

Appendix A

Review and Background Material

A.1 Expected Value, Variance and Covariance (Review)

Expected Value Let X be a random variable. If X is continuous, the expected value is defined as

$$E(X) = \int_{-\infty}^{\infty} x f_X(x) dx.$$

If X is discrete, the formula is

$$E(X) = \sum_x x p_X(x).$$

Conditional expectation uses these same formulas, only with conditional densities or probability mass functions.

Let $Y = g(X)$. The change of variables formula (a very big Theorem¹) tells us

$$E(Y) = \int_{-\infty}^{\infty} y f_Y(y) dy = \int_{-\infty}^{\infty} g(x) f_X(x) dx \tag{A.1}$$

or, for discrete random variables

$$E(Y) = \sum_y y p_Y(y) = \sum_x g(x) p_X(x).$$

One useful function $g(x)$ is the *indicator function* for a set A . $I_A(x) = 1$ if $x \in A$, and $I_A(x) = 0$ if $x \notin A$. The expected value of an indicator function is just a probability

¹The change of variables formula holds under very general circumstances; see for example Theorem 16.12 in Billingsley's *Probability and measure* [2]. It is extremely convenient and easy to apply, because there is no need to derive the probability distribution of Y . So for example the sets of values where $f_X(x) \neq 0$ and $f_Y(y) \neq 0$ (and therefore the regions over which you are integrating in expression (A.1)) may be different and you don't have to think about it. Furthermore, the function $g(x)$ is almost arbitrary. In particular, it need not be differentiable, a condition you would need to prove anything for the continuous case with ordinary calculus.

because, for discrete random variables,

$$E(I_A(X)) = \sum_x I_A(x) p_X(x) = \sum_{x \in A} p_X(x) = P(X \in A).$$

For continuous random variables, something similar happens; multiplication by $I_A(x)$ erases the density for $x \notin A$, and integration of the product from zero to infinity is just integration over the set A , yielding $P(X \in A)$.

Another useful function is a conditional expectation. If we write the conditional density

$$f_{Y|X}(y|X) = \frac{f_{X,Y}(X, y)}{f_X(X)}$$

with the capital letter X , we really mean it. X is a random variable, not a constant, and for any fixed y , the conditional density is a random variable. The conditional expected value is another random variable $g(x)$:

$$E(Y|X) = \int_{-\infty}^{\infty} y f_{Y|X}(y|X) dy.$$

This may be a strange-looking function, but still it is a function, and one can take its expected value using the change of variables formula [A.1](#).

$$E(E(Y|X)) = \int_{-\infty}^{\infty} g(x) f_X(x) dx = \int_{-\infty}^{\infty} E(Y|x) f_X(x) dx.$$

Provided $|E(Y)| < \infty$, order of integration or summation may be exchanged², and we have the double expectation formula:

$$E(Y) = E(E(Y|X)).$$

You will prove a slightly more general and useful version as an exercise.

The change of variables formula ([A.1](#)) still holds if \mathbf{X} is a vector, or even if both \mathbf{X} and \mathbf{Y} are vectors, and integration or summation is replaced by multiple integration or summation. So, for example if $\mathbf{X} = (X_1, X_2)'$ has joint density $f_{\mathbf{X}}(\mathbf{x}) = f_{X_1, X_2}(x_1, x_2)$ and $g(x_1, x_2) = a_1 x_1 + a_2 x_2$,

$$\begin{aligned} E(a_1 X_1 + a_2 X_2) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (a_1 x_1 + a_2 x_2) f_{X_1, X_2}(x_1, x_2) dx_1 dx_2 \\ &= a_1 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_1 f_{X_1, X_2}(x_1, x_2) dx_1 dx_2 + a_2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_2 f_{X_1, X_2}(x_1, x_2) dx_1 dx_2 \\ &= a_1 E(X_1) + a_2 E(X_2). \end{aligned}$$

Using this approach, it is easy to establish the linearity of expected value

$$E\left(\sum_{j=1}^m a_j X_j\right) = \sum_{j=1}^m a_j E(X_j) \tag{A.2}$$

²By Fubini's Theorem.

and other familiar properties.

The change of variables formula holds if the So, for example, since $g(x_1, x_2, \dots, x_p) = x_3$ is one possible function of x_1, x_2, \dots, x_p ,

$$\begin{aligned} \int \cdots \int x_3 f(\mathbf{x}) d\mathbf{x} &= \int \cdots \int x_3 f(x_1, \dots, x_p) dx_1 \cdots dx_p \\ &= E(X_3). \end{aligned}$$

Variance and Covariance Denote $E(X)$ by μ_X . The variance of X is defined as

$$\text{Var}(X) = E[(X - \mu_X)^2],$$

and the covariance of X and Y is defined as

$$\text{Cov}(X, Y) = E[(X - \mu_X)(Y - \mu_Y)].$$

The *correlation* between X and Y is

$$\text{Corr}(X, Y) = \frac{\text{Cov}(X, Y)}{\sqrt{\text{Var}(X)\text{Var}(Y)}}.$$

Exercises A.1 For Questions 1 through 5, please use integrals or summation. For the remaining questions, use only expected value signs.

1. Let a be a constant and let X be a random variable, either continuous or discrete (you choose). Use the change of variables formula A.1 to show that $E(a) = a$.
2. Use the change of variables formula to prove the linear property given in expression (A.2). If you assume independence, you get a zero.
3. Let X and Y be discrete random variables, with $E(|h(X)|) < \infty$. Use the change of variables formula to prove $E(h(X)) = E[E(h(X)|Y)]$. Because $E(|h(X)|) < \infty$, Fubini's Theorem says that you are free to exchange order of summation. Is the result of this problem also true for continuous random variables? Why or why not?
4. Let X and Y be continuous random variables. Prove

$$P(X \in A) = \int_{-\infty}^{\infty} P(X \in A|Y = y) f_Y(y) dy.$$

This is sometimes called the *Law of Total Probability*. Is it also true for discrete random variables? Why or why not? Hint: use indicator functions.

5. Let X and Y be continuous random variables. Prove that if X and Y are independent, $E(XY) = E(X)E(Y)$. Draw an arrow to the place in your answer where you use independence, and write “This is where I use independence.”
6. Show that $Var(X) = E[X^2] - \mu_X^2$.
7. Show that $Cov[X, Y] = E[XY] - \mu_X\mu_Y$.
8. In the following, X and Y are random variables, while a and b are fixed constants. For each pair of statements below, one is true and one is false (that is, not true in general). State which one is true, and prove it. Zero marks if you prove both statements are true, even if one of the proofs is correct.
 - (a) $Var(aX) = aVar(X)$ or $Var(aX) = a^2Var(X)$
 - (b) $Var(aX + b) = a^2Var(X) + b^2$ or $Var(aX + b) = a^2Var(X)$
 - (c) $Var(a) = 0$ or $Var(a) = a^2$
 - (d) $Cov(X + a, Y + b) = Cov(X, Y) + ab$ or $Cov(X + a, Y + b) = Cov(X, Y)$
 - (e) $Var(aX + bY) = a^2Var(X) + b^2Var(Y)$ or $Var(aX + bY) = a^2Var(X) + b^2Var(Y) + 2abCov(X, Y)$
9. Let X_1, \dots, X_n be random variables and $Y = \sum_{i=1}^n X_i$. Derive a general formula for $Var(Y)$ in terms of the variance and covariances of X_1, \dots, X_n . Show your work. Now give the useful special case that applies when X_1, \dots, X_n are independent.
10. Let X_1, \dots, X_n be independent and identically distributed random variables (the standard model of a random sample with replacement). Denoting $E(X_i)$ by μ and $V(X_i)$ by σ^2 ,
 - (a) Show $E[\bar{X}] = \mu$; that is, the sample mean is unbiased for μ .
 - (b) Find $V(\bar{X})$.
 - (c) Show $E[S^2] = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2 = \sigma^2$; that is, the sample variance is unbiased for σ^2 .

A.2 Matrix Calculations

A.2.1 Basic definitions

A matrix is a rectangular array of numbers. They are usually denoted by boldface letters like \mathbf{A} , while scalars (1×1 matrices) are lower case in italics, like a, b, c . Matrices are also written by giving their (i, j) element in brackets, like $\mathbf{A} = [a_{i,j}]$.

Let $\mathbf{A} = [a_{i,j}]$ and $\mathbf{B} = [b_{i,j}]$ be $n \times p$ matrices of constants, $\mathbf{C} = [c_{i,j}]$ be $p \times q$, and let u and v be scalars (1×1 matrices). Define

Matrix addition: $\mathbf{A} + \mathbf{B} = [a_{i,j} + b_{i,j}]$. The matrices must have the same number of rows and the same number of columns for addition (or subtraction) to be defined.

Matrix multiplication: $\mathbf{AC} = [\sum_{k=1}^p a_{i,k}c_{k,j}]$. Each element of \mathbf{AC} is the inner product of a row of \mathbf{A} and a column of \mathbf{C} . Thus, the number of columns in \mathbf{A} must equal the number of rows in \mathbf{C} . Even if $q = n$ so that multiplication in both orders is well defined, in general $\mathbf{AC} \neq \mathbf{CA}$.

Scalar multiplication: $u\mathbf{A} = [u \cdot a_{i,j}]$

Transposition: $\mathbf{A}' = [a_{j,i}]$

Symmetric matrix: A square matrix \mathbf{D} is said to be *symmetric* if $\mathbf{D} = \mathbf{D}'$.

Identity matrix: \mathbf{I} is a square matrix with ones on the main diagonal and zeros elsewhere. $\mathbf{IC} = \mathbf{C}$ and $\mathbf{AI} = \mathbf{A}$.

Diagonal matrix: A square matrix $\mathbf{D} = [d_{i,j}]$ is said to be *diagonal* if $d_{i,j} = 0$ for $i \neq j$.

Triangular matrix: A square matrix $\mathbf{D} = [d_{i,j}]$ is said to be *triangular* if $d_{i,j} = 0$ for $i < j$ or $i > j$ (or both, in which case it is also diagonal).

Distributive laws for matrix and scalar multiplication are easy to establish and are left as exercises.

A.2.2 Transpose of a product

The transpose of a product is the product of transposes, in the reverse order: $(\mathbf{AC})' = \mathbf{C}'\mathbf{A}'$.

A.2.3 Linear independence

The idea behind linear independence of a collection of vectors (say, the columns of a matrix) is that none of them can be written as a linear combination of the others. Formally, let \mathbf{X} be an $n \times p$ matrix of constants. The columns of \mathbf{X} are said to be *linearly dependent* if there exists a $p \times 1$ matrix $\mathbf{v} \neq \mathbf{0}$ with $\mathbf{Xv} = \mathbf{0}$. We will say that the columns of \mathbf{X} are *linearly independent* if $\mathbf{Xv} = \mathbf{0}$ implies $\mathbf{v} = \mathbf{0}$.

A.2.4 Row and column rank

The *row rank* of a matrix is the number of linearly independent rows. The *column rank* is the number of linearly independent columns. The rank of a matrix is the minimum of the row rank and the column rank. Thus, the rank of a matrix cannot exceed the minimum of the number of rows and the number of columns.

A.2.5 Matrix Inverse

Let \mathbf{A} and \mathbf{B} be square matrices of the same size. \mathbf{B} is said to be the *inverse* of \mathbf{A} and may be written $\mathbf{B} = \mathbf{A}^{-1}$. The definition is $\mathbf{AB} = \mathbf{BA} = \mathbf{I}$. Thus, there are always two equalities to establish when you are showing that one matrix is the inverse of another. Matrix inverses have the following properties, which may be proved as exercises.

- If a matrix inverse exists, it is unique.
- $\mathbf{A}^{-1'} = \mathbf{A}'^{-1}$
- If the scalar $u \neq 0$, $(u\mathbf{A})^{-1} = \frac{1}{u}\mathbf{A}^{-1}$.
- Suppose that the square matrices \mathbf{A} and \mathbf{B} both have inverses. Then $(\mathbf{AB})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}$.
- If \mathbf{A} is a $p \times p$ matrix, \mathbf{A}^{-1} exists if and only if the rank of \mathbf{A} equals p .

Sometimes the following formula for the inverse of a 2×2 matrix is useful:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix} \quad (\text{A.3})$$

In some cases the inverse of the matrix is its transpose. When $\mathbf{A}' = \mathbf{A}^{-1}$, the matrix \mathbf{A} is said to be *orthogonal*, because the column (row) vectors are all at right angles (zero inner product). In addition, they all have length one, because the inner product of each column (row) with itself equals one.

A.2.6 Positive definite matrices

The $n \times n$ matrix \mathbf{A} is said to be *positive definite* if

$$\mathbf{y}'\mathbf{A}\mathbf{y} > 0$$

for all $n \times 1$ vectors $\mathbf{y} \neq \mathbf{0}$. It is called *non-negative definite* (or sometimes positive semi-definite) if $\mathbf{y}'\mathbf{A}\mathbf{y} \geq 0$. Positive definiteness is a critical property of variance-covariance matrices.

A.2.7 Determinants

Let $\mathbf{A} = [a_{i,j}]$ be an $n \times n$ matrix, so that the following applies to *square* matrices. The *determinant* of \mathbf{A} , denoted $|\mathbf{A}|$, is defined as a *sum of signed elementary products*. An elementary product is a product of elements of \mathbf{A} such that there is exactly one element from every row and every column. The “signed” part is determined as follows.

Let S_n denote the set of all permutations of the set $\{1, \dots, n\}$, and denote such a permutation by $\sigma = (\sigma_1, \dots, \sigma_n)$. Each permutation may be obtained from $(1, \dots, n)$ by

a finite number of switches of numbers. If the number of switches required is even (this includes zero), let $\text{sgn}(\sigma) = +1$; if it is odd, let $\text{sgn}(\sigma) = -1$. Then,

$$|\mathbf{A}| = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n a_{i,\sigma_i}. \quad (\text{A.4})$$

Some properties of determinants are:

- $|\mathbf{AB}| = |\mathbf{A}| |\mathbf{B}|$
- $|\mathbf{A}'| = |\mathbf{A}|$
- $|\mathbf{A}^{-1}| = 1/|\mathbf{A}|$, and if $|\mathbf{A}| = 0$, \mathbf{A}^{-1} does not exist.
- If $\mathbf{A} = [a_{i,j}]$ is triangular, $|\mathbf{A}| = \prod_{i=1}^n a_{i,i}$. That is, for triangular (including diagonal) matrices, the determinant is the product of the elements on the main diagonal.
- Adding a multiple of one row to another row of a matrix, or adding a multiple of a column to another column leaves the determinant unchanged.
- Exchanging any two rows or any two columns of a matrix multiplies the determinant by -1 .
- Multiplying a single row or column by a constant multiplies the determinant by that constant, so that $|v\mathbf{A}| = v^n |\mathbf{A}|$

A.2.8 Eigenvalues and eigenvectors

Let $\mathbf{A} = [a_{i,j}]$ be an $n \times n$ matrix, so that the following applies to square matrices. \mathbf{A} is said to have an *eigenvalue* λ and (non-zero) *eigenvector* \mathbf{x} corresponding to λ if

$$\mathbf{A}\mathbf{x} = \lambda\mathbf{x}.$$

Note that λ is a scalar and $\mathbf{x} \neq \mathbf{0}$ is an $n \times 1$ matrix, typically chosen so that it has length one. It is also possible and desirable to choose the eigenvectors so they are mutually perpendicular (the inner product of any two equals zero).

To solve the eigenvalue equation, write

$$\mathbf{A}\mathbf{x} = \lambda\mathbf{x} \Rightarrow \mathbf{A}\mathbf{x} - \lambda\mathbf{x} = \mathbf{A}\mathbf{x} - \lambda\mathbf{I}\mathbf{x} = (\mathbf{A} - \lambda\mathbf{I})\mathbf{x} = \mathbf{0}.$$

If $(\mathbf{A} - \lambda\mathbf{I})^{-1}$ existed, it would be possible to solve for \mathbf{x} by multiplying both sides on the left by $(\mathbf{A} - \lambda\mathbf{I})^{-1}$, yielding $\mathbf{x} = \mathbf{0}$. But the definition specifies $\mathbf{x} \neq \mathbf{0}$, so the inverse cannot exist for the definition of an eigenvalue to be satisfied. Since $(\mathbf{A} - \lambda\mathbf{I})^{-1}$ fails to exist precisely when the determinant $|\mathbf{A} - \lambda\mathbf{I}| = 0$, the eigenvalues are the λ values that solve the determinantal equation

$$|\mathbf{A} - \lambda\mathbf{I}| = 0.$$

The left-hand side is a polynomial in λ , called the *characteristic polynomial*. If the matrix \mathbf{A} is real-valued and also symmetric, then all its eigenvalues are guaranteed to be real-valued — a handy characteristic not generally true of solutions to polynomial equations. The eigenvectors can also be chosen to be real, and for our purposes they always will be.

One of the many useful properties of eigenvalues is that **the determinant is the product of the eigenvalues**:

$$|\mathbf{A}| = \prod_{i=1}^n \lambda_i$$

A.2.9 Spectral decomposition of symmetric matrices

The *Spectral decomposition theorem* says that every square and symmetric matrix $\mathbf{A} = [a_{i,j}]$ may be written

$$\mathbf{A} = \mathbf{P}\mathbf{\Lambda}\mathbf{P}', \quad (\text{A.5})$$

where the columns of \mathbf{P} (which may also be denoted $\mathbf{x}_1, \dots, \mathbf{x}_n$) are the eigenvectors of \mathbf{A} , and the diagonal matrix $\mathbf{\Lambda}$ contains the corresponding eigenvalues.

$$\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_n \end{bmatrix}$$

Because the eigenvectors are orthonormal, \mathbf{P} is an orthogonal matrix; that is, $\mathbf{P}\mathbf{P}' = \mathbf{P}'\mathbf{P} = \mathbf{I}$.

The following shows how to get a spectral decomposition from R .

```
> help(eigen)
> A = rbind(c(-10,2),
+          c(2,5)) # Symmetric
> eigenA = eigen(A); eigenA
$values
[1]  5.262087 -10.262087

$vectors
      [,1]      [,2]
[1,] 0.1299328  0.9915228
[2,] 0.9915228 -0.1299328

> det(A)
[1] -54
> prod(eigenA$values)
[1] -54

> Lambda = diag(eigenA$values); Lambda
```



```

      [,1]      [,2]
[1,] 5.262087  0.000000
[2,] 0.000000 -10.26209

> P = eigenA$eigenvectors; P
      [,1]      [,2]
[1,] 0.1299328  0.9915228
[2,] 0.9915228 -0.1299328

> P %% Lambda %% t(P) # Matrix multiplication
      [,1] [,2]
[1,]  -10   2
[2,]   2   5

```

Another way to express the spectral decomposition is

$$\mathbf{A} = \sum_{i=1}^n \lambda_i \mathbf{x}_i \mathbf{x}_i' \quad (\text{A.6})$$

where again, $\mathbf{x}_1, \dots, \mathbf{x}_n$ are the eigenvectors of \mathbf{A} , and $\lambda_1, \dots, \lambda_n$ are the corresponding eigenvalues. It's a weighted sum of outer (not inner) products of the eigenvectors; the weights are the eigenvalues.

Continuing the R example, here is $\mathbf{x}_1 \mathbf{x}_1'$. Notice how the diagonal elements add to one, as they must.

```

> eigenA$eigenvectors[,1] %% t(eigenA$eigenvectors[,1])
      [,1]      [,2]
[1,] 0.01688253 0.1288313
[2,] 0.12883133 0.9831175

```

Reproducing (A.6) for completeness,

```

> prod1 = eigenA$eigenvectors[,1] %% t(eigenA$eigenvectors[,1])
> prod2 = eigenA$eigenvectors[,2] %% t(eigenA$eigenvectors[,2])
> eigenA$values[1]
[1] 5.262087
> eigenA$values[1]*prod1 + eigenA$values[2]*prod2
      [,1] [,2]
[1,]  -10   2
[2,]   2   5
> A
      [,1] [,2]
[1,]  -10   2
[2,]   2   5

```

A.2.10 Real symmetric matrices

For a symmetric $n \times n$ matrix \mathbf{A} (like a variance-covariance matrix), the eigenvalues are all real numbers, and the eigenvectors can be chosen to be real, perpendicular (inner product zero), and of length one. Also, for this important case the following conditions are equivalent:

- Rows linearly independent
- Columns linearly independent
- Rank = n
- Positive definite
- Non-singular (\mathbf{A}^{-1} exists)
- All eigenvalues are strictly positive

A.2.11 Trace of a square matrix

The *trace* of a square matrix $\mathbf{A} = [a_{i,j}]$ is the sum of its diagonal elements. Write

$$tr(\mathbf{A}) = \sum_{i=1}^n a_{i,i}.$$

Properties like $tr(\mathbf{A} + \mathbf{B}) = tr(\mathbf{A}) + tr(\mathbf{B})$ follow immediately from the definition. Perhaps less obvious is the following. Let \mathbf{A} be an $r \times p$ matrix and \mathbf{B} be a $p \times r$ matrix, so that the product matrices \mathbf{AB} and \mathbf{BA} are both defined. These two matrices are not necessarily equal; in fact, they need not even be the same size. But still,

$$tr(\mathbf{AB}) = tr(\mathbf{BA}). \tag{A.7}$$

To see this, write

$$\begin{aligned} tr(\mathbf{AB}) &= tr \left(\left[\sum_{k=1}^p a_{i,k} b_{k,j} \right] \right) \\ &= \sum_{i=1}^r \sum_{k=1}^p a_{i,k} b_{k,i} \\ &= \sum_{k=1}^p \sum_{i=1}^r b_{k,i} a_{i,k} \\ &= tr \left(\left[\sum_{k=1}^r b_{i,k} a_{k,j} \right] \right) \\ &= tr(\mathbf{BA}) \end{aligned}$$

Notice how the indices of summation have been changed in the line where tr appears for the second time. This is legitimate, because for example $\sum_{i=1}^r c_i$ and $\sum_{k=1}^r c_k$ both mean $c_1 + \cdots + c_r$.

Also, from the spectral decomposition (A.6), the trace is the sum of the eigenvalues:

$$tr(\mathbf{A}) = \sum_{i=1}^n \lambda_i.$$

This follows easily using (A.7), but actually it applies to *any* square matrix; the matrix need not be symmetric.

Exercises A.2

1. Which statement is true?
 - (a) $\mathbf{A}(\mathbf{B} + \mathbf{C}) = \mathbf{AB} + \mathbf{AC}$
 - (b) $\mathbf{A}(\mathbf{B} + \mathbf{C}) = \mathbf{BA} + \mathbf{CA}$
 - (c) Both a and b
 - (d) Neither a nor b

2. Which statement is true?
 - (a) $a(\mathbf{B} + \mathbf{C}) = a\mathbf{B} + a\mathbf{C}$
 - (b) $a(\mathbf{B} + \mathbf{C}) = \mathbf{Ba} + \mathbf{Ca}$
 - (c) Both a and b
 - (d) Neither a nor b

3. Which statement is true?
 - (a) $(\mathbf{B} + \mathbf{C})\mathbf{A} = \mathbf{AB} + \mathbf{AC}$
 - (b) $(\mathbf{B} + \mathbf{C})\mathbf{A} = \mathbf{BA} + \mathbf{CA}$
 - (c) Both a and b
 - (d) Neither a nor b

4. Which statement is true?
 - (a) $(\mathbf{AB})' = \mathbf{A}'\mathbf{B}'$
 - (b) $(\mathbf{AB})' = \mathbf{B}'\mathbf{A}'$
 - (c) Both a and b
 - (d) Neither a nor b

5. Which statement is true?

- (a) $\mathbf{A}'' = \mathbf{A}$
 - (b) $\mathbf{A}''' = \mathbf{A}'$
 - (c) Both a and b
 - (d) Neither a nor b
6. Suppose that the square matrices \mathbf{A} and \mathbf{B} both have inverses. Which statement is true?
- (a) $(\mathbf{AB})^{-1} = \mathbf{A}^{-1}\mathbf{B}^{-1}$
 - (b) $(\mathbf{AB})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}$
 - (c) Both a and b
 - (d) Neither a nor b
7. Which statement is true?
- (a) $(\mathbf{A} + \mathbf{B})' = \mathbf{A}' + \mathbf{B}'$
 - (b) $(\mathbf{A} + \mathbf{B})' = \mathbf{B}' + \mathbf{A}'$
 - (c) $(\mathbf{A} + \mathbf{B})' = (\mathbf{B} + \mathbf{A})'$
 - (d) All of the above
 - (e) None of the above
8. Which statement is true?
- (a) $tr(\mathbf{A} + \mathbf{B}) = tr(\mathbf{A}) + tr(\mathbf{B})$
 - (b) $tr(\mathbf{A} + \mathbf{B}) = tr(\mathbf{B}) + tr(\mathbf{A})$
 - (c) Both a and b
 - (d) Neither a nor b
9. Which statement is true?
- (a) $a tr(\mathbf{B}) = tr(a\mathbf{B})$.
 - (b) $tr(\mathbf{B})a = tr(a\mathbf{B})$
 - (c) Both a and b
 - (d) Neither a nor b
10. Which statement is true?
- (a) $(a + b)\mathbf{C} = a\mathbf{C} + b\mathbf{C}$
 - (b) $(a + b)\mathbf{C} = \mathbf{C}a + \mathbf{C}b$
 - (c) $(a + b)\mathbf{C} = \mathbf{C}(a + b)$
 - (d) All of the above

- (e) None of the above
11. Recall that \mathbf{A} symmetric means $\mathbf{A} = \mathbf{A}'$. Let \mathbf{X} be an n by p matrix. Prove that $\mathbf{X}'\mathbf{X}$ is symmetric.
 12. Recall that an inverse of the matrix \mathbf{A} (denoted \mathbf{A}^{-1}) is defined by two properties: $\mathbf{A}^{-1}\mathbf{A} = \mathbf{I}$ and $\mathbf{A}\mathbf{A}^{-1} = \mathbf{I}$. Prove that inverses are unique, as follows. Let \mathbf{B} and \mathbf{C} both be inverses of \mathbf{A} . Show that $\mathbf{B} = \mathbf{C}$.
 13. Let \mathbf{X} be an n by p matrix with $n \neq p$. Why is it incorrect to say that $(\mathbf{X}'\mathbf{X})^{-1} = \mathbf{X}^{-1}\mathbf{X}'^{-1}$?
 14. Suppose that the square matrices \mathbf{A} and \mathbf{B} both have inverses. Prove that $(\mathbf{A}\mathbf{B})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}$. You have two things to show.
 15. Let \mathbf{A} be a non-singular square matrix. Prove $(\mathbf{A}^{-1})' = (\mathbf{A}')^{-1}$.
 16. Using Question 15, prove that the inverse of a symmetric matrix is also symmetric.
 17. Let \mathbf{A} be a square matrix with the determinant of \mathbf{A} (denoted $|\mathbf{A}|$) equal to zero. What does this tell you about \mathbf{A}^{-1} ?
 18. Let \mathbf{a} be an $n \times 1$ matrix of constants. How do you know $\mathbf{a}'\mathbf{a} \geq 0$?
 19. In the following, \mathbf{A} and \mathbf{B} are $n \times p$ matrices of constants, \mathbf{C} is $p \times q$, \mathbf{D} is $p \times n$ and a, b, c are scalars. For each statement below, either prove it is true, or prove that it is not true in general by giving a counter-example. Small numerical counter-examples are best. To give an idea of the kind of proof required for most of these, denote element (i, j) of matrix \mathbf{A} by $[a_{i,j}]$.
 - (a) $\mathbf{A} + \mathbf{B} = \mathbf{B} + \mathbf{A}$
 - (b) $a(\mathbf{B} + \mathbf{C}) = a\mathbf{B} + a\mathbf{C}$
 - (c) $\mathbf{A}\mathbf{C} = \mathbf{C}\mathbf{A}$
 - (d) $(\mathbf{A} + \mathbf{B})' = \mathbf{A}' + \mathbf{B}'$
 - (e) $(\mathbf{A}\mathbf{C})' = \mathbf{C}'\mathbf{A}'$
 - (f) $(\mathbf{A} + \mathbf{B})\mathbf{C} = \mathbf{A}\mathbf{C} + \mathbf{B}\mathbf{C}$
 - (g) $(\mathbf{A}\mathbf{D})^{-1} = \mathbf{A}^{-1}\mathbf{D}^{-1}$
 20. Let \mathbf{A} and \mathbf{B} be square matrices of the same size, and \mathbf{A}^{-1} and \mathbf{B}^{-1} both exist. Show $(\mathbf{A}\mathbf{B})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}$.
 21. Let \mathbf{A} be a square symmetric matrix, and \mathbf{A}^{-1} exists. Show that \mathbf{A}^{-1} is also symmetric.
 22. Show that a matrix inverse is unique. That is, let \mathbf{B} and \mathbf{C} both be inverses of \mathbf{A} ; show $\mathbf{B} = \mathbf{C}$.

23. The *trace* of a square matrix is the sum of its diagonal elements; we write $tr(\mathbf{A})$. Let \mathbf{A} be $r \times c$ and \mathbf{B} be $c \times r$. Show $tr(\mathbf{AB}) = tr(\mathbf{BA})$.
24. Let \mathbf{a} be an $n \times 1$ matrix of constants. How do you know $\mathbf{a}'\mathbf{a} \geq 0$?
25. Recall the *spectral decomposition* of a square symmetric matrix (For example, a variance-covariance matrix). Any such matrix Σ can be written as $\Sigma = \mathbf{P}\mathbf{\Lambda}\mathbf{P}'$, where \mathbf{P} is a matrix whose columns are the (orthonormal) eigenvectors of Σ , $\mathbf{\Lambda}$ is a diagonal matrix of the corresponding (non-negative) eigenvalues, and $\mathbf{P}'\mathbf{P} = \mathbf{P}\mathbf{P}' = \mathbf{I}$.
- (a) Let Σ be a square symmetric matrix with eigenvalues that are all strictly positive.
- What is $\mathbf{\Lambda}^{-1}$?
 - Show $\Sigma^{-1} = \mathbf{P}\mathbf{\Lambda}^{-1}\mathbf{P}'$
- (b) Let Σ be a square symmetric matrix, and this time some of the eigenvalues might be zero.
- What do you think $\mathbf{\Lambda}^{1/2}$ might be?
 - Define $\Sigma^{1/2}$ as $\mathbf{P}\mathbf{\Lambda}^{1/2}\mathbf{P}'$. Show $\Sigma^{1/2}$ is symmetric.
 - Show $\Sigma^{1/2}\Sigma^{1/2} = \Sigma$.
- (c) A matrix Σ is said to be *non-negative definite* (our text calls it positive semi-definite) if $\mathbf{a}'\Sigma\mathbf{a} \geq 0$ for all vectors \mathbf{a} . Show that any symmetric matrix must be non-negative definite. Can't; this is wrong! Fix!
- (d) The (square) matrix Σ is said to be *positive definite* if $\mathbf{a}'\Sigma\mathbf{a} > 0$ for all vectors $\mathbf{a} \neq \mathbf{0}$. Show that the eigenvalues of a symmetric positive definite matrix are all strictly positive. Hint: the \mathbf{a} you want is an eigenvector.
- (e) Let Σ be a symmetric, positive definite matrix. Putting together a couple of results you have proved above, establish that Σ^{-1} exists.
26. Using the spectral decomposition (A.6) and $tr(\mathbf{AB}) = tr(\mathbf{BA})$, show that the trace of a square symmetric matrix is the sum of its eigenvalues.
27. Let \mathbf{X} be an $n \times p$ matrix of constants. Recall the definition of linear independence. The columns of \mathbf{X} are said to be *linearly dependent* if there exists $\mathbf{v} \neq \mathbf{0}$ with $\mathbf{X}\mathbf{v} = \mathbf{0}$. We will say that the columns of \mathbf{X} are *linearly independent* if $\mathbf{X}\mathbf{v} = \mathbf{0}$ implies $\mathbf{v} = \mathbf{0}$.
- Show that if the columns of \mathbf{X} are linearly dependent, then the columns of $\mathbf{X}'\mathbf{X}$ are also linearly dependent.
 - Show that if the columns of \mathbf{X} are linearly dependent, then the *rows* of $\mathbf{X}'\mathbf{X}$ are linearly dependent.
 - Show that if the columns of \mathbf{X} are linearly independent, then the columns of $\mathbf{X}'\mathbf{X}$ are also linearly independent.

- (d) Show that if $(\mathbf{X}'\mathbf{X})^{-1}$ exists, then the columns of \mathbf{X} are linearly independent.
- (e) Show that if the columns of \mathbf{X} are linearly independent, then $\mathbf{X}'\mathbf{X}$ is positive definite. Does this imply the existence of $(\mathbf{X}'\mathbf{X})^{-1}$? Locate the rule in the text, and answer Yes or No.

A.3 Random Vectors and Matrices

A *random matrix* is just a matrix of random variables. Their joint probability distribution is the distribution of the random matrix. Random matrices with just one column (say, p) may be called *random vectors*.

A.3.1 Expected Value and Variance-Covariance

Expected Value

The expected value of a matrix is defined as the matrix of expected values. Denoting the $p \times c$ random matrix \mathbf{X} by $[X_{i,j}]$,

$$E(\mathbf{X}) = [E(X_{i,j})].$$

Immediately we have natural properties like

$$\begin{aligned} E(\mathbf{X} + \mathbf{Y}) &= E([X_{i,j}] + [Y_{i,j}]) \\ &= [E(X_{i,j} + Y_{i,j})] \\ &= [E(X_{i,j}) + E(Y_{i,j})] \\ &= [E(X_{i,j})] + [E(Y_{i,j})] \\ &= E(\mathbf{X}) + E(\mathbf{Y}). \end{aligned}$$

Let $\mathbf{A} = [a_{i,j}]$ be an $r \times p$ matrix of constants, while \mathbf{X} is still a $p \times c$ random matrix. Then

$$\begin{aligned} E(\mathbf{A}\mathbf{X}) &= E\left(\left[\sum_{k=1}^p a_{i,k}X_{k,j}\right]\right) \\ &= \left[E\left(\sum_{k=1}^p a_{i,k}X_{k,j}\right)\right] \\ &= \left[\sum_{k=1}^p a_{i,k}E(X_{k,j})\right] \\ &= \mathbf{A}E(\mathbf{X}). \end{aligned}$$

Similar calculations lead $E(\mathbf{X}\mathbf{B}) = E(\mathbf{X})\mathbf{B}$, where \mathbf{B} is a matrix of constants. This yields the useful formula

$$E(\mathbf{A}\mathbf{X}\mathbf{B}) = \mathbf{A}E(\mathbf{X})\mathbf{B}. \quad (\text{A.8})$$

Variance-Covariance Matrices

Let \mathbf{X} be a $p \times 1$ random vector with $E(\mathbf{X}) = \boldsymbol{\mu}$. The *variance-covariance matrix* of \mathbf{X} (sometimes just called the *covariance matrix*), denoted by $V(\mathbf{X})$, is defined as

$$V(\mathbf{X}) = E \{ (\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})' \}. \quad (\text{A.9})$$

The covariance matrix $V(\mathbf{X})$ is a $p \times p$ matrix of constants. To see exactly what it is, suppose $p = 3$. Then

$$\begin{aligned} V(\mathbf{X}) &= E \left\{ \begin{bmatrix} X_1 - \mu_1 \\ X_2 - \mu_2 \\ X_3 - \mu_3 \end{bmatrix} \begin{bmatrix} X_1 - \mu_1 & X_2 - \mu_2 & X_3 - \mu_3 \end{bmatrix} \right\} \\ &= E \left\{ \begin{bmatrix} (X_1 - \mu_1)^2 & (X_1 - \mu_1)(X_2 - \mu_2) & (X_1 - \mu_1)(X_3 - \mu_3) \\ (X_2 - \mu_2)(X_1 - \mu_1) & (X_2 - \mu_2)^2 & (X_2 - \mu_2)(X_3 - \mu_3) \\ (X_3 - \mu_3)(X_1 - \mu_1) & (X_3 - \mu_3)(X_2 - \mu_2) & (X_3 - \mu_3)^2 \end{bmatrix} \right\} \\ &= \begin{bmatrix} E\{(X_1 - \mu_1)^2\} & E\{(X_1 - \mu_1)(X_2 - \mu_2)\} & E\{(X_1 - \mu_1)(X_3 - \mu_3)\} \\ E\{(X_2 - \mu_2)(X_1 - \mu_1)\} & E\{(X_2 - \mu_2)^2\} & E\{(X_2 - \mu_2)(X_3 - \mu_3)\} \\ E\{(X_3 - \mu_3)(X_1 - \mu_1)\} & E\{(X_3 - \mu_3)(X_2 - \mu_2)\} & E\{(X_3 - \mu_3)^2\} \end{bmatrix} \\ &= \begin{bmatrix} V(X_1) & Cov(X_1, X_2) & Cov(X_1, X_3) \\ Cov(X_1, X_2) & V(X_2) & Cov(X_2, X_3) \\ Cov(X_1, X_3) & Cov(X_2, X_3) & V(X_3) \end{bmatrix}. \end{aligned}$$

So, the covariance matrix $V(\mathbf{X})$ is a $p \times p$ symmetric matrix with variances on the main diagonal and covariances on the off-diagonals.

The matrix of covariances between two random vectors may also be written in a convenient way. Let \mathbf{X} be a $p \times 1$ random vector with $E(\mathbf{X}) = \boldsymbol{\mu}_x$ and let \mathbf{Y} be a $q \times 1$ random vector with $E(\mathbf{Y}) = \boldsymbol{\mu}_y$. The $p \times q$ matrix of covariances between the elements of \mathbf{X} and the elements of \mathbf{Y} is

$$C(\mathbf{X}, \mathbf{Y}) = E \{ (\mathbf{X} - \boldsymbol{\mu}_x)(\mathbf{Y} - \boldsymbol{\mu}_y)' \}.$$

The following rule is analogous to $Var(aX) = a^2 Var(X)$ for scalars. Let \mathbf{X} be a $p \times 1$ random vector with $E(\mathbf{X}) = \boldsymbol{\mu}$ and $V(\mathbf{X}) = \boldsymbol{\Sigma}$, while $\mathbf{A} = [a_{i,j}]$ is an $r \times p$ matrix of constants. Then

$$\begin{aligned} V(\mathbf{AX}) &= E \{ (\mathbf{AX} - \mathbf{A}\boldsymbol{\mu})(\mathbf{AX} - \mathbf{A}\boldsymbol{\mu})' \} \\ &= E \{ \mathbf{A}(\mathbf{X} - \boldsymbol{\mu})(\mathbf{A}(\mathbf{X} - \boldsymbol{\mu}))' \} \\ &= E \{ \mathbf{A}(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})' \mathbf{A}' \} \\ &= \mathbf{A} E \{ (\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})' \} \mathbf{A}' \\ &= \mathbf{A} V(\mathbf{X}) \mathbf{A}' \\ &= \mathbf{A} \boldsymbol{\Sigma} \mathbf{A}' \end{aligned} \quad (\text{A.10})$$

For scalars, $Var(X + b) = Var(X)$, and the same applies to vectors. Covariances are also unaffected by adding a constant; this amounts to shifting the whole joint distribution by a fixed amount, which has no effect on relationships among variables. So, the following rule is “obvious.” Let \mathbf{X} be a $p \times 1$ random vector with $E(\mathbf{X}) = \boldsymbol{\mu}$ and let \mathbf{b} be a $p \times 1$ vector of constants. Then $V(\mathbf{X} + \mathbf{b}) = V(\mathbf{X})$. To see this, note $E(\mathbf{X} + \mathbf{b}) = \boldsymbol{\mu} + \mathbf{b}$ and write

$$\begin{aligned} V(\mathbf{X} + \mathbf{b}) &= E\{(\mathbf{X} + \mathbf{b} - (\boldsymbol{\mu} + \mathbf{b}))(\mathbf{X} + \mathbf{b} - (\boldsymbol{\mu} + \mathbf{b}))'\} \\ &= E\{(\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})'\} \\ &= V(\mathbf{X}) \end{aligned} \tag{A.11}$$

A similar rule applies to $C(\mathbf{X} + \mathbf{b}, \mathbf{Y} + \mathbf{c})$. A direct calculation is not even necessary, though it is a valuable exercise. Think of stacking \mathbf{X} and \mathbf{Y} one on top of another, to form a bigger random vector. Then,

$$V \begin{bmatrix} \mathbf{X} \\ \mathbf{Y} \end{bmatrix} = \begin{bmatrix} V(\mathbf{X}) & C(\mathbf{X}, \mathbf{Y}) \\ C(\mathbf{X}, \mathbf{Y})' & V(\mathbf{Y}) \end{bmatrix}.$$

This is an example of a *partitioned matrix* – a matrix of matrices. At any rate, it is clear from (A.11) that adding a stack of constant vectors to the stack of random vectors has no effect upon the (partitioned) covariance matrix, and in particular no effect upon $C(\mathbf{X}, \mathbf{Y})$.

The Centering Rule

Here is a rule that will save a huge amount of effort in this course. *When calculating variances and covariances of linear combinations of random vectors, it is safe to pretend that all expected values and intercepts equal zero, even when they definitely do not.* A couple of examples will come first, followed by a general formulation.

First, combining formulas (A.10) and (A.11), it is clear that

$$V(\mathbf{A}\mathbf{X} + \mathbf{b}) = \mathbf{A}V(\mathbf{X})\mathbf{A}'.$$

Now the derivation will be repeated assuming $E(\mathbf{X}) = \mathbf{b} = \mathbf{0}$. This restriction has no effect upon the answer.

$$\begin{aligned} V(\mathbf{A}\mathbf{X} + \mathbf{b}) &= E\{\mathbf{A}\mathbf{X}(\mathbf{A}\mathbf{X})'\} \\ &= E\{\mathbf{A}\mathbf{X}\mathbf{X}'\mathbf{A}'\} \\ &= \mathbf{A}E\{\mathbf{X}\mathbf{X}'\}\mathbf{A}' \\ &= \mathbf{A}V(\mathbf{X})\mathbf{A}' \end{aligned} \tag{A.12}$$

Notice how quick that was, in addition to yielding the right answer. Also notice the final step, in which $E\{\mathbf{X}\mathbf{X}'\}$ was re-expressed as $V(\mathbf{X})$.

Now $C(\mathbf{A}\mathbf{X} + \mathbf{b}, \mathbf{C}\mathbf{Y} + \mathbf{d})$ will be calculated two ways. First, we will let $E(\mathbf{X}) = \boldsymbol{\mu}_x$, $E(\mathbf{Y}) = \boldsymbol{\mu}_y$, \mathbf{b} and \mathbf{d} all be possibly non-zero. Then we will do the calculation again assuming they are all zero, and observe that the answer is the same.

$$\begin{aligned}
C(\mathbf{AX} + \mathbf{b}, \mathbf{CY} + \mathbf{d}) &= E\{(\mathbf{AX} + \mathbf{b} - (\mathbf{A}\boldsymbol{\mu}_x + \mathbf{b}))(\mathbf{CY} + \mathbf{d} - (\mathbf{C}\boldsymbol{\mu}_y + \mathbf{d}))'\} \\
&= E\{(\mathbf{AX} - \mathbf{A}\boldsymbol{\mu}_x)(\mathbf{CY} - \mathbf{C}\boldsymbol{\mu}_y)'\} \\
&= E\{\mathbf{A}(\mathbf{X} - \boldsymbol{\mu}_x)(\mathbf{C}(\mathbf{Y} - \boldsymbol{\mu}_y))'\} \\
&= E\{\mathbf{A}(\mathbf{X} - \boldsymbol{\mu}_x)(\mathbf{Y} - \boldsymbol{\mu}_y)'\mathbf{C}'\} \\
&= \mathbf{A}E\{(\mathbf{X} - \boldsymbol{\mu}_x)(\mathbf{Y} - \boldsymbol{\mu}_y)'\}\mathbf{C}' \\
&= \mathbf{A}C(\mathbf{X}, \mathbf{Y})\mathbf{C}'
\end{aligned}$$

Now let $\boldsymbol{\mu}_x = \boldsymbol{\mu}_y = \mathbf{b} = \mathbf{d} = \mathbf{0}$, and do the calculation again.

$$\begin{aligned}
C(\mathbf{AX} + \mathbf{b}, \mathbf{CY} + \mathbf{d}) &= E\{(\mathbf{AX}(\mathbf{CY})')\} \\
&= E\{(\mathbf{AXY}'\mathbf{C}')\} \\
&= \mathbf{A}E\{\mathbf{XY}'\}\mathbf{C}' \\
&= \mathbf{A}C(\mathbf{X}, \mathbf{Y})\mathbf{C}'
\end{aligned} \tag{A.13}$$

Notice that the two answers are identical, and also notice how much simpler the calculation is when all expected values and intercepts are assumed zero. Again, the final step is to re-express $E\{\mathbf{XY}'\}$ as $C(\mathbf{X}, \mathbf{Y})$.

To state the centering rule precisely, let $\mathbf{X}_1, \dots, \mathbf{X}_m$ be random vectors, and let $\mathbf{A}_1, \dots, \mathbf{A}_m$ be constant matrices, each having r rows. The number of columns in these constant matrices may vary, as long as the product $\mathbf{A}_j\mathbf{X}_j$ can be formed for $j = 1, \dots, m$. Finally, let \mathbf{b} be an $r \times 1$ matrix of constants. Then a *linear combination* of $\mathbf{X}_1, \dots, \mathbf{X}_m$ is defined by

$$\mathbf{L} = \mathbf{A}_1\mathbf{X}_1 + \dots + \mathbf{A}_m\mathbf{X}_m + \mathbf{b}. \tag{A.14}$$

The *Centering Rule* says that to calculate $V(\mathbf{L})$ or $C(\mathbf{L}_1, \mathbf{L}_2)$, one may follow these steps:

1. Set all expected values to zero, and also set constants like \mathbf{b} (intercepts) to zero.
2. Calculate $V(\mathbf{L}) = E(\mathbf{L}\mathbf{L}')$ or $C(\mathbf{L}_1, \mathbf{L}_2) = E(\mathbf{L}_1\mathbf{L}_2')$.
3. Re-write expressions like $E(\mathbf{X}_j\mathbf{X}_j')$ as $V(\mathbf{X}_j)$. Re-write expressions like $E(\mathbf{X}_i\mathbf{X}_j')$ as $C(\mathbf{X}_i, \mathbf{X}_j)$.

Sometimes people write things like “Assuming without loss of generality that $E(X) = 0 \dots$ ” This confused me deeply the first time I saw it as a student, but it’s right. There is no loss of generality because you would get the same answer by assuming the expected value to be non-zero. But it only applies to linear combinations. Let formula (A.14) be your guide.

As another example of how helpful this strategy can be, consider the calculation of $V(\mathbf{AX} + \mathbf{BY})$, where \mathbf{X} is a $p \times 1$ random vector with $E(\mathbf{X}) = \boldsymbol{\mu}_x$, \mathbf{Y} is a $q \times 1$ random vector with $E(\mathbf{Y}) = \boldsymbol{\mu}_y$, \mathbf{A} is an $r \times p$ matrix of constants, and \mathbf{B} is an $r \times q$ matrix of constants. So, \mathbf{X} and \mathbf{Y} don’t have to be the same size, but \mathbf{AX} and \mathbf{BY} are both $r \times 1$ so they can be added.

Assuming without loss of generality that $\boldsymbol{\mu}_x = \boldsymbol{\mu}_y = \mathbf{0}$ (see why this is okay?),

$$\begin{aligned}
V(\mathbf{AX} + \mathbf{BY}) &= E(\mathbf{AX} + \mathbf{BY})(\mathbf{AX} + \mathbf{BY})' \\
&= E(\mathbf{AX} + \mathbf{BY})(\mathbf{X}'\mathbf{A}' + \mathbf{Y}'\mathbf{B}') \\
&= E(\mathbf{AXX}'\mathbf{A}' + \mathbf{AXY}'\mathbf{B}' + \mathbf{BYX}'\mathbf{A}' + \mathbf{BYY}'\mathbf{B}') \\
&= \mathbf{AE}(\mathbf{XX}')\mathbf{A}' + \mathbf{AE}(\mathbf{XY}')\mathbf{B}' + \mathbf{BE}(\mathbf{YX}')\mathbf{A}' + \mathbf{BE}(\mathbf{YY}')\mathbf{B}' \\
&= \mathbf{AV}(\mathbf{X})\mathbf{A}' + \mathbf{AC}(\mathbf{X}, \mathbf{Y})\mathbf{B}' + \mathbf{BC}(\mathbf{Y}, \mathbf{X})\mathbf{A}' + \mathbf{BV}(\mathbf{Y})\mathbf{B}' \\
&= \mathbf{AV}(\mathbf{X})\mathbf{A}' + \mathbf{BV}(\mathbf{Y})\mathbf{B}' + \mathbf{AC}(\mathbf{X}, \mathbf{Y})\mathbf{B}' + (\mathbf{AC}(\mathbf{X}, \mathbf{Y})\mathbf{B}')'.
\end{aligned}$$

This is the analogue of $Var(aX + bY) = a^2Var(X) + b^2Var(Y) + 2abCov(X, Y)$. The calculation would have been much uglier if we carried $\mathbf{A}\boldsymbol{\mu}_x + \mathbf{B}\boldsymbol{\mu}_y$ along, but the result would have been the same.

It may be worth pointing out that the trick of assuming expected values zero applies to scalar variance-covariance calculations too. For example, let

$$\begin{aligned}
Y &= \beta_0 + \beta_1X + \epsilon \\
W &= \lambda_1X + \delta_1 \\
V &= \lambda_2Y + \delta_2,
\end{aligned}$$

where $\beta_0, \beta_1, \lambda_1$ and λ_2 are constants, $Var(X) = \phi$, $Var(\epsilon) = \sigma^2$, $Var(\delta_1) = \omega_1$, $Var(\delta_2) = \omega_2$, and X, ϵ, δ_1 and δ_2 are all independent. The job is to calculate $Cov(W, V)$. Assuming without loss of generality that $E(X) = E(\epsilon) = E(\delta_1) = E(\delta_2) = \beta_0 = 0$,

$$\begin{aligned}
Cov(W, V) &= E\{WV\} \\
&= E\{(\lambda_1X + \delta_1)(\lambda_2Y + \delta_2)\} \\
&= E\{(\lambda_1X + \delta_1)(\lambda_2(\beta_1X + \epsilon) + \delta_2)\} \\
&= E\{\lambda_1\lambda_2\beta_1X^2 + \lambda_1\lambda_2X\epsilon + \lambda_1X\delta_2 + \lambda_2\beta_1\delta_1X + \delta_1\delta_2\} \\
&= \lambda_1\lambda_2\beta_1E\{X^2\} + \lambda_1\lambda_2E\{X\}E\{\epsilon\} + \lambda_1E\{X\}E\{\delta_2\} \\
&\quad + \lambda_2\beta_1E\{\delta_1\}E\{X\} + E\{\delta_1\}E\{\delta_2\} \\
&= \lambda_1\lambda_2\beta_1E\{X^2\} + 0 + 0 + 0 + 0 \\
&= \lambda_1\lambda_2\beta_1Var(X) \\
&= \lambda_1\lambda_2\beta_1\phi.
\end{aligned}$$

This calculation looks messier than it really is, because every little detail is shown. To repeat the lesson one last time, calculating variances and covariances of linear combinations is significantly easier when intercepts and expected values are assumed to be zero, and the final answer is the same — provided the final substitution step is not omitted.

Exercises A.3.1

1. Let \mathbf{X} and \mathbf{Y} be random matrices of the same dimensions. Show $E(\mathbf{X} + \mathbf{Y}) = E(\mathbf{X}) + E(\mathbf{Y})$. Recall the definition $E(\mathbf{Z}) = [E(Z_{i,j})]$.

2. Let \mathbf{X} be a random matrix, and \mathbf{B} be a matrix of constants. Show $E(\mathbf{XB}) = E(\mathbf{X})\mathbf{B}$. Recall the definition $\mathbf{AB} = [\sum_k a_{i,k}b_{k,j}]$.
3. If the $p \times 1$ random vector \mathbf{X} has variance-covariance matrix Σ and \mathbf{A} is an $m \times p$ matrix of constants, prove that the variance-covariance matrix of \mathbf{AX} is $\mathbf{A}\Sigma\mathbf{A}'$. Start with the definition of a variance-covariance matrix:

$$V(\mathbf{Z}) = E(\mathbf{Z} - \boldsymbol{\mu}_z)(\mathbf{Z} - \boldsymbol{\mu}_z)'$$

4. If the $p \times 1$ random vector \mathbf{X} has mean $\boldsymbol{\mu}$ and variance-covariance matrix Σ , show $\Sigma = E(\mathbf{XX}') - \boldsymbol{\mu}\boldsymbol{\mu}'$.
5. Let the $p \times 1$ random vector \mathbf{X} have mean $\boldsymbol{\mu}$ and variance-covariance matrix Σ , and let \mathbf{c} be a $p \times 1$ vector of constants. Find $V(\mathbf{X} + \mathbf{c})$. Show your work. This is important because it tells us we can always pretend the mean equals zero when calculating covariance matrices.
6. Let \mathbf{X} be a $p \times 1$ random vector with mean $\boldsymbol{\mu}_x$ and variance-covariance matrix Σ_x , and let \mathbf{Y} be a $q \times 1$ random vector with mean $\boldsymbol{\mu}_y$ and variance-covariance matrix Σ_y . Recall that $C(\mathbf{X}, \mathbf{Y})$ is the $p \times q$ matrix $C(\mathbf{X}, \mathbf{Y}) = E((\mathbf{X} - \boldsymbol{\mu}_x)(\mathbf{Y} - \boldsymbol{\mu}_y)')$.
 - (a) What is the (i, j) element of $C(\mathbf{X}, \mathbf{Y})$?
 - (b) Find an expression for $V(\mathbf{X} + \mathbf{Y})$ in terms of Σ_x , Σ_y and $C(\mathbf{X}, \mathbf{Y})$. Show your work.
 - (c) Simplify further for the special case where $Cov(X_i, Y_j) = 0$ for all i and j .
 - (d) Let \mathbf{c} be a $p \times 1$ vector of constants and \mathbf{d} be a $q \times 1$ vector of constants. Find $C(\mathbf{X} + \mathbf{c}, \mathbf{Y} + \mathbf{d})$. Show your work.

A.3.2 The Multivariate Normal Distribution

The $p \times 1$ random vector \mathbf{X} is said to have a *multivariate normal distribution*, and we write $\mathbf{X} \sim N(\boldsymbol{\mu}, \Sigma)$, if \mathbf{X} has (joint) density

$$f(\mathbf{x}) = \frac{1}{|\Sigma|^{\frac{1}{2}}(2\pi)^{\frac{p}{2}}} \exp \left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})'\Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu}) \right],$$

where $\boldsymbol{\mu}$ is $p \times 1$ and Σ is $p \times p$ symmetric and positive definite. Positive definite means that for any non-zero $p \times 1$ vector \mathbf{a} , we have $\mathbf{a}'\Sigma\mathbf{a} > 0$.

- Since the one-dimensional random variable $Y = \sum_{i=1}^p a_i X_i$ may be written as $Y = \mathbf{a}'\mathbf{X}$ and $Var(Y) = V(\mathbf{a}'\mathbf{X}) = \mathbf{a}'\Sigma\mathbf{a}$, it is natural to require that Σ be positive definite. All it means is that every non-zero linear combination of \mathbf{X} values has a positive variance.
- Σ positive definite is equivalent to Σ^{-1} positive definite.

The multivariate normal reduces to the univariate normal when $p = 1$. Other properties of the multivariate normal include the following.

1. $E(\mathbf{X}) = \boldsymbol{\mu}$
2. $V(\mathbf{X}) = \boldsymbol{\Sigma}$
3. If \mathbf{c} is a vector of constants, $\mathbf{X} + \mathbf{c} \sim N(\mathbf{c} + \boldsymbol{\mu}, \boldsymbol{\Sigma})$
4. If \mathbf{A} is a matrix of constants, $\mathbf{A}\mathbf{X} \sim N(\mathbf{A}\boldsymbol{\mu}, \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}')$
5. All the marginals (dimension less than p) of \mathbf{X} are (multivariate) normal, but it is possible in theory to have a collection of univariate normals whose joint distribution is not multivariate normal.
6. For the multivariate normal, zero covariance implies independence. The multivariate normal is the only continuous distribution with this property.
7. The random variable $(\mathbf{X} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\mathbf{X} - \boldsymbol{\mu})$ has a chi-square distribution with p degrees of freedom.
8. After a bit of work, the multivariate normal likelihood may be written as

$$L(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = |\boldsymbol{\Sigma}|^{-n/2} (2\pi)^{-np/2} \exp -\frac{n}{2} \left\{ \text{tr}(\widehat{\boldsymbol{\Sigma}}\boldsymbol{\Sigma}^{-1}) + (\bar{\mathbf{x}} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\bar{\mathbf{x}} - \boldsymbol{\mu}) \right\}, \quad (\text{A.15})$$

where $\widehat{\boldsymbol{\Sigma}} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})'$ is the sample variance-covariance matrix (it would be unbiased if divided by $n - 1$).

Exercises A.3.2

1. Let X_1 be Normal(μ_1, σ_1^2), and X_2 be Normal(μ_2, σ_2^2), independent of X_1 . What is the joint distribution of $Y_1 = X_1 + X_2$ and $Y_2 = X_1 - X_2$? What is required for Y_1 and Y_2 to be independent?
2. Let $\mathbf{X} = (X_1, X_2, X_3)'$ be multivariate normal with

$$\boldsymbol{\mu} = \begin{bmatrix} 1 \\ 0 \\ 6 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Sigma} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Let $Y_1 = X_1 + X_2$ and $Y_2 = X_2 + X_3$. Find the joint distribution of Y_1 and Y_2 .

3. Let $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$, where \mathbf{X} is an $n \times p$ matrix of known constants, $\boldsymbol{\beta}$ is a $p \times 1$ vector of unknown constants, and $\boldsymbol{\epsilon}$ is multivariate normal with mean zero and covariance matrix $\sigma^2 \mathbf{I}_n$, where $\sigma^2 > 0$ is a constant. In the following, it may be helpful to recall that $(\mathbf{A}^{-1})' = (\mathbf{A}')^{-1}$.

- (a) What is the distribution of \mathbf{Y} ?

- (b) The maximum likelihood estimate (MLE) of β is $\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$. What is the distribution of $\hat{\beta}$? Show the calculations.
- (c) Let $\hat{\mathbf{Y}} = \mathbf{X}\hat{\beta}$. What is the distribution of $\hat{\mathbf{Y}}$? Show the calculations.
- (d) Let the vector of residuals $\mathbf{e} = (\mathbf{Y} - \hat{\mathbf{Y}})$. What is the distribution of \mathbf{e} ? Show the calculations. Simplify both the expected value (which is zero) and the covariance matrix.
4. Prove the formula (A.15) for the multivariate normal likelihood. Show all the calculations.
5. Prove that for *any* positive definite Σ , the likelihood (A.15) is maximized when $\bar{\mathbf{x}} = \boldsymbol{\mu}$. How do you know this maximum must be unique? Cite the necessary matrix facts from Section A.2 of this Appendix.

A.4 Maximum Likelihood

We will use the following framework for estimation and inference. The data are D_1, \dots, D_n (the letter D stands for data). The distribution of these independent and identically distributed random variables depends on the parameter θ , which is an element of the parameter space Θ . That is,

$$D_1, \dots, D_n \stackrel{i.i.d.}{\sim} P_\theta, \theta \in \Theta.$$

Both the data values and the parameter may be vectors, even though they are not written in boldface. The data have likelihood function

$$L(\theta) = \prod_{i=1}^n f(d_i; \theta),$$

where $f(d_i; \theta)$ is the density or probability mass function evaluated at d_i .

A.4.1 Some Very Basic Math

I have noticed that a major obstacle for many students when doing likelihood calculations is a set of basic mathematical operations they actually know. But the mechanics are rusty, or the notation used in Statistics is troublesome. So, with sincere apologies to those who don't need this, here are some basic rules. Feel free to skip this.

- The distributive law: $a(b + c) = ab + ac$. You may see this in a form like

$$\theta \sum_{i=1}^n x_i = \sum_{i=1}^n \theta x_i$$

- Power of a product is the product of powers: $(ab)^c = a^c b^c$. You may see this in a form like

$$\left(\prod_{i=1}^n x_i \right)^\alpha = \prod_{i=1}^n x_i^\alpha$$

- Multiplication is addition of exponents: $a^b a^c = a^{b+c}$. You may see this in a form like

$$\prod_{i=1}^n \theta e^{-\theta x_i} = \theta^n \exp\left(-\theta \sum_{i=1}^n x_i\right)$$

- Powering is multiplication of exponents: $(a^b)^c = a^{bc}$. You may see this in a form like

$$\left(e^{\mu t + \frac{1}{2}\sigma^2 t^2}\right)^n = e^{n\mu t + \frac{1}{2}n\sigma^2 t^2}$$

- Log of a product is sum of logs: $\ln(ab) = \ln(a) + \ln(b)$. You may see this in a form like

$$\ln \prod_{i=1}^n x_i = \sum_{i=1}^n \ln x_i$$

- Log of a power is the exponent times the log: $\ln(a^b) = b \ln(a)$. You may see this in a form like

$$\ln(\theta^n) = n \ln \theta$$

- The log is the inverse of the exponential function: $\ln(e^a) = a$. You may see this in a form like

$$\ln \left(\theta^n \exp\left(-\theta \sum_{i=1}^n x_i\right) \right) = n \ln \theta - \theta \sum_{i=1}^n x_i$$

Exercises A.4.1

1. Choose the correct answer.

(a) $\prod_{i=1}^n e^{x_i} =$

i. $\exp\left(\prod_{i=1}^n x_i\right)$

ii. e^{nx_i}

iii. $\exp\left(\sum_{i=1}^n x_i\right)$

(b) $\prod_{i=1}^n \lambda e^{-\lambda x_i} =$

i. $\lambda e^{-\lambda^n x_i}$

ii. $\lambda^n e^{-\lambda n x_i}$

iii. $\lambda^n \exp\left(-\lambda \sum_{i=1}^n x_i\right)$

iv. $\lambda^n \exp\left(-n\lambda \sum_{i=1}^n x_i\right)$

v. $\lambda^n \exp\left(-\lambda^n \sum_{i=1}^n x_i\right)$

(c) $\prod_{i=1}^n a_i^b =$
 i. na^b
 ii. a^{nb}
 iii. $(\prod_{i=1}^n a_i)^b$

(d) $\prod_{i=1}^n a^{b_i} =$
 i. na^{b_i}
 ii. a^{nb_i}
 iii. $\sum_{i=1}^n a^{b_i}$
 iv. $a^{\prod_{i=1}^n b_i}$
 v. $a^{\sum_{i=1}^n b_i}$

(e) $(e^{\lambda(e^t-1)})^n =$
 i. $ne^{\lambda(e^t-1)}$
 ii. $e^{n\lambda(e^t-1)}$
 iii. $e^{\lambda(e^{nt}-1)}$
 iv. $e^{n\lambda(e^t-n)}$

(f) $(\prod_{i=1}^n e^{-\lambda x_i})^2 =$
 i. $e^{-2n\lambda x_i}$
 ii. $e^{-2\lambda \sum_{i=1}^n x_i}$
 iii. $2e^{-\lambda \sum_{i=1}^n x_i}$

2. True, or False?

(a) $\sum_{i=1}^n \frac{1}{x_i} = \frac{1}{\sum_{i=1}^n x_i}$
 (b) $\prod_{i=1}^n \frac{1}{x_i} = \frac{1}{\prod_{i=1}^n x_i}$
 (c) $\frac{a}{b+c} = \frac{a}{b} + \frac{a}{c}$
 (d) $\ln(a+b) = \ln(a) + \ln(b)$
 (e) $e^{a+b} = e^a + e^b$
 (f) $e^{a+b} = e^a e^b$
 (g) $e^{ab} = e^a e^b$
 (h) $\prod_{i=1}^n (x_i + y_i) = \prod_{i=1}^n x_i + \prod_{i=1}^n y_i$
 (i) $\ln(\prod_{i=1}^n a_i^b) = b \sum_{i=1}^n \ln(a_i)$
 (j) $\sum_{i=1}^n \prod_{j=1}^n a_j = n \prod_{j=1}^n a_j$
 (k) $\sum_{i=1}^n \prod_{j=1}^n a_i = \sum_{i=1}^n a_i^n$
 (l) $\sum_{i=1}^n \prod_{j=1}^n a_{i,j} = \prod_{j=1}^n \sum_{i=1}^n a_{i,j}$

3. Simplify as much as possible.

- (a) $\ln \prod_{i=1}^n \theta^{x_i} (1 - \theta)^{1-x_i}$
- (b) $\ln \prod_{i=1}^n \binom{m}{x_i} \theta^{x_i} (1 - \theta)^{m-x_i}$
- (c) $\ln \prod_{i=1}^n \frac{e^{-\lambda} \lambda^{x_i}}{x_i!}$
- (d) $\ln \prod_{i=1}^n \theta (1 - \theta)^{x_i-1}$
- (e) $\ln \prod_{i=1}^n \frac{1}{\theta} e^{-x_i/\theta}$
- (f) $\ln \prod_{i=1}^n \frac{1}{\beta^\alpha \Gamma(\alpha)} e^{-x_i/\beta} x_i^{\alpha-1}$
- (g) $\ln \prod_{i=1}^n \frac{1}{2^{\nu/2} \Gamma(\nu/2)} e^{-x_i/2} x_i^{\nu/2-1}$
- (h) $\ln \prod_{i=1}^n \frac{1}{\sigma \sqrt{2\pi}} e^{-\frac{(x_i-\mu)^2}{2\sigma^2}}$
- (i) $\prod_{i=1}^n \frac{1}{\beta-\alpha} I(\alpha \leq x_i \leq \beta)$ (Express in terms of the minimum and maximum y_1 and y_n .)

A.4.2 Point Estimation

Let $\hat{\theta}$ denote the usual Maximum Likelihood Estimate (MLE). That is, it is the parameter value for which the likelihood function is greatest, over all $\theta \in \Theta$. Because the log is an increasing function, maximizing the likelihood is equivalent to maximizing the log likelihood, which will be denoted

$$\ell(\theta) = \ln L(\theta).$$

In elementary situations where the support of the distribution does not depend on the parameter, you get the MLE by closing your eyes, differentiating the log likelihood, setting the derivative to zero, and solving for θ . Then if you are being careful, you carry out the second derivative test; if $\ell''(\hat{\theta}) < 0$, the log likelihood is concave down at your answer, and you have found the maximum.

This is all very well, and you will be asked to do it in some of the exercises. But in this course, as in much of applied statistics, you will find that you can write the log likelihood and differentiate it easily enough, but when you set the derivatives to zero, you obtain a set of equations that are impossible to solve explicitly. This means that the problem needs to be solved numerically. That is, you use a computer to calculate the value of the log likelihood for a set of parameter values, and you search until you have found the biggest one.

But how do you search? It's easy in one or two dimensions, but structural equation models can easily involve dozens, scores or even hundreds of parameters. It's a bit like being dropped by helicopter onto a mountain range, and asked to find the highest peak blindfolded. All you can do is walk uphill. The gradient is the direction of steepest increase, so walk that way. How big a step should you take? That's a good question. When you come to a place where the surface is level, or approximately level, stop. How level is level enough? That's another good question. Once you find a "level" place, you can check to see if the surface is concave down there. If so, you're at a maximum. Is it

the global maximum (the real MLE), or just a local maximum? It's usually impossible to tell for sure. You can get the helicopter to drop you in several different places fairly far apart, and if you always arrive at the same maximum you will feel more confident of your answer. But it could still be just a local maximum that is easy to reach. The main thing to observe is that where you start is *very* important. Another point is that for realistically big problems, you need high-grade, professionally written software.

The following example is one that you can do by hand, though maybe not with your eyes closed. But it will serve to illustrate the basic ideas of numerical maximum likelihood.

Example A.4.2.1

Let D_1, \dots, D_n be a random sample from a normal distribution with mean θ and variance θ^2 . A sample of size 50 yields:

```

5.85 -15.02 -13.24 -1.63 -0.07 -2.40 -3.02 -3.19 -5.16  0.79 -1.03 -10.69
-12.96 -4.55  0.57 -7.94 -6.80  2.95 -9.01 -9.33 -11.93 -7.13 10.34 -1.01
-4.18 -1.30 -7.56 -1.25 -4.64 -4.88 -4.06 -1.91 -1.81 -6.92 -13.27 -5.52
 4.40 -12.17 -4.55 -5.82 -0.81 -19.28 -4.97 -7.78 -5.07 -5.45 -4.27 -4.98
-9.56 -9.33

```

Find the maximum likelihood estimate of θ . You only need an approximate value; one decimal place of accuracy will do.

Again, this is a problem that can be solved explicitly by differentiation, and the reader is invited to give it a try before proceeding. Have the answer? Is it still the same day you started? Now for the numerical solution. First, write the log likelihood as

$$\begin{aligned} \ell(\theta) &= \ln \prod_{i=1}^n \frac{1}{|\theta| \sqrt{2\pi}} e^{-\frac{(d_i - \theta)^2}{2\theta^2}} \\ &= -n \ln |\theta| - \frac{n}{2} \ln(2\pi) - \frac{\sum_{i=1}^n d_i^2}{2\theta^2} + \frac{\sum_{i=1}^n d_i}{\theta} - \frac{n}{2}. \end{aligned}$$

We will do this in R. The data are in a file called `norm1.data`. Read it. Remember that `>` is the R prompt.

```

> D <- scan("norm1.data")
Read 50 items

```

Now define a function to compute the log likelihood.

```

loglike1 <- function(theta) # Assume data are in a vector called D
{
  sumdsq <- sum(D^2); sumd <- sum(D); n <- length(D)
  loglike1 <- -n * log(abs(theta)) - (n/2)*log(2*pi) - sumdsq/(2*theta^2) +
    sumd/theta - n/2
  loglike1 # Return value of function
} # End definition of function loglike1

```

Just to show how the function works, compute it at a couple of values, say $\theta = 2$ and $\theta = -2$.

```
> loglike1(2)
[1] -574.2965
> loglike1(-2)
[1] -321.7465
```

Negative values of the parameter look more promising, but it is time to get systematic. The following is called a *grid search*. It is brutal, inefficient, and usually effective. It is too slow to be practical for large problems, but this is a one-dimensional parameter and we are only asked for one decimal place of accuracy. Where should we start? Since the parameter is the mean of the distribution, it should be safe to search within the range of the data. Start with widely spaced values and then refine the search. All we are doing is to calculate the log likelihood for a set of (equally spaced) parameter values and see where it is greatest. After all, that is the *idea* behind the MLE.

```
> min(D); max(D)
[1] -19.28
[1] 10.34
> Theta <- -20:10
> cbind(Theta, loglike1(Theta))
      Theta
[1,]   -20 -211.5302
[2,]   -19 -208.6709
[3,]   -18 -205.6623
[4,]   -17 -202.4911
[5,]   -16 -199.1423
[6,]   -15 -195.6002
[7,]   -14 -191.8486
[8,]   -13 -187.8720
[9,]   -12 -183.6580
[10,]  -11 -179.2022
[11,]  -10 -174.5179
[12,]   -9 -169.6565
[13,]   -8 -164.7513
[14,]   -7 -160.1163
[15,]   -6 -156.4896
[16,]   -5 -155.6956
[17,]   -4 -162.7285
[18,]   -3 -193.8796
[19,]   -2 -321.7465
[20,]   -1 -1188.0659
[21,]    0      NaN
[22,]    1 -1693.1659
```

```
[23,]    2 -574.2965
[24,]    3 -362.2463
[25,]    4 -289.0035
[26,]    5 -256.7156
[27,]    6 -240.6729
[28,]    7 -232.2734
[29,]    8 -227.8888
[30,]    9 -225.7788
[31,]   10 -225.0279
```

First, we notice that at $\theta = 0$, the log likelihood is indeed Not a Number. For this problem, the parameter space is all the real numbers except zero – unless one wants to think of a normal random variable with zero variance as being degenerate at μ ; that is, $P(D = \mu) = 1$. (In his case, what would the data look like?)

But the log likelihood is greatest around $\theta = -5$. We are asked for one decimal place of accuracy, so,

```
> Theta <- seq(from=-5.5,to=-4.5,by=0.1)
> Loglike <- loglike1(Theta)
> cbind(Theta,Loglike)
      Theta  Loglike
[1,] -5.5 -155.5445
[2,] -5.4 -155.4692
[3,] -5.3 -155.4413
[4,] -5.2 -155.4660
[5,] -5.1 -155.5487
[6,] -5.0 -155.6956
[7,] -4.9 -155.9136
[8,] -4.8 -156.2106
[9,] -4.7 -156.5950
[10,] -4.6 -157.0767
[11,] -4.5 -157.6665
> thetahat <- Theta[Loglike==max(Loglike)]
>           # Theta such that Loglike is the maximum of Loglike
> thetahat
[1] -5.3
```

To one decimal place of accuracy, the maximum is at $\theta = -5.3$. It would be easy to refine the grid and get more accuracy, but that will do. This is the last time we will see our friend the grid search, but you may find the approach useful in homework.

Now let's do the search in a more sophisticated way, using R's `nlm` (non-linear minimization) function.³ The `nlm` function has quite a few arguments; try `help(nlm)`. The

³The `nlm` function is good but generic. See Numerical Recipes for a really good discussion of routines for numerically minimizing a function. They also provide source code. The *Numerical Recipes* books have

ones you always need are the first two: the name of the function, and a starting value (or vector of starting values, for multiparameter problems).

Where should we start? Since the parameter equals the expected value of the distribution, how about the sample mean? It is often a good strategy to use Method of Moment estimators as starting values for numerical maximum likelihood. Method of Moments estimation is reviewed in Section ??.

One characteristic that `nlm` shares with most optimization routines is that it likes to *minimize* rather than maximizing. So we will minimize the negative of the log likelihood function. For this, it is necessary to define a new function, `loglike2`.

```
> mean(D)
[1] -5.051
> loglike2 <- function(theta) { loglike2 <- -loglike1(theta); loglike2 }
> nlm(loglike2,mean(D))
$minimum
[1] 155.4413

$estimate
[1] -5.295305

$gradient
[1] -1.386921e-05

$code
[1] 1

$iterations
[1] 4
```

By default, `nlm` returns a list with four elements; `minimum` is the value of the function at the point where it reaches its minimum, `estimate` is the value at which the minimum was located; that's the MLE. `Gradient` is the slope in the direction of greatest increase; it should be near zero. `Code` is a diagnosis of how well the optimization went; the value of 1 means everything seemed okay. See `help nlm` for more detail.

We could have gotten just the MLE with

```
> nlm(loglike2,mean(D))$estimate
[1] -5.295305
```

That's the answer, but the numerical approach misses some interesting features of the problem, which can be done with paper and pencil in this simple case. Differentiating the

versions for the Pascal, Fortran and Basic languages as well as C. This is a case where a book definitely delivers more than the title promises. It may be a cookbook, but it is a very good cookbook written by expert Chemists.

log likelihood separately for $\theta < 0$ and $\theta > 0$ to get rid of the absolute value sign, and then re-uniting the two cases since the answer is the same, we get

$$\ell'(\theta) = -\frac{n}{\theta} + \frac{\sum_{i=1}^n d_i^2}{\theta^3} - \frac{\sum_{i=1}^n d_i}{\theta^2}.$$

Setting $\ell'(\theta) = 0$ and re-arranging terms, we get

$$n\theta^2 + \left(\sum_{i=1}^n d_i\right)\theta - \left(\sum_{i=1}^n d_i^2\right) = 0.$$

Of course this expression is not valid at $\theta = 0$, because the function we are differentiating is not even defined there. The quadratic formula yields two solutions:

$$\frac{-\sum_{i=1}^n d_i \pm \sqrt{(\sum_{i=1}^n d_i)^2 + 4n \sum_{i=1}^n d_i^2}}{2n} = \frac{1}{2} \left(-\bar{d} \pm \sqrt{\bar{d}^2 + 4 \frac{\sum_{i=1}^n d_i^2}{n}} \right), \quad (\text{A.16})$$

where \bar{d} is the sample mean.

Let's calculate these for the given data.

```
> meand <- mean(D) ; meandsq <- sum(D^2)/length(D)
> (-meand + sqrt(meand^2 + 4*meandsq) )/2
[1] 10.3463
> (-meand - sqrt(meand^2 + 4*meandsq) )/2
[1] -5.2953
```

The second solution is the one we found with the numerical search. What about the other one? Is it a minimum? Maximum? Saddle point? The second derivative test will tell us. The second derivative is

$$\ell''(\theta) = \frac{n}{\theta^2} - \frac{3 \sum_{i=1}^n d_i^2}{\theta^4} + \frac{2 \sum_{i=1}^n d_i}{\theta^3}.$$

Substituting [A.16](#) into this does not promise to be much fun, so we will be content with a numerical answer for this particular data set. Call the first root `t1` and the second one (our MLE) `t2`.

```
> t1 <- (-meand + sqrt(meand^2 + 4*meandsq) )/2 ; t1
[1] 10.3463
> t2 <- (-meand - sqrt(meand^2 + 4*meandsq) )/2 ; t2
[1] -5.2953
> n <- length(D)
> # Now calculate second derivative at t1 and t2
> n/t1^2 - 3*sum(D^2)/t1^4 + 2*sum(D)/t1^3
[1] -0.7061484
> n/t2^2 - 3*sum(D^2)/t2^4 + 2*sum(D)/t2^3
[1] -5.267197
```

The second derivative is negative in both cases; they are both local maxima! Which peak is higher?

```
> loglike1(t1)
[1] -224.9832
> loglike1(t2)
[1] -155.4413
```

So the maximum we found is higher, which makes sense because it's within the range of the data. But we only found it because we started searching near the correct answer.

Let's plot the log likelihood function, and see what this thing looks like. We know that because the natural log function goes to minus infinity as its (positive) argument approaches zero, the log likelihood plunges to $-\infty$ at $\theta = 0$. A plot would look like a giant icicle and we would not be able to see any detail where it matters. So we will zoom in by limiting the range of the y axis. Here is the R code.

```
Theta <- seq(from=-15,to=20,by=0.25); Theta <- Theta[Theta!=0]
Loglike <- loglike1(Theta)
# Check where to break off the icicle
max(Loglike); Loglike[Theta==3]; Loglike[Theta==3]

plot(Theta,Loglike,type='l',xlim=c(-15,20),ylim=c(-375,-155),
     xlab=expression(theta),ylab="Log Likelihood")
# This is how you get Greek letters.
```

Here is the picture. You can see the local maxima around $\theta = -5$ and $\theta = 10$, and also that the one for negative θ is a higher.

Presumably we would have reached the bad answer if we had started the search in a bad place. Let's try starting the search at $\theta = +3$.


```
> nlm(loglike2,3)
$minimum
[1] 283.7589

$estimate
[1] 64.83292

$gradient
[1] 0.701077

$code
[1] 4

$iterations
[1] 100
```

Figure A.1: Log Likelihood for Example [A.4.2.1](#)

What happened?! The answer is way off, nowhere near the positive root of 10.3463. And the **minimum** (of *minus* the log likelihood) is over 283, when it would have been 224.9832 at $\theta = 10.3463$.

What happened was that the slope of the function was very steep at our starting value of $\theta = 3$, so `nlm` took a huge step in a positive direction. It was too big, and landed in a nearly flat place. Then `nlm` wandered around until it ran out of its default number of iterations (notice `iterations=100`). The exit code of 4 means maximum number of iterations exceeded.

It should be better if we start close to the answer, say at $\theta = 8$.

```
> nlm(loglike2,8)
```

```
$minimum
```

```
[1] 224.9832
```

```
$estimate
```

```
[1] 10.34629
```

```
$gradient
```

```
[1] -4.120564e-08
```

```
$code
```

```
[1] 1
```

```
$iterations
```

```
[1] 6
```


That's better. The moral of this story is clear. Good starting are *very* important.

Now let us look at an example of a multi-parameter problem where an explicit formula for the MLE is impossible, and numerical methods are required.

Example A.4.2.2

Let D_1, \dots, D_n be a random sample from a Gamma distribution with parameters $\alpha > 0$ and $\beta > 0$. The probability density function is

$$f(x; \alpha, \beta) = \frac{1}{\beta^\alpha \Gamma(\alpha)} e^{-x/\beta} x^{\alpha-1}$$

for $x > 0$, and zero otherwise. Here is a random sample of size $n = 50$. For this example, the data are simulated using R, with known parameter values $\alpha = 2$ and $\beta = 3$. The seed for the random, number generator is set so the pseudo-random numbers can be recovered if necessary.

```
> set.seed(3201); alpha=2; beta=3
> D <- round(rgamma(50,shape=alpha, scale=beta),2); D
 [1] 20.87 13.74  5.13  2.76  4.73  2.66 11.74  0.75 22.07 10.49  7.26  5.82 13.08
[14]  1.79  4.57  1.40  1.13  6.84  3.21  0.38 11.24  1.72  4.69  1.96  7.87  8.49
[27]  5.31  3.40  5.24  1.64  7.17  9.60  6.97 10.87  5.23  5.53 15.80  6.40 11.25
[40]  4.91 12.05  5.44 12.62  1.81  2.70  3.03  4.09 12.29  3.23 10.94
> mean(D); alpha*beta
 [1] 6.8782
 [1] 6
> var(D); alpha*beta^2
 [1] 24.90303
 [1] 18
```

The parameter vector $\boldsymbol{\theta} = (\alpha, \beta)$, and the parameter space Θ is the first quadrant of \mathbb{R}^2 .

$$\Theta = \{(\alpha, \beta) : \alpha > 0, \beta > 0\}$$

The log likelihood is

$$\begin{aligned} \ell(\alpha, \beta) &= \ln \prod_{i=1}^n \frac{1}{\beta^\alpha \Gamma(\alpha)} e^{-d_i/\beta} d_i^{\alpha-1} \\ &= \ln \left[\beta^{-n\alpha} \Gamma(\alpha)^{-n} \exp\left(-\frac{1}{\beta} \sum_{i=1}^n d_i\right) \left(\prod_{i=1}^n d_i\right)^{\alpha-1} \right] \\ &= -n\alpha \ln \beta - \ln \Gamma(\alpha) - \frac{1}{\beta} \sum_{i=1}^n d_i + (\alpha - 1) \sum_{i=1}^n \ln d_i. \end{aligned}$$

The next step would be to partially differentiate the log likelihood with respect to α and β , set both partial derivatives to zero, and solve two equations in two unknowns. But even if you are confident that the gamma function is differentiable (it is), you will be unable to solve the equations. It has to be done numerically.

Define an R function for the minus log likelihood. Notice the `lgamma` function, a direct numerical approximation of $\ln \Gamma(\alpha)$. The plan is to numerically minimize the minus log likelihood function over all (α, β) pairs, for this particular set of data values.

```
> # Gamma minus log likelihood: alpha=a, beta=b
> gmll <- function(theta,datta)
+   {
+     a <- theta[1]; b <- theta[2]
+     n <- length(datta); sumd <- sum(datta); sumlogd <- sum(log(datta))
+     gmll <- n*a*log(b) + n*lgamma(a) + sumd/b - (a-1)*sumlogd
+     gmll
+   } # End function gmll
```

Where should the numerical search start? One approach is to start at reasonable estimates of α and β — estimates that can be calculated directly rather than by a numerical approximation. As in Example A.4.2.1, Method of Moments estimators are a convenient, high-quality choice.

For a gamma distribution, $E(D) = \alpha\beta$ and $Var(D) = \alpha\beta^2$. So,

$$\alpha = \frac{E(D)^2}{Var(D)} \quad \text{and} \quad \beta = \frac{Var(D)}{E(D)}.$$

Replacing population moments by sample moments and writing $\tilde{\alpha}$ and $\tilde{\beta}$ for the resulting Method of Moments estimators, we obtain

$$\tilde{\alpha} = \frac{\bar{D}^2}{S_D^2} \quad \text{and} \quad \tilde{\beta} = \frac{S_D^2}{\bar{D}},$$

where \bar{D} is the sample mean and S_D^2 is the sample variance. For these data, the Method of Moments estimates are reasonably close to the correct values of $\alpha = 2$ and $\beta = 3$, but they are not perfect. Parameter estimates are not the same as parameters!

```
> momalpha <- mean(D)^2/var(D); momalpha
[1] 1.899754
> mombeta <- var(D)/mean(D); mombeta
[1] 3.620574
```

Now for the numerical search. This time, we will request that the `nlm` function return the *Hessian* at the place where the search stops. The Hessian is defined as follows. Suppose we are minimizing a function $g(\theta_1, \dots, \theta_k)$ — say, a minus log likelihood. The

Hessian is a $k \times k$ matrix of mixed partial derivatives. It may be written in terms of its (i, j) element s

$$\mathbf{H} = \left[\frac{\partial^2 g}{\partial \theta_i \partial \theta_j} \right]. \quad (\text{A.17})$$

When the determinant of the Hessian matrix is positive at a point, the function is concave up there. It's like a second derivative test.

In the following, notice how the `nlm` function assumes that the first argument of the function being minimized is a vector of arguments over which we should minimize, and any other arguments (in this case, the name of the data vector) can be specified by name in the `nlm` function call.

```
> gammasearch = nlm(gmll,c(momalpha,mombeta),hessian=T,datta=D); gammasearch
$minimum
[1] 142.0316

$estimate
[1] 1.805930 3.808674

$gradient
[1] 2.847002e-05 9.133932e-06

$hessian
      [,1]      [,2]
[1,] 36.68932 13.127271
[2,] 13.12727  6.222282

$code
[1] 1

$iterations
[1] 6

> det(gammasearch$hessian)
[1] 55.96605
```

The `nlm` object `gammasearch` is a linked list. The item `minimum` is the value of the minus log likelihood function where the search stops. The item `estimate` is the point at which the search stops, so $\hat{\alpha} = 1.805930$ and $\hat{\beta} = 3.808674$. The `gradient` is

$$\left(-\frac{\partial \ell}{\partial \alpha}, -\frac{\partial \ell}{\partial \beta} \right)'$$

Besides being the direction of steepest decrease, it's something that should be zero at the MLE. And indeed it is, give or take a bit of numerical inaccuracy.

The Hessian at the stopping place is in `gammasearch$hessian`. Its determinant is positive, something that `nlm` checks too. So we have a minimum. Finally, `code=1` means normal termination of the search, and `iterations=6` means the function took 6 steps downhill to reach its target.

It is very helpful to have the true parameter values $\alpha = 2$ and $\beta = 3$ for this example. $\hat{\alpha} = 1.8$ seems pretty close, while $\hat{\beta} = 3.8$ seems farther off. This is a reminder of how informative confidence intervals and tests can be.

The Invariance Principle

The Invariance Principle of maximum likelihood estimation says that *the MLE of a function is that function of the MLE*. An example comes first, followed by formal details.

Example A.4.2.3

Let D_1, \dots, D_n be a random sample from a Bernoulli distribution (1=Yes, 0=No) with parameter $\theta, 0 < \theta < 1$. The parameter space is $\Theta = (0, 1)$, and the likelihood function is

$$L(\theta) = \prod_{i=1}^n \theta^{d_i} (1 - \theta)^{1-d_i} = \theta^{\sum_{i=1}^n d_i} (1 - \theta)^{n - \sum_{i=1}^n d_i}.$$

Differentiating the log likelihood with respect to θ , setting the derivative to zero and solving yields the usual estimate $\hat{\theta} = \bar{d}$, the sample proportion.

Now suppose that instead of the probability, we write this model in terms of the *odds* of $D_i = 1$, a re-parameterization that is often useful in categorical data analysis. Denote the odds by θ' . The definition of odds is

$$\theta' = \frac{\theta}{1 - \theta} = g(\theta). \tag{A.18}$$

As θ ranges from zero to one, θ' ranges from zero to infinity. So there is a new parameter space: $\theta' \in \Theta' = (0, \infty)$.

To write the likelihood function in terms of θ' , first solve for θ , obtaining

$$\theta = \frac{\theta'}{1 + \theta'} = g^{-1}(\theta').$$

The likelihood in terms of θ' is then

$$L(g^{-1}(\theta')) = \frac{\theta'^{\sum_{i=1}^n d_i}}{(1 + \theta')^n}. \tag{A.19}$$

At this point one could differentiate the log of (A.19) with respect to θ' , set the derivative to zero, and solve for θ' . The point of the invariance principle is that this is unnecessary. The maximum likelihood estimator of $g(\theta)$ is $g(\hat{\theta})$, so one need only look at (A.18) and write

$$\hat{\theta}' = \frac{\hat{\theta}}{1 - \hat{\theta}} = \frac{\bar{d}}{1 - \bar{d}}.$$

It is sometimes convenient to parameterize a statistical model in more than one way. The invariance principle can save a lot of work in practice, because it says that you only have to maximize the likelihood function once. It is useful theoretically too.

In Example A.4.2.3, the likelihood function has only one maximum and the function g linking θ' to θ is one-to-one, which is why we can write g^{-1} . This is the situation where the invariance principle is clearest and most useful. Here is a proof.

Let the parameter θ (which could be a vector even though it is not in bold face) be an element of the parameter space Θ , and let the function $g : \Theta \rightarrow \Theta'$ be one-to-one, where $\Theta' = \{\theta' : \theta' = g(\theta), \theta \in \Theta\}$. The meaning of g one-to-one is that there exists a function $g^{-1} : \Theta' \rightarrow \Theta$ such that $g^{-1}(g(\theta)) = \theta$ for all $\theta \in \Theta$.

As usual, denote the likelihood function by $L(\theta)$. In addition, suppose it has a unique maximum at $\hat{\theta} \in \Theta$, so that for all $\theta \in \Theta$ with $\theta \neq \hat{\theta}$, $L(\theta) < L(\hat{\theta})$. Now,

$$L(g^{-1}(g(\theta))) = L(g^{-1}(\theta')), \quad \theta' \in \Theta'.$$

Maximizing $L(g^{-1}(\theta'))$ over all $\theta' \in \Theta'$ yields $\hat{\theta}'$ satisfying $L(g^{-1}(\hat{\theta}')) \geq L(g^{-1}(\theta'))$ ⁴ for all $\theta' \in \Theta'$. Because $\theta' \in \Theta'$, there exists $\theta \in \Theta$ with $g(\theta) = \theta'$. The objective is to show that this θ equals $\hat{\theta}$.

Suppose on the contrary that $\theta \neq \hat{\theta}$. Then $L(g^{-1}(\hat{\theta}')) = L(\theta) \geq L(\hat{\theta})$. But the maximum over Θ is unique; i.e., $L(\hat{\theta}) > L(\theta)$. This contradiction establishes $\theta = \hat{\theta}$, completing the proof.

Exercises A.4.2

- For each of the following distributions, derive a general expression for the Maximum Likelihood Estimator (MLE). You don't have to do the second derivative test. Then use the data to calculate a numerical estimate.
 - $p(x) = \theta(1 - \theta)^x$ for $x = 0, 1, \dots$, where $0 < \theta < 1$. Data: 4, 0, 1, 0, 1, 3, 2, 16, 3, 0, 4, 3, 6, 16, 0, 0, 1, 1, 6, 10. Answer: 0.2061856
 - $f(x) = \frac{\alpha}{x^{\alpha+1}}$ for $x > 1$, where $\alpha > 0$. Data: 1.37, 2.89, 1.52, 1.77, 1.04, 2.71, 1.19, 1.13, 15.66, 1.43. Answer: 1.469102
 - $f(x) = \frac{\tau}{\sqrt{2\pi}} e^{-\frac{\tau^2 x^2}{2}}$, for x real, where $\tau > 0$. Data: 1.45, 0.47, -3.33, 0.82, -1.59, -0.37, -1.56, -0.20. Answer: 0.6451059
 - $f(x) = \frac{1}{\theta} e^{-x/\theta}$ for $x > 0$, where $\theta > 0$. Data: 0.28, 1.72, 0.08, 1.22, 1.86, 0.62, 2.44, 2.48, 2.96. Answer: 1.517778
- Let X_1, \dots, X_5 be a random sample from a Gamma distribution with parameters $\alpha > 0$ and $\beta = 1$. That is, the density is

$$f(x; \alpha) = \frac{1}{\Gamma(\alpha)} e^{-x} x^{\alpha-1}$$

⁴Actually the maximum over Θ' is unique. The proof of this will be left as an exercise.

for $x > 0$, and zero otherwise.

The five data values are 2.06, 1.08, 0.96, 1.32, 1.53. Find an approximate numerical value of the maximum likelihood estimate of α . Your final answer is one number. For this question you will hand in a one-page printout. On the back, you will write a brief explanation of what you did.

3. For each of the following distributions, try to derive a general expression for the Maximum Likelihood Estimator (MLE). Then, use R's `nlm` function to obtain the MLE numerically for the data supplied for the problem. The data are in a separate HTML document, because it saves a lot of effort to copy and paste rather than typing the data in by hand, and PDF documents can contain invisible characters that mess things up. NOTE! Put them here as well as in assignment HTML document.

(a) $f(x) = \frac{1}{\pi[1+(x-\theta)^2]}$ for x real, where $-\infty < \theta < \infty$.

(b) $f(x) = \frac{1}{2}e^{-|x-\theta|}$ for x real, where $-\infty < \theta < \infty$.

(c) $f(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}x^{\alpha-1}(1-x)^{\beta-1}$ for $0 < x < 1$, where $\alpha > 0$ and $\beta > 0$.

For each distribution, be able to state (briefly) why differentiating the log likelihood and setting the derivative to zero does not work. For the computer part, bring to the quiz one sheet of printed output for each of the 3 distributions. The three sheets should be separate, because you may hand only one of them in. Each printed page should show the following, *in this order*.

- Definition of the function that computes the likelihood, or log likelihood, or minus log likelihood or whatever.
 - How you got the data into R – probably a `scan` statement.
 - Listing of the data for the problem.
 - The `nlm` statement and resulting output.
4. Let $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$, where \mathbf{X} is an $n \times p$ matrix of known constants, $\boldsymbol{\beta}$ is a $p \times 1$ vector of unknown constants, and $\boldsymbol{\epsilon}$ is multivariate normal with mean zero and covariance matrix $\sigma^2\mathbf{I}_n$, with $\sigma^2 > 0$ an unknown constant.
 - (a) What is the distribution of \mathbf{Y} ? There is no need to show any work.
 - (b) Assuming that the columns of \mathbf{X} are linearly independent, show that the maximum likelihood estimate of $\boldsymbol{\beta}$ is $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$. Don't use derivatives. The trick is to add and subtract $\hat{\boldsymbol{\beta}}$, distribute the expected value, and simplify. Does your answer apply for any value of σ^2 ? Why or why not?
 - (c) Given the MLE of $\boldsymbol{\beta}$, find the MLE of σ^2 . Show your work. This time you may differentiate.

A.4.3 Interval Estimation and Testing

All the tests and confidence intervals here are based on large-sample approximations, primarily the Central Limit Theorem. See Appendix ?? for basic definitions and results. They are valid as the sample size $n \rightarrow \infty$, but frequently perform well for samples that are only fairly large. How big is big enough? This is a legitimate question, and the honest answer is that it depends upon the distribution of the data. In practice, people often just apply these tools almost regardless of the sample size, because nothing better is available. Some do it with their eyes closed, some squint, and some have their eyes wide open.

The basic result comes from the research of Abraham Wald (give a source) in the 1950s. *As the sample size n increases, the distribution of the maximum likelihood estimator $\hat{\boldsymbol{\theta}}_n$ approaches a multivariate normal with expected value $\boldsymbol{\theta}$ and variance-covariance matrix $\mathbf{V}_n(\boldsymbol{\theta})$.* It is quite remarkable that anyone could figure this out, given that it includes cases like the Gamma, where no closed-form expressions for the maximum likelihood estimators are possible. The theorem in question is not true for every distribution, but it is true if the distribution of the data is not too strange. The precise meaning of “not too strange” is captured in a set of technical conditions called *regularity conditions*. Volume 2 of *Kendall’s advanced theory of statistics* [7] is a good textbook source for the details.

If $\boldsymbol{\theta}$ is a $k \times 1$ matrix, then $\mathbf{V}_n(\boldsymbol{\theta})$ is a $k \times k$ matrix, called the *asymptotic covariance matrix* of the estimators. It’s not too surprising that it depends on the parameter $\boldsymbol{\theta}$, and it also depends on the sample size n . Using the asymptotic covariance matrix, it is possible to construct a variety of useful tests and confidence intervals.

Fisher Information

The fact that $\mathbf{V}_n(\boldsymbol{\theta})$ depends on the unknown parameter will present no problem; substituting $\hat{\boldsymbol{\theta}}_n$ for $\boldsymbol{\theta}$ yields an *estimated* asymptotic covariance matrix. So consider the form of the matrix \mathbf{V} .

Think of a one-parameter maximum likelihood problem, where we differentiate the log likelihood, set the derivative to zero and solve for θ ; the solution is $\hat{\theta}$. The log likelihood will be concave down at $\hat{\theta}$, but the exact way it looks will depend on the distribution as well as the sample size. In particular, it could be almost flat at $\hat{\theta}$, or it could be nearly a sharp peak, with extreme downward curvature. In the latter case, clearly the log likelihood is more informative about θ . It contains more information. One of the many good ideas of R. A. F. Fisher was that the second derivative reflects curvature, and can be viewed as a measure of the information provided by the sample data. It is called the *Fisher Information* in his honour.

Now with increasing sample size, nearly all log likelihood functions acquire more and more downward curvature at the MLE. This makes sense – more data provide more information. But how about the information from just one observation? If you look at the second derivative of the log likelihood function,

$$\frac{\partial^2 \ell}{\partial \theta^2} = \frac{\partial^2}{\partial \theta^2} \ln \prod_{i=1}^n f(d_i; \theta) = \sum_{i=1}^n \frac{\partial^2}{\partial \theta^2} \ln f(d_i; \theta),$$

you see that it is the sum of n quantities. Each observation is contributing a piece to the downward curvature. But how much? Well, it depends on the particular data value x_i . But the data are a random sample, so in fact the contribution is a random quantity: $\frac{\partial^2}{\partial \theta^2} \ln f(X_i; \theta)$. How about the information one would *expect* an observation to contribute? Okay, take the expected value. Finally, note that because the curvature is down at the MLE, the quantity we are discussing is negative. But we want to call this “information,” and it would be nicer if it were a positive number, so higher values meant more information. Okay, multiply by -1 . This leads to the definition of the Fisher Information in a single observation:

$$I(\theta) = E \left[-\frac{\partial^2}{\partial \theta^2} \ln f(D_i; \theta) \right]. \quad (\text{A.20})$$

The information is the same for $i = 1, \dots, n$, and the Fisher Information in the entire sample is just $nI(\theta)$.

It was clear that Fisher was onto something good, because for many problems where the variance of $\hat{\theta}$ can be calculated exactly, it is one divided by the Fisher Information. Subsequently Cramér and Rao discovered the *Cramér-Rao Inequality*, which says that for *any* statistic T that is an unbiased estimator of θ ,

$$\text{Var}(T) \geq \frac{1}{nI(\theta)}.$$

That’s impressive, because to have a small variance is a great property in an estimator; it means precise estimation. The Cramér-Rao inequality tells us that in terms of variance, one cannot do better than an unbiased estimator whose variance equals one over the Fisher Information, and many MLEs do that. Subsequently, Wald⁵ showed that (under some regularity conditions) as $n \rightarrow \infty$, the variances of maximum likelihood estimators in general attain the Cramér-Rao lower bound. Thus, to learn the asymptotic variance of $\hat{\theta}$, you do not need an explicit formula for $\hat{\theta}$. All you need is the Fisher Information. Also, in terms of variance nothing can beat maximum likelihood estimation, at least for large sample. So one needs a good reason to choose anything else, if the distribution of the data is known so you can write down the likelihood.

Calculating the expected value in (A.20) is often not too hard because taking the log and differentiating twice results in some simplification; it’s a source of many fun homework problems. But still it can be a chore, especially for multiparameter problems, which will be taken up shortly. For larger sample sizes, the Law of Large Numbers (Appendix ??) guarantees that the expected value can be approximated quite well by a sample mean, so that

$$I(\theta) = E \left[-\frac{\partial^2}{\partial \theta^2} \ln f(D_1; \theta) \right] \approx \frac{1}{n} \sum_{i=1}^n -\frac{\partial^2}{\partial \theta^2} \ln f(D_i; \theta).$$

This is sometimes called the *observed* Fisher Information.

⁵Need a reference

Multiplying the observed Fisher Information by n to get the approximate information in the entire sample yields

$$\sum_{i=1}^n -\frac{\partial^2}{\partial\theta^2} \ln f(D_i; \theta) = \frac{\partial^2}{\partial\theta^2} \sum_{i=1}^n -\ln f(D_i; \theta) = \frac{\partial^2}{\partial\theta^2} \left(-\ln \prod_{i=1}^n f(D_i; \theta) \right).$$

That's just the second derivative of the minus log likelihood.

The parameter θ is unknown, so to get the *estimated* Fisher Information in the whole sample, substitute $\hat{\theta}$. The result is

$$\frac{\partial^2}{\partial\theta^2} \left(-\ln \prod_{i=1}^n f(D_i; \hat{\theta}) \right).$$

That's the second derivative of minus the log likelihood, evaluated at the maximum likelihood estimate. And, it's a function of the sample data that is not a function of any unknown parameters; in other words it is a statistic. If you have already carried out the second derivative test to check that you really had a maximum, all you need to do to estimate the variance of $\hat{\theta}$ is take the reciprocal of the second derivative and multiply by -1 . It is truly remarkable how neatly this all works out.

Generalization to the multivariate case is very natural. Now the parameter is $\boldsymbol{\theta} = (\theta_1, \dots, \theta_k)'$ and the Fisher Information *Matrix* is a $k \times k$ matrix of (expected) mixed partial derivatives, defined by

$$\mathcal{I}(\boldsymbol{\theta}) = \left[-E \left(\frac{\partial^2}{\partial\theta_i \partial\theta_j} f(\mathbf{D}_1; \boldsymbol{\theta}) \right) \right],$$

where the boldface \mathbf{D}_i is an acknowledgement that the data might also be multivariate.

In the estimated observed Fisher Information evaluated at the MLE (which will simply be called the ‘‘Fisher Information Matrix’’ unless other wise noted), expected value is replaced by a sample mean and $\boldsymbol{\theta}$ is replaced by $\hat{\boldsymbol{\theta}}$. The formula is

$$\mathcal{J}(\hat{\boldsymbol{\theta}}) = \left[\frac{\partial^2}{\partial\theta_i \partial\theta_j} \left(-\ln \prod_{q=1}^n f(\mathbf{D}_q; \hat{\boldsymbol{\theta}}) \right) \right] = \left[\frac{\partial^2}{\partial\theta_i \partial\theta_j} (-\ell(\hat{\boldsymbol{\theta}})) \right]. \quad (\text{A.21})$$

In the one-dimensional case, one over the estimated Fisher Information is the (estimated) asymptotic variance of the maximum likelihood estimator. *In the multi-parameter case, the Fisher Information is a matrix, and the estimated asymptotic variance-covariance matrix is its inverse.* Denoting the estimated asymptotic covariance matrix by $\hat{\mathbf{V}}_n$, we have

$$\hat{\mathbf{V}}_n = \mathcal{J}(\hat{\boldsymbol{\theta}}_n)^{-1}. \quad (\text{A.22})$$

Now comes the really good part. Comparing Formula (A.21) for the Fisher Information to Formula (A.17) for the Hessian, we see that they are exactly the same. And *the Hessian evaluated at $\hat{\boldsymbol{\theta}}$ is a by-product of the numerical search*⁶.

⁶At least for generic numerical minimization routines like R's `nlm`. Some specialized methods like iterative proportional fitting of log-linear models and Fisher scoring (iteratively re-weighted least squares) for generalized linear models maximize the likelihood indirectly and do not require calculation of the Hessian.

So to get the asymptotic covariance matrix, minimize minus the log likelihood, tell the software to give you the Hessian, and calculate the inverse by computer. The theoretical story may be a bit long here, but what you have to do in practice is quite simple.

Continuing with the Gamma distribution Example A.4.2.2, the Hessian is

```
> gammasearch$hessian
      [,1]      [,2]
[1,] 36.68932 13.127271
[2,] 13.12727  6.222282
```

and the asymptotic covariance is just

```
> Vhat = solve(gammasearch$hessian); V
      [,1]      [,2]
[1,]  0.1111796 -0.2345577
[2,] -0.2345577  0.6555638 .
```

The diagonal elements of $\widehat{\mathbf{V}}$ are the estimated variances of the sampling distributions of $\widehat{\alpha}$ and $\widehat{\beta}$ respectively, and their square roots are the standard errors.

```
> SEalphahat = sqrt(Vhat[1,1]); SEbetahat = sqrt(Vhat[2,2])
```

In general, let θ denote an element of the parameter vector, let $\widehat{\theta}$ be its maximum likelihood estimator, and let the standard error of $\widehat{\theta}$ be written $S_{\widehat{\theta}}$. Then Wald's central Limit Theorem for maximum likelihood estimators tells us that

$$Z = \frac{\widehat{\theta} - \theta}{S_{\widehat{\theta}}} \tag{A.23}$$

has an approximate standard normal distribution. In particular, for the Gamma example

$$Z_1 = \frac{\widehat{\alpha} - \alpha}{S_{\widehat{\alpha}}} \quad \text{and} \quad Z_2 = \frac{\widehat{\beta} - \beta}{S_{\widehat{\beta}}}$$

may be treated as standard normal.

Confidence Intervals

These quantities may be used to produce both tests and confidence intervals. For example, a 95% confidence interval for the parameter θ is obtained as follows.

$$\begin{aligned} 0.95 &\approx Pr\{-1.96 \leq Z \leq 1.96\} \\ &= Pr\left\{-1.96 \leq \frac{\widehat{\theta} - \theta}{S_{\widehat{\theta}}} \leq 1.96\right\} \\ &= Pr\left\{\widehat{\theta} - 1.96 S_{\widehat{\theta}} \leq \theta \leq \widehat{\theta} + 1.96 S_{\widehat{\theta}}\right\} \end{aligned}$$

This could also be written $\hat{\theta} \pm 1.96 S_{\hat{\theta}}$.

If you are used to seeing confidence intervals with a \sqrt{n} and wondering where it went, recall that $S_{\bar{X}} = \frac{S}{\sqrt{n}}$. The \sqrt{n} is also present in the confidence interval for θ , but it is embedded in $S_{\hat{\theta}}$.

Here are the 95% confidence intervals for the Gamma distribution example:

```
> alphahat = gammasearch$estimate[1]; betahat = gammasearch$estimate[2]
> Lalpha = alphahat - 1.96*SEalphahat; Ualpha = alphahat + 1.96*SEalphahat
> Lbeta = betahat - 1.96*SEbetahat; Ubeta = betahat + 1.96*SEbetahat
> cat("\nEstimated alpha = ",round(alphahat,2)," 95 percent CI from ",
+     round(Lalpha,2)," to ",round(Ualpha,2), "\n\n")
```

```
Estimated alpha = 1.81 95 percent CI from 1.15 to 2.46
```

```
> cat("\nEstimated beta = ",round(betahat,2)," 95 percent CI from ",
+     round(Lbeta,2)," to ",round(Ubeta,2), "\n\n")
```

```
Estimated beta = 3.81 95 percent CI from 2.22 to 5.4
```

Notice that while the parameter estimates may not seem very accurate, the 95% confidence intervals do include the true parameter values $\alpha = 2$ and $\beta = 3$.

Z-tests

The standard normal variable in (A.23) can be used to form a Z -test of $H_0 : \theta = \theta_0$ using

$$Z = \frac{\hat{\theta} - \theta_0}{S_{\hat{\theta}}}.$$

So for example, suppose the data represent time intervals between events occurring in time, and we wonder whether the events arise from a Poisson process. In this case the distribution of times would be exponential, which means $\alpha = 1$. To test this null hypothesis at the 0.05 level,

```
> Z = (alphahat-1)/SEalphahat; Z
[1] 2.417046
> pval = 2*(1-pnorm(abs(Z))); pval # Two-sided test
[1] 0.01564705
```

So, the null hypothesis is rejected, and because the value is positive, the conclusion is that the true value of α is greater than one⁷.

When statistical software packages display this kind of large-sample Z -test, they usually just divide $\hat{\theta}$ by its standard error, testing the null hypothesis $H_0 : \theta = 0$. For parameters like regression coefficients, this is usually a good generic choice.

⁷The following basic question arises from time to time. Suppose a null hypothesis is rejected in favour of a two-sided alternative. Are we then “allowed” to look at the sign of the test statistic and conclude that $\theta < \theta_0$ or $\theta > \theta_0$, or must we just be content with saying $\theta \neq \theta_0$? The answer is that directional conclusions are theoretically justified as well as practically desirable. Think of splitting up the two-

Wald Tests

The approximate multivariate normality of the MLE can be used to construct a larger class of hypothesis tests for *linear* null hypotheses. A linear null hypothesis sets a collection of linear combinations of the parameters to zero. Suppose $\boldsymbol{\theta} = (\theta_1, \dots, \theta_k)'$ is a $k \times 1$ vector. A linear null hypothesis can be written

$$H_0 : \mathbf{C}\boldsymbol{\theta} = \mathbf{h},$$

where \mathbf{C} is an $r \times k$ matrix of constants, with rank r , $r \leq k$. As an example let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_7)'$, and the null hypothesis is

$$\theta_1 = \theta_2, \quad \theta_6 = \theta_7, \quad \frac{1}{3}(\theta_1 + \theta_2 + \theta_3) = \frac{1}{3}(\theta_4 + \theta_5 + \theta_6).$$

This may be expressed in the form $\mathbf{C}\boldsymbol{\theta} = \mathbf{h}$ as follows:

$$\begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 \\ 1 & 1 & 1 & -1 & -1 & -1 & 0 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \\ \theta_4 \\ \theta_5 \\ \theta_6 \\ \theta_7 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

Recall from Appendix A.3.2 that if $\mathbf{X} \sim N_k(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, and \mathbf{C} is an $r \times k$ constant matrix of rank r , then

$$\mathbf{C}\mathbf{X} \sim N_r(\mathbf{C}\boldsymbol{\mu}, \mathbf{C}\boldsymbol{\Sigma}\mathbf{C}')$$

and

$$(\mathbf{C}\mathbf{X} - \mathbf{C}\boldsymbol{\mu})'(\mathbf{C}\boldsymbol{\Sigma}\mathbf{C}')^{-1}(\mathbf{C}\mathbf{X} - \mathbf{C}\boldsymbol{\mu}) \sim \chi^2(r).$$

Similar facts hold asymptotically — that is approximately, as the sample size n approaches infinity. Because (approximately) $\hat{\boldsymbol{\theta}}_n \sim N_k(\boldsymbol{\theta}, \hat{\mathbf{V}}_n)$,

$$\mathbf{C}\hat{\boldsymbol{\theta}}_n \sim N_r(\mathbf{C}\boldsymbol{\theta}, \mathbf{C}\hat{\mathbf{V}}_n\mathbf{C}')$$

sided level α test (call it the *overall test*) into two one-sided tests with significance level $\alpha/2$. The null hypotheses of these tests are $H_{0,a} : \theta \leq \theta_0$ and $H_{0,b} : \theta \geq \theta_0$. Exactly one of these null hypotheses will be rejected if and only if the null hypothesis of the overall test is rejected, so the set of two one-sided tests is fully equivalent to the overall two-sided test. And directional conclusions from the one-sided tests are clearly justified.

On a deeper level, notice that the null hypothesis of the overall test is the intersection of the null hypotheses of the one-sided tests, and its critical region (rejection region) is the union of the critical regions of the one-sided tests. This makes the two one-sided tests a set of *union-intersection multiple comparisons*, which are always simultaneously protected against Type I error at the significance level of the overall test. Performing the two-sided test and then following up with a one-sided test is very much like following up a statistically significant ANOVA with Scheffé tests. Indeed, Scheffé tests are another example of union-intersection multiple comparisons. See [3] for details.

and

$$(\mathbf{C}\hat{\boldsymbol{\theta}}_n - \mathbf{C}\boldsymbol{\theta})'(\mathbf{C}\hat{\mathbf{V}}_n\mathbf{C}')^{-1}(\mathbf{C}\hat{\boldsymbol{\theta}}_n - \mathbf{C}\boldsymbol{\theta}) \sim \chi^2(r).$$

So, if $H_0 : \mathbf{C}\boldsymbol{\theta} = \mathbf{h}$ is true, we have the Wald test statistic

$$W_n = (\mathbf{C}\hat{\boldsymbol{\theta}}_n - \mathbf{h})'(\mathbf{C}\hat{\mathbf{V}}_n\mathbf{C}')^{-1}(\mathbf{C}\hat{\boldsymbol{\theta}}_n - \mathbf{h}) \sim \chi^2(r), \quad (\text{A.24})$$

where again,

$$\hat{\mathbf{V}}_n = \mathcal{J}(\hat{\boldsymbol{\theta}})^{-1} = \left[\frac{\partial^2}{\partial\theta_i\partial\theta_j} \left(-\ell(\hat{\boldsymbol{\theta}}) \right) \right]^{-1}.$$

Here is a test of $H_0 : \alpha = \beta$ for the Gamma distribution example. A little care must be taken to ensure that the matrices in (A.24) are the right size.

```
> # H0: C theta = 0 is that alpha = beta <=> alpha-beta=0
> # Name C is used by R
> CC = rbind(c(1,-1)); is.matrix(CC); dim(CC)
[1] TRUE
[1] 1 2
> thetahat = as.matrix(c(alphahat,betahat)); dim(thetahat)
[1] 2 1
> W = t(CC%%thetahat) %% solve(CC%%Vhat%%t(CC)) %% CC%%thetahat
> W = as.numeric(W) # it was a 1x1 matrix
> pval2 = 1-pchisq(W,1)
> cat("Wald Test: W = ", W, ", p = ", pval2, "\n")
Wald Test: W = 3.245501 , p = 0.07161978
```

We might as well define a function to do Wald tests in general. In the function `WaldTest`, the null hypothesis is $\mathbf{L}\boldsymbol{\theta} = \mathbf{h}$, but that's just because the name `C` is used by R for contrasts. The function returns a pair of quantities, the Wald test statistic and the p -value.

```
> WaldTest = function(L,thetahat,h=0) # H0: L theta = h
+ {
+   WaldTest = numeric(2)
+   names(WaldTest) = c("W","p-value")
+   dfree = dim(L)[1]
+   W = t(L%%thetahat-h) %% solve(L%%Vhat%%t(L)) %% (L%%thetahat-h)
+   W = as.numeric(W)
+   pval = 1-pchisq(W,dfree)
+   WaldTest[1] = W; WaldTest[2] = pval
+   WaldTest
+ } # End function WaldTest
```

Here is the same test of $H_0 : \alpha = \beta$ done immediately above, just to test out the function. Notice that the default value of \mathbf{h} in $H_0 : \mathbf{L}\boldsymbol{\theta} = \mathbf{h}$ is zero, so it does not have to be specified. The matrix `CC` has already been created, and the computed values are the same as before, naturally.

```
> WaldTest(CC,as.matrix(c(alphahat,betahat)))
      W      p-value
3.24550127 0.07161978
```

Here is a test of $H_0 : \alpha = 2, \beta = 3$, which happen to be the true parameter values. The null hypothesis is not rejected.

```
> C2 = rbind(c(1,0),
+           c(0,1) )
> WaldTest(C2,as.matrix(c(alphahat,betahat)),c(2,3))
      W      p-value
1.3305497 0.5141322
```

Finally, here is a test of $H_0 : \alpha = 1$, which was done earlier with a Z -test.

```
> WaldTest(t(c(1,0)),as.matrix(c(alphahat,betahat)),1)
      W      p-value
5.84210645 0.01564708
> Z; pval
[1] 2.417045
[1] 0.01564708
> Z^2
[1] 5.842106
```

The results of the Wald and Z tests are identical, with $W_n = Z^2$. In general, suppose the matrix \mathbf{C} in $H_0 : \mathbf{C}\boldsymbol{\theta} = \mathbf{h}$ has just a single row, and that row contains one 1 in position j and all the rest zeros. Take a look at Formula (A.24) for the Wald test statistic. Pre-multiplying by \mathbf{C} in $\mathbf{C}\widehat{\mathbf{V}}_n$ picks out row j of $\widehat{\mathbf{V}}_n$, and post-multiplying by \mathbf{C}' picks out column j of the result, so that $\mathbf{C}\widehat{\mathbf{V}}_n\mathbf{C}' = \widehat{v}_{j,j}$, and inverting it puts it in the denominator. In the numerator, $(\mathbf{C}\widehat{\boldsymbol{\theta}}_n - \mathbf{h})'(\mathbf{C}\widehat{\boldsymbol{\theta}}_n - \mathbf{h}) = (\widehat{\theta}_j - \theta_{j,0})^2$, so that $W_n = Z^2$. Thus, squaring a large-sample Z -test gives a Wald chisquare test with one degree of freedom.

A.4.4 Likelihood Ratio Tests

Likelihood ratio tests fall into two categories, exact and large-sample. The main examples of exact likelihood ratio tests include are the standard F -tests and t -tests associated with regression and the analysis of variance for normal data. Here, we concentrate on the large-sample likelihood ratio tests. Like the Wald tests, they are very flexible and are distributed approximately as chi-square under the null hypothesis for large samples. In fact, they are *asymptotically equivalent* to the Wald test under H_0 , meaning that if the null hypothesis is true, the difference between the likelihood ratio statistic and the Wald statistic goes to zero in probability as the sample size approaches infinity.

Since the Wald and likelihood ratio tests are equivalent, does it matter which one you use? The answer is that usually, Wald tests and likelihood ratio tests lead to the same conclusions and their numerical values are close. But the tests are only equivalent as

$n \rightarrow \infty$. When there is a meaningful difference, the likelihood ratio tests usually perform better, especially in terms of controlling Type I error rate for relatively small sample sizes.

Table A.1 below contains the most extreme example I know. For a particular structural equation model with normal data (details don't matter for now), ten thousand data sets were randomly generated so that the null hypothesis was true. This was done for several sample sizes: $n = 50, 100, 250, 500$ and $1,000$. Using each of the 50,000 resulting data sets, the null hypothesis was tested with a Wald test and a likelihood ratio test at the $\alpha = 0.05$ significance level. If the asymptotic results held, we would expect both tests to reject H_0 500 times at each sample size.

Table A.1: Wald versus likelihood ratio: Type I error in 10,000 simulated datasets

Test	n				
	50	100	250	500	1000
Wald	1180	1589	1362	0749	0556
Likelihood Ratio	0330	0391	0541	0550	0522

So for this deliberately nasty example, the Wald test requires $n = 1,000$ before it settles down to something like the theoretical 0.05 significance level. The likelihood ratio test needs $n = 250$, and for smaller sample sizes it is conservative, with a Type I error rate somewhat *lower* than 0.05⁸. In general, when the Wald and likelihood ratio tests have a contest of this sort, it is usually a draw. When there is a winner, it is always the likelihood ratio test, but the margin of victory is seldom as large as this.

Consider the following hypothesis-testing framework. The data are D_1, \dots, D_n . The distribution of these independent and identically distributed random variables depends on the parameter θ , and we are testing a null hypothesis H_0 using a large sample likelihood ratio test.

$$D_1, \dots, D_n \stackrel{i.i.d.}{\sim} D_\theta, \theta \in \Theta,$$

$$H_0 : \theta \in \Theta_0 \text{ v.s. } H_A : \theta \in \Theta \cap \Theta_0^c,$$

The data have likelihood function

$$L(\theta) = \prod_{i=1}^n f(d_i; \theta),$$

where $f(d_i; \theta)$ is the density or probability mass function evaluated at d_i .

Let $\hat{\theta}$ denote the usual Maximum Likelihood Estimate (MLE). That is, it is the parameter value for which the likelihood function is greatest, over all $\theta \in \Theta$. And, let $\hat{\theta}_0$

⁸This suggests that the power will not be wonderful for smaller sample sizes, in this example. But keeping Type I error rates below 0.05 is the first priority.

denote the *restricted* MLE. The restricted MLE is the parameter value for which the likelihood function is greatest, over all $\theta \in \Theta_0$. This MLE is *restricted* by the null hypothesis $H_0 : \theta \in \Theta_0$. It should be clear that $L(\hat{\theta}_0) \leq L(\hat{\theta})$, so that the *likelihood ratio*.

$$\lambda = \frac{L(\hat{\theta}_0)}{L(\hat{\theta})} \leq 1.$$

The likelihood ratio will equal one if and only if the overall MLE $\hat{\theta}$ is located in Θ_0 . In this case, there is no reason to reject the null hypothesis.

Suppose that the likelihood ratio is strictly less than one. If it's a *lot* less than one, then the data are a lot less likely to have been observed under the null hypothesis than under the alternative hypothesis, and the null hypothesis is questionable. This is the basis of the likelihood ratio tests.

If λ is small (close to zero), then $\ln(\lambda)$ is a large negative number, and $-2 \ln \lambda$ is a large positive number.

Tests will be based on

$$\begin{aligned} G &= -2 \ln \left(\frac{\max_{\theta \in \Theta_0} L(\theta)}{\max_{\theta \in \Theta} L(\theta)} \right) \\ &= -2 \ln \left(\frac{L(\hat{\theta}_0)}{L(\hat{\theta})} \right) \\ &= -2 \ln L(\hat{\theta}_0) - [-2 \ln L(\hat{\theta})] \\ &= 2 \left(-\ell(\hat{\theta}_0) - [-\ell(\hat{\theta})] \right). \end{aligned} \tag{A.25}$$

Thus, the test statistic G is the *difference* between two $-2 \log$ likelihood functions. This means that to carry out a test, you can minimize $-\ell(\theta)$ twice, first over all $\theta \in \Theta$, and then over all $\theta \in \Theta_0$. The test statistic is the difference between the two minimum values, multiplied by two.

If the null hypothesis is true, then the test statistic G has, if the sample size is large, an approximate chisquare distribution, with degrees of freedom equal to the difference of the *dimension* of Θ and Θ_0 . For example, if the null hypothesis is that 4 elements of θ equal zero, then the degrees of freedom are equal to 4. More generally, if the null hypothesis imposes r linear restrictions on θ (as in $H_0 : \mathbf{C}\theta = \mathbf{h}$), then the degrees of freedom equal r . Another way to obtain the degrees of freedom is by counting the equal signs in the null hypothesis.

Think of a multiple regression model with p regression coefficients. Here, $\theta = (\boldsymbol{\beta}, \sigma^2)$. Consider the null hypothesis $H_0 : \mathbf{L}\boldsymbol{\beta} = \boldsymbol{\gamma}$, where \mathbf{L} is an $r \times p$ matrix. This null hypothesis imposes r linear restrictions on the parameter, one for each row of \mathbf{L} . The dimension of Θ is $p + 1$; the dimension of Θ_0 is $p + 1 - r$. The degrees of freedom equal r .

The p -value associated with the test statistic G is $Pr\{X > G\}$, where X is a chisquare random variable with r degrees of freedom. If $p < \alpha$, we reject H_0 and call the results “statistically significant.” The standard choice is $\alpha = 0.05$.

To illustrate the likelihood ratio tests, consider (one last time) the Gamma distribution Example A.4.2.2. For comparison, the likelihood ratio method will be used test the same three null hypotheses that were tested earlier using Wald tests. They are

- $H_0 : \alpha = 1$
- $H_0 : \alpha = \beta$
- $H_0 : \alpha = 2, \beta = 3$

For $H_0 : \alpha = 1$, the restricted parameter space is $\Theta_0 = \{(\alpha, \beta) : \alpha = 1, \beta > 0\}$. Because the Gamma distribution with $\alpha = 1$ is exponential, the restricted MLE is $\hat{\theta}_0 = (1, \bar{d})$. It is more informative, though, to use numerical methods.

To maximize the likelihood function (or minimize minus the log likelihood) over Θ_0 , it might be tempting to impose the restriction on θ , simplify the log likelihood, and write the code for a new function to minimize. But this strategy is *not* recommended. It's time consuming, and mistakes are possible. In the R work shown below, notice how the function `gmll1` is just a “wrapper” for the unrestricted minus log likelihood function `gmll`. It is a function of β (and the data, of course), but all it does is call `gmll` with α set to one and β free to vary.

```
> gmll1 <- function(b,datta) # Restricted gamma minus LL with alpha=1
+   { gmll1 <- gmll(c(1,b),datta)
+     gmll1
+   } # End of function gmll1
> mean(D) # Resticted MLE of beta, just to check
[1] 6.8782
```

The next step is to invoke the nonlinear minimization function `nlm`. The second argument is a (vector of) starting value(s). Starting the search at $\beta = 1$ turns out to be unfortunate.

```
> gsearch1 <- nlm(gmll1,1,datta=D); gsearch1
$minimum
[1] 282.6288

$estimate
[1] 278.0605

$gradient
[1] 0.1753689

$code
[1] 4

$iterations
[1] 100
```

The answer `g1search$estimate=278.0605` is way off the correct answer of $\bar{d} = 6.8782$, it took 100 steps, and the exit code of 4 means the function ran out of the default number of iterations. Starting at the unrestricted $\hat{\beta}$ works better.

```
> gsearch1 <- nlm(gmll1,betahat,datta=D); gsearch1
$minimum
[1] 146.4178

$estimate
[1] 6.878195

$gradient
[1] -1.768559e-06

$code
[1] 1

$iterations
[1] 7
```

That's better. Good starting values are important! Now the test statistic is easy to calculate.

```
> G = 2 * (gsearch1$minimum-gammasearch$minimum); pval = 1-pchisq(G,df=1)
> G; pval
[1] 8.772448
[1] 0.003058146
```

Let us carry out the other two tests, and then compare the Wald and likelihood ratio test results together in a table.

For $H_0 : \alpha = \beta$, the restricted parameter space is $\Theta_0 = \{(\alpha, \beta) : \alpha = \beta > 0\}$.

```
> gmll2 <- function(ab,datta) # Restricted gamma minus LL with alpha=1
+   { gmll2 <- gmll(c(ab,ab),datta)
+     gmll2
+   } # End of function gmll2
> abstart = (alphahat+betahat)/2
> gsearch2 <- nlm(gmll2,abstart,datta=D); gsearch2
Warning messages:
1: NaNs produced in: log(x)
2: NA/Inf replaced by maximum positive value
$minimum
[1] 144.1704

$estimate
```

```
[1] 2.562369
```

```
$gradient
```

```
[1] -4.991384e-07
```

```
$code
```

```
[1] 1
```

```
$iterations
```

```
[1] 4
```

```
> G = 2 * (gsearch2$minimum-gammasearch$minimum); pval = 1-pchisq(G,df=1)
```

```
> G; pval
```

```
[1] 4.277603
```

```
[1] 0.03861777
```

This seems okay; it only took 4 iterations and the exit code of 1 is a clean bill of health. But the warning messages are a little troubling. Probably they just indicate that the search tried a negative parameter value, outside the parameter space. The R function `nlminb` does minimization with bounds. Let's try it.

```
> gsearch2b <- nlminb(start=abstart,objective=gml12,lower=0,datta=D); gsearch2b
```

```
$par
```

```
[1] 2.562371
```

```
$objective
```

```
[1] 144.1704
```

```
$convergence
```

```
[1] 0
```

```
$message
```

```
[1] "relative convergence (4)"
```

```
$iterations
```

```
[1] 5
```

```
$evaluations
```

```
function gradient
```

```
7 8
```

Since `nlminb` gives almost the same restricted $\hat{\alpha} = \hat{\beta} = 2.5624$ (and no warnings), the warning messages from `nlm` were probably nothing to worry about.

Finally, for $H_0 : \alpha = 2, \beta = 3$ the restricted parameter space Θ_0 is a single point and no optimization is necessary. All we need to do is calculate the minus log likelihood there.

```
> G = 2 * (gml1(c(2,3),D)-gammasearch$minimum); pval = 1-pchisq(G,df=1)
> G; pval
[1] 2.269162
[1] 0.1319713
```

The top panel of Table A.2 shows the Wald and likelihood ratio tests that have been done on the Gamma distribution data. But this is $n = 50$ and the asymptotic equivalence of the tests has barely begun to show. In the lower panel, the same tests were done for a sample of $n = 200$, formed by adding another 150 cases to the original data set. The results are typical; the χ^2 values are much closer except where they are far out on the tails, and both test lead to the same conclusions (though not always to the truth).

Table A.2: Tests on data from a gamma distribution with $\alpha = 2$ and $\beta = 3$

$n = 50$				
	Wald		Likelihood Ratio	
H_0	χ^2	p -value	χ^2	p -value
$\alpha = 1$	5.8421	0.0156	8.7724	0.0031
$\alpha = \beta$	3.2455	0.0762	4.2776	0.0386
$\alpha = 2, \beta = 3$	1.3305	0.5141	2.2692	0.1320
$n = 200$				
$\alpha = 1$	34.1847	5.01e-09	58.2194	2.34e-14
$\alpha = \beta$	0.9197	0.3376	0.9664	0.3256
$\alpha = 2, \beta = 3$	1.5286	0.4657	1.2724	0.2593

Exercises A.4.4

- Let Y_1, \dots, Y_n be a random sample from a distribution with density $f(y) = \frac{1}{\theta} e^{-\frac{y}{\theta}}$ for $y > 0$, where the parameter $\theta > 0$. We are interested in testing $H_0 : \theta = \theta_0$.
 - What is Θ ?
 - What is Θ_0 ?
 - What is Θ_1 ?
 - Derive a general expression for the large-sample likelihood ratio statistic $G = -2 \log \frac{\ell(\hat{\theta})}{\ell(\theta_0)}$.
 - A sample of size $n = 100$ yields $\bar{Y} = 1.37$ and $S^2 = 1.42$. One of these quantities is unnecessary and just provided to irritate you. Well, actually it's a mild substitute for reality, which always provides you with a huge pile of information you don't need. Anyway, we want to test $H_0 : \theta = 1$. You can do this with a calculator. When I did it a long time ago I got $G = 11.038$.

- (f) At $\alpha = 0.05$, the critical value of chisquare with one degree of freedom is 3.841459. Do you reject H_0 ? Answer Yes or No.
2. The label on the peanut butter jar says peanuts, partially hydrogenated peanut oil, salt and sugar. But we all know there is other stuff in there too. In the United States, the Food and Drug administration requires that a shipment of peanut butter be rejected if it contains an average of more than 8 rat hairs per pound (well, I'm not sure if it's exactly 8, but let's pretend). There is very good reason to assume that the number of rat hairs per pound has a Poisson distribution with mean λ , because it's easy to justify a Poisson process model for how the hairs get into the jars. We will test $H_0 : \lambda = \lambda_0$.
- (a) What is Θ ?
- (b) What is Θ_0 ?
- (c) What is Θ_1 ?
- (d) Derive a general expression for the large-sample likelihood ratio statistic.
- (e) We sample 100 1-pound jars, and observe a sample mean of $\bar{Y} = 8.57$. Should we reject the shipment? We want to test $H_0 : \lambda = 8$. What is the value of G ? You can do this with a calculator. When I did it a long time ago I got $G = 3.97$.
- (f) Do you reject H_0 at $\alpha = 0.05$? Answer Yes or No.
- (g) Do you reject the shipment of peanut butter? Answer Yes or No.
3. The normal distribution has density

$$f(y) = \frac{1}{\sigma\sqrt{2\pi}} \exp \left\{ -\frac{(y - \mu)^2}{2\sigma^2} \right\}.$$

Find an explicit formula for the MLE of $\theta = (\mu, \sigma^2)$. This example is in practically every mathematical statistics textbook, so the full solution is available. But please try it yourself first.

4. Write an R function that performs a large-sample likelihood ratio test of $H_0 : \sigma^2 = \sigma_0^2$ for data from a single normal random sample. The function should take the sample data and σ_0^2 as input, and return 3 values: G , the degrees of freedom, and the p -value. Run your function on the data in `var.dat`, testing $H_0 : \sigma^2 = 2$; see link to the data on the course web page.

For this question, you need to bring a printout with a listing of your function (showing how it is defined), and also part of an R session showing execution of the function, and the resulting output.

5. For k samples from independent normal distributions, the usual one-way analysis of variance tests equality of means assuming equal variances. Now you will construct

a large-sample likelihood ratio test for equality of means, except that you will *not* assume equal variances. Write an R function to do it.

Input to the function should be the sample data, in the form of a matrix. The first column should contain group membership (the independent variable). It is okay to assume that the unique values in this column are the integers from 1 to k . The second column should contain values of the normal random variates – the dependent variable.

The function should return 3 values: G , the degrees of freedom, and the p -value. Run your function on the sample in `kars.dat`; see link to the data on the course web page. This data set shows country of origin and gas mileage for a sample of automobiles.

6. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be a random sample from a multivariate normal population with mean $\boldsymbol{\mu}$ and variance-covariance matrix $\boldsymbol{\Sigma}$. Using the MLEs

$$\hat{\boldsymbol{\mu}} = \bar{\mathbf{X}} \text{ and } \hat{\boldsymbol{\Sigma}} = \frac{1}{n} \sum_{i=1}^n (\mathbf{X}_i - \bar{\mathbf{X}})(\mathbf{X}_i - \bar{\mathbf{X}})',$$

derive the large-sample likelihood ratio test G for testing whether the components of the random vectors \mathbf{X}_i are independent. That is, we want to test whether $\boldsymbol{\Sigma}$ is diagonal. It is okay to use material from the class notes without proof.

7. Using R, write a program to compute the test you derived in the preceding question. Your program should return 3 values: G , the degrees of freedom, and the p -value. Run it on the sample in `fourvars.dat`; see link to the data on the course web page. Bring a printout listing your program and illustrating the run on `fourvars.dat`. Of course it would be nice if your program were general, but it is not required. Note that for this problem, numerical maximum likelihood is not needed. Both your restricted and your unrestricted MLEs can and should be in explicit form.

A.5 Large Sample Theory (Preliminary)

For this part, it helps to start by going down to the basement and taking a look at the foundations of the building. There is an underlying sample space Ω , consisting of sample points $\omega \in \Omega$ ⁹. The specific nature of a point ω in applications depends on what is being observed. For example, if we were observing whether a single individual is male or female, Ω might be $\{F, M\}$. If we selected a pair of individuals and observed their genders in order, Ω might be $\{(F, F), (F, M), (M, F), (M, M)\}$. If we selected n individuals and just *counted* the number of females, Ω might be $\{0, \dots, n\}$. For limits problems, the points in Ω are infinite sequences.

⁹Throughout most of this book, Ω is a covariance matrix. The symbol will briefly have its usual meaning here, just for the discussion of almost sure convergence

Let \mathcal{A} be a class of subsets of Ω (that is, a set of *events*), and let \mathcal{P} be a probability function that assigns numbers between zero and one inclusive to the elements of \mathcal{A} . A *random variable* $X = X(\omega)$ is a function that maps Ω into some other space, typically \mathbb{R} or \mathbb{R}^k . Think of taking a measurement: if Ω is a set of students, $X(\omega)$ might be the cumulative grade point average of student ω .

Suppose the random variable X maps Ω into the set of real numbers \mathbb{R} . Then X induces a probability measure on a class¹⁰ \mathcal{B} of subsets of \mathbb{R} , by means of

$$Pr\{X \in B\} = \mathcal{P}(\{\omega \in \Omega : X(\omega) \in B\})$$

for $B \in \mathcal{B}$.

Suppose we have a sample of data $X_1(\omega), \dots, X_n(\omega)$, and we calculate a function of the sample data $T = T(X_1, \dots, X_n)$. For example T could be a *statistic* like the sample mean \bar{X} . It is helpful to write $T = T_n(\omega)$, to indicate that T is a random variable (a function from Ω into \mathbb{R}) that depend upon the sample size n .

Frequently it is useful to let $n \rightarrow \infty$, because when the sequence T_1, T_2, \dots converges, it is an indication of what happens when the sample is large enough. But this is not just a sequence of numbers; it is a sequence of functions. Several different types of convergence are meaningful.

A.5.1 Modes of Convergence

Throughout, let T_1, T_2, \dots be a sequence of random variables, and let T be another random variable. It is quite possible and often useful for $T = T(\omega)$ to be a constant — that is, a constant function of ω . In that case T is a “degenerate” random variable, with $P\{T = c\} = 1$ for some constant c .

Almost Sure Convergence

We say that T_n converges *almost surely* to T , and write $T_n \xrightarrow{a.s.} T$ if

$$\mathcal{P}\{\omega : \lim_{n \rightarrow \infty} T_n(\omega) = T(\omega)\} = 1.$$

That is, except possibly for $\omega \in A$ with $\mathcal{P}(A) = 0$, $T_n(\omega)$ converges to the random variable $T(\omega)$ like an ordinary limit, and all the usual rules apply — for example, the limit of a continuous function is the continuous function of the limit, L’Hôpital’s rule and so on. Almost sure convergence is also called *convergence with probability one*, or sometimes *strong convergence*.

Almost sure convergence may be the most technically “advanced” mode of convergence, but it is also perhaps the easiest to work with, because you treat the sequence T_1, T_2, \dots like numbers, find the limit, and then mention that the result applies “except possibly on a set of probability zero.”

The main entry point to establishing almost sure convergence is the *Strong Law of Large Numbers*, which involves almost sure convergence to a constant. Let X_1, \dots, X_n be

¹⁰I’m thinking of the Borel σ -algebra, but there is no need to go that far.

independent and identically distributed random variables with expected value μ . Denote the sample mean as usual by $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$. The Strong Law of Large Numbers (SLLN) says

$$\bar{X}_n \xrightarrow{a.s.} \mu. \quad (\text{A.26})$$

The only condition required for this to hold is the existence of the expected value.

Let X_1, \dots, X_n be independent and identically distributed random variables; let X be a general random variable from this same distribution, and $Y = g(X)$. The change of variables formula (A.1) can be combined with the Strong Law of Large Numbers to write

$$\frac{1}{n} \sum_{i=1}^n g(X_i) = \frac{1}{n} \sum_{i=1}^n Y_i \xrightarrow{a.s.} E(Y) = E(g(X)). \quad (\text{A.27})$$

This means that sample moments converge almost surely to population moments:

$$\frac{1}{n} \sum_{i=1}^n X_i^k \xrightarrow{a.s.} E(X^k)$$

It even yields rules like

$$\frac{1}{n} \sum_{i=1}^n U_i^2 V_i W_i^3 \xrightarrow{a.s.} E(U^2 V W^3).$$

Convergence in Probability

We say that T_n converges *in probability* to T , and write $T_n \xrightarrow{P} T$ if for all $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} P\{|T_n - T| < \epsilon\} = 1.$$

Convergence in probability is implied by almost sure convergence, so corresponding to the Strong Law of Large Numbers is the Weak Law of Large Numbers (WLLN). Let X_1, \dots, X_n be independent and identically distributed random variables with expected value μ . Then the sample mean converges in probability to μ :

$$\bar{X}_n \xrightarrow{P} \mu. \quad (\text{A.28})$$

A change of variables rule like expression (A.27) holds, and sample moments converge in probability to population moments. These rules follow from the corresponding facts about almost sure convergence.

Another way of establishing convergence in probability to a constant without using the definition is the *Variance Rule*. Let θ be a constant. Then if $\lim_{n \rightarrow \infty} E(T_n) = \theta$ and $\lim_{n \rightarrow \infty} \text{Var}(T_n) = 0$, it follows that $T_n \xrightarrow{P} \theta$. But convergence in probability does not imply the conditions of the Variance Rule.

Convergence in Distribution

Denote the cumulative distribution functions of T_1, T_2, \dots by $F_1(t), F_2(t), \dots$ respectively, and denote the cumulative distribution function of T by $F(t)$. We say that T_n converges *in distribution* to T , and write $T_n \xrightarrow{d} T$ if for every point t at which F is continuous,

$$\lim_{n \rightarrow \infty} F_n(t) = F(t).$$

The main entry point to convergence in distribution is the *Central Limit Theorem*. X_1, \dots, X_n be independent and identically distributed random variables with mean μ and variance σ^2 . Then

$$Z_n = \frac{\sqrt{n}(\bar{X}_n - \mu)}{\sigma} \xrightarrow{d} Z \sim N(0, 1).$$

In applications, the sample standard deviation may be substituted for σ , and the result still holds.

A useful tool is provided by the univariate *delta method*¹¹. Let $\sqrt{n}(X_n - \theta) \xrightarrow{d} X$, and let $g(x)$ be a function with $g'(\theta) \neq 0$ and $g''(x)$ continuous at $x = \theta$. Then

$$\sqrt{n}(g(X_n) - g(\theta)) \xrightarrow{d} g'(\theta)X.$$

In particular, $\sqrt{n}(g(\bar{X}_n) - g(\mu)) \xrightarrow{d} Y \sim N(0, g'(\mu)^2\sigma^2)$.

Connections among the Modes of Convergence

- $T_n \xrightarrow{a.s.} T \Rightarrow T_n \xrightarrow{P} T \Rightarrow T_n \xrightarrow{d} T$.
- If a is a constant, $T_n \xrightarrow{d} a \Rightarrow T_n \xrightarrow{P} a$.

A.5.2 Consistency

For this application, T_1, T_2, \dots are not just random variables: They are *statistics*¹² that estimate some parameter θ . The statistic T_n is said to be *consistent* for θ if $T_n \xrightarrow{P} \theta$ for all $\theta \in \Theta$.

Let us take a closer look at this important concept. Using the definition of convergence in probability, saying that T_n is consistent for θ means that for any tiny positive constant ϵ , no matter *how* tiny,

$$\lim_{n \rightarrow \infty} P\{|T_n - \theta| < \epsilon\} = 1.$$

So, take an arbitrarily small interval around the true parameter value. For any given sample size n , a certain amount of the probability distribution of T_n falls between $\theta - \epsilon$ and $\theta + \epsilon$. Consistency means that in the limit, *all* the probability falls in this interval, no

¹¹The delta method is named after the way it is proved; it uses Taylor's theorem, and the "delta" part is connected to the definition of a derivative. We will just use it.

¹²A statistic is a function of the sample data that does not depend functionally upon any unknown parameter. That is, symbol for the parameter does not appear in the formula for the statistic.

matter how small the interval is. Basically, consistency is saying that for a large enough sample size, the statistic (estimator) will probably be close to parameter it is estimating — regardless of how strict your definitions of “probably” and “close” might be.

Even better than ordinary consistency is *strong consistency*, which means $T_n \xrightarrow{a.s.} \theta$. Instead of saying T_n will probably be close to θ , strong consistency says that for a large enough sample size, the probability that it *will* be close equals one. Because almost sure convergence implies convergence in probability, strong consistency implies ordinary consistency.

One last remark is that while consistency is an important property in an estimator, in a way it is the least we should expect. Consistency means that with an infinite amount of data, we would know the truth. If this is *not* the case, something is seriously wrong¹³.

Exercises A.5.2

1. Let X_1, \dots, X_n be a random sample from a Poisson distribution with parameter λ , and let X be a general random variable with that distribution. You know that $E(X) = Var(X) = \lambda$; there is no need to prove it.

From the Strong Law of Large Numbers, it follows immediately that \bar{X}_n is strongly consistent for λ . Let

$$\hat{\lambda} = \frac{\sum_{i=1}^n (X_i - \bar{X}_n)^2}{n - 4}.$$

Is $\hat{\lambda}$ also consistent for λ ? Answer Yes or No and prove your answer.

2. Let X_1, \dots, X_n be a random sample from a Binomial distribution with parameters 3 and θ . That is,

$$P(X_i = x_i) = \binom{3}{x_i} \theta^{x_i} (1 - \theta)^{3-x_i},$$

for $x_i = 0, 1, 2, 3$. Find a reasonable estimator of θ , and prove that it is strongly consistent. Where you get your estimator does not really matter, but please state how you thought of it.

3. Let X_1, \dots, X_n be a random sample from a continuous distribution with density

$$f(x; \tau) = \frac{\tau^{1/2}}{\sqrt{2\pi}} e^{-\frac{\tau x^2}{2}},$$

where the parameter $\tau > 0$. Let

$$\hat{\tau} = \frac{n}{\sum_{i=1}^n X_i^2}.$$

Is $\hat{\tau}$ consistent for τ ? Answer Yes or No and prove your answer. Hint: You can just write down $E(X^2)$ by inspection. This is a very familiar distribution; have confidence!

¹³In structural equation models, a parameter that is not identifiable cannot be estimated consistently. This is why model identification is such an important topic.

4. Independently for $i = 1, \dots, n$, let

$$Y_i = \beta X_i + \epsilon_i,$$

where $E(X_i) = E(\epsilon_i) = 0$, $Var(X_i) = \sigma_X^2$, $Var(\epsilon_i) = \sigma_\epsilon^2$, and ϵ_i is independent of X_i . Let

$$\hat{\beta} = \frac{\sum_{i=1}^n X_i Y_i}{\sum_{i=1}^n X_i^2}.$$

Is $\hat{\beta}$ consistent for β ? Answer Yes or No and prove your answer.

5. Let X_1, \dots, X_n be a random sample from a Gamma distribution with $\alpha = \beta = \theta > 0$. That is, the density is

$$f(x; \theta) = \frac{1}{\theta^\theta \Gamma(\theta)} e^{-x/\theta} x^{\theta-1},$$

for $x > 0$. Let $\hat{\theta} = \bar{X}_n$. Is $\hat{\theta}$ consistent for θ ? Answer Yes or No and prove your answer.

A.5.3 Convergence of random vectors

1. Definitions (All quantities in boldface are vectors in \mathbb{R}^m unless otherwise stated)

★ $\mathbf{T}_n \xrightarrow{a.s.} \mathbf{T}$ means $P\{\omega : \lim_{n \rightarrow \infty} \mathbf{T}_n(\omega) = \mathbf{T}(\omega)\} = 1$.

★ $\mathbf{T}_n \xrightarrow{P} \mathbf{T}$ means $\forall \epsilon > 0, \lim_{n \rightarrow \infty} P\{\|\mathbf{T}_n - \mathbf{T}\| < \epsilon\} = 1$.

★ $\mathbf{T}_n \xrightarrow{d} \mathbf{T}$ means for every continuity point \mathbf{t} of $F_{\mathbf{T}}$, $\lim_{n \rightarrow \infty} F_{\mathbf{T}_n}(\mathbf{t}) = F_{\mathbf{T}}(\mathbf{t})$.

2. $\mathbf{T}_n \xrightarrow{a.s.} \mathbf{T} \Rightarrow \mathbf{T}_n \xrightarrow{P} \mathbf{T} \Rightarrow \mathbf{T}_n \xrightarrow{d} \mathbf{T}$.

3. If \mathbf{a} is a vector of constants, $\mathbf{T}_n \xrightarrow{d} \mathbf{a} \Rightarrow \mathbf{T}_n \xrightarrow{P} \mathbf{a}$.

4. Strong Law of Large Numbers (SLLN): Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be independent and identically distributed random vectors with finite first moment, and let \mathbf{X} be a general random vector from the same distribution. Then $\bar{\mathbf{X}}_n \xrightarrow{a.s.} E(\mathbf{X})$.

5. Central Limit Theorem: Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be i.i.d. random vectors with expected value vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Then $\sqrt{n}(\bar{\mathbf{X}}_n - \boldsymbol{\mu})$ converges in distribution to a multivariate normal with mean $\mathbf{0}$ and covariance matrix $\boldsymbol{\Sigma}$.

6. Slutsky Theorems for Convergence in Distribution:

(a) If $\mathbf{T}_n \in \mathbb{R}^m$, $\mathbf{T}_n \xrightarrow{d} \mathbf{T}$ and if $f : \mathbb{R}^m \rightarrow \mathbb{R}^q$ (where $q \leq m$) is continuous except possibly on a set C with $P(\mathbf{T} \in C) = 0$, then $f(\mathbf{T}_n) \xrightarrow{d} f(\mathbf{T})$.

(b) If $\mathbf{T}_n \xrightarrow{d} \mathbf{T}$ and $(\mathbf{T}_n - \mathbf{Y}_n) \xrightarrow{P} \mathbf{0}$, then $\mathbf{Y}_n \xrightarrow{d} \mathbf{T}$.

(c) If $\mathbf{T}_n \in \mathbb{R}^d$, $\mathbf{Y}_n \in \mathbb{R}^k$, $\mathbf{T}_n \xrightarrow{d} \mathbf{T}$ and $\mathbf{Y}_n \xrightarrow{d} \mathbf{c}$, then

$$\begin{pmatrix} \mathbf{T}_n \\ \mathbf{Y}_n \end{pmatrix} \xrightarrow{d} \begin{pmatrix} \mathbf{T} \\ \mathbf{c} \end{pmatrix}$$

7. Slutsky Theorems for Convergence in Probability:

(a) If $\mathbf{T}_n \in \mathbb{R}^m$, $\mathbf{T}_n \xrightarrow{P} \mathbf{T}$ and if $f : \mathbb{R}^m \rightarrow \mathbb{R}^q$ (where $q \leq m$) is continuous except possibly on a set C with $P(\mathbf{T} \in C) = 0$, then $f(\mathbf{T}_n) \xrightarrow{P} f(\mathbf{T})$.

(b) If $\mathbf{T}_n \xrightarrow{P} \mathbf{T}$ and $(\mathbf{T}_n - \mathbf{Y}_n) \xrightarrow{P} 0$, then $\mathbf{Y}_n \xrightarrow{P} \mathbf{T}$.

(c) If $\mathbf{T}_n \in \mathbb{R}^d$, $\mathbf{Y}_n \in \mathbb{R}^k$, $\mathbf{T}_n \xrightarrow{P} \mathbf{T}$ and $\mathbf{Y}_n \xrightarrow{P} \mathbf{Y}$, then

$$\begin{pmatrix} \mathbf{T}_n \\ \mathbf{Y}_n \end{pmatrix} \xrightarrow{P} \begin{pmatrix} \mathbf{T} \\ \mathbf{Y} \end{pmatrix}$$

8. Delta Method (Theorem of Cramér, Ferguson p. 45): Let $g : \mathbb{R}^d \rightarrow \mathbb{R}^k$ be such that the elements of $\dot{g}(\mathbf{x}) = \left[\frac{\partial g_i}{\partial x_j} \right]_{k \times d}$ are continuous in a neighborhood of $\boldsymbol{\theta} \in \mathbb{R}^d$. If \mathbf{T}_n is a sequence of d -dimensional random vectors such that $\sqrt{n}(\mathbf{T}_n - \boldsymbol{\theta}) \xrightarrow{d} \mathbf{T}$, then $\sqrt{n}(g(\mathbf{T}_n) - g(\boldsymbol{\theta})) \xrightarrow{d} \dot{g}(\boldsymbol{\theta})\mathbf{T}$. In particular, if $\sqrt{n}(\mathbf{T}_n - \boldsymbol{\theta}) \xrightarrow{d} \mathbf{T} \sim N(\mathbf{0}, \boldsymbol{\Sigma})$, then $\sqrt{n}(g(\mathbf{T}_n) - g(\boldsymbol{\theta})) \xrightarrow{d} \mathbf{Y} \sim N(\mathbf{0}, \dot{g}(\boldsymbol{\theta})\boldsymbol{\Sigma}\dot{g}(\boldsymbol{\theta})')$.

Exercises A.5

- 1.
- 2.

A.6 Method of Moments (Very Preliminary)

Let the joint distribution of the data X_1, \dots, X_n depend on the parameters $\theta_1, \dots, \theta_r$. Following standard notation, we define the k th sample moment as

$$m'_k = \frac{1}{n} \sum_{i=1}^n X_i^k$$

Clearly, $E(m'_k)$ is a function of $\theta_1, \dots, \theta_r$. So pick r values of k , and write the system of equations

$$\begin{aligned} E(m'_{k_1}) &= g_1(\theta_1, \dots, \theta_r) \\ &\vdots \\ E(m'_{k_r}) &= g_r(\theta_1, \dots, \theta_r) \end{aligned}$$

Once this has been done, follow these steps.

1. Remove the expected values on the left hand side of the equations.
2. Add hats to the θ values on the right hand side.
3. Solve for the $\hat{\Theta}$ values. These are the Method of Moments (MOM) estimators.

Here are some comments.

- If the data are a random sample (independent and identically distributed) then $E(m'_k) = \mu'_k$, and we have the usual Method of Moments estimators.
- Of course it does not matter if you solve the equations first and then put the hats on, or if you do it the other way around.
- You can hope that the values of k are just $1, \dots, r$, but there is no guarantee. In particular, if the distribution of the X_i values is symmetric about zero, then all the odd-numbered moments are zero, and only even values of k will be useful.
- If the data are not a random sample (think regression), don't depend on formulas like $E(\bar{X}) = \mu$; they were derived for the *i.i.d.* case, and might not be correct. Do the calculation.

Exercises [A.6](#)

- 1.
- 2.