tilde{\mathcal{T}}) - \mathbf{D}(\mathcal{T})$  is positive semidefinite for any unbiased estimator  $\tilde{\mathcal{T}}$ . That is*

$$\mathbf{a}'[\mathbf{D}(\tilde{\mathcal{T}}) - \mathbf{D}(\mathcal{T})]\mathbf{a} \geq 0 \quad \forall \mathbf{a} \in \mathbb{R}^p$$

*Because the dispersion matrices of unbiased estimators are the variance-covariance matrices, this condition is equivalent to*

$$\mathbf{a}' \text{Var}(\tilde{\mathcal{T}})\mathbf{a} - \mathbf{a}' \text{Var}(\mathcal{T})\mathbf{a} \geq 0 \Rightarrow \text{Var}(\mathbf{a}'\tilde{\mathcal{T}}) - \text{Var}(\mathbf{a}'\mathcal{T}) \geq 0$$

**Theorem 8** (Gauss-Markov). *In the full-rank case (i.e.  $\text{rank}(\mathbf{X}) = p$ ) the minimum dispersion linear unbiased estimator of  $\boldsymbol{\beta}$  is  $\hat{\boldsymbol{\beta}}_{LSE}$  with dispersion matrix  $\sigma^2(\mathbf{X}'\mathbf{X})^{-1}$ . It is also called the best linear unbiased estimator or BLUE of  $\boldsymbol{\beta}$ .*

*Proof.* Any linear estimator of  $\boldsymbol{\beta}$  can be written as  $\mathbf{A}\mathcal{Y}$  for some  $p \times n$  matrix  $\mathbf{A}$ . (That's what it means to be a linear estimator.) To be an unbiased linear estimator we must have

$$\boldsymbol{\beta} = E[\mathbf{A}\mathcal{Y}] = \mathbf{A}E[\mathcal{Y}] = \mathbf{A}\mathbf{X}\boldsymbol{\beta} \quad \forall \boldsymbol{\beta} \in \mathbb{R}^p \Rightarrow \mathbf{A}\mathbf{X} = \mathbf{I}_p$$

The variance-covariance matrix such a linear unbiased estimator,  $\mathbf{A}\mathcal{Y}$ , is

$$\text{Var}(\mathbf{A}\mathcal{Y}) = \mathbf{A} \text{Var}(\mathcal{Y})\mathbf{A}' = \mathbf{A}\sigma^2\mathbf{I}_n\mathbf{A}' = \sigma^2\mathbf{A}\mathbf{A}'$$

Now we must show that

$$\text{Var}(\mathbf{a}'\mathbf{A}\mathcal{Y}) - \text{Var}(\mathbf{a}'\widehat{\boldsymbol{\beta}}_{LSE}) = \sigma^2 \mathbf{a}' (\mathbf{A}\mathbf{A}' - (\mathbf{X}'\mathbf{X})^{-1}) \mathbf{a} \geq 0, \forall \mathbf{a} \in \mathbb{R}^p.$$

In other words, the symmetric matrix,  $(\mathbf{A}\mathbf{A}' - (\mathbf{X}'\mathbf{X})^{-1})$ , must be positive semi-definite. Consider

$$\begin{aligned} \mathbf{A}\mathbf{A}' &= [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'] [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']' \\ &= [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'] [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']' + [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'] [(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']' + \\ &\quad (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X} [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']' + [(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'] [\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}] \\ &= [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'] [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']' + (\mathbf{X}'\mathbf{X})^{-1}, \end{aligned}$$

showing that  $\mathbf{A}\mathbf{A}' - (\mathbf{X}'\mathbf{X})^{-1}$  is the positive semi-definite matrix  $[\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'] [\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']'$ . Therefore  $\widehat{\boldsymbol{\beta}}_{LSE}$  is the BLUE for  $\boldsymbol{\beta}$ .  $\square$

**Corollary 7.** *If  $\text{rank}(\mathbf{X}) = p < n$ , the best linear unbiased estimator of  $\mathbf{a}'\boldsymbol{\beta}$  is  $\mathbf{a}'\widehat{\boldsymbol{\beta}}_{LSE}$ .*

To extend the Gauss-Markov theorem to the rank-deficient case we must define

**Definition 6** (Estimable linear function). *An estimable linear function of the parameters  $\boldsymbol{\beta}$  in the linear model,  $\mathcal{Y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I}_n)$ , is any function of the form  $\mathbf{l}'\boldsymbol{\beta}$  where  $\mathbf{l}$  is in the row span of  $\mathbf{X}$ . That is,  $\mathbf{l}'\boldsymbol{\beta}$  is estimable if and only if there exists  $\mathbf{c} \in \mathbb{R}^n$  such that  $\mathbf{l} = \mathbf{X}'\mathbf{c}$ .*

The coefficients of the estimable functions form a  $\text{rank}(\mathbf{X}) = k$ -dimensional linear subspace of  $\mathbb{R}^p$ . In the full-rank this subspace is all of  $\mathbb{R}^p$  so any linear combination  $\mathbf{l}'\boldsymbol{\beta}$  is estimable.

In the rank-deficient case (i.e.  $\text{rank}(\mathbf{X}) = k < p$ ), consider the singular value decomposition  $\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}'$  with  $\mathbf{D}$  a diagonal matrix having non-negative, non-increasing diagonal elements, the first  $k$  of which are positive and the last  $p - k$  are zero. Let  $\mathbf{U}_k$  be the first  $k$  columns of  $\mathbf{U}$ ,  $\mathbf{D}_k$  be the first  $k$  rows and  $k$  columns of  $\mathbf{D}$ , and  $\mathbf{V}_k$  be the first  $k$  columns of  $\mathbf{V}$ . The coefficients  $\mathbf{l}$  for an estimable linear function must lie in the column span of  $\mathbf{V}_k$  because

$$\mathbf{l} = \mathbf{X}'\mathbf{c} = \mathbf{V}_k \underbrace{\mathbf{D}_k \mathbf{U}_k' \mathbf{c}}_{\mathbf{a}} = \mathbf{V}_k \mathbf{a}$$

We will write the  $p \times (p - k)$  matrix formed by the last  $p - k$  columns of  $\mathbf{V}$  as  $\mathbf{V}_{p-k}$  so that

$$\boldsymbol{\beta} = \mathbf{V}\mathbf{V}'\boldsymbol{\beta} = [\mathbf{V}_k \quad \mathbf{V}_{(p-k)}] \begin{bmatrix} \mathbf{V}_k' \\ \mathbf{V}_{(p-k)}' \end{bmatrix} \boldsymbol{\beta} = \mathbf{V}_k \boldsymbol{\gamma} + \mathbf{V}_{p-k} \boldsymbol{\delta}$$

where  $\boldsymbol{\gamma} = \mathbf{V}_k' \boldsymbol{\beta}$  and  $\boldsymbol{\delta} = \mathbf{V}_{p-k}' \boldsymbol{\beta}$  are the estimable and inestimable parts of the parameter vector in the  $\mathbf{V}$  basis.

Now any estimable function is of the form

$$\mathbf{l}'\boldsymbol{\beta} = \mathbf{a}'\mathbf{V}_k' \boldsymbol{\beta} = \mathbf{a}'\boldsymbol{\gamma} + \mathbf{0} = \mathbf{a}'\boldsymbol{\gamma},$$

where  $\boldsymbol{\gamma}$  is the parameter in the full-rank model  $\mathcal{Y} \sim \mathcal{N}(\mathbf{D}_k \mathbf{U}_k \boldsymbol{\gamma}, \sigma^2 \mathbf{I}_n)$ .

So anything we say about estimable functions of  $\boldsymbol{\beta}$  can be transformed into a statement about  $\boldsymbol{\gamma}$  in the full rank model and anything we say about the fitted values,  $\mathbf{X}\boldsymbol{\beta}$ , or the residuals can be expressed in terms of the full-rank  $\mathbf{D}_k \mathbf{U}_k \boldsymbol{\gamma}$ . In particular, the hat matrix,  $\mathbf{H} = \mathbf{U}_k \mathbf{U}_k'$ , and has  $\text{rank}(\mathbf{H}) = k$  and the projection into the orthogonal (residual) space is  $\mathbf{I}_n - \mathbf{H}$ .

**Corollary 8** (Gauss-Markov extension to rank-deficient cases).  $\mathbf{l}'\hat{\boldsymbol{\beta}}_{LSE} = \mathbf{a}'\hat{\boldsymbol{\gamma}}_{LSE}$  is the BLUE for any estimable linear function,  $\mathbf{l}'\boldsymbol{\beta}$ , of  $\boldsymbol{\beta}$ .

*Proof.* By the Gauss-Markov theorem  $\hat{\boldsymbol{\gamma}}_{LSE}$  is the BLUE for  $\boldsymbol{\gamma}$  and  $\mathbf{l}'\boldsymbol{\beta} = \mathbf{a}'\boldsymbol{\gamma}$  is a linear function of  $\boldsymbol{\gamma}$ .  $\square$

**Theorem 9.** Suppose that  $k = \text{rank}(\mathbf{X}) \leq p$ . Then an unbiased estimator of  $\sigma^2$  is

$$S^2 = \frac{\|\mathcal{Y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2}{n - k} = \frac{\|\hat{\boldsymbol{\epsilon}}\|^2}{n - k} = \frac{\sum_{i=1}^n \hat{\epsilon}_i^2}{n - k}.$$

*Proof.* The simple proof is to observe that this estimator is the unbiased estimator of  $\sigma^2$  for the full-rank version of the model,  $\mathcal{Y} \sim \mathcal{N}(\mathbf{D}_k \mathbf{U}_k \boldsymbol{\gamma}, \sigma^2 \mathbf{I}_n)$ .  $\square$